Practical Neutron Scattering at a Steady State Neutron Source

Tom Heitmann

The Missouri Research Reactor, University of Missouri, Columbia Missouri, United States

Wouter Montfrooij

Department of Physics and Astronomy and the Missouri Research Reactor, University of Missouri, Columbia Missouri, United States

This book is dedicated to Mihai Popovici

Preface

The primary purpose of this booklet is to serve as the companion to the IGERT (Integrative Graduate Education and Research Traineeship) course on neutron scattering at a steady state nuclear reactor. However, it is our hope that this booklet will serve the broader purpose of a quick reference guide for those who will perform neutron scattering experiments once the course is over, or for those who have an interest in performing neutron scattering experiments for the first time.

We aim to provide a conceptual understanding of the technique of neutron scattering, an understanding that should allow a neutron scattering user to successfully plan, carry out and analyze scattering experiments in the fields of physics, chemistry, biology, engineering and geology. As such, this booklet will be very light on equations; in fact, equations serve the purpose of providing background information rather than being an essential part of teaching the technique. Because of our aim, we will sometimes have to make our arguments in a rather hand-waving manner, and we have no intention of being complete in the materials we present. We understand that this may make some experienced users- especially those with a physics background- cringe, however, this cannot be avoided.

Given the different backgrounds and interests of the IGERT participants, we start out this booklet at a rather basic level, perhaps unbecoming of a graduate course. These introductory chapters can easily be skipped by those already familiar with the essentials of modern physics, such as wave particle duality and interference effects. Notwithstanding these limitations in scope, we hope that the reader will enjoy this booklet and we certainly welcome any and all feedback. Finally, everyone is welcome to reproduce any part of this book, as long as it is for educational purposes.

Tom Heitmann and Wouter Montfrooij

August 2012,

Columbia, Missouri.

The participants of the 2012 IGERT neutron scattering course taught at the University of Missouri (instructors: Wouter Montfrooij and Tom Heitmann). Posing, from left to right, are Jagath Gunasekera, Kevin Tarwater, Andrew Miskowiec, Matt Connolly, Joe Schaeperkoetter, David Stalla, Elmar Dohnke, Tom Zhang, Peiwu Qin, Tyler Rash and Dongmei Yu.

Acknowledgements

Contents

1
Introduction and Outline

Neutron scattering is an experimental technique whereby we can learn about the microscopic structures of our materials and the movements of the atoms and molecules within these structures. Some of the questions that we might have about the structure and dynamics of our materials can very easily be answered by doing neutron scattering experiments; the answers to some other questions require some more elaborate experiments, while the answers to yet some other questions should not be sought by using this technique. We will provide a practical guide as to what can be measured, and what should be left alone.

Neutron scattering is an interference based technique, similar to light scattering. The principles behind interference are easy to visualize, and once these principles are understood, neutron scattering will no longer be a 'black-box' technique but rather it becomes a versatile tool for learning about the microscopic behavior of one's materials. At this level, the neutron scattering user can benefit fully from the unique capabilities of this scattering technique, capabilities that allow for the uncovering of details that cannot be learned through any other technique. As such, we start out this booklet with a brief review of what interference entails, and how it relates to scattering experiments.

As the title of this booklet promises, the presented materials are meant to be a practical guide. Practical guides and user manuals already exist for various spectrometers in various places, but generally it is difficult to come by all the required information before planning an experiment. This booklet should help somewhat in filling in this gap. In addition, most basic texts on neutron scattering deal with the connection between the idealized scattering by a sample, and the underlying microscopic behavior. In general, this information is not all that helpful in planning a successful experiment, in selecting the best instrument for answering particular questions, and in being able to correct the resulting data for unwanted scattering events and artifacts. Once again, this booklet should be somewhat helpful in overcoming these difficulties.

The outline of this booklet is the following. The first part of the book discusses some generalities about neutron scattering. First we discuss the relationship between neutrons, interference, and what is measured in a scattering experiment. This is done in Chapter 2. In the next chapter, we review some of the basic properties of the neutron and we combine this with the interference technique to review what can be measured. In Chapter 4 we delve into a bit more detail and discuss what actually is measured, and how this relates to what information we would like to uncover.

In the second part of this booklet we continue with a more detailed discussion of the various most commonly used spectrometers that are available for doing neutron scattering experiments. There are dedicated chapters for diffractometers (Chapter 5), powder diffractometers (Chapter 6), SANS diffractometers (Chapter 7), reflectometers (Chapter 8) and triple-axis spectrometers (Chapter 9). These chapters describe the (basic) operations of the spectrometers, how to use them in a sensible manner, and potential pitfalls in acquiring good data. We conclude this part of the booklet with a brief review of some other, less frequently used spectrometers (Chapter 10).

The third part of the booklet deals with the nuts and bolts of planning and analyzing a successful neutron scattering experiment. Chapter 11 not only deals with selecting the best instrument for the job, but also deals with the sample environment that needs to be selected, and it lists the various supporting measurements that need to be performed in order to have he best chance that the measured data will be able to be fully analyzed, free from unwanted signals and experimental artifacts. Chapter 12 describes how to analyze data during and after an experiment, and how to avoid certain pitfalls.

The appendices that follow these parts detail common procedures and provide a bit of a physics background. The procedures are most useful in planning and analyzing experiments, but they have been relegated to appendices as they stand in the way of understanding the concepts behind the neutron scattering technique. The last part of this booklet contains a description of experimental modules particular to the IGERT course. These modules are not helpful to readers that are not course participants.

We would also like to draw the readers attention to multiple texts and textbooks that are available to learn more about neutron scattering. The following is our selection of these texts. Our selection is based upon how closely these texts match the (anticipated) interest of the IGERT participants; as such, the many omissions from this list do not reflect our opinions of these books, no slight was intended in any of our choices.

Physics angle:
- G.L. Squires: Introduction to the Theory of Thermal Neutron Scattering (Dover publications, 1996)
- Steven W. Lovesey, Theory of Neutron Scattering from Condensed Matter, Vols. 1 and 2 (Clarendon Press, 1984)
- Gen Shirane, Stephen M. Shapiro, John M. Tranquada: Neutron Scattering with a Triple-Axis Spectrometer (Cambridge, 2002)
- G.E. Bacon: Fifty Years of Neutron Diffraction (Hilger, 1987)
- Albert Furrer, Joel Mesot and Thierry Strässle: Neutron Scattering in Condensed Matter Physics (World Scientific, 2009)

Biology angle:

- Chick Wilson: Single Crystal Neutron Diffraction from Molecular Materials (World Scientific, 2000)
- S.J. Perkins: High Flux X-ray and Neutron Solution Scattering, in Methods in Molecular Biology Volume 22, C. Jones; B. Mulloy; A.H. Thomas, editors (Humana Press, 1994)
- Neutrons in Biology, edited by Schoenborn and Knott (Plenum 1996)

Chemistry angle:

- G.E. Bacon: Neutron Scattering in Chemistry (Butterworths, 1977)
- B.J.M. Willis: Chemical Applications of Neutron Scattering (Oxford University Press, 1973)
- R.H. Ottewill, Small Angle Neutron Scattering, in Colloidal Dispersions, J.W. Goodwin (editor), Special Publication No. 43, Royal Society of Chemistry, 1982

Books accompanying neutron scattering schools:

- Neutron and Synchroton Radiation for Condensed Matter Studies, Baruchel ed. (Hercules school, Springer Verlag 1991)
- R.J. Newport; B.D. Rainford; R. Cywinski (editors): Neutron Scattering at a Pulsed Source (Adam Hilger, 1988)

Popular reference articles:

- Fundamentals of Neutron scattering by Condensed Matter, R. Scherm, 1972.
- Neutron Scattering: A Primer, Roger Pynn.
- The SANS Toolbox, Boualem Hammouda.

Part I

The basics

This part of the booklet establishes the connection between interference phenomena and neutron scattering. Chapter 2 is intended for readers who do not have a physics background. Chapter 3 describes the some of the properties in the neutron that are relevant to neutron scattering. Chapter 4 deals with how to make the connection between the number of scattering events (counts) in a detector and the microscopic structure and motion of the sample that we stuck in the beam of neutrons.

2
Interference and Neutron Scattering

Interference takes place when two waves are in the same location at the same time. These two waves can help each other, resulting in an amplified effect, or they can negate the effects of each other, resulting in a diminished effect. All waves do this, whether they be sound waves, light waves, waves in the ocean or waves in your bathtub. Under special circumstances, two waves can completely cancel each other; we refer to this situation as destructive interference. At the other extreme, two waves can help each other to the fullest extent possible; we refer to this situation as constructive interference. Thus, the usage of the term interference in waves differs from its usage in our everyday language where it is always meant to signify a reduced outcome.

Looking at waves in water provides a visualization of what we just mentioned in words. Imagine a surface wave that is rolling up against a wall with two openings in it as shown in Fig. 2.1. The water in the two openings is put in motion when the incoming wave hits it. As a result, the two openings will now act as the sources of two new waves, and these two new waves will spread outwards while putting the surface of the water in motion. However, these two new waves will occupy the same region of the water, and therefore, we will observe an interference pattern. Note that one can easily create one's own interference pattern by sitting in a bathtub, sticking one's hands in the water and moving them about in a repeated fashion.

The resulting interference pattern is easy to understand. Pick any point in the water behind the two openings. The water at this point will be subject to one wave trying to move it in a particular fashion, and another wave trying to do what the other wave wants it to do. If both waves are trying to do the same thing at the same time, then we will observe an amplified effect. Imagine that one wave was trying to move the water up, and the other wave was doing the same thing. Then, as a result, the water will be moved up quite a bit more. The two waves' combined action results in a higher amplitude of the up and down motion than what an individual wave can accomplish. We mention up and down motion because the same would hold for when one wave is trying to create an indentation in the water, and the second wave is doing the identical thing, resulting in a deeper indentation. Thus, this type of interference is characterized by an increased amplitude of the overall motion.

Whether two waves help each other, or try to cancel each other depends on how long it took the individual waves to get to the same point in the water. For instance, if we look at water that is located at an identical distance from both inlets, then the two

Fig. 2.1 Water waves making it through two openings will create an interference pattern as demonstrated on this photo from *Google* Earth, and as explained by Fabrizio Logiurato in Physics Education 47, 73 (2012).

waves will take identical times to reach this point. As a result, the two waves will try to move the water in an identical manner, and we witness constructive interference. Conversely, we can have points in the water that are slightly closer to one opening than to the other, and we can even find points where the second wave took so much longer to arrive that when it finally gets there and tries to push the water up, that the first wave has already moved to the point where it is trying to push the water down. At these points, the water will be pushed up just as much as it will be pushed down. We call this destructive interference since the two forces exactly cancel each other, and the water will not move at all. This situation is shown in Fig. 2.2.

There is more than one region where we can witness constructive and destructive interference. For instance, there will be points that the second wave took so long to reach, that by the time this second wave gets ready to try to push the water up, the first wave is once again trying to do the same thing. The first wave already went through a cycle of (pushing the water) up and down, and now it is ready to push the water up again. Clearly, this is also a situation where the two waves will amplify each other. The same holds true for when the second wave arrives two cycles later, or three, or four etc.

By the same token, it can happen that the second wave arrives 1/2 a cycle later, or 1 1/2 cycles later, or 2 1/2 etc. These are all situations where we would observe complete cancelation of the two waves, regions of destructive interference. All this can be seen in Fig. 2.2. Of course, there are also regions that are neither here nor there; at these regions the water will neither be tranquil, nor will it show full amplification.
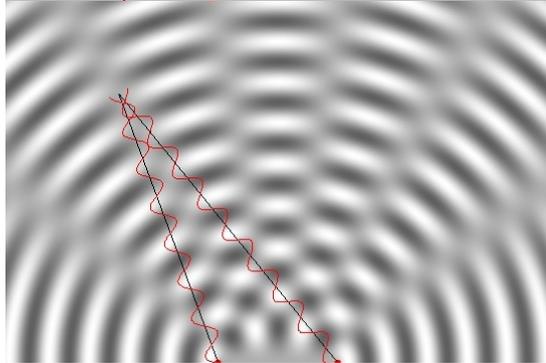
Fig. 2.2 Water waves originating from two sources interfere with each other and create a distinct pattern. The resulting wave pattern has regions where the waves help each other (dark for regions where the water is pushed down, and light for regions are where the water is pushed up), as well as regions where the waves cancel each other (the radial spokes where the water does not move at all). The red lines depict how two waves can exactly cancel each other at a point (where the black lines intersect for this example). Source: 'Ripple Tank' simulation software.

The locations at which we find complete amplification (and cancelation) are determined by how much longer one wave traveled to reach this location than the other. This depends both on how fast the wave is going, and on how far apart the two openings are since the latter determines how much farther one point in the water is from one opening than it is from the other. These are pretty obvious statements, but they form the basis of the theory of scattering.

In scattering experiments one observes (measures) the interference pattern, and from this one tries to figure out what the spacing is between the two openings. For instance, one can well imagine that by observing (Fig. 2.2) the regions where there is maximum amplification of the two waves, that one can identify which of these regions of maximum amplification correspond to the two waves having traveled an identical amount of time, which regions correspond to the second wave lagging behind the first wave by one cycle, etc. Then provided one knows how fast these water waves were traveling, one can deduce- with the aid of some math- what the spacing must have been between the two openings.

In real scattering experiments one does not use water waves, of course, but rather one uses light waves or- as we shall see- neutrons. The distance between the two openings in the wall then becomes the spacing between two atoms, a number that tells us something about the structure of our material. We can even have many atoms arranged in some periodic structure, the equivalent of multiple holes in the wall. As long as we study the ensuing interference pattern carefully, we can deduce all the relevant spacings between the atoms. It requires some more math, but the principle behind

this deduction is identical to figuring out the distance between the two inlets in Fig. 2.2.

Before we look at the interference pattern in even more detail, we have to ask ourselves the question: what do neutrons have to do with this? We can see an interference pattern emerging when one uses light waves, or even very special light waves such as xrays, but neutrons are particles, so what is the relevance of interference to neutron scattering? The answer to this is the most remarkable discovery of the 20^{th} century, namely that neutrons (and all particles) can, and do, behave like waves.

Neutrons are undeniably particles. Whenever we detect one, it looks and smells like a particle. They arrive one by one, and we can even make our detectors produce a satisfactory clicking sound when they arrive. The same holds for all little particles in nature, from electrons to protons, from atoms to buckyballs. Yet, according to an uncountable number of experiments, whenever we do NOT look at a particle, it behaves like a wave. It turns out to be a law of nature: **when we look at a particle, then it is a particle, and when we do not look, then it is a wave**. Admittedly, this sounds bizarre, but it is what nature does. In here, we just take this law as an established experimental fact and we refer the reader to the popular literature to learn more about this highly counterintuitive and unsettling 20^{th} century law of physics.

According to this law of nature, neutrons behave like waves when we do not look, such as when they are flying through a material. And when they behave like waves, they will suffer from interference effects like any old wave, and we can use the interference pattern to learn something about the material itself, such as the spacings between atoms. This is exactly what we do when we perform neutron scattering experiments.

Neutrons are waves. In experiments we can control the speed of these neutron waves, which in turn will make our interference pattern sensitive to small distances (such as interatomic distances for fast waves), or to large distances (such as the size of a protein for slow waves). It is this range of length scales that make neutrons such a versatile probe of matter. In addition, the neutron- having no electric charge- has no problem penetrating deeply into all sorts of materials. This makes the technique very useful for studying the bulk properties of materials rather than probing what is happening at the surface. Also, how easily a neutron is scattered by the atoms depends on what the atom is; this makes it possible to identify not only how far apart the atoms are, but also to identify which atom sits where. Of course, there will be some math involved, but this is an issue that has been resolved and should not concern us too terribly much (at this point).

The converse to our new law also holds true: waves are particles. In particular, when we detect light, we detect it one little bit at a time. Therefore, light is made up of particles when we look, and these particles are called photons. This is at the basis of the photo-electric effect, which in turn is at the basis of every automatic toilet flusher and door opener. The upside of all this is that whatever we learn about
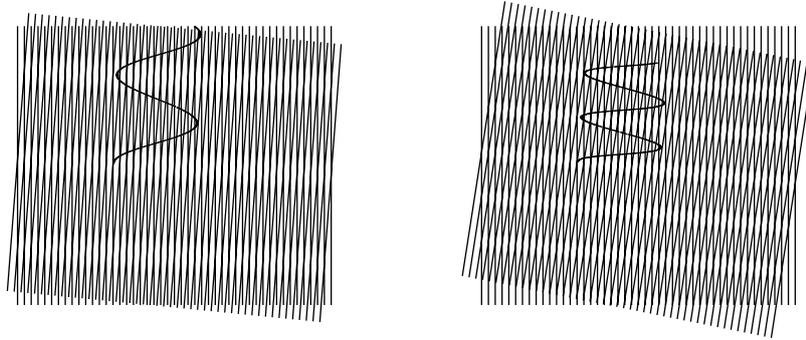
Fig. 2.3 In scattering experiments, we deal with an incoming wave, and an outgoing wave. The wave length associated with the waves is given by the distance between the vertical lines, the wave crests. The incoming (moving from left to right) and outgoing (moving at an angle) waves give rise to an interference pattern that shows up as the white, almost horizontal bands. The distance between successive white bands corresponds to what is called the probing wavelength and which has been plotted in the figure as the sinusoidal shape. The probing wave length λ_{probe} depends both on the wave length of the plane waves, and on the scattering angle. The latter is illustrated in the figure by comparing a scattering angle of 5^{o} (left half of figure) and 10^{o} (right half).

neutron scattering in terms of interference, we can also directly apply it to light scattering, and vice versa. All the basic concepts are the same between the two techniques.

2.1 Interference and Scattering

In a scattering experiment, a wave can even interfere with itself. This can be seen as follows. Picture a wave that is about to be scattered by the individual atoms in a piece of material. The wave actually occupies a region of space, it is not restricted to be at a particular point. As such, the wave that has not been scattered yet will occupy the same space at the same time as the wave that is being scattered by the atoms. This incoming wave and the scattered wave will now interfere with each other, ultimately determining under what angle the scattered wave takes off.

We do not have to look into the details of the scattering process to understand what is probed in a scattering experiment; we can simply look at pictures of waves. Since a neutron behaves like a wave, both when it is incident on the sample and when it is scattered by the sample, we can depict a scattering event by drawing an incoming and a scattered wave. This is shown in Fig. 2.3.

As a note up front, this sketch leaves out the details of the scattering process as far as it concerns individual atoms. However, these details are not important for understanding the basics of a scattering experiment. In the left panel of Fig. 2.3 we see a wave traveling from the left to the right. The wavelength is given by the distance

between the vertical lines. Also shown is a wave that (somehow) scattered and that is now heading for a detector that is placed at an angle of 5^o with respect to the direction of the incoming wave. This scattered wave has an identical wavelength as the incoming wave. This is the situation that corresponds to the neutron being scattered by the sample without gaining or losing any of its energy; it merely changes direction.

The incoming and the scattered wave create an interference pattern, resulting in the bright and dark regions shown in Fig. 2.3 and known as a moiré pattern. The orientation of the bright and dark bands is such that they are neither parallel to the incoming nor to the scattered direction; rather, it is somewhere in between. For the case shown in the figure where the neutron neither gains nor loses energy, the direction of the bands is exactly half way in between the incoming and the scattered direction.

The distance between the bands, such as the prominent white bands visible in the figure, is what we call the probing wavelength λ_{probe}. It is on this length scale that we are probing our sample in the scattering experiment. During the course of the scattering event depicted in this figure, the wave crests of the incoming and scattered wave will travel along the direction of propagation of the neutron, and as a result, the white bands will also be moving. The direction of propagation of these bands is perpendicular to the bands themselves. All this is indicated in the figure by the sinusoidal shape which shows that the probing wave length is associated with a probing wave that travels through the sample in the direction shown in the figure.

Note that the probing wavelength is not the same as the wavelength of the incoming wave; in the case shown in the figure, it is much larger. The probing wavelength does depend on the incoming wavelength- the larger the incoming wave length, the larger the probing wavelength- but it also depends on the angle between the incoming and the scattered wave. This is illustrated in the right hand panel of figure 2.3 where we doubled the angle between the incoming and the scattered wave. As can be seen in this particular example, the probing wavelength shortened as a result.

When we are studying a particular material with the aid of neutron scattering, we try to ensure that the probing wavelength (roughly) matches the length scale of the objects we are interested in. We will require exact matches for materials that have a periodic (repeating) structure, such as crystals, and we require somewhat looser matches for materials that only have a repeating structure on short length scales. Examples of the latter are liquids where the atoms do not all have the same distance to their neighbors, or polymers spread out in a film.

These sketches already illustrate that there are limitations to the technique of neutron scattering when it comes to studying structures that require very large probing wavelengths. Either one has to look for scattered neutrons at very small scattering angles, or one has to use incident waves that have very long wavelengths. The former runs into problems because not all neutrons are being scattered, and one cannot distinguish easily anymore between those neutrons that have been scattered, and those that
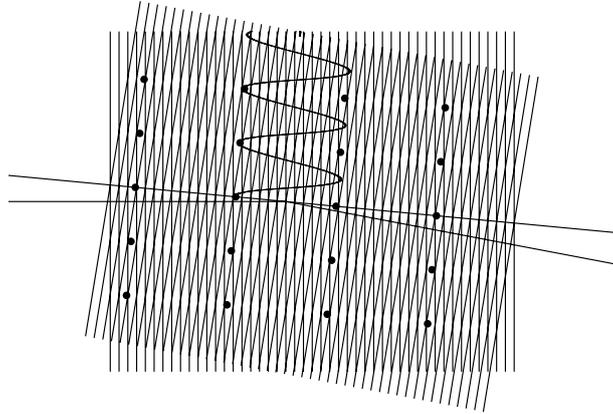
Fig. 2.4 When the interference pattern produced by the incoming wave and outgoing wave corresponds to a natural distance in the sample, then this will result in constructive interference. In here, if the probing wave length λ_{probe} exactly matches the lattice spacing d of the solid, and if the sample is oriented in just the right way as shown in the figure, then the scattering process depicted might actually occur and a detector placed in the direction of the outgoing wave can detect a scattered neutron (or photon in xray scattering). The solid lines depict that under these scattering conditions the radiation appears to bounce off of a mirror-like surface formed by the array of atoms.

simply made it through the sample unaffected. The latter ends up being problematic because when neutrons are produced inside reactors, they typically have wave lengths of about 1 Å. It is possible to cool down these neutrons, thereby stretching their wave length, but this is not a minor undertaking. Because of these two limitations, neutrons are not very good at probing length scales above $\sim$ 100-1,000 Å. Luckily, this upper limit represents such a large length scales that we can use different (non-neutron scattering) techniques to probe our materials.

Leaving these limitations aside for the moment, let us look at an example of neutrons being scattered by a crystal. A crystal is characterized by atoms lined up in a periodic structure. When a crystal is hit by a neutron wave (which is NOT the same as a wave of neutrons), then every atom will act as a center from which a scattered wave emerges. This is very similar to the case of the two holes in the wall in Fig. 2.1 where each opening acts as the source of a new wave. If all these atoms are exposed to the neutron wave in the same way, then all these individual scattered waves *can* add up to a viable total wave, and the neutron will actually end up being scattered.

The above mystery language is clarified in Fig. 2.4. In this figure, the distance d between the atoms and the orientation of the crystal is such that the atoms are exactly located corresponding to the spacings between the white bands of the interference pattern:

$$\lambda_{\text{probe}} = d.$$

This equation is referred to as Bragg's law. What this equation implies is that if one atom emits a scattered wave, then any of the other atoms also emits a scattered wave that will be trying to do the same thing as the other scattered waves when the waves meet up. Compare this again to the case of interference in water. When two waves meet up that are trying to push the water in the same direction, then these two waves help each other. The same is shown in Fig. 2.4. The distances and orientations of everything involved is such that when scattered waves originating from individual atoms meet up, then they help each other and they will give rise to the nice big scattered wave that can be seen leaving the sample scattered under a 10^o angle.

Even if this language (or the math behind it) is not clear, it should be clear that one can expect something special to happen when all the atoms shown in Fig. 2.4 line up with the interference bands. In this figure, a detector located at an angle of 10^o with respect to the direction of propagation of the incoming waves will detect many scattered neutrons. Should this detector be placed at an angle of 9^o, then it will not detect any scattered neutron waves. It will not detect any, because at such an angle the white interference bands would not have corresponded exactly to the distance between the atoms.

The above description tells us how we can actually determine the interatomic distances between atoms. We can stick our crystal in a beam of neutrons, and get ourselves a detector. Then we can move our detector around, positioning it by moving it from small angles to large angles as shown in Fig. 2.5, to see if we detect (m)any neutrons being scattered. Likely, we will not find any. But then we can rotate our crystal a little bit and try again by moving our detector around. We repeat this procedure until we find an orientation of the crystal and a position of the detector for which we do notice a lot of scattering events. When this happens, then we must have the situation depicted in Fig. 2.4. Given this, we can then calculate the distance between the atoms based on the angle where our detector measured the most scattering events, and based upon the wavelength λ of the neutron (the distance between wave crests in Fig. 2.4). We can capture this in the following representation of Bragg's law (see also exercise 2.1):

$$\lambda_{\text{probe}} = \frac{\lambda}{2\sin(\theta/2)} = d; \lambda = 2d\sin(\theta/2). \tag{2.1}$$

This is scattering in its most basic and most readily interpretable form. We can take it a step further by, for instance, having the crystal consist of more than one type of atom, or by melting the crystal so that the periodic structure still looks ordered on short length scales, but it will look like the familiar liquid disorder on longer length scales. Under such conditions the neutrons will still be scattered, but not as well as before and not all at the same angle; therefore, the results will be harder to interpret. But they still can be interpreted and from a conceptual point of view, they are not all that different from the pure case shown in Fig. 2.4. For future reference, the case
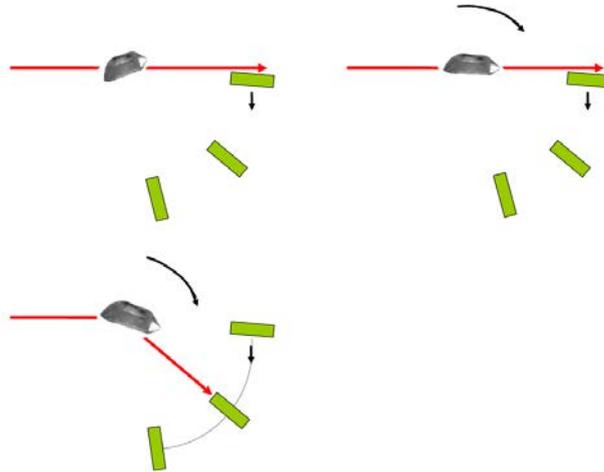
Fig. 2.5 In order to determine the interatomic spacing of a crystal, we stick the sample in the neutron beam (red arrow). We then move our detector around (green box in the top left panel) and see if we count any neutrons anywhere. When we fail to find the expected counts (the neutrons can be 'seen' to fly through the sample), then we rotate the crystal a little bit (top right panel) and we move the detector around again. We repeat this procedure until we have rotated the crystal to just the right angle (bottom left panel) that would scatter neutrons (red arrow) into the detector when we position it at just the right angle (in this example at a scattering angle of approximately 40^o). The situation in this bottom panel must then correspond to satisfying the constructive interference condition sketched in Fig. 2.4.

shown in Fig. 2.4 is referred to as Bragg scattering.

There is one more thing that we can already learn about scattering by looking at interference patterns such as the one in Fig. 2.4, and this has to do with the amount of momentum transferred from the neutron to the sample in a scattering event. Looking at Fig. 2.4, we see that the neutron changed direction just as if the neutron bounced of a wall formed by the atoms within the crystal. And, as is the case when bouncing a ball of a wall (at an angle or not), the momentum that will be transferred to the wall is in the direction perpendicular to the wall. This will be important when we study the details of the movement of the atoms within samples. This picture (Fig. 2.4) shows how the neutron will impart momentum onto these atoms, in what direction it will apply a force so as to make these atoms move.

The amount of momentum transferred depends on the scattering angle. Using the analogue of a ball bouncing off a wall after being thrown at the wall under a certain angle, it is easy to visualize that the least amount of momentum will be transferred when the scattering angle is very small, and that the most momentum is being trans-

ferred when the scattering is such that the neutron bounces back in the direction it came. This should be intuitively clear; if you want to knock over an object, you throw the ball straight at it, not under a glancing angle.

Combining these (obvious) considerations with how the interference pattern changes (Fig. 2.3) with changes in scattering angle, we see that a large probing wave length corresponds to small amounts of momentum being transferred. Decreasing the probing wavelength goes hand in hand with increasing the amount of momentum transferred. Therefore, small angle scattering usually involves small amounts of momentum transfer. This represents a limitation of the technique of neutron scattering as sometimes we would like to transfer large amounts of momentum while still scattering under small angles. This is simply not possible with neutrons.

In summary, when doing and interpreting a neutron scattering experiment, it is useful to keep pictures such as Fig. 2.3 in mind in order to keep track of what is happening in real space and in real time. As we will see in later chapters, in neutron scattering one typically only keeps track of how much momentum is transferred, and of how much energy is transferred. This bookkeeping works well in most cases, however, a mental image of what is going on in terms of interference patterns makes it easier to understand what will happen when we change something to the experiment, such as increasing the neutron wavelength or trying to obtain scattering data at small scattering angles.

2.2   Exercises

Exercise 2.1

Show that the distance λ_{probe} between the white bands in Fig. 2.3 is given by:

$$\lambda_{\mathrm{probe}} = \lambda_{\mathrm{incoming}}/2\sin(\theta/2).$$

In here, θ is the angle under which the neutron is scattered, and $\lambda_{\mathrm{incoming}}$ is the wave length of the incoming neutron given by the distance between the vertical wave crests in the figure.

Exercise 2.2

What is the maximum amount of momentum that can be transferred from the neutron to the sample (when the incoming neutron carries 1 unit of momentum)?

Exercise 2.3

The water waves shown in Fig. 2.1 effectively go around the corner once they make it through the openings. This is a common characteristic of all waves, not just water waves. What are some other examples of waves going around the corner, such as light

waves and sound waves?

Exercise 2.4

When we use two tuning forks that are tuned to almost the same frequency (ν_1 and ν_2), then the two sound waves will interfere. We will hear an overall sound wave with frequency $(\nu_1 + \nu_2)/2$, and on top of it we will hear a beat. This beat affects the amplitude of the combined sound wave, and the beat frequency is given by $(\nu_1 - \nu_2)/2$. What would be the equivalent in scattering of hearing a beat to the signal?

3

Characteristics of the Neutron and What Can Be Measured

This chapter and the next chapter (Chapter 4) contain very similar information as presented in the popular articles by Roger Pynn and Reinhard Scherm (listed in Chapter 1). We strongly encourage the reader to look through these excellent treatises.

The neutron interacts with the nucleus of an atom, so it can measure where the nuclei of the atoms are in a properly designed experiment. This is called neutron diffraction. If nuclei move around, such as when the atoms are participating in a sound wave, then the neutron will experience this changing in position of the nuclei provided the neutron hangs around for long enough to notice any perceptible change in position. This is called inelastic neutron scattering.

The neutron is also a little magnet, so when it encounters other magnets such as atomic magnetic moments, it will react accordingly because it feels a magnetic force. In a well designed diffraction experiment, the positions of the atomic magnets can be determined by looking at the scattering pattern of the neutron. If the atomic magnets are changing orientation, such as in a spin wave, then the neutron is affected by this movement, again provided the neutron hangs around for long enough.

These basic notions allow us to calculate what can be measured, and how to design an appropriate experiment. Let us look at a few examples. We will start with the easiest thing to visualize, and that is the motion of the atoms. A good rule of thumb is to try to ensure that the speed of the neutrons is of the same order of magnitude as the speed of the motion one is trying to map out. Thus, if one is studying the behavior of sound waves in a material, then the neutron should be traveling at a speed comparable to the sound wave velocity. If one tries to measure the slow diffusion of a particle through a liquid, then one should use neutrons that do not go very fast at all.

The reason for this (approximate) speed matching is not too hard to see. A neutron is a wave, and what we see is an interference pattern. Therefore, if we want to measure changes (in atomic positions) that take place over a certain amount of time, then we have to ensure that the neutron is present both at the start of the event, and at the end of it.

For instance, if we look at a sound wave, then the neutron has to be around for a
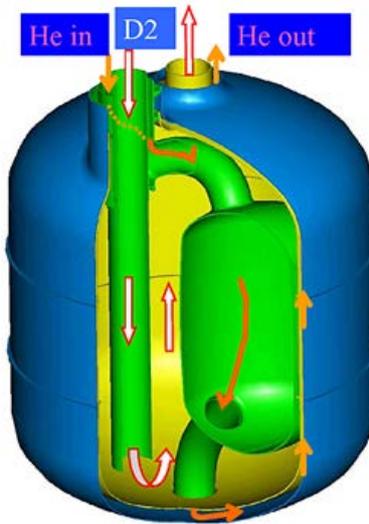
Fig. 3.1 The deuterium cold source at the OPAL reactor in Sydney, Australia (Source: www.ansto.gov.au). Helium is used for cooling the deuterium, the deuterium is used to slow down the neutrons.

time comparable to half the period of the sound wave (the time it takes to go from maximum to minimum amplitude). If the neutron hangs around for much longer, then we would simply measure the average of many up and downs and all the in-betweens, washing out the details of the information we are interested in. Conversely, if the neutron only sticks around for a much shorter amount of time, then there hardly will have been any changes in the positions of the atoms, and no interference pattern will be produced that carries any information about changes in positions that take place after the neutron has left the area.

The above immediately explains why the study of the motion of large structures (such as present in biological systems) takes place using cold source neutron spectrometers. Large objects tend to move slowly, and any motion on biological time scales takes forever to be noticeable (from the point of view of a neutron). Typical neutrons that come out of a reactor travel at speeds of several thousands of meters per second (corresponding to the temperature of the water that surrounds the core of a reactor), implying that they will only spend about 1 picosecond in an area the size of 50 Å. Should something not appreciably change until one waits for at least 10 nanoseconds, then nothing will be measured on a 1 ps timescale. The trick around this is to slow the neutrons down.

In order to slow neutrons down, one makes the neutrons collide with cold atoms so they can give up some of their energy. Only after they have given up their energy will they be sent down a beam tube to be used in experiments. This giving up of energy
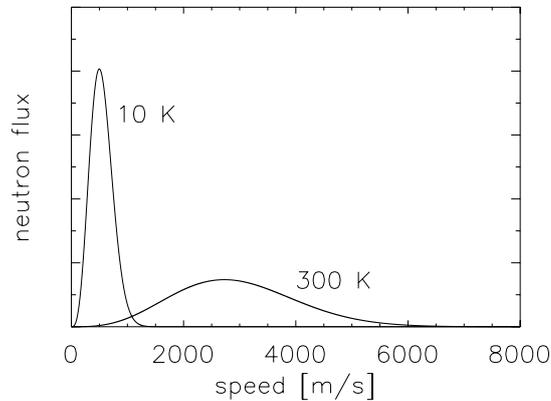
Fig. 3.2 The distribution of speeds v of neutrons [a.k.a. flux $\phi(v)$] that make it down a beam tube is strongly dependent on the temperature of the moderator. Shown are the same number of neutrons for the temperatures indicated in the figure, calculated using a Maxwellian velocity distribution $\phi(v) \sim v^3 e^{-mv^2/2k_B T}$.

is done in a cold source, which simply is a container with hydrogen or deuterium containing material kept at a cold temperature. When neutrons collide with light nuclei, they give up a substantial fraction of their kinetic energy with the result that after about ten collisions the neutrons will have the same temperature as the surrounding material. We say that the neutrons have been moderated. An example of such a moderator is shown in Fig. 3.1.

As a concrete example, when the hydrogen material is held at 10 K, then a typical neutron will emerge traveling at a speed of 500 m/s (see Appendix A and Eq. A.1). This is still fast, but considerable better than 2,000 m/s. Moreover, the neutrons will emerge from the cold source with a range of energies and speeds, and there will be plenty of neutrons that travel at speeds less than 200 m/s. This is good, because it implies that it would take such a neutron 25 picoseconds to traverse an area the size of 50 Å. Therefore, we can follow much slower processes in cold neutron scattering than we could with thermal neutron scattering.

There is an additional consideration. Namely, the scattering angle also (partially) determines the time scale of the measurement. Looking back at the interference patterns shown in Fig. 2.3, we see that the probing wave length depends on both the wave length (speed) of the neutron, and on the scattering angle. The smaller the scattering angle, the larger the probing wavelength. This probing length is directly linked to the size of the objects that a neutron can see.

Perhaps this size business is best phrased as a question, namely: how big is a neutron? We cannot think of the neutron as a tiny ball in a scattering event since a tiny ball would not give us an interference pattern. We have to think of the neutron as

Table 3.1 Approximate wave lengths and speeds of neutron. Use Eq. A.1 for exact conversions.

speed [m/s]	wavelength [Å]	Energy [meV]	time to traverse 5 nm
200	20	0.2	25 ps
400	10	0.8	12.5 ps
1000	4	5	5 ps
2000	2	20	2.5 ps
4000	1	80	1.25 ps

an incoming wave and a scattered wave with an interference pattern originating from the region where the scattering takes place. The size of this region is effectively the size of a neutron during a scattering event. The neutron cannot glean information about particles that are larger than this region, and it cannot tell us anything about the dynamics of the particles in this region if the neutron has already left the region before noticeable changes happen.

The size of this region of overlap between the incident and final neutron wave can be estimated through geometry, or by applying Heisenberg's uncertainty principle. Either way, it yields that the region of overlap (and therefore the size of the neutron during a scattering event) is proportional to $\sim \lambda / \sin(\theta/2)$, with θ the angle by which the neutron is scattered, and λ the wavelength of the neutron. For small angles this yields that the size of the neutron is proportional to the wave length of the neutron divided by the scattering angle λ/θ. This immediately tell us that if you want to study large molecules, you have to do small angle scattering, preferably using neutrons with long wave lengths (cold neutrons).

The bottom line is that the larger the objects you are interested in, the smaller your scattering angles will have to be, combined with using neutrons of long wave lengths (cold neutrons). The slower the dynamics you are interested in studying, the slower the neutrons have to be moving that you are using.

As to the specific characteristics of the neutron, we can summarize them in a table, but they are just numbers that give us the prefactors in how big a neutron is when it is scattered at a certain angle, and how likely it is that a neutron is scattered by a particular atom. These numbers can be put in recipes, but we do dwell on them here as they do not add to the basic understanding of neutron scattering. Instead, we look at the overall concepts that go into neutron scattering.

3.1   Why We ♡ Neutrons

The neutron has quite a few characteristics that make it very well suited to be used as a probe for the investigation of the structure and dynamics on atomic length scales. In no particular order, they are the following. The neutron, having no electric charge, only interacts very weakly with other materials. This means that a neutron is more likely than not to pass through a material. This is an advantage because this implies that materials are almost uniformly illuminated by neutrons; neutrons can scatter from anywhere in the sample, not just from the surface as is the case in conventional light scattering. For this reason, we call neutrons a bulk probe.

A major advantage of neutrons interacting weakly with materials is that they can easily penetrate a cryostat, furnace, or pressure cell used to house the sample and keep it at the desired environmental conditions. This allows for studying of the materials under non-ambient conditions. Clearly, this is essential if one would like to study, for instance, the superconducting state of a metal that only manifests itself at low temperatures.

A major disadvantage of neutrons interacting weakly with materials is that in general one requires a substantial amount of material, typically of the order of 0.1 cm^3. It is not always possible to obtain this much material, or obtain the amount of material in the desired state (such as a single crystal as opposed to a powder, a collection of very small grains).

Another advantage of neutron scattering is that the likelihood that a neutron is scattered by a particular atom depends not only on the type of atom, but also on the type of isotope. Moreover, the likelihood is not such that heavier atoms are more likely to scatter neutrons than light elements. The latter is the case in xray scattering where it is very easy to scatter photons by the heavier elements, but much more difficult to scatter from the lighter elements. In particular, one can easily study hydrogen containing samples using neutron scattering, but hardly at all by using xray scattering. This makes studying (some details of) biological materials using xray scattering very difficult.

The fact that each individual isotope has a different scattering power for neutrons allows for more versatility. For instance, if one is interested in the behavior of a particular type of atom in a material, but not so much in the behavior of another type, one can choose suitable isotopes to bring one type of atom to the fore at the expense of the other. This is most important in biological materials where we can vary the scattering strength of the hydrogen by substituting deuterium in place of hydrogen. We can even choose mixtures of hydrogen and deuterium in such a way that all the other non-hydrogenous atoms are emphasized. This is called contrast matching and it is what makes neutron scattering such a powerful tool for studying biological materials.

In (non-magnetic) neutron scattering, the neutrons are actually scattered by the nuclei of the atoms, not by the clouds of electrons as is the case for light scattering.

By extension, if we combine neutron and xray scattering experiments, then we learn complementary information, such as the shape of the electronic clouds in combination with the positions of the nuclei. This tells us about how well atoms are bounded together, and in what way.

3.2   Two Types of Scattering: Coherent and Incoherent

There is one more difference between neutron scattering and light scattering, a difference which sometimes is a big headache, and which at other times can be a big advantage. In our interference pictures in the previous chapter we assumed that all the atoms react the same way to being struck by a neutron. This is not always correct. For instance, when we have a crystal made up of one species of atoms, but of different isotopes of this atom, then from a neutron scattering point of view we have a mixture of scattering centers. When this happens, not all atoms react in the same way to being struck by the neutron wave, there will be some variation. When isotopes are distributed randomly (as is the case in most materials), then there will also be a random component to the way a neutron is scattered. This has a noticeable effect on the scattering.

When there is a random component, then we will lose some of our interference pattern. After all, if the atoms that were placed at identical spots in the interference pattern of Fig. 2.4 do not respond in identical ways, then the individual scattered waves cannot add up in a perfect way either. As a result, our interference pattern (and as a result our measured scattering) will be washed out a bit. How much it will be washed out depends on the degree of randomness introduced. For instance, if the natural distribution of isotopes is such that 99% of all atoms are made up of an identical isotope, then our interference pattern will only be mildly affected. By the same token, if we have many different isotopes, but all isotopes more or less scatter neutrons very similarly, then also not much will happen.

However, if we have many different isotopes that scatter quite differently, we could potentially reach the case where we do not see a reflection of the interference pattern. It is as if the atoms were not located at the positions corresponding to the white bands. In the case of such randomness, we cannot say much about what one atom is doing in response to its neighbor is doing some up and down motion. Such information would no longer be contained in the interference pattern.

This randomness does not mean that neutrons are no longer being scattered, it only means that we cannot learn about the relative positions and motions of neighboring atoms. However, we can still learn what individual atoms are doing. After all, when an individual atom scatterers a neutron, then changes position and again interacts with the neutron wave, then both interactions will be done by the same atom, and therefore, we can expect the ensuing interference pattern to contain information about the position of such an individual atom relative to the positions of the same atom a short time earlier: we can follow the movement of individual atoms in this way.

When all atoms scatter in the same way, we call it coherent scattering. When we can only follow what individual atoms are doing, we call it incoherent scattering. Thus, coherent scattering tells us about relative positions of atoms and collective motions such as sound waves, incoherent scattering tells us about the individual motions of atoms such as diffusion through a liquid. A simplified way to remember the difference between the two is the following. When someone opens a bottle of perfume, you hear the unscrewing of the cap. This sound wave is a collective effect and it can be studied using coherent neutron scattering which would tell us what neighboring atoms are doing in response to being pushed. Conversely, when you smell it, it means that a molecule actually made it to your nose. This individual motion can be studied by incoherent neutron scattering where the neutron would scatter from the aromatic molecule in one position, and some time later by the same molecule in a different position. Clearly, to follow this you must ensure that the neutron goes slowly enough so that it sticks around for a long enough period of time for the motion of the stinky molecule to actually take place.

Interestingly enough, even collections of identical isotopes can give rise to incoherent scattering. This has to do with the actual mechanism through which a neutron is being scattered by a nucleus. Without going into details at this point, when a neutron is scattered by a nucleus, it fleetingly forms a compound nucleus with the protons and neutrons that are already there. The likelihood that such a compound nucleus will be formed determines how likely it is that a neutron will be scattered. In turn, the likelihood that such a compound nucleus can be formed, even fleetingly, depends on how many neutrons are already there (different isotopes), and on what these neutrons are doing. The latter has to do with the Pauli exclusion principle, but we can contort it in here to state that the neutron itself is a tiny magnet, and that the nucleus of an atom also can be a tiny magnet. Depending whether these magnets are lined up, at an angle, or pointing in the opposite direction, it will be more difficult or more easy for a compound nucleus to form.

How does this give rise to incoherent scattering even if there is only one isotope? If the magnets formed by the nuclei point in random directions (and this is what nuclei actually do), then sometimes these magnets will be lined up with the neutron's magnetic poles, and sometimes not. This introduces a randomness, and applying the same reasoning as we did earlier, this will result in incoherent scattering. After all, if the magnets corresponding to neighboring nuclei are not lined up, then neighboring nuclei will not scatter the neutron in identical ways.

Not all isotopes give rise to incoherent scattering. Namely, if the nucleus that the neutron is trying to scatter from is not a magnet, then there is no preferred orientations, and therefore all neutrons will be scattered in an equal manner. Nuclei that are not magnetic have that the magnetic moments of the protons and neutrons in them exactly cancel each other. For instance, helium-4 has two neutrons and two protons in its nucleus. The two protons have opposite spin (direction of magnetic poles) to each other and thereby cancel each other, and the same holds true for the two neutrons.

As a result, the nucleus of helium-4 is non-magnetic and all scattering is 100% coherent.

Xray scattering is always coherent scattering. This is because it is the electrons around the nuclei that scatter the photons, and one electron has an identical scattering power to any other electron. Thus, xray scattering will neither show an isotope effect (since all isotopes have the same number of electrons), nor any effect due to what happens inside the nucleus (since the nucleus is not what scatters the photons). Therefore, incoherent scattering is solely restricted to neutron scattering experiments.

Why can incoherent scattering be both an advantage and a disadvantage? On the one hand, if you are interested in the collective motion of atoms, then you do not want the sought-after signal to be overwhelmed by incoherent scattering that contains no information about collective motions. On the other hand, if you are interested in what individual molecules are doing, for instance when they are locked up in the cage formed by their neighbors and they are trying to escape this confinement, then incoherent scattering is the way to learn about it.

3.3   Neutron Wavelength and Conversions

Neutrons behave like waves when we do not look, this is a fact of life. And neutrons show diffraction effects just like any other wave when their wave length is comparable to the size of the opening they are trying to get through. This is all very well, but how big is the wave length λ of a neutron? This number can be determined by experiment. We can simply measure the wavelength for a certain setup, like one where we already know the size of the opening the neutron is trying to get through. Then we can change the speed of the neutron to see how the wave length depends on the speed. And, we can even repeat this experiment not just for neutrons, but for electrons and buckyballs.

The outcome of all these experiments is that the wave length of a particle is inversely proportional to the speed and to the mass of the particle. In other words, it is inversely proportional to the momentum $p = mv$ of the particle. The constant of proportionality is called h, Planck's constant. Thus, $\lambda = h/p$. This relationship was first put forward by Louis de Broglie, and has been verified ever since for any particle we can do experiments on. Interestingly, it also works the opposite way, we can calculate the momentum of a photon of light given the wavelength of light. We rewrite the de Broglie relationship as follows:

$$p = \hbar q = mv = \frac{h}{\lambda}; q = \frac{2\pi}{\lambda}. \tag{3.1}$$

In this equation, we have rewritten the momentum of a neutron as $\hbar q$ with $\hbar = h/2\pi$. The units of q are inverse length, and q is often referred to as the wave number of the neutron. When the results of most neutron scattering experiments are presented, they are presented as a function of q, which we should read as a short hand notation of the amount of momentum transferred by the neutron to the sample.

Now that we have the relationship between momentum, speed and wave length of the neutron, we can also express the kinetic energy of a neutron ($E = mv^2/2$) in terms of wave length etc. We show the conversions in Appendix A. One important factoid is that if we were to double the wave length of a neutron, then we halve its speed. At the same time, we cut down its energy by a factor of four upon doubling the wave length. These considerations will become important when we consider higher-order contamination of our (nominally) mono-energetic neutron beam.

3.4   Neutron Detection and Scattering Lengths

We can detect with great efficiency whether a neutron is scattered in a particular direction by placing a detector at that position. A standard detector contains helium-3; there are also other types of materials used in neutron detectors, such as lithium-6 and boron-10.

All the detection mechanisms in these detectors involve the capture of the neutron by a nucleus. In other words, we detect the presence of a neutron through a nuclear reaction. Nuclear reactions typically involve energy scales of MeV (million electron volts). During a neutron scattering experiment, the neutron will emerge with a kinetic energy of the order of meV (milli electron volts). This immediately tells us that our detectors are not capable of telling us what the energy is of the scattered neutron, we simply cannot detect a few meV in a reaction that involves a few MeV. For this reason, if we are interested in the energy of the scattered neutron, we have to determine this energy through other means before the neutron is detected. We will discuss this further in the context of triple-axis spectrometers.

In the most commonly used detectors, tubes filled with helium-3, the neutron is captured. The helium-3 nucleus does not actually turn into a helium-4 nucleus, rather tritium is formed and a proton is released from the nucleus. This charged proton is then detected by means of its ionization track. As a rule of thumb, we can more or less treat a neutron detector as a black box and not worry about it any further. There are some caveats which are further discussed in Chapter 9.

Neutron detectors employ isotopes which have a high cross section for neutron capture. Helium-3 is a very good candidate, and the most commonly used detector is a tube of 1" diameter filled with 10 bar of helium-3 gas. However, there is a worldwide shortage of helium-3, and the price of it has increased tenfold over the past few years. Alternatives are being investigated such as boron-10 detectors. It is likely that the European Spallation Source (ESS) that will be constructed in Sweden will be using alternative detectors by the time this new source becomes operational (in 2020).

This discussion brings us to the parameter b that we use to indicate how easily a nucleus scatters a neutron. This parameter is referred to as the scattering length, and we have tabulated them for most commonly used isotopes in Appendix B. The scattering length is determined by a quantum mechanical scattering process involving the nuclear force. Luckily for us, we do not have to understand the details of this

process, we can simply take the scattering length as a number that has been measured experimentally.

The scattering length is typically of the order of the radius of the nucleus r_N, but not identical to it. The fact that it is close to the radius of the nucleus makes sense when we think of the nucleus presenting a target to the neutron of size $4\pi r_N^2$, the quantum mechanical cross-section for a hard-sphere of radius r_N. [1] But let's not dwell on details, our interpretation is such that the meaning of the scattering length is that it is a measure of how effectively a neutron is scattered by a nucleus.

We have some conventions for calculating the cross-section σ that a nucleus presents to a neutron based on the tabulated scattering lengths: $\sigma = 4\pi b^2$. In Appendix C we show, using quantum mechanics, why the scattering length has to be of the order of the nuclear radius, and we also give an approximation on how to calculate scattering lengths using a square-well potential. This clearly is only intended for those that might be interested in such a topic, for everyday neutron scattering is suffices to simply look up the scattering lengths in a Table B.1.

We show the scattering lengths for naturally occurring mixtures of isotopes in Fig. 3.3. We immediately notice a few things. First, the scattering length is indeed of the order of r_N for most nuclei (r_N is given by the dotted line in this figure). We also see quite a bit of variation amongst the nuclei. In fact, we even see the occurrence of negative scattering lengths. How negative scattering lengths should be interpreted is explained in Appendix C, but it is important to realize that this is not a trick. When analyzing neutron scattering data for samples that contain isotopes with negative scattering lengths, one had better put the minus signs in the appropriate places.

Another thing that can be seen in this figure is that light and heavy nuclei have a similar scattering length and, therefore, a similar scattering cross-section. Compare this to xray scattering where the scattering cross-section increases proportional to the number of electrons; xrays are not very good at all in scattering off of light atoms, in particular hydrogen is practically invisible. Having badmouthed xrays in this paragraph, we should also mention that xrays have a big advantage over neutrons: the flux of photons greatly exceeds the flux of neutrons.

Scattering lengths can even be complex numbers. A complex number indicates that there is a chance that the nucleus will absorb the neutron rather than re-emit it. While it is possible to calculate incoherent scattering lengths and relate this to an incoherent cross-section, there is not much point in doing so since in practice we use these four numbers (Table B.1) for any given isotope:
b_{coh}, $\sigma_{\mathrm{coh}} = 4\pi b_{\mathrm{coh}}^2$, σ_{inc} and σ_{abs}.

---

[1] The factor of 4 is counterintuitive since classically it should be a factor of 1 since πr_N^2 represents the size of the target when you are trying to hit it with a ball. However, according to quantum mechanics a neutron is a wave, and therefore, it will envelop the entire target, so that the size of the target is the surface of the sphere, $4\pi r_N^2$.
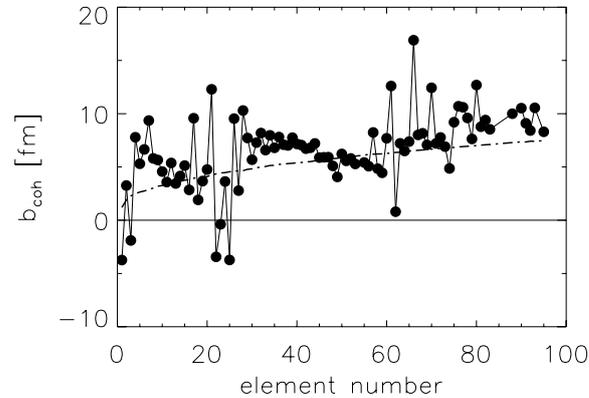
Fig. 3.3 The coherent scattering length as a function of element number for natural mixtures of isotopes. The dashed line is given by $b = r_N = 1.2A^{1/3}$, with A the number of nucleides (neutrons and protons).

The way we use these numbers is the following. We use the coherent scattering length b_{coh} when we model the scattered intensity in a powder pattern in the process of trying to determine the arrangement of the atoms, or when we scatter off of a multi-component mixture. We use the coherent scattering cross-section σ_{coh} when we have systems consisting of identical atoms, or when we want to calculate how thick we should make our sample for usage in a scattering experiment (see Chapter 11 on planning a neutron scattering experiment).

We use the incoherent scattering cross-section σ_{inc} to calculate the strength of the incoherent signal; this signal contains information about individual atoms, we do not have to deal with incoherent scattering lengths b_{inc} since we will not add two scattered waves (each with amplitude b_{inc}) originating from neighboring atoms. Finally, we use the absorbtion cross-section σ_{abs} to calculate how many neutrons will not make it through the sample, or out of the sample after having been scattered. For some isotopes, such as lithium-6 or Cd, the absorbtion cross-section is so large that one tends to avoid using these isotopes in neutron scattering samples: in order to give the neutron a chance to make it out of the sample, the sample should be very small. However, small samples do not scatter a lot, so a neutron scattering experiment on a strongly absorbing isotope is going to be a difficult one to perform.

The listed values for absorbtion cross-sections are for neutrons that travel at a speed of 2200 m/s. The reason for this is that the cross-section is inversely proportional to the speed of the neutron: the faster the neutron goes, the less likely it is to be absorbed by the nucleus. A useful, but incorrect, picture for this is that the longer a neutron takes to traverse a nucleus, the more chance it has of being absorbed. When calculating the absorbtion cross-section for one's sample, make sure to use the number evaluated at the actual speed v of the neutron:

$$\sigma_{\text{abs}}(v) = \sigma_{\text{abs}}(\text{tabulated})\frac{2200}{v}; \sigma_{\text{abs}}(E) = \sigma_{\text{abs}}(\text{tabulated})\sqrt{25.3/E}. \qquad (3.2)$$

This equation holds when speed v is measured in m/s, and the energy E of the neutron in meV.

As an example of all these numbers, we will calculate the efficiency of a 1 inch diameter cyclinder filled with 10-bar helium-3 for detecting neutrons that have an energy of 14 meV. A helium-3 detector is operated at room temperature, and even 10 bar of the stuff still behaves as an ideal gas. Therefore, we can calculate the number density $n = N/V$ of the gas inside the detector using the ideal gas law: $n = N/V = p/(k_B T)$ For a pressure $p=$ 10 bar= 10^6 N/m^2 we find n= 2.4 x 10^{26}atoms/m^3= 2.4 x 10^{20}atoms/cm^3. We will need this density to figure out how many nuclei there are in the path of the neutron when it is fired at a detector, so that we can calculate the likelihood that the neutron will not be captured by any of them.

An important measure characterizing our problem is the following dimensionless product: $nL\sigma_{\text{abs}}(E)$. Here, L is the length of material the neutron has to travel through, which would be 1" for the center of the detector tube. This product is almost self explanatory. The longer L, the less likely a neutron will not be captured. The higher the absorption cross-section, the less likely a neutron will not be captured. The more atoms there are per volume, the less likely a neutron will not be captured.

If we look at the problem as a transmission problem, then the chance that a neutron will make it through the next portion of the tube without being captured depends on how likely it was that the neutron made it to the start of this section. In nature, whenever something depends on how much there is at the outset, we can expect an exponential function. The same holds here, the transmission probability T is given by:

$$T = e^{-nL\sigma_{\text{abs}}(E)}. \qquad (3.3)$$

By extension, the detection efficiency is given by 1-T. When we look up σ_{abs} in a table, we find $\sigma_{\text{abs}}=$ 5333 barn= 5333 x 10^{-24} cm^2. A small target area indeed, but nuclei are small. Bearing in mind that this tabulated cross-section is for neutrons with an energy of 25.3 meV, we find (eqn 3.2) that the capture cross-section for a single helium-3 nucleus to be 7169 barn for neutrons that have an energy of 14 meV. Therefore, for a length L of 2.54 cm we find that the transmission is given by $T = 0.0126$. In other words, 98.7 % of the neutrons do not make it through the center of the tube, and therefore, these neutrons are being detected.

We would actually need to make a correction to this, as the tube is round and not square; some neutrons will try to make it through the much thinner sides. We can introduce an effective length L_{eff} that captures this effects, and it can be calculated to be (roughly): $L_{\text{eff}} = L\pi/4$. Using this effective length we find 96.8 % as the overall detection efficiency for neutrons with an energy of 14 meV. When we do the same calculation for neutrons of energy 80 meV we find 76%. Finally, for a 1/2" detector tube filled with 10-bar helium-3 gas (as in use on TRIAX) we find the detector efficiency

for a single tube to be 82 % for 14 meV neutrons.

3.5   Exercises

Exercise 3.1

The physics of neutron scattering is described by the wave function

$$\psi = e^{i\vec{k}.\vec{r}} + f(\theta)\frac{e^{ikr}}{r}.$$

The first part represents the incoming plane wave, the second part represents the spherical scattered wave. We want to verify that each part of this expression satisfies the Schrödinger equation for when the neutron is not in the sample.

a) Show that the plane wave $\psi(k) = e^{i\vec{k}.\vec{r}}$ is a solution of the free particle Schrödinger equation. Use spherical coordinates if you want to complete this exercise before the end of the semester.

b) Show that the scattered wave $\psi_{\text{scat}}(k) = f(\theta)\frac{e^{ikr}}{r}$ is a solution of the free particle Schrödinger equation, provided that $f(\theta)$ is independent of θ.

Exercise 3.2

The technique of contrast matching is mostly used for the study of biological systems by means of neutron scattering. This technique relies on the different scattering lengths that elements and isotopes have. When we talk about contrast matching in biology, we refer to partially deuterating our samples.

a) Make a table of the neutron scattering lengths for common biological elements, and make sure to include hydrogen as well as deuterium.

b) According to the attached figure, the scattering length of (roughly) a 9-91% mixture of H_2O-D_2O has a scattering length of zero. Do a calculation to find the exact percentage.

c) What is the meaning of having an H_2O-D_2O solution that has zero scattering length? In other words, what would the neutrons do when they encounter this mixture?

d) The figure shows scattering length density as opposed to scattering length. The scattering length density is calculated by taking the scattering length, and multiplying it by the number density of the system. Why is scattering length density the appropriate measure (as opposed to simply the scattering length) when we want to do contrast matching?
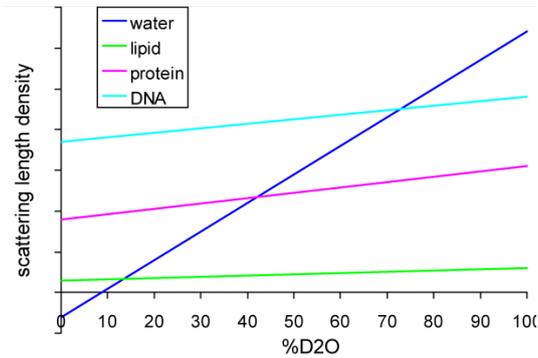
Fig. 3.4 The scattering length density for biological relevant materials as a function of deuteration percentage. Source: Wikipedia, Small-angle neutron scattering

e) Between a deuteration percentage of 40-45%, the scattering length (density) of a partially deuterated protein matches that of the H_2O-D_2O solution. So both the solution and the proteins scatter just as much. Why is this advantageous for certain studies?

Exercise 3.3

Consider a sample in the (slab) geometry shown in the figure. Typically, for neutron scattering experiments we want a sample that scatters about 5-10% of the incoming neutrons. If it is less, then we waste a lot of neutrons and might not find any discernable signal above the background level; if it is much more then many neutrons will be scattered multiple times and we no longer can deduce what is going on inside of the sample. In this problem, we will use 10 % as our target number.
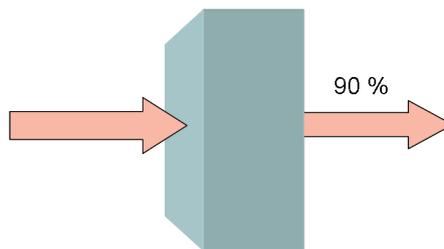
Fig. 3.5 The typical thickness of a sample is such that 90% of the neutrons (traveling from the left to the right) are transmitted (a.k.a. wasted) without scattering.

When we calculate how thick a sample should be, we use $\sigma = 4\pi b^2$ for the cross-section of an individual atom. We treat our sample as a collection of individual atoms

in our calculation. For instance, the total cross-section of a sample with N atoms would be $N\sigma = N4\pi b^2$.

We will calculate how thick our sample should be so that 90% of the neutrons are transmitted (and, therefore, 10% will be scattered).

a) Why do we only need to know the thickness of the sample? Why do the width and height not matter?

b) Argue that we are trying to solve the following equation,

$$0.9 = e^{-nL\sigma_{\text{coh}}}.$$

with L the thickness of the sample, and n the number density of the atoms in the sample.

c) What should be the thickness for a sample of lead? The density of lead is 11.342 grams per cubic centimeter, and the mass of one lead atom is 207.2 atomic mass units. Make sure that you do the unit conversions correctly, so that the argument of the exponent in eqn 3.4 is a number without units.

d) What should be the thickness of (for the physicists) a sample of activated carbon with 10% H_2 loading, or (for the biologists) a sample of heavy water (D_2O). If you are a mechanical engineer, calculate how thick the wall of your heat pipe should be so that 90% of the neutrons can fly through it. In order to calculate this, you will have to match the number density for each element to the cross-section of each element, or you can calculate the cross-section for one formula unit, and figure out how many formula units there are per volume.

4

What Is Measured (Ideally and in Reality)?

We have seen that neutron scattering is an interference technique and that it is used to map out the relative positions of atoms and to follow their motion, but we have not discussed exactly what it is that is measured and what the connection is with what we want to know about our samples.

What we measure is an interference pattern. This interference pattern is the combination of various scattered waves, and the scattered waves originate on individual atoms when these are exposed to an incoming neutron wave. Interference can occur when scattered waves from neighboring atoms start to share the same physical space, or it can occur when the wave that is scattered from an individual atom interacts with the wave that was scattered by the same atom moments earlier.

In all cases, we measure correlation functions. We measure the correlation between where one atom is at a certain time and where another atom is a fraction of a second later. It can even be the correlation where one atom is at a certain time, and where the same atom is a little bit later. Or, it can be a correlation between where one atom is at a certain time, and where another atom is at the same time. But in all cases, we measure the correlation between two (or more) scattering events.

Depending whether we are looking at more than one atom, at the same time or not, these scattering events and correlation functions go by different names. We show the names of these functions in Table 4.1. This table is for future reference, not for studying at this point. The important thing to realize is that beneath the nomenclature and the plethora of symbols, there is always an interference pattern originating from waves being emitted by scattering centers (the nuclei of atoms in most cases) when they are being struck by a neutron wave.

If neutron scattering were a black box technique, we would only need to throw a sample into the sample pit, and out would come the functions listed in Table 4.1 that tell us what we want to know about the material we are studying. Unfortunately, neutron scattering is not a black box technique, it is brown at best so we should have a look at what is measured, and how to get from there to the functions listed in this table.

In a neutron scattering experiment, we send a bunch of neutrons at the sample,

Table 4.1 Overview of the various correlation functions, and their interconnectedness. This list is incomplete, we refer the reader to standard textbooks for additional functions. $L[...]$ stands for the Laplace transform, and $Ft[...]$ for the Fourier transform (both from $\vec{r} \to \vec{q}$ and from $t \to \omega$). A and B stand for some microscopic quantity, such as the microscopic density n.

name	notation		definition
van Hove correlation function	$G_{AB}(\vec{r} - \vec{r}', t - t')$	$=$	$< A(\vec{r}, t)B(\vec{r}', t') >_{eq}$
response function	$\chi"_{AB}(\vec{r} - \vec{r}', t - t')$	$=$	$< [A(\vec{r}, t), B(\vec{r}', t')] >_{eq}$
relaxation function	$C_{AB}(\vec{r} - \vec{r}', t - t')$	$=$	$< [A(\vec{r}, t)]^* B(\vec{r}', t') >$
	with		$i\beta\partial_t C_{AB}(\vec{r}, t)/2 = \chi"(\vec{r}, t)$
intermediate scattering function	$F_{AB}(q, t)$	$=$	$Ft[G_{AB}(\vec{r}, t)]$
susceptibility	$\chi_{AB}(q, z)$	$=$	$L[\chi"_{AB}(\vec{r}, t)]$
		$=$	$\frac{1}{\pi} \int_\infty^\infty d\omega \chi"_{AB}(q, \omega) \frac{1}{\omega - z}$
dynamic susceptibility	$\chi_{AB}(q, \omega)$	$=$	$\chi_{AB}(q, z = \omega + i0^+)$
imaginary part of	$\chi"_{AB}(q, \omega)$	$=$	$\frac{1 - e^{-\beta\hbar\omega}}{2\hbar} S_{AB}(q, \omega)$
dynamic susceptibility		$=$	$Ft[\chi"_{AB}(\vec{r}, t)]$
dynamic structure factor	$S_{AB}(q, \omega)$	$=$	$Ft[F_{AB}(q, t)]$
symmetrized	$S_{AB}^{sym}(q, \omega)$	$=$	$Ft[C_{AB}(\vec{r}, t)]$
dynamic structure factor		$=$	$\frac{1 - e^{-\beta\hbar\omega}}{\beta\hbar\omega} S_{AB}(q, \omega)$
current-current correlation function	$C_L(q, \omega)$	$=$	$\frac{\omega^2}{q^2} S_{nn}(q, \omega)$
static structure factor	$S_{AB}(q)$	$=$	$F_{AB}(q, t = 0)$
static susceptibility	$\chi_{AB}(q)$	$=$	$\chi_{AB}(q, z = 0)$
		$=$	$\int_\infty^\infty \frac{d\omega}{\pi} \frac{\chi"_{AB}(q, \omega)}{\omega}$

and we measure how many are being scattered, under what angle, and sometimes even whether they gained or lost some energy in their interaction with the sample. The amount of energy lost by the neutron, $E = \hbar\omega$, is easy to picture, and fairly easy to measure. Especially when we secretly picture the neutron as a little ball in a collision. The amount of momentum transferred p to the sample, $p = \hbar q$, is also easy to picture and to measure. So what is the problem?

We list a few of the problems that stand in the way of translating the number of detector clicks to the functions listed in Table 4.1. This list is not complete (sorry), and it is not listed in order of importance as the order is determined by the sample, the spectrometer, and the problem that is being studied. The list also contains some jargon that will be explained in subsequent chapters; the jargon should not stand in the way of appreciating the multiple steps involved in going from counts to desired functions.

- Background problem: it can be that the sample is so small that it does not scatter a great many of the neutrons that we send in its direction. Perhaps in this situation when we position our detector somewhere we would only observe one count every minute or so corresponding to a neutron that was scattered by the sample. At the same time, our detector might pick up lots of unwanted neutrons consisting of neutrons that come from the sample container, or perhaps even neutrons that came from somewhere else altogether such as the spectrometer next door. These unwanted counts can drown the signal we are interested in- and, as a result- we cannot accurately determine the correlation functions from the scattering pattern.

- Multiple scattering: it can be that the sample is so effective at scattering neutrons that the neutrons that are trying to make it out of the sample after a scattering event will be scattered once more, or even twice more. If these multiply scattered neutrons make it to a detector, then we do not know how much energy and momentum was transferred in a single collision, we just know the overall amount transferred. This muddles the picture to such an extent that we cannot determine the correlation functions. We would not know how many neutrons had been scattered once (those are the events we want to count), how many twice, etc.

- Attenuation of the sample: the back of the sample can be exposed to fewer neutrons because it is in the shadow of the front of the sample. The shape of the sample can be such that it is easier for the neutrons to make it out of the sample in one direction than in another; this would skew our interference pattern and we would mistakenly believe that the sample has an intrinsic property which in reality is cause by an extrinsic one (the shape of the sample, not what the atoms are doing inside it).

- Higher order contamination: when a crystal is used as a monochromator or analyzer to select neutrons of a specific energy and wave length, then neutrons with half that wave length, or a third etc. will also be reflected in the same direction.

This is a direct consequence of the neutron being a wave: if some wavelength satisfies the Bragg condition [$\lambda = 2d\sin(\theta/2)$], than half this wave length will also satisfy it. Therefore, neutrons with 4 times, or 9 times the energy of the neutron we are interested in will also make their way to the sample or detector. This is called higher order contamination, and the quickest way around it is to use filters to remove the higher orders. However, this does not entirely remove all of the problems as explained in the next point.

- Monitor contamination: even when higher order contamination is removed before it reaches the detector, it can still have been counted by the incident beam monitor which is used to normalize the data. This can result in misjudging the number of neutrons with a particular energy that are impingent on the sample, with some energies affected more than others. Ultimately, this would lead to a skewing of line shapes.

- Instrument resolution: when we are doing experiment, we never just allow a neutron of one particular energy to hit our sample, we always have a range of allowed energies. This is unavoidable, otherwise we simply would not have enough neutrons to do experiments. We also allow neutrons within a certain range of angles to hit our sample and be detected. Again, this is unavoidable, all our pieces of equipment are finite sized, and crystals are not perfect but they have an intrinsic mosaicity. The result is that our scattered signal gets smeared out. If this smearing out exceeds the width of the features we are trying to study, then we have lost essential information. All together, the degree of smearing is called the instrumental resolution function, as it determines which features can still be resolved. If the resolution function only smeared out our signals a little bit, then we can still correct for it.

Despite the length of this list, it should not stand in the way of inferring the sought after correlation functions from an experiment, provided the experiment is well executed and the data are properly corrected. The high level of government investment in neutron scattering facilities should serve as proof that such careful experimental planning can indeed be accomplished.

4.1   The Connection Between Counts and What We Want To Know

In this section we will have a qualitative look at what is being measured. In an *ideal* world we have the following relationship:

Measured counts $\sim$ property of the sample.

The symbol $\sim$ means "directly proportional to". This is a very powerful statement. It says that the number of counts- something which has to do with the spectrometer, the overall setup, and with using neutrons as a probe- gives us direct information of

what is happening inside the sample, independent of the probe we were using or the details of the spectrometer. This implies that the results are not merely easy to interpret, it implies that there is no ambiguity as to what is being measured. Clearly, this is a good thing.

The above relationship (in bold) is valid in an ideal world, one that neutron scatterers do not live in. But it remains approximately valid, and by carrying out careful (but straightforward) corrections, we can restore the validity. For instance, unwanted background counts would mess up the relationship. However, by measuring the background, and by subtracting it from the measured counts, we can correct for this. The relationship also only holds for neutrons that were scattered only once. But we can model what fraction of neutrons were scattered more than once, and also subtract this number from the measured counts. The same holds for all other effects that could affect the above relationship; we can carry out corrections so that the relationship will hold true (again).

In order to use this relationship to our advantage, we need to be a little bit more specific about exactly what the left hand and the right hand sides entail. The left hand side is as easy as just sticking your detector somewhere and tallying up the number of counts, the right hand side requires some more thought.

The left hand side must be a measure of the likelihood that the sample wanted to scatter the neutrons. The higher the likelihood, the higher the number of counts. To make sense of this likelihood, we need to include the conditions under which the counts were gathered. First, we do not just measure the neutrons that were scattered in one particular direction only (the center of the detector), we measure all the neutrons that hit the detector. The detector spans a solid angle $\Delta\Omega$, so all neutrons that are scattered into this solid angle are counted. This is shown in Fig. 4.1

When we also measure what the final energy E_{final} of the neutron is, then we only accept neutrons at that energy plus some range ΔE around it. So what we would get out of it is the likelihood that a neutron is being scattered into a solid angle $\Delta\Omega$ with a range of energies in the interval $E_{\text{final}} - \Delta E/2 < E < E_{\text{final}} + \Delta E/2$. When we use the symbol σ to denote this likelihood (or cross-section as the terminology in scattering goes) then we can put some symbols to our relation:

$$\text{Counts} \sim \frac{d^2\sigma}{d\Omega dE}\Delta\Omega\Delta E \sim \frac{d^2\sigma}{d\Omega dE} \sim \text{ property of the sample},$$

that is:

$$\frac{d^2\sigma}{d\Omega dE} \sim \text{ property of the sample}.$$

The fraction on the left hand side is called the double differential cross-section. It is called double because we differentiate (distinguish) both with respect to solid angle, as well as to energy range. This is just terminology and math language, it still says
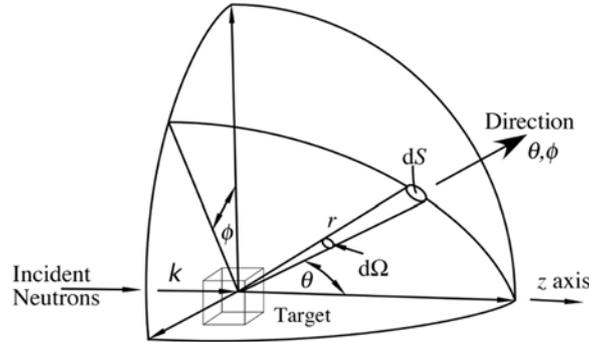
Fig. 4.1 Figure adapted from Squires' "Introduction to the Theory of Thermal Neutron Scattering" (Dover publications, 1996). The neutrons that are being scattered by the sample and that are within a solid angle ($d\Omega$ for an infinitesimal small solid angle) will be detected by a finite size detector with surface area ΔS that covers a solid angle $\Delta\Omega$.

'counts'. Note that we did not worry about $\Delta\Omega$ and ΔE. These are numbers, related for instance to the size of the detector; during the course of the experiment, these numbers will not suddenly change so we do not have to worry about them too much at this point.

We also use the differential cross-section. This is the appropriate measure when we do not care about what the final energy of the scattered neutron is; instead, we simply measure all of the neutrons that are scattered into a particular solid angle. In such a setup, without energy analysis, we simple have:

$$\text{Counts} \ \sim \ \frac{d\sigma}{d\Omega}\Delta\Omega \sim \frac{d\sigma}{d\Omega} \sim \ \text{some other property of the sample.}$$

By the way, if we simply measure all scattered neutrons, no matter what energy they have or what direction they were scattered in, then we measure the total cross-section σ. We are almost ready to look at what properties of the sample are being measured, we should just have a quick think about the constant of proportionality since, in fact, it is not a constant. Even though it is not a constant, we lumped it with the $\sim$ part because it does not tell us anything about our sample.

The constant of proportionality contains a term with the ratio of the speed of the scattered neutron and that of the incoming neutron: $v_{\text{final}}/v_{\text{initial}} = k_{\text{final}}/k_{\text{initial}}$. Thus, the expression for the double differential cross-section reads:

$$\boxed{\text{Counts} \ \sim \ \frac{d^2\sigma}{d\Omega dE}\Delta\Omega\Delta E \sim \frac{d^2\sigma}{d\Omega dE} \sim \frac{k_{\text{final}}}{k_{\text{initial}}} \ (\text{property of the sample})}. \qquad (4.1)$$

The origin of this ratio requires some math and quantum mechanics, which can be found in any neutron scattering textbook. We will not reproduce this here, we merely

mention from where this term originates. When we discussed what the meaning is of this cross-section quantity, we called it the ratio of how many neutrons will be scattered compared to how many neutrons hit the sample. In fact, we should have taken the ratio of how many neutrons per second are scattered by the sample, compared to how many neutrons per second are hitting it. This explains the term v_{initial}: it is a measure of how fast the incoming neutrons fly. The faster they go, the higher the incoming neutron flux and the more neutrons per second will hit the sample.

The term proportional to v_{final} (or k_{final} if you prefer) is trickier. Ultimately, we are interested in how prone the sample is to being kicked into an excited state, we are interested in the likelihood that such a transition will happen. For instance, we are interested in how willing the sample is to absorb 1 meV in energy. What we actually measure in an experiment is the likelihood of this transition, *times* the number of possible ways that this transition can be materialized. We are interested in the former, not in the latter. However, it can be shown (by Léon van Hove) that the number of possible ways is proportional to the final speed of the neutron. Hence, by explicitly including this factor in the expression of the cross-section, we are ensured that what we end up with is a property of the sample, not whether we probed the sample with high energy or low energy radiation.

4.2   Equal Time Correlation Functions

The connection between detected neutron counts and the intrinsic properties of the sample was figured out in 1954 by Léon van Hove. His original paper remains a very good starting point for learning about correlation functions and we certainly encourage the reader to study his manuscript. In this section, we will follow a somewhat opposite approach. We will look at some fairly simple systems, and wave our hands into what is being measured when we do neutron scattering experiments on these simple systems. We will then make our systems slightly more complicated, and we will arrive at the main equation for diffraction experiment at the end of this section.

In fact, if you want to try to skip this section, here is the main equation we will 'derive'. If it already makes sense to you, then skip the remainder of this section.

$$\boxed{\frac{d\sigma}{d\Omega} = \langle \sum_{i,j=1}^{N} b_i b_j e^{i\vec{q}.(\vec{R}_i - \vec{R}_j)} \rangle}. \tag{4.2}$$

We will start with a diatomic molecule, such as O_2 or N_2. We assume that this molecule is in a gas of similar molecules, but that the molecules do not influence each other. That is, the motion and position of one molecule has no effect on the other, and therefore, we will not see such an effect in neutron scattering experiments. By extension, we might as well look at a single molecule during a scattering process. We picture such a molecule in Fig. 4.3.

In this figure we are using neutrons of a fixed wave length, and we have placed our detector at a scattering angle of 10^o. All these numbers are simply to illustrate

Fig. 4.2 Léon van Hove provided the field on neutron scattering with the theoretical impetus. Photo source: CERN courier, April 30, 2001.

the connection between the interference pattern, and a property of the system we are studying. In addition, we have put our diatomic molecule in our neutron beam at some random angle. The neutron will scatter from both atoms of the molecule, but because of the separation between them, the two atoms will not be located at similar positions in the interference pattern. In this figure, we see that one atom is located just below a white band, while the other is slightly above a band. Since these bands correspond to identical points in the up and down cycle of the probing wave, we see that the neutron waves re-emitted by the atoms are not going to be exactly 'in phase'.

'Phase' is actually a number that tells us what fraction of a cycle (of a wave) has been completed. For example, if we look at a wave that starts with 'up', proceeds to 'down' and goes back 'up', then a phase of 1/2 will correspond to 'down', a phase of 1 will correspond to 'up', a phase of '1 1/2' corresponds to down, a phase of 1/4 corresponds to zero. This is shown in Fig 4.4. If we multiply this number by 2π, the number of radians in a circle, then we call it the **phase angle**. In other words, if we look at a sine or cosine wave, the phase angle would simply be the argument x in $\sin(x)$.

Back to Fig. 4.3. When we want to see whether the scattering from both atoms adds up constructively, or destructively, or somewhere in between, then we have to figure out what the difference is in phase between the two positions. In the figure, the difference in phase between the two atoms is a little over 3 (cycles), which can be seen immediately by counting the white bands. How do we do this counting in terms of math?

What we have to count is how many white bands, or probing wavelengths λ_{probe}, fit in the distance between the two atoms. Actually, not the distance *per sé*, but rather the projection of the distance onto the direction of propagation of the white bands
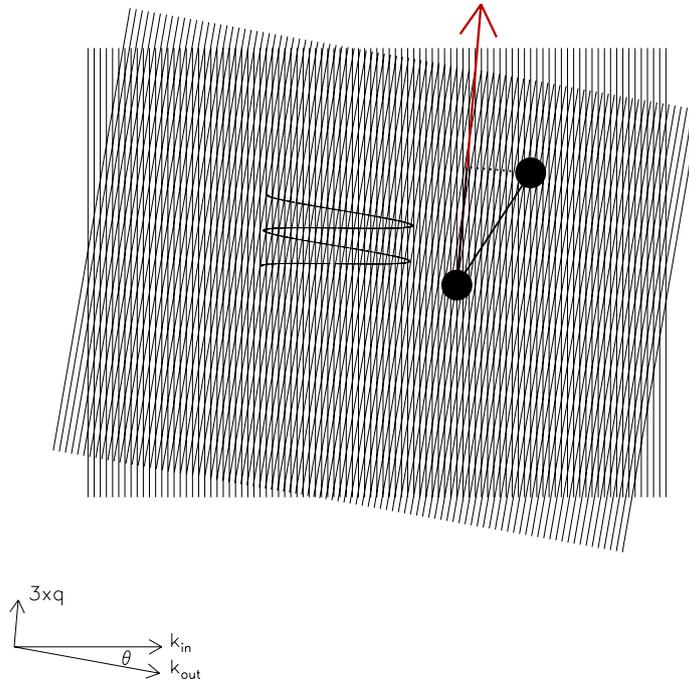
Fig. 4.3 A diatomic molecule as seen during one particular scattering event. The direction of momentum transfer of the neutron to the sample is given by the big red arrow in the figure. The incident and scattered neutron directions (scattering angle $\theta = 10^o$) are indicated below the figure by the arrows labeled k_{in} and k_{out}. The momentum transfer $\hbar q$ is also shown (to scale) in this part of the figure. In order to figure out the relative phase between two atoms that are being probed using a probing wave length indicated by the sinusoidal shape in the figure, we need to find the projection of the separation between the two atoms onto the direction of momentum transfer. This projection is shown as the black line on top of the red arrow.

(indicated by the big red arrow in Fig. 4.3). Hence, an expression for the difference in phase angle would be:

$$\text{phase angle difference} = \frac{2\pi}{\lambda_{\mathrm{probe}}} \text{ (projected distance)} = q \text{ (projected distance)}.$$

In here we have used eqn 3.1 to relate λ_{probe} to q[1]. Since the direction of momentum transfer is perpendicular to the white bands, we can rewrite the above equation using

---

[1]Note that q in this expression stands for the amount of momentum transferred, it does not stand for the momentum of the neutron. Eqn 3.1 is in fact very general: anything that has momentum has an associated wave length, and vice versa, including λ_{probe}.
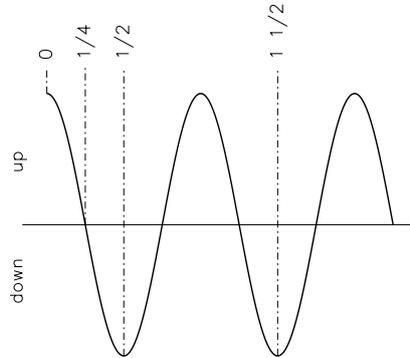
Fig. 4.4 The phase of a wave tells us what the amplitude of a wave is at a particular point. When we multiply the phases shown in the figure by 2π, then we get the phase angles. Note that a phase of '1/2' or '1 1/2' would have the same effect on an object subject to the wave pictured here. Hence it does not matter how many cycles have already passed, the only thing that matters is the fraction of a full cycle.

the scalar product $\vec{A}.\vec{B} = AB\cos\Theta$ with Θ the angle between the vectors $\vec{A}$ and $\vec{B}$. Thus,

$$\text{phase angle difference} = \vec{q}.(\vec{R}_2 - \vec{R}_1).$$

$\vec{R}_2$ and $\vec{R}_1$ are the position vectors of the two atoms, $|\vec{R}_2\text{-}\vec{R}_1|$ is the separation between the two atoms, and $|\vec{R}_2\text{-}\vec{R}_1|$ times the cosine between the two vectors $\vec{q}$ and $(\vec{R}_2 - \vec{R}_1)$ is the projected distance.

How would this phase angle difference enter an expression for what we can learn about the system? As can be seen from Fig. 4.3, we would get the same result whether the phase difference is 2.4 or 3.4 cycles. Atoms would respond in the same way, so we would not be able to see any difference. Therefore, our mathematical expressions should reflect this in the sense that the functions we encounter should repeat themselves when we change the argument of these function by 2π. Such functions are sines and cosines and combinations thereof. General combinations of sines and cosines are the exponential function with a complex argument (see Fig. 4.5):

$$\text{scattered signal} \sim \frac{d\sigma}{d\Omega} \sim e^{i\,(\text{phase angle difference})} = e^{i\vec{q}.(\vec{R}_2 - \vec{R}_1)}.$$

The strength of the scattered signal should (also) depend on how effective each atom is at scattering a neutron wave. We will use the symbol b_i to describe how effective atom i is at scattering the neutron. We call these b's the scattering length of an atom. These numbers were discussed in the previous chapter and are listed in Table B.1, but as a reminder, they are typically of the order of the radius of the nucleus (that is, pretty small), so that the square of this number represents a surface area. Think of
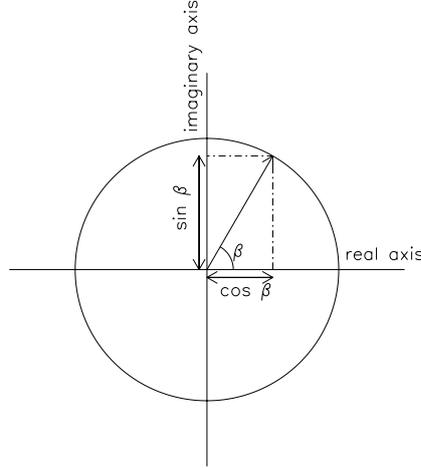
Fig. 4.5 The exponent of a complex number, $e^{i\beta}$, is a convenient way to describe a combination of sines and cosines, since $e^{i\beta} = \cos\beta + i\sin\beta$.

it as how big a target a nucleus represents when seen by a neutron. Putting in these measures of cross section, we are edging closer to our final expression (not there yet):

$$\frac{d\sigma}{d\Omega} \sim b_1 b_2 e^{i\vec{q}\cdot(\vec{R}_2 - \vec{R}_1)}.$$

There are actually three terms missing from this expression. Let's put them in, and we can see that they make sense:

$$\frac{d\sigma}{d\Omega} \sim b_1 b_1 e^{i\vec{q}\cdot(\vec{R}_1 - \vec{R}_1)} + b_2 b_2 e^{i\vec{q}\cdot(\vec{R}_2 - \vec{R}_2)} + b_1 b_2 e^{i\vec{q}\cdot(\vec{R}_1 - \vec{R}_2)} + b_2 b_1 e^{i\vec{q}\cdot(\vec{R}_2 - \vec{R}_1)}.$$

We can see where the three new terms come from. We had been looking at the difference in phase angle between atom 1 and 2, but the difference between atom 2 and 1 would be subject to the same reasoning. Also, the neutron wave originating from atom 1 can easily interfere with the neutron wave originating from atom 1, an instant later. The same holds for atom 2. In short, we should have run over all possible combinations. We can write the above expression in a more compact form, which we can then easily generalize to include more atoms for different systems:

$$\frac{d\sigma}{d\Omega} \sim \sum_{i=1}^{2} b_i e^{i\vec{q}\cdot\vec{R}_i} \sum_{j=1}^{2} b_j e^{-i\vec{q}\cdot\vec{R}_j} = \sum_{i,j=1}^{2} b_i b_j e^{i\vec{q}\cdot(\vec{R}_i - \vec{R}_j)}. \tag{4.3}$$

Suppose that instead we had started this section by simply giving this equation (actually we did just that). We still could have made sense of it. Instead, we would have read it as follows: the number of counts we measure is given by the sum over all possible interference patterns created by two atoms, including an atom with itself.

Therefore, we expect to find the difference in position in the argument, and we expect to find how effective individual atoms are at scattering the neutrons in there too. We get complex exponential functions (combinations of sines and cosines) because not only is it the relative phase that matters, but on top of it we do not care if our wave is in an upswing for the first time, or for the thirtieth time. We merely care whether it is in an upswing or in a downswing.

Eqn 4.3 is correct for the situation depicted in Fig. 4.3. However, this figure is very unlikely to represent an actual scattering experiment. After all, in how many scattering experiments do we hold a diatomic molecule at a given angle? In a real experiment, we would just have a gas of those molecules. All these molecules are zooming around, and their orientation with respect to the direction of momentum transfer can be anything. In short, we will measure an average over all possible orientations. Therefore, our expression for an actual scattering experiment should reflect such an average. This is no problem:

$$\frac{d\sigma}{d\Omega} \sim \langle \sum_{i,j=1}^{2} b_i b_j e^{i\vec{q}\cdot(\vec{R}_i - \vec{R}_j)} \rangle, \tag{4.4}$$

where the brackets $\langle...\rangle$ denote such an average. In physics averages go by many names, in this booklet we shall use the term ensemble average, even when this is not appropriate.

When we perform the average over all possible orientations of a diatomic molecule made up of identical atoms ($b_1 = b_2$), we get the following expression (see exercise 4.2):

$$\frac{d\sigma}{d\Omega} \sim b^2[1 + \frac{\sin(qa)}{qa}],$$

with $a = |\vec{R}_2 - \vec{R}_1|$, the separation between the two atoms in the molecule. In Fig. 4.6 we show the measured counts for a collection of nitrogen and oxygen molecules. We can easily see the oscillation corresponding to $\sin(qa)$ reaching successive maxima. Denoting the distance between maxima by k^*, we must have that $k^*a = 2\pi$, or $a = 2\pi/k^*$. Reading of k^* from the figure ($k^* = 5.7$ Å^{-1} for N_2 and $k^* = 5.2$ Å^{-1} for O_2), we find the distance between two atoms in a nitrogen molecule as $a = 1.1$ Å and $a = 1.2$ Å in an oxygen molecule.

Let us now generalize eqn 4.4 to a system containing N identical atoms. We drop the restriction that the atoms form diatomic molecules, or that they do not interact with each other. The generalization is rather straightforward:

$$\boxed{\frac{d\sigma}{d\Omega} = Nb^2 S(\vec{q})}, \tag{4.5}$$

with

$$\boxed{S(\vec{q}) = \frac{1}{N}\langle \sum_{i,j=1}^{N} e^{i\vec{q}\cdot(\vec{R}_i - \vec{R}_j)} \rangle}. \tag{4.6}$$
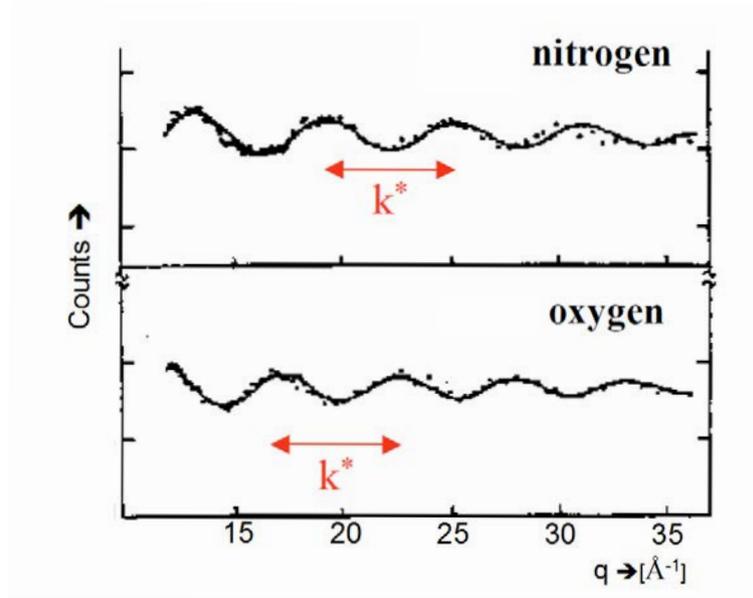
Fig. 4.6 Neutron scattering experiments on diatomic gases or liquids immediately reveal the fixed bond distance between the two atoms that make up the molecule. This distance a gives rise to the observed oscillation, with the period of the oscillation given by $qa = 2\pi$. Source: neutron scattering course by Ignatz de Schepper, source of original data unknown.

$S(\vec{q})$ is a correlation function that is called the static structure factor. It is called a correlation function because it correlates the position of atom i with that of atom j, and it is called 'static' because this expression represents a snapshot in the sense that we did not include any time dependence in the position vectors. For this reason, we also refer to this class of correlation functions as equal-time or snapshot correlation functions. The reason why it is called structure factor is because it tells us something about the atomic structure within our system, such as the distance between the atoms in a diatomic molecule in our example of Fig. 4.6.

Note that eqn 4.6 makes no reference to the neutron. It is a property of the system, it does not contain any references to the probe that was used to distill this information. This was our aim in this chapter: we wanted to show that the number of counts in the detector was directly proportional to a property of the sample.

The static structure factor is always defined by eqn 4.6, for liquids, gases and solids alike. There will be differences between these three phases of matter, of course. For instance, when we scatter by a liquid we do not get any directional information back simply because there are no preferred directions in a liquid. Therefore, in a liquid the averaging contained in the brackets $\langle ... \rangle$ will average over all possible angles. For single crystals we would not perform this directional averaging since crystals do have

specific orientations. This can easily be seen in a neutron scattering experiment where a crystal in the correct orientation will scatter almost all of the neutrons, while it will hardly scatter any when we rotate the crystal by a few degrees.

We will look into examples of the static structure factor for crystals when we discuss diffractometers (Chapter 5), and for a collection of many small crystals (called a powder) when we discuss powder diffractometers (Chapter 6). We now briefly look into the general shape of the static structure factor for a mono-atomic liquid, such as argon as this will be illustrative for how measurements as a function of q relate to what is going on in a system as a function of the distance between atoms.

In Fig. 4.7, we show the measured counts for neutrons scattered by liquid argon at 85 K. In fact, the count rate has been converted according to eqn 4.6 to give us the static structure factor $S(q)$. As can be seen in this figure, there is quite a bit of detail to be discerned in this figure, which should tell us something about the structure in the liquid.

First of all, we observe that the oscillations in $S(q)$ die out with increased q. Recalling that q is inversely proportional to the probing wavelength, we see that we lose our interference pattern when the probing wave length becomes very small. The reason for this is illustrated in Fig. 4.8.

For very small probing wave lengths, we see that the phase difference resulting from the separation between two atoms varies rapidly as a function of their separation. Since the separation between two atoms in a liquid is not fixed, we can expect to find a wide distribution of phase differences. Sometime we find that a pair of atoms generates a phase difference corresponding to an integer number of cycles (constructive interference), sometimes to half an integer (destructive interference), and most of the times to some number in between these two extremes. When we carry out the averaging over all distances and over all angles between the direction of momentum transfer and the direction of this pair of atoms, then we are averaging over all possible (almost random) numbers and we cannot distinguish any particular interference pattern. This is why we observe the oscillations to die out, we simply observe scattering by N individual atoms, just as if the positions of the atoms are not correlated with each other. When this scattering is normalized by dividing the total signal by N to yield $S(q)$, then $S(q)$ will be seen to reach a level of 1 (as indeed is observed in the figure). Note that this is different from our example of diatomic molecules where we had that one distance did not vary, namely the distance between two atoms in a molecule.

While a liquid does not exhibit a periodic structure, there is a considerable amount of short range order. This is caused by the fact that atoms do not sit on top of each other. Especially in a dense liquid, like argon in our example, we see that we have a fairly well defined average distance between neighboring atoms. Therefore, when we probe the liquid with a wavelength corresponding to this average distance ($\lambda_{\mathrm{probe}} \approx d_{\mathrm{avg}}$), then we can expect to see constructive interference (see Fig. 4.8). This
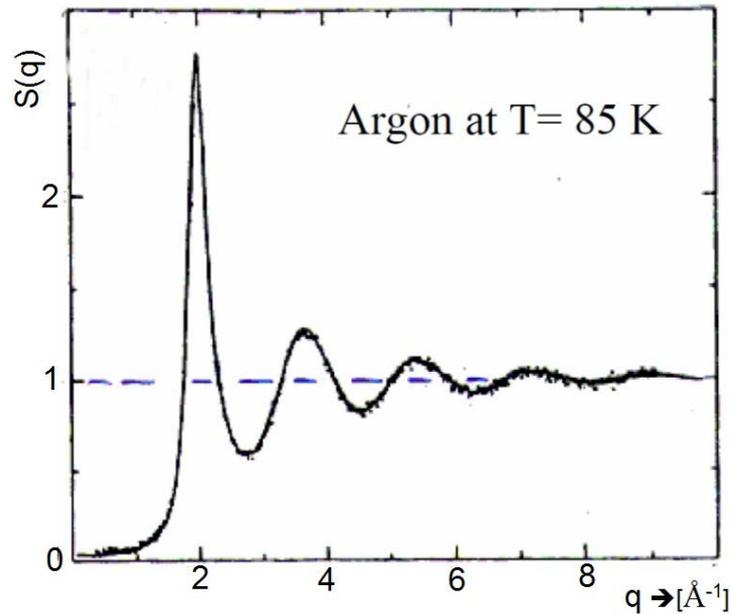
Fig. 4.7 Neutron scattering experiments on liquid argon at $T=$ 85 K are a direct measure of the static structure factor $S(q)$ for this mono-atomic system. We can see that for large momentum transfers the structure factor approaches 1, the limit where we no longer distinguish any interference between neighboring atoms. This is also known as the incoherent limit since we only observe individual atoms at these high q-values. The fairly sharp peak at $q_{\max} \approx 2\text{Å}^{-1}$ is (roughly) given by the average interatomic distance d_{avg} as $q_{\max} = 2\pi/d_{\text{avg}}$. The solid line through the points is the result for a molecular dynamics computer simulation using a Lennard-Jones interaction potential. Source: neutron scattering course by Ignatz de Schepper, source of original data unknown.

is the origin of the main peak at $q \approx 2$ Å^{-1} in Fig. 4.7. The second and third and so forth peaks can also be traced back to this average nearest neighbor distance.

When the liquid becomes more and more dense, then the average distance becomes more and more similar between different pairs of atoms. As a result, the peaks in the structure factor will become sharper: they increase in height and they narrow in width. We are seeing that the probing wave length has to match the interatomic separation more and more accurately in order to get constructive interference. Conversely, when we decrease the density of the liquid, then we allow for more variation in interatomic separations, and the peak becomes less pronounced and broader. In the limiting case of a very low density gas, the peak disappears and the static structure factor starts resembling a constant: $S(q)=1$.
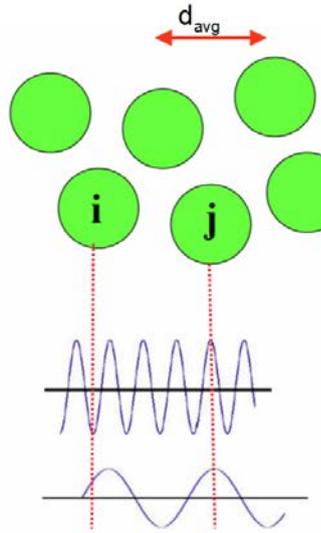
Fig. 4.8 In a liquid that consists of atoms displaying short range order while lacking long range order at the same time, we can expect the following. If we probe the liquid using a very short probing wave length (rapidly varying blue squiggle) then the phase difference between neighboring atoms will almost become random, and the interference pattern disappears. When we probe the liquid on a wave length corresponding to the average atomic distance d_{avg} (bottom blue squiggle) then we expect to see a certain level of constructive interference survive, much like the case of the diatomic molecule where we had one well-defined length. Source: neutron scattering course by Ignatz de Schepper.

We can also easily visualize what would happen to the static structure factor if the liquid were to freeze into a solid. In this case, all interatomic separations would be fixed, and the peaks of the static structure factor become very narrow (in q) and very tall indeed. Such peaks are referred to as Bragg peaks. All constructive interference already visible in Fig. 4.7 becomes very pronounced indeed, and all destructive interference (the dips in between the peaks in Fig. 4.7) also becomes very pronounced. In fact, it becomes so pronounced that all scattering in between the peaks disappears. We will deal with such structure factors in our chapters on diffraction and powder diffraction.

There is more information contained in the static structure factor of liquid argon shown in Fig. 4.7, such as the intercept $S(q = 0)$, but we refer the reader to specialized textbooks on liquids for these details. In fact, it is possible to carry out what is called a Fourier transformation of $S(q)$ to obtain a correlation function that directly deals with distances. This function is called the pair correlation function $g(r)$, and we give details of this transformation in Appendix D.

One final word in this section. When a system consists of more than one type of atom, then we cannot separate the neutron from the property of the sample quite as easily as we did in eqns 4.5 and 4.6. We can still do it for single crystals through analysis, but for liquids it would require additional measurements. Here is the problem for the general case:

$$\frac{d\sigma}{d\Omega} = N\frac{1}{N}\langle \sum_{i,j=1}^{N} b_i b_j e^{i\vec{q}.(\vec{R}_i - \vec{R}_j)}\rangle. \tag{4.7}$$

We can no longer take the terms $b_i b_j$ out of the summation, however, for the case of a single crystal we can still find a configuration (that is, the collection of position vectors $\vec{R}_i$ and $\vec{R}_j$) that will reproduce the measured differential cross section (given our knowledge of the tabulated values for the isotope dependent scattering lengths b). For liquids the situation is a little more tricky because the intrinsic directional averaging that accompanies scattering by a liquid no longer allows us to reconstitute the static structure factor without ambiguity. The solution in this case would be to repeat the measurement, but now with different isotopes, and hence, with different scattering lengths. Oh well.

4.3 Time Dependent Correlation Functions

When we measure the energy gain or loss of a neutron during a scattering experiment, we are in a position to follow the motion of atoms. Again, Léon van Hove figured out the precise connection between the (double differential) cross-section and what goes on in our system, the correlation function. Since we are talking about motion as a function of time on the one hand, and energy transfer on the other hand, a Fourier or Laplace transform will show up in the connection. So the generalization of eqn 4.7 will look something like this (FT stands for Fourier Transform):

$$\boxed{\frac{d^2\sigma}{d\Omega dE} = \frac{k_f}{k_i}\text{FT}[\langle \sum_{i,j=1}^{N} b_i b_j e^{i\vec{q}.(\vec{R}_i - \vec{R}_j(t))}\rangle]}. \tag{4.8}$$

It is a somewhat useful equation to picture what is going on inside the system. In practice, it is not a useful equation to analyze one's data. While it looks like a deceptively simple equation, the time evolution of the position of a particle actually provides us with a visualization problem because the position of a particle, and the position of the same particle a little while later do not commute. This is quantum mechanics speak for 'you need sharp eyes to see it' to quote Reinhard Scherm (see popular Scherm text mentioned in the first chapter).

In practice, we analyze our system directly in terms of how much energy is required to excite it, so that perhaps the neutron causes a phonon to appear. We will show some examples before too long, let's first make sense of eqn 4.8 the same way we did in the

preceding section, by looking at the special case of a mono-atomic system where all b's are identical. For this special case we rewite eqn 4.8 as

$$\boxed{\frac{d^2\sigma}{d\Omega dE} = N\frac{k_f}{k_i}b^2 S(\vec{q}, E)}, \tag{4.9}$$

with

$$\boxed{S(\vec{q}, E) = \frac{1}{N}\text{FT}[\langle \sum_{i,j=1}^{N} e^{i\vec{q}.(\vec{R}_i - \vec{R}_j(t))}\rangle]}. \tag{4.10}$$

The above equation shows a nice separation between properties of the neutron (the b's and the ratio k_f/k_i) on the one hand, and the intrinsic property of the system ($S(\vec{q}, E)$) on the other hand. As before, this strict separation is only possible for a mono-atomic system, however, the interpretation of the results is the same for a multi-atomic system.

$S(\vec{q}, E)$ is referred to as the dynamic structure factor of the system, and it is a measure of how easily spontaneous fluctuations in a system arise and decay. For completeness, the function that is Fourier transformed in eqn 4.10 to yield the dynamic structure factor is called the intermediate scattering function. There are a few useful relationships that pertain to $S(\vec{q}, E)$ that can serve as a check on the data analysis procedure. We list those relationships in Appendix E, here we only mention one, namely the detailed balance relationship:

$$S(\vec{q}, -E) = e^{-E/k_B T}S(\vec{q}, E). \tag{4.11}$$

This relationship is not too hard to derive, but we refer the reader to the neutron scattering literature for this. In words, the relationship states that the neutron is more likely to lose energy than to gain energy when being scattered by the sample. In fact, it holds for any type of probe, not just the neutron. In the extreme case of $T = 0$ K, the neutron can only lose energy. This of course makes perfect sense given that a system at zero Kelvin cannot give any energy. It is, however, a relationship that has practical consequences: when studying a system at low temperatures, do not bother with measuring on the energy gain side for the neutron, there is hardly any signal (counts) there.

There are correlation functions that are related to the dynamic structure factor, such as the dynamic susceptibility and the relaxation function. We will come back to these functions in Chapter 12 when we look into our data analysis, and what can be misinterpreted when scrutinizing one's data. For now, let's look at some examples for liquids and solids to gain a better understanding of what $S(\vec{q}, E)$ entails.

4.3.1   Liquids

Liquids are always a good starting point since it is (almost) possible to visualize the motion of atoms, and compare the expectations for these motions with what we observe in the dynamic structure factor. Since this section is meant to illustrate the
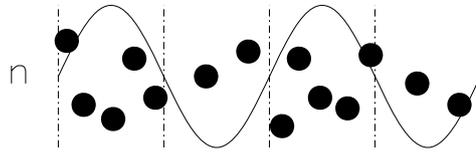
Fig. 4.9 A sketch representing microscopic fluctuations of the number density n. Essentially, what is depicted is a small scale sound wave with regions of increased density (compared to the average density of about 3.5 atoms between the vertical lines) followed by regions of decreased density. The wave length of this fluctuation is indicated by the sinusoidal line shape.

meaning of $S(\vec{q}, E)$, we will assume a mono-atomic liquid so that eqns 4.9 and 4.10 are valid. Also, since it is a liquid, we do not need to keep the vector symbol in $S(\vec{q}, E)$.

Liquids are capable of sustaining sound waves as anyone knows who has ever screamed under water. These sound waves also persist on a much shorter length scale, as shown in Fig. 4.9. A neutron can excite such a sound wave or absorb it, should it already be present. In order to do so, the probing wavelength λ_{probe} associated with the interference pattern has to match the wavelength of the fluctuation. This determines how much momentum the neutron should transfer to the sample. Also, a sound wave costs a certain amount of energy to create; the neutron has to give up this amount of energy.

When we put these consideration together, we expect to observe a peak in $S(q, E)$ at an energy transfer E, where E corresponds to the energy of the sound wave that has a wave length of $\lambda = 2\pi/q$. We show such a peak in Fig. 4.10. This figure only shows $S(q, E = \hbar\omega)$ for the neutron losing some of its energy, but one can clearly see a peak being present at non-zero energy transfers. The four panels in Fig. 4.10 show the results for increasingly shorter wave lengths. We observe, by noticing that the peak moves out to higher energy transfers for increased q, that is costs more and more energy to create a sound wave of shorter and shorter wave length. This observation is in line with everyday experience that low tones (long wave lengths) carry less energy than high tones (shorter wave lengths).

We can also see in this figure that there is a feature around $E = 0$, the elastic channel. This corresponds to diffusive motion. When we (locally) increase the density in a liquid, it does not necessarily lead to a propagating pressure wave (a.k.a. a sound wave). The excess density can also simply diffuse away. Such a process sometimes makes a neutron gain a little bit of energy, sometimes it loses a bit of energy depending on what direction the diffusing particles were heading. On average, there is no net energy gain. Therefore, diffusive features show up as broad (in energy transfer) lines around the elastic channel ($E = 0$). This type of scattering is also referred to as quasi-elastic scattering, and it is important in liquids, polymers and biological systems.
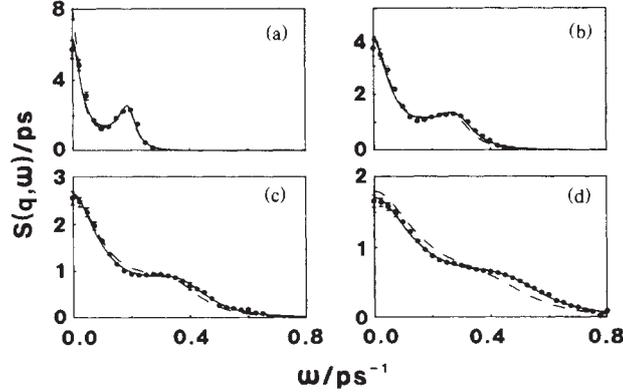
Fig. 4.10 Neutron scattering data (dots) for $S(q, \omega)$ for Ar at $T = 301$ K and at a particle density $n = 0.0504$ Å^{-3}. The figure has been adapted with permission from U. Bafile et *al.*, Phys. Rev. Lett. 65, 2394 (1990). The q-values are 0.050 Å^{-1}, 0.075 Å^{-1}, 0.10 Å^{-1} and 0.125 Å^{-1} for panels (a), (b), (c) and (d), respectively.

Let's also look at the other extreme, where the probing wave length is very short indeed, much shorter than the distance between the atoms. So we are again looking at the situation shown in Fig. 4.8. Even though we are now investigating how much energy is gained or lost when the liquid is probed at a certain wave length λ_{probe}, the same overall arguments- related to length scale- apply as before. At very short probing wave lengths, we will only get constructive interference between one atom and itself. Moreover, we will only be able to follow the short-time motion of this atom. This is easy to see, especially if we treat the problem classically by stating that $\vec{R}_i(t) = \vec{R}_i(t = 0) + \vec{v}_i(t)t$ and evaluating the large q limit:

$$
\begin{aligned}
S(q, E) &= \tfrac{1}{N}\text{FT}[\langle \sum_{i,j=1}^{N} e^{i\vec{q}.(\vec{R}_i - \vec{R}_j(t))}\rangle] \\
&\approx \tfrac{1}{N}\text{FT}[\langle \sum_{i=1}^{N} e^{i\vec{q}.(\vec{R}_i - \vec{R}_i(t))}\rangle] = \tfrac{1}{N}\text{FT}[\langle \sum_{i=1}^{N} e^{-i\vec{q}.\vec{v}_i t}\rangle].
\end{aligned}
\tag{4.12}
$$

If we wait too long (large t) then the argument $\vec{q}.\vec{v}t$ will be very large so we would be looking at a wildly oscillating function. Such functions do not give rise to constructive interference when the averaging $\langle ... \rangle$ is performed. Therefore, the only contribution we can expect is for small values of $\vec{q}.\vec{v}t$, which occur for short times. So what exactly are we measuring in this limit of (very) short probing wavelengths? The behavior of individual particles followed over very short times. Essentially, times so short that the atoms do not even have time to make it over to a neighbor and collide. This is very much like an ideal gas, where the particles do not interact with each other. This limit is therefore called the ideal-gas limit, and it has indeed been observed in many neutron scattering experiments on liquids.

Thus, whenever we measure at very high momentum transfers $\hbar q$ in liquids, then we follow the behavior of individual atoms over short times. In experiments, this shows up as a broad line shape centered around the recoil energy. While according to classical arguments ideal gas behavior shows up around $E = 0$, when we include quantum mechanical arguments, the scattering will be centered around the recoil energy $E_{\text{recoil}} = (\hbar q)^2/(2m)$ with m the mass of the atom in the liquid that causes the scattering.

4.3.2 Solids

What can we expect for solids? Clearly, we can expect to find sound waves. Therefore, as was the case in liquids, we expect to see peaks in $S(\vec{q}, E)$ corresponding to the energy required to excite a sound wave of wavelength $\lambda = 2\pi/q$. In fact, these sound waves should be much more orderly than the sound waves in a liquid, since all the atoms sit on well defined positions. We can expect a sound wave in solids to be able to persist for longer, and by extension, travel further. This brings up the following question: how would we observe- in neutron scattering experiments- how long something lasts for until it decays back to equilibrium?

The answer to this is that the sound waves will produce much sharper (in energy) peaks than is the case in liquids (such as the peaks shown in liquid argon in Fig. 4.10). The mathematical relationship between how long excitations persist for, and what we see in neutron scattering is that the width Γ (in energy) of a peak is given by the decay time τ of an excitation as follows: $\Gamma = 1/\tau$.

We can give a hand waving argument why the peaks sharpen up the longer the excitation lasts for. Looking back at the interference patterns shown in Chapter 2, we see that an interference pattern shows up as bands. But we have to bear in mind that these bands are not stationary. If there is a sound wave present in the crystal, or if we want to create one, then we have to make sure that the probing wave length matches the wave length of the bands. But we also have to ensure that our probing wave travels just as fast as the sound wave. Only when this is the case can we ensure that we still get constructive interference: when the points of high density of the sound wave have moved by half a wave length, our probing wave should have moved the same distance so as to maintain the condition of constructive interference. If we had not matched this movement, then we no longer would have had a perfect match.

Suppose a sound wave persists for quite a long time. During this time, we must make sure that the speed of our probing wave length maintains its match with the speed of the sound wave. The speed of our probing wavelength is determined (in part) by the amount of energy transferred. For a perfect match, we must be very precise in the amount of energy we transfer. A slight mismatch can result in going from a constructive interference condition to a destructive one after a few cycles. For instance, a mismatch of 10 % becomes a mismatch of 50 % after 5 cycles, which would give us perfect destructive interference.
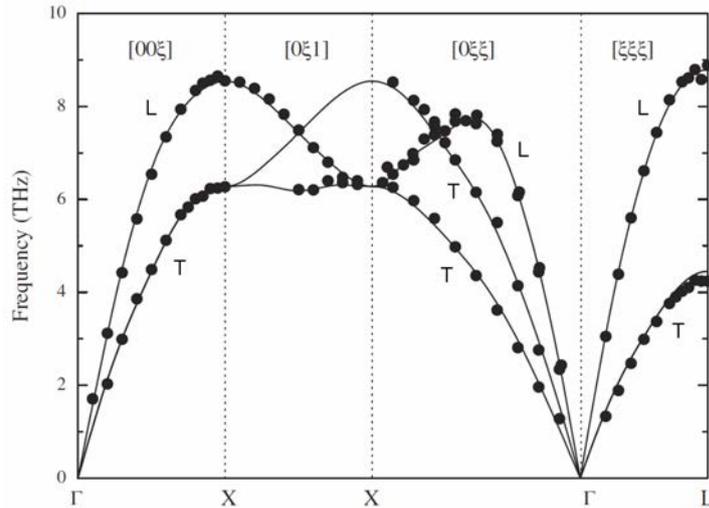
Fig. 4.11 Figure reproduced with permission from Z.-Y. Zeng *et al.*, Physica B 405, 3665 (2010). Shown is the dispersion relation for nickel measured along various crystallographic directions. The label 'L' stands for longitudinal, the label 'T' for transverse. The points are obtained by reading off the peak positions in inelastic scattering experiments (performed by Birgenau in 1964), the solid lines are theoretical calculations.

In contrast to this, if an excitation does not persist for very long, then we can get away with a less than perfect match. Suppose an excitation had decayed back to equilibrium by the time it has propagated over a distance comparable to its wave length. If we had probed this wave with a 10 % mismatch in propagation speed, then we still would have gotten constructive interference since in this case nothing exists anymore after 5 cycles to give us destructive interference. So we would not get perfect constructive interference, but we would get some. As a result, we would see some scattered signal not just at the exact match, but also for energy transfers close to it. In this way, the width (in energy) for which we still see a signal, is a measure of how long it took for the excitation to decay. A long-lived excitation required a close match over many cycles, so this would result in a sharp peak in our neutron scattering spectra. A fast decaying excitation would result in a broad peak. Thus, $\Gamma = 1/\tau$.

When there are multiple types of atoms in a solid, then we can actually have many more types of vibrations. Not all the atoms necessarily have to be moving in the same direction; we can also have that some atoms move to the left, and some move to the right. This latter type of out-of-phase vibration costs much more energy than a vibration where all atoms move in the same direction during the back and forth motion. When all atoms move in the same direction, we call them acoustic phonons (because you can hear them when the wave length is very long). When they move in opposite directions, we call them optical phonons. The latter term comes from the fact that
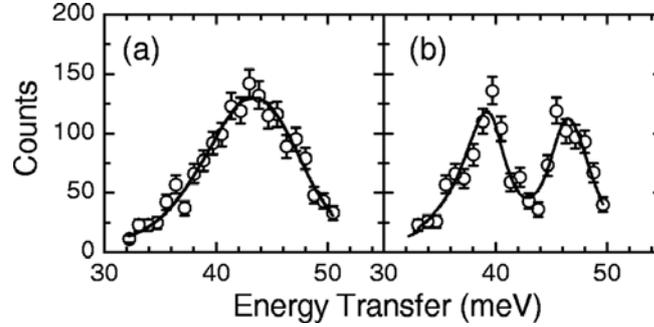
Fig. 4.12 Figure reproduced with permission from R. J. McQueeney *et al.*, Phys. Rev. B73, 174409 (2006). A magnetic spin wave in magnetite. The panel on the left shows the excitation just above the charge-order transition (the Verwey transition); the panel on the right shows the results just below the transition. This demonstrates that neutrons are sensitive to small changes in the magnetic excitation spectrum. Note the actual value of the energy transferred by the neutron, these magnetic excitations require quite a bit of energy to create.

they are visible in light scattering experiments.

Neutrons can easily excite both types of phonons. We can measure how much energy is required to excite these phonons as a function of their wave lengths, and as a function of direction of momentum transfer in the crystal. Some phonons propagate (vibrate) faster along one direction then along another direction. The excitation energies as a function of wave length and direction of propagation are referred to as the dispersion relations, an example of which is shown in Fig. 4.11 where the acoustic phonons are displayed as a function of momentum transfer along the three main crystallographic directions. Note that there are three acoustic phonon branches, one called a longitudinal branch, and the other two called transverse branches. We will revisit this topic when we deal with inelastic neutron spectrometers.

4.4   Incoherent and Magnetic Scattering

Phonons are an example of collective behavior, and it will be observed in neutron scattering experiments through the coherent cross-section. As mentioned, neutrons can also be scattered incoherently by the sample, and they are also sensitive to magnetic excitations through the magnetic cross-section. As such, eqn 4.9 is a little oversimplified, and it makes sense to write out the coherent, incoherent and magnetic contributions separately. We do this for a mono-atomic system, and use the relationship that $\sigma_{\mathrm{coh}} = 4\pi b_{\mathrm{coh}}^2$:

$$\frac{d^2\sigma}{d\Omega dE} = N \frac{k_f}{k_i} \left[\frac{\sigma_{\mathrm{coh}}}{4\pi} S_{\mathrm{coh}}(q, E) + \frac{\sigma_{\mathrm{inc}}}{4\pi} S_{\mathrm{inc}}(q, E) + \frac{\sigma_{\mathrm{mag}}}{4\pi} S_{\mathrm{mag}}(q, E) \right]. \qquad (4.13)$$

In other words, we can treat all three contributions separately, and we simply refer to them as the coherent dynamic structure factor, the incoherent dynamic structure

factor and the magnetic dynamic structure factor. The magnetic dynamic structure factor will contain information about the magnetic excitations in the system, such as the magnons (also called the spin waves, the magnetic equivalent of the phonons). We show an example of magnetic excitations in Fig. 4.12 for magnetite, the oldest known magnetic material.

Actually, we should have carried out the same separation (as shown in eqn 4.13) for the equal time correlation functions. We did not do it for the incoherent static structure factor, since that function would not make any sense. It would be independent of momentum transfer as it does not contain any information on structural effects, since it is related to scattering by a single entity. Anything incoherent has to do with individual atoms, with local excitations. But we definitely should differentiate between the magnetic and non-magnetic parts of the structure factor. The latter is called the nuclear structure factor; the former the magnetic structure factor. One routinely carries out these separations during the analysis stages of powder diffraction patterns (Chapter 6).

4.5 Exercises

Exercise 4.1

When we do neutron scattering on biological samples, we (frequently) use thin slabs rather than cylindrical shapes. Why is this?

Exercise 4.2

The scattering cross-section for diatomic molecules where the atoms are separated by a distance a is given by

$$\frac{d\sigma}{d\Omega} \sim b^2[1 + \frac{\sin(qa)}{qa}].$$

Derive this equation by carrying out the appropriate directional averaging that is contained within eqn 4.4.

Exercise 4.3

In eqn 4.12, repeated below, we showed (in the classical approximation) that for large momentum transfers the dynamic structure factor $S(q, E)$ for liquids looks like that of an ideal gas. At these high momentum transfers, we only observe the behavior of individual atoms over short times.

$$\begin{aligned}
S(q, E) &= \tfrac{1}{N}\mathrm{FT}[\langle \textstyle\sum_{i,j=1}^{N} e^{i\vec{q}\cdot(\vec{R}_i - \vec{R}_j(t))}\rangle] \\
&\approx \tfrac{1}{N}\mathrm{FT}[\langle \textstyle\sum_{i=1}^{N} e^{i\vec{q}\cdot(\vec{R}_i - \vec{R}_i(t))}\rangle] \approx \tfrac{1}{N}\mathrm{FT}[\langle \textstyle\sum_{i=1}^{N} e^{-i\vec{q}\cdot\vec{v}_i t}\rangle] \\
&= \mathrm{FT}[\langle e^{-i\vec{q}\cdot\vec{v}_1 t}\rangle]
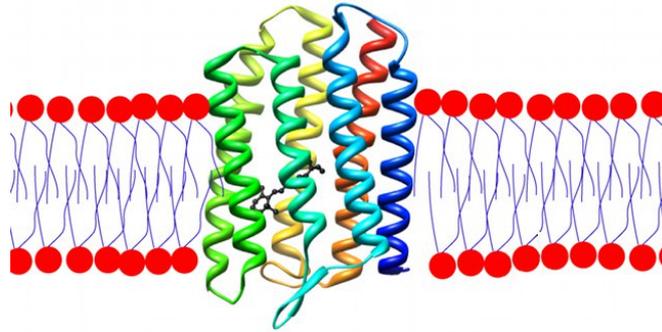\end{aligned}$$

.

Fig. 4.13 Rhodopsin embedded in a lipid bilayer. Source: Wikipedia, Rhodopsin.

In here, we will not worry about any pre-factors, but we will calculate $S(q, E)$ by evaluating this equation.

a) Use a Maxwellian distribution of molecular velocities to evaluate the right hand side of this equation before doing a Fourier transform:

$$[\langle e^{-i\vec{q}\cdot\vec{v}_1 t}\rangle] = \int e^{-(1/2)mv^2/k_B T} e^{-i\vec{q}\cdot\vec{v}t} d\vec{v} / \int e^{-(1/2)mv^2/k_B T} d\vec{v}.$$

b) Carry out the fourier transform to arrive at an expression for $S(q, E)$:

$$S(q, E) = \int_{-\infty}^{\infty} e^{iEt/\hbar} [\text{answer from a)}] \, dt.$$

Note that this is the classical expression for an ideal gas. When we include quantum mechanics, we find that the peak positions are shifted to $(\hbar q)^2/2m$, but their widths remain the same. This latter result is derived in many standard texts on neutron scattering such as Lovesey's.

Exercise 4.4

Looking at figure 4.10, and using the conversion table in Appendix A when needed, estimate the speed of sound for the argon gas when it was probed by the neutrons in this scattering experiment.

Exercise 4.5

A membrane shares some characteristics with a diatomic molecule, namely, we have a well defined separation between the two sides of the membrane. Therefore, when we do neutron scattering at high momentum transfers, we might expect a similar oscillation in the scattering cross-section as shown in Fig. 4.13. Estimate, or look up, the bilayer separation and make a sketch of the expected scattering profile for large momentum transfers for the membrane shown in the figure below (Fig. 4.13). What is

the expected period of oscillation in Å^{-1}?

Exercise 4.6

If you are interested in how tightly bound a particular atom is to its surroundings, such as a hydrogen atom inside of a carbon nanotube, you can learn this information by looking at the motion of this atom. For instance, you could imagine measuring the characteristic energy (frequency) of the harmonic oscillator potential that is trapping this atom.
Where in $q-$space should you be looking if you want to measure this bonding by doing a neutron scattering experiment? Why?

Part II

The instruments

This part of the booklet deals with the specifics of the most commonly used spectrometers at reactor sources. Every chapter discusses a particular spectrometer, not merely at the instrumental level, but it also discusses what types of experiments can be performed using these spectrometers, and what type of science can be done on them.

5
Diffractometers

Diffractometers are used to investigate the structure of materials, and changes therein when these materials are subject to a change in temperature or pressure. Powder diffractometers are a specialized version of diffractometers and they will be discussed separately in Chapter 6. In here, we look at the most basic version of the diffractometer: the 2-axis diffractometer at a reactor source, and the one-axis diffractometer at a pulsed neutron source.

The basic 2-axis diffractometer was developed at Oak Ridge National Laboratory by Clifford Shull and Ernie Wollan. The operation of this type of spectrometer is identical to an xray diffractometer that one might find in the chemistry department. Neutrons of one particular energy (wavelength) are scattered out of the main neutron beam by using the Bragg reflection of a monochromator crystal. Examples of such crystals are silicon, copper, and pyrolytic graphite (PG). This scattered beam of neutrons is directed at the sample we wish to investigate. The neutrons scattered by the sample are counted by moving a detector around the sample. This process is shown in Fig. 5.2.

A diffractometer at a pulsed neutron source does not require a monochromating crystal since the energy (wavelength) of the neutrons can be determined from the time it takes the neutrons to travel the distance between the source and the detector. As such, the only thing we would need for such a diffractometer is a single detector placed at some angle. We would not even have to move the detector around the sam-

Fig. 5.1 The inventors of the diffractometer, Ernie Wollan (left) and Cliff Shull.
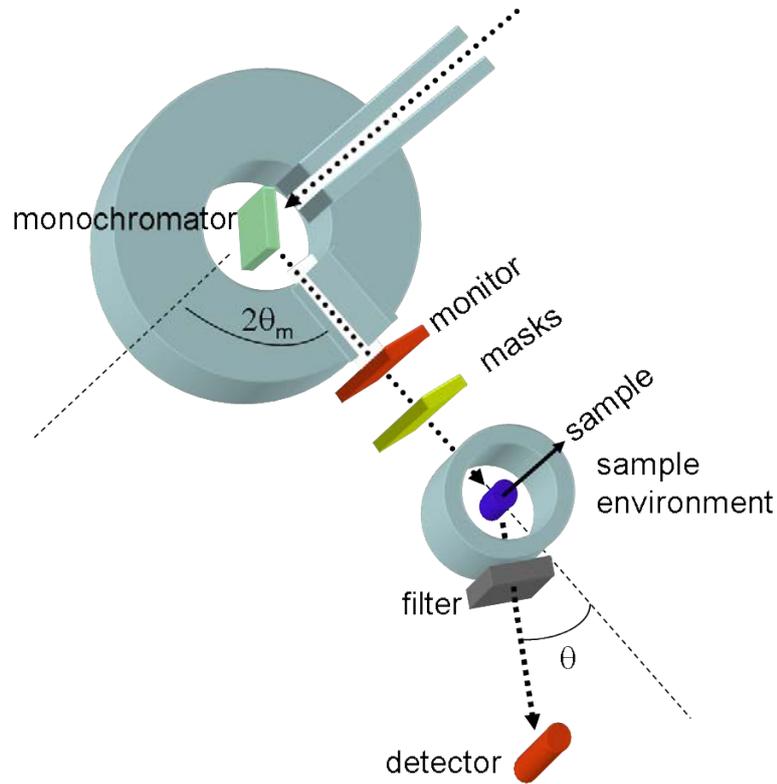
Fig. 5.2 A diffractometer at a reactor source requires a device for selecting a certain wave length out of all the neutrons that come down a beam tube. The monochromator crystal that selects these neutrons Bragg reflects these neutrons onto the sample. A detector can then be moved around the sample to measure the number of scattered neutrons as a function of the amount of momentum that has been transferred to the sample.

ple. Because every neutron incident on the sample that is scattered into the detector would have a different incoming wavelength, every scattered neutron would transfer a different amount of momentum to the sample. In this way, we can still measure how likely it is for the sample to scatter a neutron as a function of how much momentum is transferred between the neutron and the sample.

In practice, one is very unlikely to encounter a single detector diffractometer (even though one might well (ab)use a more sophisticated spectrometer as a single detector 2-axis diffractometer). An obvious improvement on the single detector 2-axis diffractometer is to use multiple detectors- located at different angles- or to use a position sensitive detector. For the case of a diffractometer at a reactor source, one would then measure multiple momentum transfers using a single incident neutron wave length. For the case of a multi-detector diffractometer at a pulsed source, one could greatly

Fig. 5.3 A straightforward diffractometer. In fact, it is one step above straightforward as this diffractometer (2XC at MURR) uses five detectors simultaneously. The openings in front of the five detectors can easily be identified in this photo, and the detectors themselves are tucked away inside of shielding material. Also visible is the sample mounted on the sample table.

reduce the counting time by measuring all the momentum transfers in multiple detectors simultaneously. An example of a multi-detector at a reactor source is shown in Fig. 5.3, an example of such an improved diffractometer at a pulsed neutron source is shown in Fig. 5.4.

For the remainder of this chapter, we will focus on the most primitive of diffractometers; namely, the 2-axis diffractometer at a reactor source. As mentioned, the powder diffractometer will be discussed in Chapter 6, whereas the single crystal four-circle diffractometer will not be mentioned at all. This latter instrument is no longer in very high use at reactor neutron sources now that single-crystal diffractometers at pulsed neutron sources have been developed. Another diffractometer which will not be discussed in great detail is the stress diffractometer used in engineering studies, although we will briefly describe this instrument in the last section of this chapter.

We will first discuss the basics of the 2-axis spectrometer, including its limitations. After that, we will have a detailed look at the components of the instrument.

5.1   The Skinny

The scattering profile that we measure on a 2-axis spectrometer can be interpreted without ambiguity when the neutrons are being scattered elastically, or almost elasti-
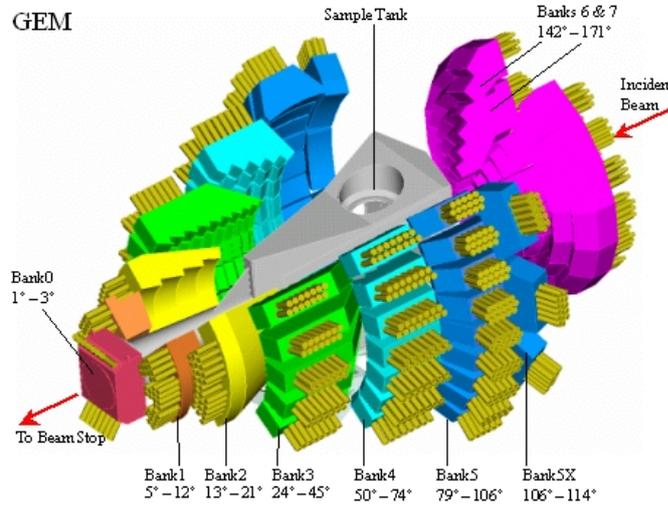
Fig. 5.4 Shown is a schematic of the General Materials Diffractometer GEM at the ISIS spallation source in the United Kingdom (wwwisis2.isis.rl.ac.uk). In principle, each of the many detectors is capable of measuring the full diffraction pattern, so the gain in neutron count rate is evident.

cally. In contrast, experiments on materials where the neutrons can gain or lose a large fraction of their energy during a scattering event require quite a bit of data massaging. There are two reasons for this, as we will detail next.

When we do a 2-axis scattering experiment, we wish to measure the static structure factor $S(\vec{q})$ of a material, with $\hbar\vec{q}$ the amount of momentum transferred from the neutron to the sample. We are not interested in any energy transferred from the neutron to the sample; thus, we are after (see Appendix E)

$$S(\vec{q}) = \int_{-\infty}^{\infty} S(\vec{q}, E) dE, \qquad (5.1)$$

with E the amount of energy transferred from the neutron to the sample, the difference between the incoming energy E_i and the final energy E_f. In an ideal world, we (would like to) use the neutron detector to do this energy integration for us. Especially when we use a neutron detector that is highly efficient in detecting all neutrons- independent of their energy- then we stand a good chance of accomplishing our goal.

In practice, we will only do a good job when the amount of energy transferred to the sample (if any), is small compared to the energy of the neutron. First, if we transfer a large amount of energy, then this implies that there will be a large difference between the incoming neutron wave number k_i and the scattered neutron wave number k_f. This messes up our nominal value of q as this is determined- for the case of diffraction- by

assuming that all scattering events are elastic ($k_i = k_f$):

$$|\vec{q}| = [k_i^2 + k_f^2 - 2k_i k_f \cos(\theta)]^{1/2} = 2k_i \sin(\theta/2).$$

Therefore, we would make increasingly more substantial errors in q when the difference between k_i and k_f increases.

An even more serious problem is that the exact integration of eqn 5.1 is not actually carried out by the detector. What we measure instead is related to our differential cross-section in the following approximation (stated here for the case of a mono-atomic sample so as to keep our notation from being too cumbersome):

$$\left(\frac{d\sigma}{d\Omega}\right)_{\text{eff}} = \int_0^\infty \epsilon(E_f) \frac{d^2\sigma}{d\Omega dE_f} dE_f = N \frac{\sigma_{\text{coh}}}{4\pi} \int_{-\infty}^{E_i} \epsilon(E_f) \frac{k_f}{k_i} S(\vec{q}, E_i - E_f) d(E_i - E_f)$$

In here, we have used the symbol $\epsilon(E_f)$ to indicate the energy dependent detection efficiency of our detector. Even in the case where $\epsilon(E_f)$=1 for all neutron energies, we still would not measure the integral contained in eqn 5.1. This is because in our setup we are integrating over the function $(k_f/k_i)S(\vec{q}, E_i - E_f)$ rather than over $S(\vec{q}, E_i - E_f)$. In addition, the upper limit of our energy integration is limited to the neutron giving up all of its energy.

This makes it clear why we do a good job at measuring $S(\vec{q})$ when we use very large incident energies compared to what the neutron transfers. In those instances, the upper limit of our integral would be large enough, and the ratio k_f/k_i is very close to 1 for all scattering events. When this is satisfied (and $\epsilon(E_f)$=1) we find

$$\left(\frac{d\sigma}{d\Omega}\right)_{\text{eff}} = N \frac{\sigma_{\text{coh}}}{4\pi} \int_{-\infty}^{E_i} \frac{k_f}{k_i} S(\vec{q}, E_i - E_f) d(E_i - E_f)$$

$$\approx N \frac{\sigma_{\text{coh}}}{4\pi} \int_{-\infty}^{\infty} S(\vec{q}, E') dE' = N \frac{\sigma_{\text{coh}}}{4\pi} S(\vec{q}).$$

This discussion also implies that the problem does not really exist for samples that scatter almost elastically, but that we can expect problems in liquids and other 'soft' materials. For those 'soft' materials there exists a correction scheme known as Placzek corrections. We refer the reader to the literature for more information on these corrections.

Given these limitations, let's have a look at some of the science that can be done on a 2-axis diffractometer. In fact, let's look at some physics that is done best on a 2-axis diffractometer rather than on a 1-axis diffractometer at a pulsed source.

For scattering events that are almost elastic, we can measure at an exact q-value of our choosing. We simply move our detector to the appropriate scattering angle.

For the cases where we are only interested in this particular q-value, the advantage of measuring multiple *q*-values at the same time is no longer relevant. In those cases, multi-detector instruments do not offer an advantage over a single detector one. While multi-detector instruments at a pulsed source would still allow us to measure this particular q-value in all detectors, we would actually not have as much intensity as on a straightforward 2-axis spectrometer. This is because we would not be using all the other wave lengths that are produced in the spallation process. When this is the case, we actually start to notice that a pulsed source is off most of the time, whereas a reactor source is on all of the time.

Under what circumstances would we be interested in a single q-value, or in a narrow range of q-values? This is the case when we are interested in the evolution of a particular feature upon changing an external parameter such as pressure, temperature, hydration level, stress-causing load, etc. For instance, close to a phase transition where the system changes from a paramagnet to a ferromagnet, we can see the correlations between neighboring atoms build up. Or, we can follow the changes in conformation of polymer molecules on a surface when we add an additional layer of molecules, or raise the temperature. We can follow the folding of a protein as a function of temperature or hydration level by looking for the appearance of a specific inter amino-acid distance that tells us that the folding process has been completed. Diffractometers have even been used to watch cement dry.

In short, whenever we are interested in specific structural changes, we can use a 2-axis diffractometer. Not only will we measure at what temperature or pressure these changes have been completed, but we can also follow the approach to ordering as well as changes in ordered structure. We illustrate this using an example in hard condensed matter physics where the approach to ordering in an unusual quantum magnet was studied using 2-axis diffraction. The details of the physics in this example are not important, rather this example should be viewed as a nice demonstration of the richness of information that is available to the experimenter when doing neutron scattering experiments. Readers who are not quite ready, or not interested, in the physics of some particular system, can skip to the start of the next section.

The system studied was a so-called quantum critical point system. These are metals that harbor magnetic ions. When this metal is cooled down, the magnetic ions try to align themselves with their neighbors. What makes these systems interesting though is that the conduction electrons are getting progressively better at shielding these magnetic moments by aligning their own magnetic moments (originating from the intrinsic spin of the electron) opposite to the magnetic moment of the ions. Thus, we have two effects competing with each other: when the moments are shielded, they cannot align with their neighbors. Conversely, when the moments are aligned, then the conduction electrons can no longer shield them. If this competition stays in place all the way down to zero Kelvin, then we can expect new physics to emerge, since the approach to ordering will be determined by quantum fluctuations rather than by thermal fluctuations.

2-axis diffractometry sheds light on how this competition between ordering and shielding gives rise to a new ground state, at least in some of these quantum critical systems. We show the results of the approach to ordering for one particular temperature in Fig. 5.5. There are a few noteworthy aspects to these results.

First, scattering can be observed at positions (arrows) that do not correspond to the Bragg reflections based on the chemical unit cell. This implies that the distance over which the magnetic structure repeats itself, the magnetic unit cell, is larger than the distance over which the chemical structure repeats itself. The data in this figure were obtained using 2" of graphite as a filter to prevent all higher order contamination from reaching the detector. Such precautions were necessary in this case since partially shielded magnetic moments scatter far fewer neutrons than second or third order nuclear reflections do. This filter will be discussed in more detail in the next section.

Second, the peaks that are visible in this figure are not resolution limited; instead, they are spread out in q-space (reciprocal space). As we intimated in earlier chapters, this implies that the (magnetic) ordering is not long-range. That is, ordering does not span the breadth of the crystal. From the width Γ of the peaks in q-space, we can determine the length L in real space over which the moments are correlated: $L = 1/\Gamma$, with L measured in lattice units and Γ in reciprocal lattice units. The reciprocal lattice is discussed in Chapter 6, here we just need that one reciprocal lattice unit corresponding to transfer of momentum along the crystallographic $\vec{c}$-direction is $2\pi/c$, where c is the length of the c-axis of the nuclear unit cell. Therefore, the figure (Fig. 5.5) tells us that the magnetic moments are lined up with their nearest neighbors, but that the ordering is no longer recognizable by the time we are looking at two magnetic moments separated over a large distance. Therefore, we are looking at the approach to ordering, not at the resulting ordered structure itself.

Third, we can see that the intensity of the peaks diminishes with increasing amounts of momentum transfer. This is an interference effect. In magnetic scattering, the neutron is scattered by the electronic cloud of unpaired electrons around a nucleus, but not by the nucleus itself. When we probe the system at higher and higher momentum transfers, we are using shorter and shorter probing wavelengths. When we reach probing wavelengths comparable to the size of the electronic cloud, then we will start to see that the scattered wave originating from one part of the cloud has traveled a different distance than the scattered wave originating from the opposite size of the cloud, as shown in Fig. 5.6. When this happens, we no longer get perfect constructive interference, and our scattered signal will be weaker. This is a good thing, because we can calculate the size (and shape) of the electronic cloud from the demise of our scattering intensity. In our example, we could conclude that the electronic clouds corresponded to f-orbitals.

Fourth, the data shown are actually the raw data minus the data taken at elevated temperature. This procedure removes any unwanted nuclear scattering, such
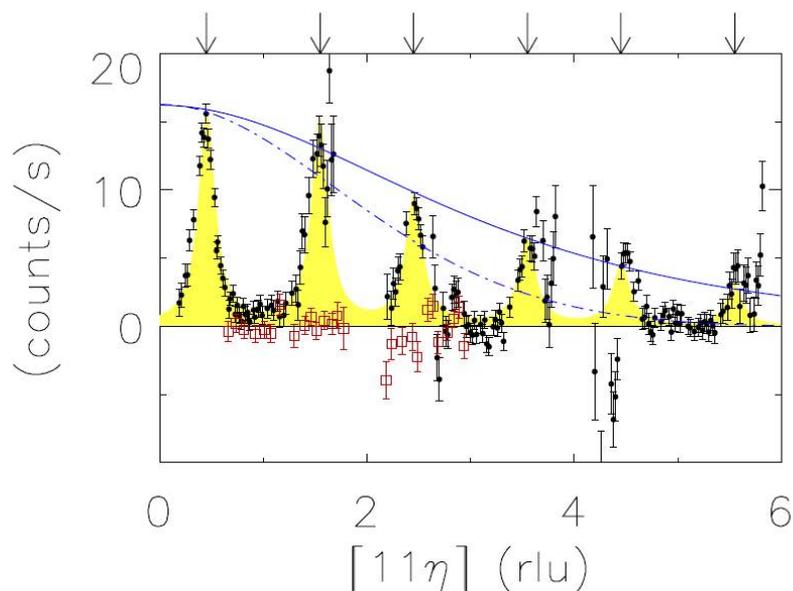
Fig. 5.5 Shown is the magnetic scattered intensity obtained from subtracting data taken at elevated temperature (56 K)- where there are no magnetic correlations present- from data taken at 2 K. η is a measure of how much momentum is transferred to the sample along the c-axis direction of the crystal. (rlu) stands for reciprocal lattice units, and only the amount of momentum transferred along the c-direction was varied in obtaining these data. The black dots and red dots are for momentum transfers along different directions, revealing the direction in which the magnetic moments are pointing. The solid blue line is the expected demise in scattered intensity for scattering originating from f-orbitals, the dashed-dotted line is for d-orbitals. Clearly, the scattering cannot be attributed to some transition metal ions that were also present in the sample. Note that there are some gaps in the data corresponding to strong nuclear reflections such as the (112) and the (114) reflections. Even subtracting data taken at an elevated temperature will not entirely get rid off nuclear peaks. The change in temperature will have affected the lattice parameters sufficiently so that these very intense peaks will have shifted slightly in q-space rendering a straightforward subtraction impossible. Figure reproduced from W. Montfrooij *et al.*, Phys. Rev. B76, 052404 (2007).

as an incoherent signal that could possibly obfuscate the results. Such a background subtraction is not limited to magnetic scattering, one could also imagine a biological system with various levels of deuteration. For instance, subtracting two different deuteration levels (by and large) leaves the signal due to the hydrogen (or deuterium) in the sample.

Fifth, the red points in this figure were measured under identical conditions, but along a different crystallographic direction. Unlike the black data points, no signal is observed. This absence of signal tells us in what direction the magnetic moments
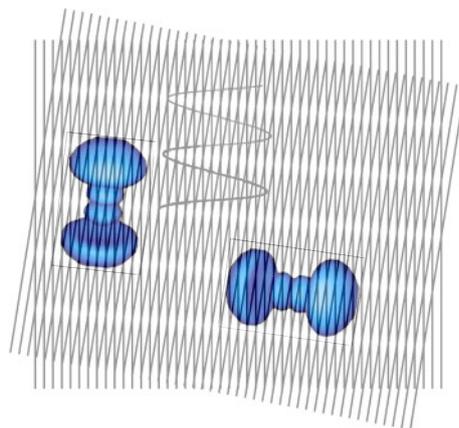
Fig. 5.6 When the scattering originates from the electronic cloud rather than from the much smaller nucleus, then we will see a demise in scattered intensity because not all parts of the electronic cloud (orbital) are probed with the same phase. Studying this demise very carefully as a function of probing wave length and direction of momentum transfer, we should be able to not only distinguish the overall shape and size of the orbital, we should also be able to determine whether the occupied orbitals have any preferred orientations. This latter possibility is known as orbital ordering.

are pointing. This is because unlike the case of nuclear scattering, the magnetic cross-section depends on the relative angle between the direction of momentum transfer and the orientation of the magnetic moment the neutron is being scattered by.

While Fig. 5.5 told us a few things about the magnetic moments in this system, it did not tell us anything ground breaking about this type of materials, the quantum critical magnets. The new knowledge came when it was measured how many magnetic ions had lined up with their neighbors along various crystallographic directions. On very general grounds it was expected that this number would be different along every direction as the distances between the moments was not the same along all directions, and therefore, the magnetic moments should align more strongly with each other along directions where the moments were more closely spaced. This is not at all what was observed, however, as shown in Fig. 5.7.

This figure (Fig. 5.7) shows the evolution (with cooling down) of the scattered signal, reflecting that more and more moments are lining up with their neighbors, thereby giving rise to constructive interference. Note that magnetic moments pointing in random directions would not give rise to scattered waves that would add up constructively, because the phases of the scattered waves would be random. We would only observe scattered waves originating from individual moments under those circumstances, similar to why we would get incoherent nuclear scattering when we have a collection of different isotopes, or a collection of identical isotopes with a nuclear spin.
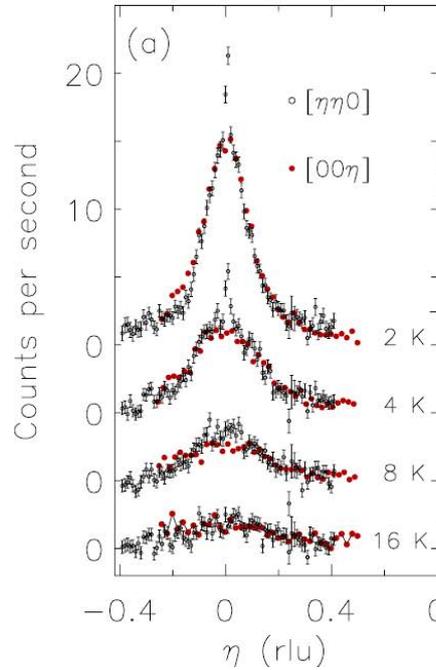
Fig. 5.7 The temperature dependence of the magnetic scattering of the first peak shown in Fig. 5.5. One observes that the scattering starts to emerge at 16 K, and increases in intensity while narrowing in width in q-space when we cool down more. Since the inverse of this width is directly proportional to the distance over which moments are correlated, we see that (for all temperatures) there are just as many moments lined up along the c-direction (red symbols, momentum transfer along the (00η)-direction) as there are along another high symmetry direction (black symbols, momentum transfer along the $(\eta\eta0)$-direction. Figure reproduced from W. Montfrooij *et al.*, Phys. Rev. B76, 052404 (2007).

In our example, since we actually see peaks in the scattered intensity, we know that moments must be lining up with each other in some way.

We see the peaks first emerging at 16 K, implying that above this temperature the thermal energy is still high enough for the moments to lose their orientation with respect to their neighbors. But once we cool down more, we see that the thermal energy available for destroying order becomes less than the energy gain associated with lining up two atomic magnets, and we observe more and more moments lining up with their neighbors. This is reflected both in the intensity of the peaks (reflecting the number of moments partaking in the constructive interference process) and the characteristic width Γ of the peaks: the smaller the width, the larger the distance L in real space over which the moments are correlated.

The most remarkable feature that can be seen in Fig. 5.7 is that the width of the

peaks is identical along different crystallographic directions. This can be observed at any of the temperatures that the data were taken. This implies that independent of what direction one looks, one will always find that (on average) the same number of moments have lined up. As mentioned, this is unexpected given the difference in interatomic distances between the magnetic ions in this (tetragonal) structure. After some thinking, which is the most fun part when confronted with very accurate and very puzzling neutron scattering data, it was concluded that increasingly larger magnetic clusters must be forming in this material upon cooling.

Clusters nicely explain the direction independence of the measured signal, provided the cluster size is determined by a random process. This makes perfect sense for a system where there is a competition between moments aligning, and moments being shielded. Cool such a system down, and random moments will be shielded. Cool down more, and more moments will be shielded. This is known as a percolation problem. Removing more and more moments will result in isolated clusters of surviving moments. The moments within such clusters then line up with their neighbors because finite size effects (a consequence of quantum mechanics) dictate them to do so. These ordered clusters will then scatter neutrons, revealing how many moments there are within the clusters. Then when the temperature is lowered further, larger and larger clusters will become isolated, yielding a stronger scattered signal characterized by a smaller width in q-space.

Such microscopic detail of what happens in a system that is trying to satisfy competing interest helps explain other unusual features of these types of systems that have to do with resistivity and specific heat. What is really interesting for our purposes though, is that all this was deduced from a straightforward diffraction experiment. All that was needed for the success of this experiment was a careful setup (employing graphite filters) and a plan of how to get rid off an unwanted background caused by (incoherent) nuclear scattering events that could easily have overwhelmed the magnetic signal.

We will show some more examples of what is possible by doing diffraction in the last section of this chapter. Next, however, we will look at some of the components that are used on 2-axis spectrometers, such as the aforementioned graphite filters.

5.2 Components of a 2-Axis Spectrometer

The components of a 2-axis spectrometer can also be found on a triple-axis spectrometer, which will be discussed in detail in Chapter 9. Since we do not change the incident or final energy of the neutron, and since we do not measure as a function of energy transfer, the choice of what option to use for which components is not nearly as important for a 2-axis spectrometer as it is for a 3-axis one. We therefore refer the reader to Chapter 9 for most details, in this section we merely discuss those components that determine the difference between a successful and a failed experiment.

For a successful experiment, we must ensure that our interpretation of the data is not plagued by inelastic effects as explained in the introduction. We can minimize inelastic effects by choosing neutrons with high incoming energy, or equivalently, neutrons that have short wavelengths. If we cannot get around the hurdle that the inelasticity problem represents, then we should do our experiment using a spectrometer that measures energy transfers as well as momentum transfers.

For a 2-axis spectrometer, we can select neutrons with a sufficiently short wave length using a monochromator crystal. Such a crystal is used to Bragg reflect neutrons with wavelength λ out of the beam of neutrons that are traveling down the beam tube. Thus, we are using Bragg's law:

$$\lambda = 2d \sin \theta_m. \tag{5.2}$$

In here, $2\theta_m$ is the angle between the neutron beam and the beam of redirected neutrons of wavelength λ. We can think of θ_m as the angle under which the 'mirror' (shown in Fig. 2.4) is positioned. d is the distance between the planes of atoms that make up the monochromator crystal. As we can see right away, a monochromator with a small lattice spacing (such as copper) will select neutrons with a shorter wavelength than when we use a monochromator with a larger lattice spacing, such as graphite.

Some diffractometers allow for the monochromator angle θ_m to be varied so that the wavelength can be varied; in other instances, the monochromator 'take-off' angle is fixed. While the latter might seem unnecessarily restrictive, it does come with the advantage that there are fewer moving parts, and therefore, one can do a better job at shielding the instrument from unwanted background counts. Some diffractometers do not allow the user to either change the monochromator crystal for another one, or vary the take-off angle. It is up to the user to figure out if the diffractometer can nonetheless be used for one's experiment. Note that when using a triple-axis spectrometer, one can vary both d (by dropping a different monochromator into the beam or by rotating the existing monochromator) and θ_m.

For wavelengths λ that satisfy Bragg's law, wave lengths $\lambda/2$, $\lambda/3$ etc. will also satisfy our constructive interference criterium. Therefore, neutrons with shorter wavelengths will also be reflected out of the beam and directed at the sample. We call these unwanted neutrons higher-order contamination, and they must be taken care of. If we allow them to be detected, then these scattering events would correspond to double the amount of momentum transfer that we are interested in, and we mess up our data.

Higher-order contamination is particularly problematic when we are studying antiferromagnetic ordering. Higher order contamination gives rise to nuclear peaks being detected at the same scattering angles as antiferromagnetic peaks, rendering it neigh impossible to study (the onset of) antiferromagnetic ordering. For cases like this, we use filters to remove the unwanted higher order neutrons, we use monochromator crystals that do not allow for second order contamination to be reflected, or we use a

velocity selector.

There are monochromator crystals that do not allow for second order reflection because of their diamond-like crystal structure. Of this family of monochromotors, silicon is most widely used. In particular, the (111), (311) and (511) reflections are popular choices. Note, however, that third order reflections are still allowed; these third order reflections come in at 9 times the neutron energy of the uncontaminated beam. Unless energies this high are cut out by other filters (such as a sapphire filter in the beam tube), we would still have a contamination problem.

Filters can be very effective at getting rid off second and third order contamination. The filter most commonly used is pyrolytic graphite. This filter can remove $\lambda/2$ and $\lambda/3$ contamination from the beam, *provided* λ is suitably chosen such as λ= 2.4 Å. Chapter 9 has a graph in it that can help the user in selecting the appropriate wavelength. Typically, one uses 1 or 2 inch thick graphite filters when studying magnetic excitations that would otherwise be contaminated by nuclear scattering. There exist other filters similar to graphite, but we will not detail those here.

Another good filter option that operates in a slightly different way is employing Be, or BeO. These materials only allow neutrons with wave lengths longer than 4 Å(Be) or 3.7 Å(BeO) to pass through. Shorter wave lengths are scattered out of the beam. This is very useful for doing diffraction experiments that utilize long wave lengths, such as experiments on biological materials or other systems where one is either interested in long lattice spacings or in large structures. The downside of these filters is that they have to be cooled down to liquid nitrogen temperatures in order to be effective in transmitting the sought-after neutrons.

Yet another option for getting rid off higher order contamination is placing a velocity selector in the beam. A velocity selector, such as the one shown in Fig. 5.8, is essentially a set of blades- covered in neutron absorbing material such as Gd- that rotate around an axis at just the right speed to allow neutrons of the desired wavelength to pass through the channel between the blades. Neutrons that go much faster or slower will end up being absorbed by the blades. Such velocity selectors are frequently found on SANS instruments where they are also used to select the incoming wave length without the aid of a monochromator, they are found at the end of neutron guide tubes, and sometimes they are used on triple axis spectrometers such as on Dualspec at CINS (Chalk River, Ontario).

One final set of components that merits our attention are collimators and masks. Masks are moveable, or adjustable holes in a neutron absorbing material. These masks are very good at cutting down on unwanted background that might be caused by the neutron beam scattering off of the cryostat, humidity chamber, or furnace. The function of the mask is to ensure that only the sample is exposed to the neutron beam. Masks are always good to use, we do not lose any useful neutrons, but we are able to cut down greatly on the unwanted neutrons. Masks are normally placed on either side

Fig. 5.8 Shown is a velocity selector as used at the small angle neutron scattering (SANS) spectrometer Quokka at ANSTO in Australia (www.ansto.gov.au). When spinning, only neutrons of a certain velocity will be able to make it past the curved blades.

of the sample; that is, in between the monochromator and the sample, and in between the sample and the detector.

The function of collimators is to cut down on the angular divergence of the incoming and scattered beam. The collimators most frequently used are Soller collimators: these consists of parallel blades that are made of Cd, or some other material painted with Gd. Only neutrons that travel parallel, or close to the direction of the blades can make it through (Fig. 5.9). The separation between the blades dictates the angular divergence that the emerging beam can have upon exit. Since angle is directly related to wave length and momentum transfer, collimators essentially determine the resolution in q-space. That is, they determine what range of q-values will be detected around the chosen q-value. Collimators have blades that are spaced apart in such a way that the divergence of the neutron beam will be, typically, 80', 60', 40', 20' or 10' upon exiting. Clearly, the larger the divergence, the more neutrons will make it through, giving us more counts. The choice of collimator therefore deserves careful consideration and it should be based upon the desired q-resolution.

We will give a rough guide on how to calculate the q-resolution. This is only a rough guide, free software exists to help the user with this task. The q-resolution can easily be visualized for elastic scattering on a 2-axis diffractometer. Starting with the equation for momentum transfer $\hbar q$ when neutrons of incident wave number k_i are scattered over an angle θ by the sample,

$$q = 2k_i \sin(\theta/2)$$

we find the spread Δq in q-values through differentiation:

$$\Delta q = 2 \sin(\theta/2)\Delta k_i + k_i \cos(\theta/2)\Delta\theta.$$

Here, $\Delta\theta$ is the acceptance angle of our collimation. Dividing these two equations leads directly to

Fig. 5.9 Soller collimators, as shown on the left (80', commercially available), are straightforward devices. The divergence of the emerging neutron beam is directly related to the spacing between the blades. Radial collimators such as the one made by the Missouri Physics Department machine shop (right panel, used on CNCS at the SNS) serve the same function, but they are designed for instruments that have multiple detectors, or position sensitive detectors. The sample will be placed at the center of curvature of the blade assembly. For some instruments the collimator is made to oscillate so that no single part of the detector(s) will be permanently in the direct shadow of the blade. For example, the powder diffractometer at MURR is equipped with an oscillating radial collimator.

$$\frac{\Delta q}{q} = \frac{\Delta k_i}{k_i} + \frac{\Delta\theta}{2\tan(\theta/2)}.$$

Normally, the two terms on the right hand side are independent of each other, so that we should add them quadratically:

$$\left(\frac{\Delta q}{q}\right)^2 = \left(\frac{\Delta k_i}{k_i}\right)^2 + \left(\frac{\Delta\theta}{2\tan(\theta/2)}\right)^2. \tag{5.3}$$

This equation tells us that if we are using collimators to reduce the second term on the right hand side, then we do not need to worry about collimators for back scattering ($\theta/2 = 90^o$). However, this equation also tells us that collimators will be important for scattering at small or moderate angles. Alternatively, if you have a choice, position your detector at large scattering angles; this way, you will not have to use very tight collimators and you will count more neutrons.

The collimation also (partially) determines the spread in k_i. In fact, the spread in k_i is determined both by the intrinsic mosaicity $\Delta\theta_{\text{crystal}}$ of the monochromator crystal, and by the collimators acceptance angle $\Delta\theta$. This follows directly from Bragg's law:

$$k_i = \frac{\pi}{d\sin(\theta_m)}; \frac{\Delta k_i}{k_i} = \frac{\sqrt{(\Delta\theta_{\text{crystal}})^2 + (\Delta\theta)^2}}{2\tan(\theta_m)}. \tag{5.4}$$

As long as one ensures that the angles are entered in radians, then these calculations are not too terribly hard to perform. Should one also use a collimator between the

reactor source and the monochromator (and one will), then a good approximation for finding the effective acceptance angle for the pair of collimators is obtained by reciprocally adding the two acceptance angles:

$$\left(\frac{1}{\Delta\theta_{\text{effective}}}\right)^2 = \left(\frac{1}{\Delta\theta_{\text{before}}}\right)^2 + \left(\frac{1}{\Delta\theta_{\text{after}}}\right)^2.$$

So far we have treated q-resolution as if it is one number, however, momentum transfer is a vector, and therefore, the q-resolution will also have three components. The full three-dimensional q-resolution is referred to as the resolution ellipse, and this ellipse can be visualized using software such as RESLIB. Without using the software, it helps to look at some actual data to develop some sense of this resolution ellipse.

We can think of the q-resolution as being separated in the resolution perpendicular to the scattering plane (one component), and within the scattering plane (two components). The resolution perpendicular to the scattering plane is not very good on a 2-axis diffractometer. We tend not to use the words 'not very good', instead we call it relaxed. However, this relaxed resolution does not affect the interpretation of our data (in most cases), since our data are measured as a function of momentum transfer within the scattering plane, not perpendicular to it. The only instance we have to worry about the vertical resolution is when we are studying the onset of order *and* we need to know the absolute amount of scattering that is taking place so that we can deduce the size of a magnetic moment that is subject to ordering. This does not happen very often, so normally we do not have any vertical collimation in place.

The in-plane resolution is always important. In Fig. 5.10 we show the measured q-resolution along the direction parallel to the direction of momentum transfer (longitudinal), and perpendicular to it (transverse). These data were taken on a Bragg peak of a single crystal. Nominally, Bragg reflection only happens when both the wavelength of the incoming neutron is correctly chosen, and when the detector is just at the right scattering angle. This should result in a very sharp delta-peak of the scattering: either there is a lot of scattering, or none at all. However, in practice, this all-or-nothing type of scattering is smeared out over a range of scattering angles and incident wavelengths, giving us our resolution ellipse.

Looking at these actual data, we can read off from the figures by how much this delta-function has been smeared out in the longitudinal and in the transverse direction. We find 0.034Å^{-1} and 0.023Å^{-1} for the longitudinal and transverse resolution, respectively. We can compare this to the overall size of the resolution ellipse given by eqn 5.3. For this particular experiment, 40' collimators were used on either side of the sample and on either side of the monochromator crystal, the incident neutron wave length was $\lambda=2.4$ Å ($k_i=2.6\text{Å}^{-1}$), and the amount of momentum transferred for the (002) reflection was 2 x $2\pi/c = 4\pi/10.5$ Å$= 1.2\text{Å}^{-1}$. The average scattering angle θ is given by $\theta=2$ x $\sin^{-1}(2.4$ Å$/(2$ x $10.5/2)$ Å$) = 26.4^o$. The final piece of information we need is that the PG002 reflection (with lattice spacing $d=3.35\text{Å}$) was used
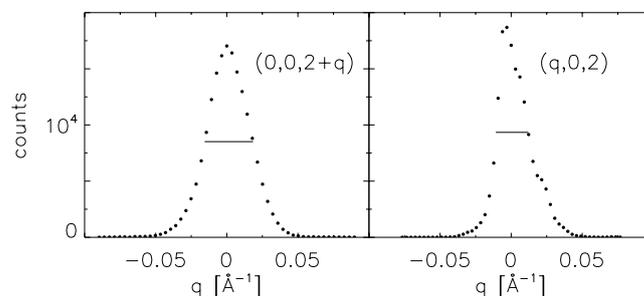
Fig. 5.10 The longitudinal (left panel) and transverse (right panel) resolution function for the (002) reflection of a body centered system. The characteristic width (horizontal line) of the scattering shown in the left panel gives us our longitudinal resolution, the right panel corresponds to the transverse resolution. Note that the two resolution widths are not the same. Also note the shoulder in the scattering in the right panel. This is caused by the single crystal actually having some grains in it that are partially misaligned. This is what is called the intrinsic mosaicity of a crystal. Transverse scans are a good way of probing this mosaicity.

to select the wave length of the incident neutrons. From this we calculate that $\theta_m = 21^o$.

Let's plug in these numbers into eqn 5.4 and eqn 5.3. For the effective acceptance angle $\Delta\theta$ that our monochromator crystal 'sees' we use $40'/\sqrt{2}$, and we take 30' for the intrinsic mosaicity of our PG monochromator crystal. With these numbers we find that $\Delta k_i/k_i = 0.022$. Putting these numbers into eqn 5.3, and using that the acceptance angle that our sample sees ($\Delta\theta$) is $40'/\sqrt{2}$, since all collimators that were used were 40' collimators, we find that $\Delta q/q = 0.028$. Since $q = 1.2\text{Å}^{-1}$ we find that $\Delta q = 0.034$ Å^{-1}. This number appears to be a reasonable measure for the resolution that is shown in Fig. 5.10, at least for the purpose of doing quick calculations. Had we used collimators that were 20', then our resolution would have been twice as good (roughly). Not quite twice as good since the intrinsic mosaicity of both the monochromator crystal and our sample crystal would have played a more prominent role. In practice, one chooses similar collimators at all positions except between the sample and the detector where one uses a considerably more relaxed collimation.

Lastly, a very effective way of reducing Δk_i is to reduce k_i itself. If we pick neutrons that only go half as fast, then our Δk_i will be reduced, and provided we can still get to our desired q-value, all scattering angles will increase with the accompanying increase in the tangents that enter eqns 5.3 and 5.4.

5.3 Diffractometers With Area Sensitive Detectors

We mentioned that diffractometers with multiple detectors can collect data much faster when the problem is suited to benefit from measuring multiple q-values at the same time. We do not actually need separate detectors, we can also use position or area sensitive detectors. Position sensitive detectors are long detector tubes, while area
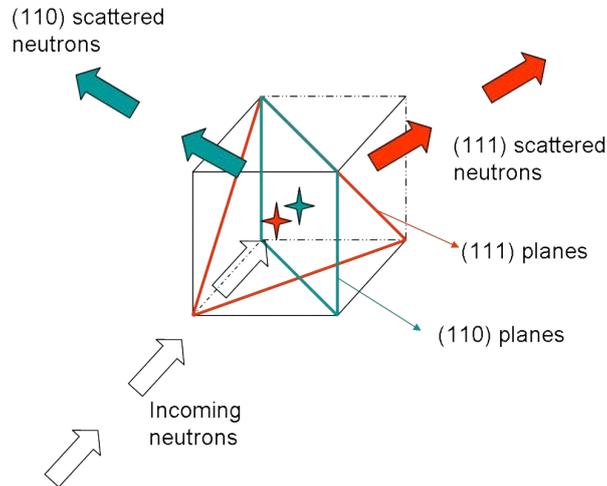
Fig. 5.11 When a crystal is oriented with a high symmetry direction parallel to the incoming beam, then the various sets of lattice planes will select a neutron wave length out of a white beam and diffract it at the corresponding Bragg angle. This figure shows neutrons coming in (open arrows) parallel to the a-axis, and being diffracted by the (110) planes (green and green arrows), and by the (111) planes (red, red arrows).

sensitive detectors are square or disk-shaped boxes. Such detectors not only detect that a neutron hit the detector, but also where it hit the detector to within an accuracy of a few millimeters. Position sensitive detectors tell us where along the tube the neutron arrives; area detectors tell us both the horizontal and vertical position of where the neutron entered the detector. So, from a data analysis point of view, such detectors work as if we have a continuous assembly of smaller detectors, each a few millimeters wide.

In our chapter on powder diffraction (Chapter 6) we we look at a position sensitive detector. In this section, we will look at some special diffractometers that employ area detectors. Area detectors are typically found on SANS instruments (Chapter 7), stress diffractometers, and Laue cameras. Position sensitive detectors tend to be used on powder diffractometers. We first describe the Laue camera that is used to determine the structure and atomic positions of materials that have many atoms in a unit cell, such as proteins.

When we stick a crystal in a neutron or xray beam, oriented in such a way that the direction of incident radiation coincides with a direction of symmetry of the crystal, then we can expect Bragg reflection to occur provided the incident wavelength is just right. This can be seen as follows. For Bragg reflection to occur off of a specific set of crystal planes separated by a distance d oriented at an angle $\theta/2$ with respect to the incoming beam (See Fig. 5.11), we must have that the required incident wavelength
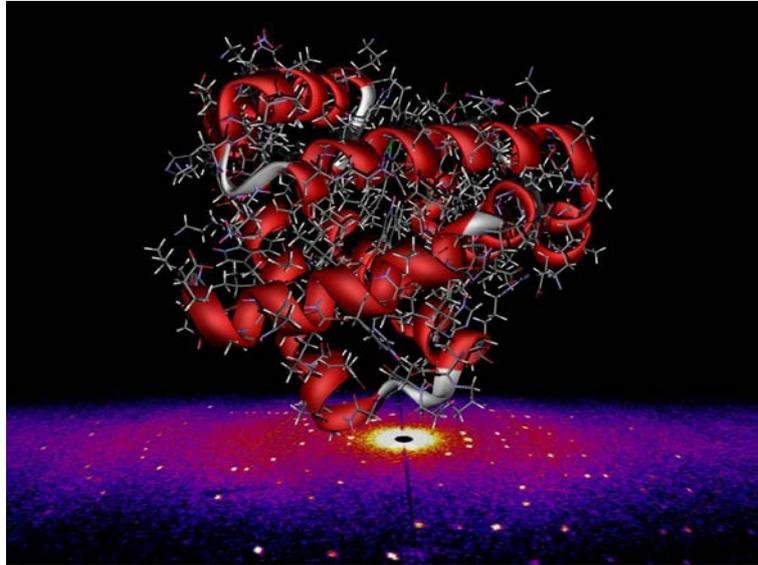
Fig. 5.12 Source: www.nmi3.eu. The structure of myoglobin is becoming better characterized. Thanks to xray and neutron diffraction experiments performed by researchers at München (Andreas Östermann) and at the Japan Atomic Energy Research Institute (Ryota Kuroki), even the positions of many hydrogen atoms are now known. The top shows the spatial structure of the myoglobin protein; the bottom the neutron Laue picture from which the positions of 157 bound hydrogen D_2O molecules could be determined.

is given by $\lambda = 2d \sin(\theta/2)$. When this is satisfied, then the neutron will be scattered into a particular direction as if the lattice planes form a mirror.

Since every neutron that is scattered will be scattered under an angle that depends on its wave length, it actually makes sense to use a whole bunch of wave lengths at the same time. This way, we will measure multiple Bragg reflections at the same time, but since all Bragg reflections will correspond to different scattering angles, we can still tell them apart by doing some thinking. The statement 'some thinking' applies to the fact that the scattering angle only yields the ratio $\lambda/2d$, however, we will be able to measure so many Bragg reflections that we can deduce the size of the unit cell of the crystal, and from there, determine all possible lattice spacings d. It should make sense from this description that an area sensitive detector is required for a Laue camera.

Laue cameras are used for electron diffraction, xray diffraction, and neutron diffraction. We show an example of a neutron Laue pattern in Fig. 5.12. For this particular example, proteins have been grown into a single crystal protein structure. In this case, we are not actually interested in the lattice that proteins form as this can hardly be of biological relevance, rather we are interested in the intensity of the spots. The intensity of the spots is directly related to the structure factor of that particular reflection as
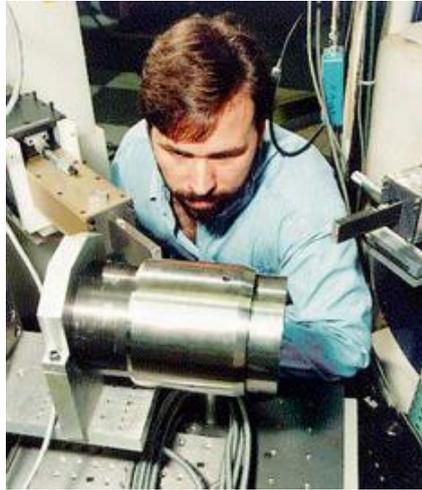
Fig. 5.13 Ron Rogge of the Canadian Institute for Neutron Scattering CINS is setting up an experiment to determine the thermal evolution of stress in a 3-component (Inconel, Zr-2.5%Nb, & Stainless Steel) rolled joint. Note the two apertures that are used to define the scattering volume (the long rectangular shapes, one of which is close to Ron's left shoulder), and the 90^o detector angle. For this instrument (L3 at the NRU reactor), the spatial resolution (defined by a volume element) ranges from 0.2 to 1,000 mm^3. The resolution in lattice strain is of the order 0.01%. Source: www.cins.ca.

given by eqn 4.6. Provided we collect enough reflections, we should be able to determine the positions of all atoms within a protein, as well as their type. The protein structure shown in Fig. 5.12 is based on a combination of xray and neutron Laue results. This modern use of a Laue camera is different from its original use where the spots were actually made visible on a photographic plate (hence the name camera) and the crystal symmetry could be deduced from that; however, (at the time) the intensity of the spots was not known to sufficient accuracy to determine which atom was sitting where.

Another important type of diffractometer is the residual stress diffractometer. Such a spectrometer measures the strain in a material when it has been subject to stress(es). Strain manifests itself as a change in lattice parameter d at a local level. Some areas of a material might show compression, others might show expansion. The way strain ($\Delta d/d$) is actually measured depends on whether the instrument is at a reactor source, or at a spallation source. In both cases, Bragg's diffraction law and the differentiation thereof forms the basis of this type of measurement:

$$\lambda = 2d\sin(\theta/2); \frac{\Delta\lambda}{\lambda} = \frac{\Delta d}{d} + \frac{\Delta\theta}{2\tan(\theta/2)}. \tag{5.5}$$

For a reactor source, we have that $\Delta\lambda$ equals zero, so that strain is measured by the shift $\Delta\theta$ of the Bragg peak. For a pulsed source diffractometer, we have that the detector position is fixed ($\Delta\theta$=0), so that strain will be proportional to the change in wave

Fig. 5.14 Tom Holden (second from left) and his team get ready to get out of the way so that they can measure strains in this turbine when it is actually spinning. This spectrometer (L3 at CINS) has since been equipped with a rig capable of applying biaxial stress. Source: www.cins.ca.

length, which in turn will be proportional to the change in time of flight of the neutron.

It is possible to measure the strain in a small sub-volume of a large piece of material, in situ. The trick here is to narrow down the incoming beam of neutrons, and to put an equally restrictive aperture in the scattered beam. This way, the volume of the material that the detected neutrons have scattered from (called the gage volume) is given by the cross-section of these two beams. An idea of what these apertures look like can be seen by scrutinizing Fig. 5.13. For most stress measurements, the detector is placed at a 90^o angle so that the volume being examined has an easy-to-analyze geometric shape.

Stress diffractometers have been employed widely to study metal fatigue and how metals break. Studies have been carried out on the wheels of trains, pieces from the space shuttle, and even a Pratt and Whitney Gas turbine while operational (Fig. 5.14). Also, many engineered materials harbor smaller grains. These grains can exhibit preferred orientations, which can enhance some of the properties of the materials. With neutrons, these preferred orientations can be mapped out and spatial variations of this
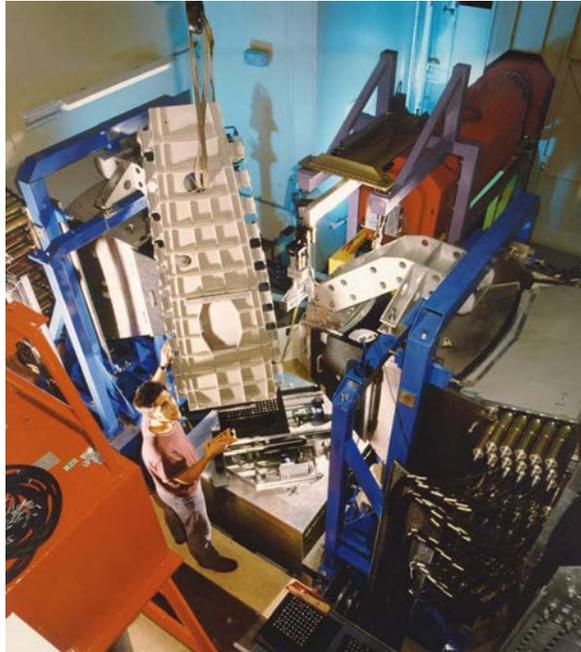
Fig. 5.15 The ENGIN-X instrument while it is measuring the residual stress within friction stir welds on an Airbus prototype wing rib. This image, and the description, was taken from the ISIS website. Source: www.isis.stfc.ac.uk/instruments/gem/

texture can be evaluated in most materials commonly used in industry, such as steel, aluminum, titanium, zirconium, nickel-based alloys, ceramics and composites.

Multiple detectors, or position sensitive detectors can be helpful in the following way. Since it is not clear whether a Bragg peak will have shifted to a smaller or larger scattering angle, or even have shifted perpendicular to the scattering plane, it makes sense to simply have a position sensitive detector cover a bunch of scattering angles. Also, it is a big time saver, since one does not have to do the full scan over the diffraction peak (like the one shown in Fig. 5.10) in order to get the width and intensity of it. The width is an important parameter when characterizing individual grains, while the intensity is used to map out any texture in the sample.

One final instrument to mention in this section is ENGIN-X. This is the stress diffractometer at the ISIS pulsed source. Its operational principle is based on a white beam of neutrons coming in, and measuring the diffraction pattern in a detector at 90^o. To make this diffractometer very efficient, multiple detectors have been placed roughly at 90^o spread over two banks, as shown in Fig. 5.15. This instrument has been designed with accessibility in mind, and it can be used to study very large structures.

As mentioned, the most widely used diffractometer is the Small Angle Neutron Scattering (SANS) diffractometer. This instrument will be discussed in detail in Chapter 7.

5.4 Exercises

Exercise 5.1

A multi-detector spectrometer such as the GEM spectrometer shown in Fig. 5.4 can collect data on liquids and amorphous materials very fast, but it is not that useful for collecting data on single crystals. Explain why.

Exercise 5.2

When one is interested in following the phase transition of a magnetic system that orders antiferromagnetically at low temperatures, and when the magnetic unit cell is very large, what is the best experimental setup in order to obtain accurate data, free from higher-order contamination?

Exercise 5.3

A piece of aluminum has been subject to stress, and it now has a strain of 0.1 %. This means that $\Delta d/d = 0.001$. What minimal angular resolution must a reactor source based diffractometer have for us to be able to observe a peak shift in the (111)-Bragg reflections of aluminum? ($d_{(111)} = 2.33$ Å)
If we use neutrons of incident wave length of 2.4 Å, and assuming that our monochromator is a perfect crystal, discuss what collimators (acceptance angles) you would be using for this experiment.

6
Powder Diffractometers

The term powder diffraction refers to measuring the scattering pattern from a powdered sample, as opposed to obtaining a diffraction pattern off of a single crystal. Powder diffractometers are dedicated instruments optimized for collecting such diffraction patterns; the term optimized refers to the ability to collect the pattern as fast as possible, or with as high a q-resolution as possible. Powder diffractometers are found both at reactor sources and at spallation sources. They are the workhorse instruments when it comes to determining the structure of materials.

A powder is a collection of very small crystallites, oriented randomly. Neutrons can be Bragg relected from individual crystallites, provided these crystallites are oriented in such a way as to meet the Bragg reflection condition (eqn 2.1 for neutrons of wavelength λ and lattice planes spaced a distance d apart. As long as there are enough crystallites, and as long as they cover all orientations equally, we can use the reflection pattern to determine the structure of the material.

The advantage over doing single crystal diffraction is that the data collection is much quicker, and we do not require our material to be a single crystal. The latter is an obvious advantage as it is very difficult, sometimes even impossible, to obtain single crystal samples. The disadvantage of powder diffraction over single crystal diffraction, namely that reflection peaks originating from different crystallographic planes can overlap in scattering angle, was essentially resolved in the late sixties by Hugo Rietveld. Rietveld came up with a method for fitting the scattered intensity that virtually separates overlapping peaks. This method is now known as the Rietveld refinement method, and we will say more about this method in section 6.4.

Any self-respecting powder diffractometer has multiple detectors, or position sensitive detectors. The advantage of this is clear: we wish to measure the scattered intensity at all momentum transfers, so we wish to cover the entire range of scattering angles between $0 < \theta < 180^o$. What better way to do this than to have multiple detectors cover as much of this range as possible, or to utilize position sensitive detectors? We show a schematic of a powder diffractometer at a reactor source in Fig. 6.2. We already showed the GEM diffractometer at ISIS spallation source in the previous chapter.

Not all powder diffractometers are the same. The user may have an option to select the incident neutron wave length (such as is the case for D20 shown in Fig. 6.2), but frequently the incident wave length is fixed. Typical incident wave lengths are in the

Fig. 6.1 Hugo Rietveld, shown here on the beamport floor of the Petten research reactor in the Netherlands, developed an algorithm for analyzing powder patterns that put powder diffraction on par with single crystal diffraction.

range of 1-1.5 Å, allowing us to get up to momentum transfers of around 10 Å^{-1}. These are just ball park numbers, but there is no reason why we would want to go out to much higher momentum transfers as the scattering at these high q-values mostly yields information about individual atoms, not about the relative positions of atoms. We discussed this effect in Chapter 4, and we will return to it in the section when we discuss the Debye-Waller factor. Sometimes there is a need to actually go to lower momentum transfers when one is interested in magnetic structures that have large unit cells. Such spectrometers will be found at cold sources.

An important characteristic of a powder diffractometer is its q-resolution. This tells us how easy it is to separate two peaks that occur at almost the same scattering angle. The q-resolution is closely related to the angular resolution, but it also depends on the scattering angle itself. Collimators help define the angular resolution, but the type and shape of monochromator crystal, as well as the sample size are also important elements in determining the overall resolution. Later on in this chapter we will describe the powder diffractometer at MURR as an example of what can be achieved in terms of optimizing the resolution without sacrificing too much in intensity.

We will first discuss some of the basics of powder diffractometers and show some examples of what has been measured. This is followed by a look into the components of these instruments in more detail, and into how to analyze the diffraction pattern so that we glean the most about our samples.
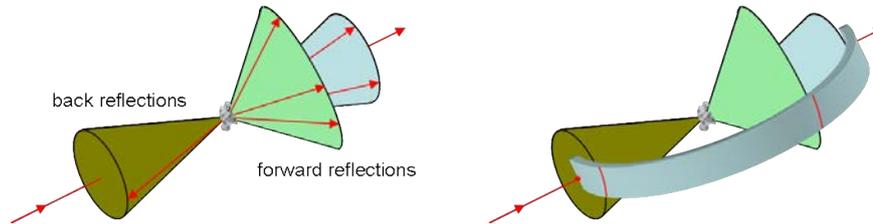
Fig. 6.2 Schematic of a powder diffractometer at a reactor source. The one shown here goes by the name of D20, and it is located at the Institute Laue-Langevin (ILL) in France. Source: www.ill.eu.

6.1   The Skinny

When neutrons are Bragg reflected from a single crystal, we observe that the neutrons are being scattered into one particular direction or- when we take the instrumental resolution into account- into a particular solid angle $d\Omega$. When neutrons are being Bragg reflected by a powder, we observe that the scattered intensity is spread out over a cone, called the Debye-Scherrer cone. The apex angle of this cone is double the scattering angle θ of the neutron. We show this in Fig. 6.3. We can visualize the scattering as originating from a great many individual crystallites whose crystal planes all meet the reflection condition. This is sketched in Fig. 6.4.

In a scattering experiment, we move a detector of fixed height over a range of scattering angles (or we simply cover the entire range of angles with detectors so that we do not have to move them), thereby intercepting part of the cone (Fig. 6.3). The scattered intensity we measure is proportional to what fraction of the cone is being intercepted by the detector, and to the strength of the reflection σ_{cone}. We wish to determine the latter from our data as this number is a direct measure of the relative

Fig. 6.3 Scattering for a particular reflection occurs at a Bragg angle $\theta_B = \theta/2$, which shows up as scattered intensity spread out over a cone with apex angle 2θ (left panel). A detector of fixed height, shown by the blue band in the panel on the right, will intercept a fraction of this cone. This fraction depends on the scattering angle.

positions of the atoms in the crystal, and of what type of atoms they are:

$$\sigma_{\text{cone}} \sim \langle \sum_{i,j=1}^{N} b_i b_j e^{i\vec{q}\cdot(\vec{R}_i - \vec{R}_j)} \rangle. \tag{6.1}$$

Note that we do not actually measure the dependence of σ on vector $\vec{q}$, we only measure its dependence on the length q of $\vec{q}$ given by $q = 4\pi/\lambda \sin(\theta/2)$. We rely on the Rietveld refinement algorithm to provide us with a model of the atomic structure, from which the average over all crystallite orientations as contained within the brackets $< ... >$ of eqn 6.1 can be calculated and compared to the measured intensity.

As can be seen from Fig. 6.3, the amount of scattering that is intercepted by the detector of fixed height depends on the scattering angle. If the detector of height h
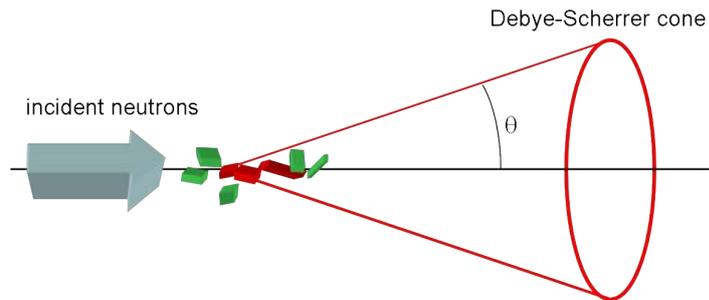
Fig. 6.4 When crystallites are in the right orientation (the red ones), they will scatter the neutrons. There are 2π orientations for a crystallite that will meet the reflection condition without changing the angle of incidence of the neutron beam with respect to the lattice planes that are Bragg reflecting the neutrons. When we view the problem in spherical coordinates, the inclination will be fixed (the inclination is the angle with the zenith direction, and we pick the direction of the incoming neutrons as the zenith direction), but that still leaves the azimuth angle unspecified, giving us 2π degrees of rotation.

is a distance R from the sample, then the fraction of scattering intercepted is given by $h/(2\pi R \sin\theta)$. Thus the observed signal in our detector will be proportional to $\sigma_{\mathrm{cone}}(\theta)/\sin\theta$. This angular dependence is in fact part of the full Lorentz factor, which is given by the denominator of the following expression:

$$\text{counts} \sim \frac{\sigma_{\mathrm{cone}}}{\sin\theta \sin(\theta/2)}. \tag{6.2}$$

The origin is the term $\sin(\theta/2)$ is not as easy to visualize as the $\sin\theta$ term, since it is a combination of three effects: how much of the sample is in the intersection of the incident and scattered beam, over what range of angles will the scattering occur (intrinsic resolution effect), and how many orientations are possible for a crystallite in Bragg reflection. The third term yields the $\sin(\theta/2)$ term, the first and second term cancel each other out. The interested reader can look this up a standard text on neutron or xray diffraction, or just have faith in the wisdom of our colleagues and take this term for granted. The net effect of the Lorentz factor is to enhance the scattering at low (near zero) and high (near 180^o) angles compared to the scattering at 90^o. And the good news is that this factor does not complicate our analysis.

We show a typical powder diffraction pattern in Fig. 6.5. One of the first things that can be observed in this figure is that at small q-values (low scattering angles) we observe nice sharp peaks, but these peaks can be seen to diminish in intensity with increased scattering angle. The second thing we notice is that the spacing between the peaks gets smaller as we go from smaller to larger scattering angles.

The demise of the peaks with increased scattering angle is an interference effect known as the Debye-Waller factor. As we have argued in preceding chapters, when we probe the system at increasingly shorter probing wave lengths, we are more and more susceptible to small deviations from a perfectly ordered structure. In liquids, we notice this right away as there is no periodic structure and we only measure (to good approximation) nearest neighbor distances. In solids, however, we have a similar effect caused by the thermal and zero-point motion of the atoms around their equilibrium positions. Denoting the position of atom j by $\vec{R}_j$, we can decompose this position into

$$\vec{R}_j = \vec{R}_{j,\mathrm{eq}} + \vec{u}_j,$$

with $\vec{R}_{j,\mathrm{eq}}$ the time-averaged lattice position, and $\vec{u}_j$ the deviation of atom j from this average. When we look at the intensities of scattered waves, we will encounter arguments of the exponential function such as:

$$e^{i\vec{q}.(\vec{R}_i - \vec{R}_j)} = e^{i\vec{q}.(\vec{R}_{i,\mathrm{eq}} - \vec{R}_{j,\mathrm{eq}})} e^{i\vec{q}.(\vec{u}_i - \vec{u}_j)} = e^{i\vec{q}.(\vec{u}_i - \vec{u}_j)}.$$

The last equality holds for q-vectors corresponding to a Bragg reflection as the product $\vec{q}.(\vec{R}_{i,\mathrm{eq}} - \vec{R}_{j,\mathrm{eq}})$ will be an integer multiple of 2π in that case. As can be seen from this equation, when q is large, the fluctuations around the equilibrium positions will diminish the strength of the interference pattern, especially so when there is no correlation between the fluctuations around the equilibrium positions of the atoms.
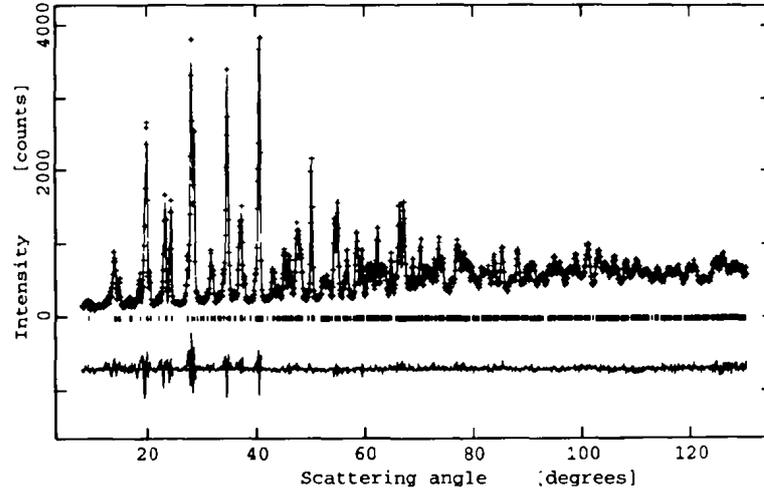
Fig. 6.5 An example of a powder pattern. The pattern shown, collected on YBa$_2$Cu$_3$O$_7$ at Risø, is displayed as the number of counts as a function of the angle over which neutrons were detected. The data themselves are shown by points, the fit according to the Rietveld refinement algorithm is shown by the solid line. The difference between the two is shown separately at the bottom of the graph. All Bragg reflections that are allowed by the symmetry of the unit cell are denoted by the vertical tick marks. The agreement of the fit is expressed in two quality factors: R_w= 8.02 %, χ^2= 3.88. Figure reproduced from W. Montfrooij *et al.*, Journ. of Appl. Cryst. 29, 285 (1996).

This would be the case when we look at random thermal fluctuations of atoms separated by large distances: when $\vec{u}_i$ and $\vec{u}_j$ are uncorrelated, upon carrying out the averaging contained in the brackets $< ... >$, we will see the average effects of how far individual atoms have wondered away from the equilibrium positions. Thus, we would only be averaging over the fluctuations of individual atoms. This is captured in the Debye-Waller factor DWF:

$$\text{DWF} =< e^{i\vec{q}\cdot(\vec{u}_i - \vec{u}_j)} >=< e^{i\vec{q}\cdot\vec{u}} >^2= e^{-q^2(<u_x^2>+<u_y^2>+<u_z^2>)}. \tag{6.3}$$

This Debye-Waller factor is the reason why we do not measure out to very high q-values: the interference pattern will have changed into the pattern of scattering by individual atoms.

Our second observation from Fig. 6.5, namely that the spacing between reflections gets smaller when the scattering angle increases, is fairly straightforward to explain. We illustrate this for the case of scattering by a sample where the unit cell has a cubic structure with side length a. The spacing between lattice planes d_{hkl} for the Bragg reflections corresponding to the Miller indices (hkl) is given by $d_{hkl} = a/\sqrt{h^2 + k^2 + l^2}$.
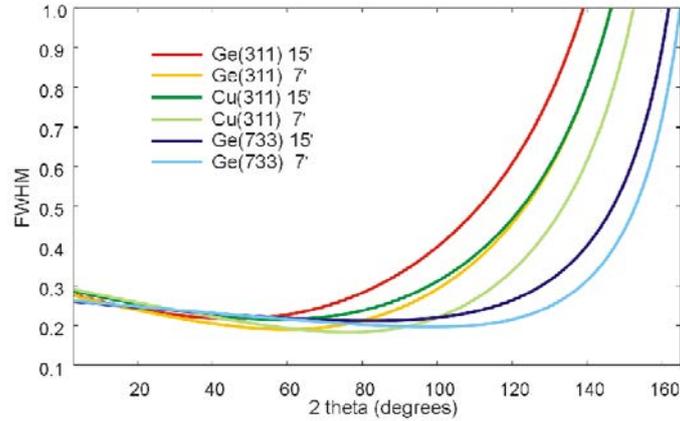
Fig. 6.6 Measured examples of the angular resolution of the BT-1 spectrometer at NIST as a function of scattering angle for various choices of monochromator. FWHM stands for Full Width at Half Maximum, with maximum the maximum intensity in the smeared out Bragg peak. Note that the axis labeling of '2 theta' refers to what we call θ in this write-up. Source: www.ncnr.nist.gov.

Miller indices will be explained in section 6.3. Thus, the scattering angles θ_{hkl} are given by

$$\theta_{hkl} = 2\sin^{-1}(\lambda/2d_{hkl}) = 2\arcsin(\sqrt{(h^2 + k^2 + l^2)}\lambda/2a). \tag{6.4}$$

Thus, the (100) and (101) reflection will be well separated in angle since there is a large difference between $\sqrt{1}$ and $\sqrt{2}$, but the (600) and (601) reflection will be quite close since $\sqrt{36}$ and $\sqrt{37}$ are quite close.

Before we show some examples of what has been measured in powder diffraction, we briefly mention the angular resolution of a diffractometer and the unwanted or 'background' scattering. The angular resolution depends on the collimation that is in place, but also on the scattering angle. In order to properly account for the measured intensity of a Bragg reflection, we need to know over what angular range $\Delta\theta$ the Bragg peak (cone really) has been smeared out. In practice, we fit the angular resolution to a three parameter model. These parameters are normally called u, v and w, and they yield an angular dependent resolution function of the form:

$$[\Delta\theta(\theta)]^2 = u\tan^2(\theta) + v\tan(\theta) + w. \tag{6.5}$$

We show examples of spectrometer resolution functions in Fig. 6.6. As can be seen, the resolution is a smoothly varying function of scattering angle and therefore, it is well suited to being fitted. More on the spectrometer resolution function can be found in section 6.4.
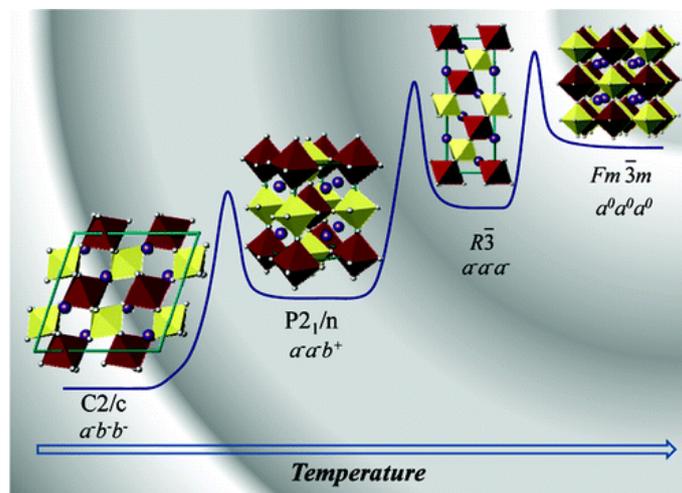
Fig. 6.7 This figure has been copied from the Paul Scherrer Institute (PSI) highlights of 2010, and it was published by S.A. Larregola *et al.* in the Journal of the American Chemical Society 132, 14470 (2010). The room-temperature crystal structure of Pb_2TmSbO_6 was determined from neutron powder diffraction and synchrotron xray powder diffraction to be the monoclinic C2/c space group. The material undergoes a series of structural changes upon heating, culminating in the high symmetry cubic structure Fm-3m.

The word 'background' in powder diffraction refers to all scattering that does not give us information about the relative positions of the atoms. Thus, incoherent scattering is lumped in with the background, and even the scattering by individual atoms at high momentum transfers gets labeled as background. This is not entirely correct of course, but if we are only interested in the positions and species of atoms in a unit cell, and in the amplitude of the fluctuations around these positions, then from an analysis point of view the rest is indeed 'background'. Or almost. In section 6.4 we show an example of how we can slightly misjudge the amplitude of the fluctuations when we label scattering by individual atoms- that has its origin in the coherent cross-section- as background. Notwithstanding this caveat, it is not a bad approach to lump all other scattering together as background.

When we analyze the scattered intensity as a function of scattering angle- or, scattering profile for short- we model (fit) this background by a polynominal function in θ. How many terms we need to include depends on the spectrometer and the spectrometer settings. Typically, we start off our refinement using only a few parameters, and we keep adding more as long as we get a significantly better agreement with the data. This is more or less a case of trial and error, but it is good practice to make a separate plot of the background after it has been fitted to make sure that our polynominal looks sensible.

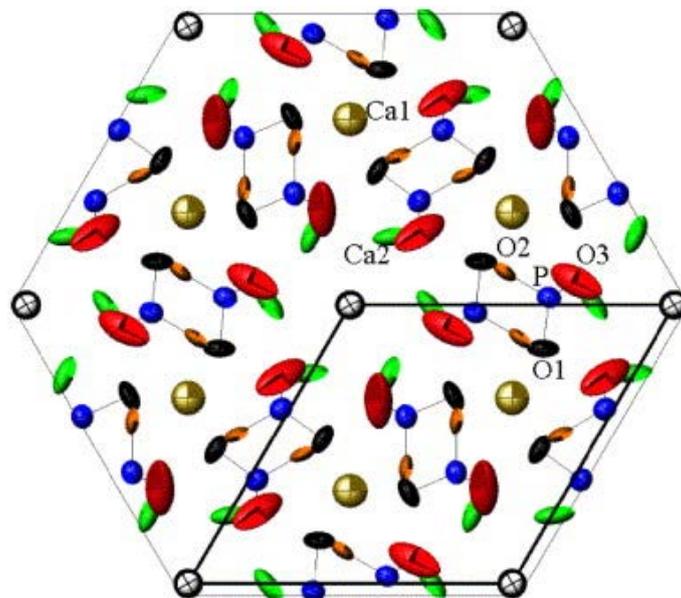So let us look at some examples. In Fig. 6.7, we show the unit cell structure of

Fig. 6.8 Figure reprinted from Th. Leventouri, Biomaterials 27, 3339 (2006), with permission from Elsevier. Shown is a projection down the c-axis of a synthetic CHAp containing 7.2% CO_2, as calculated from a Rietveld refinement of room temperature neutron powder diffraction patterns. The solid lines depict the bonds between the atoms labeled P, O1, O2 and O3 of the phosphate tetrahedra, the ellipses represent the (anisotropic) Debye-Waller factors indicating how far, and in what direction, atoms can wander from their equilibrium positions.

Pb_2TmSbO_6. This is a newly synthesized material and it was found to undergo a series of structural phase transitions upon cooling. Powder diffraction is the method of choice to determine the crystal structure for new compounds.

We show an example of a synthetic material in Fig. 6.8 that is similar in crystal structure to biological hydroxyapatites. Small differences in structure greatly affect the bioactivity of these synthetic materials. Detailed measurements and Rietveld refinement allow for the testing (and refuting) of theoretical calculations indicating that the hexagonal space groups are energetically unfavorable for hydroxyapatite crystallization.

The Debye-Waller factors themselves present us with a wealth of information, especially when the amplitude of the fluctuations is different along different crystallographic directions. For instance, when we cool down a system through a structural phase transition, we might already be able to see above the transition temperature that the amplitude of fluctuation of a particular atom gets much larger along some direction. This heralds the shift in position of that particular atom, and it shows that
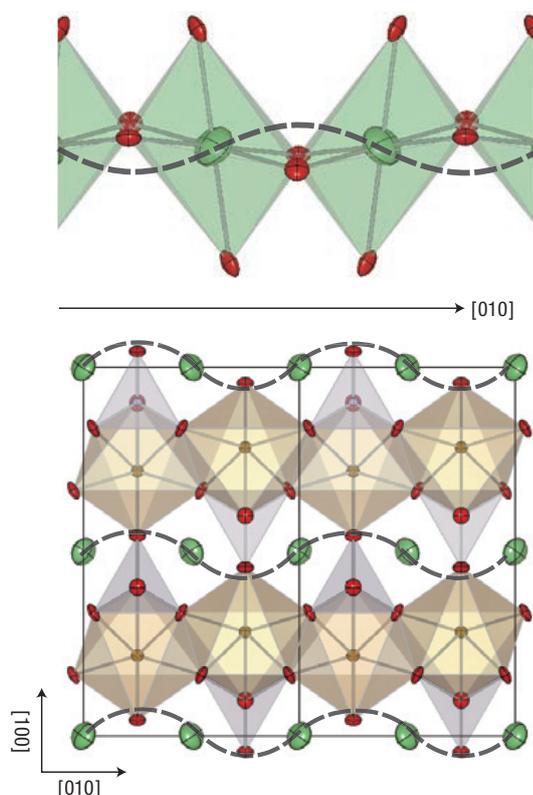
Fig. 6.9 This figure has been reproduced with permission from Macmillan Publishers Ltd: Nature materials (Shin-ichi Nishimura *et al.*, Nature materials 7, 707 (2008)), copyright 2008. Powder diffraction can illustrate how vibrations of lithium atoms evolve into diffusion paths. This was demonstrated by Shin-ichi Nishimura for the battery material $LiFePO_4$. The Rietveld refinement of the powder data indicated the tilting of the Debye-Waller ellipsoids (green candies). The expected diffusive paths that characterize the motion of the lithium atoms while depleting the battery are depicted by the dashed lines.

this structural phase transition is driven by a soft-mode. We can even observe that these Debye-Waller ellipsoids will show some orientation relative to their neighbors. An example of this effect is shown in Fig. 6.9 and it elucidates the diffusion pathway of lithium atoms in a battery material.

6.2 Components of a Powder Spectrometer

The components of a powder diffractometer are very similar to those of a 2-axis diffractometer, so this will be a short section. We briefly discuss two powder diffractometer designs: BT-1 at NIST and PSD at MURR. Both powder diffractometers are located
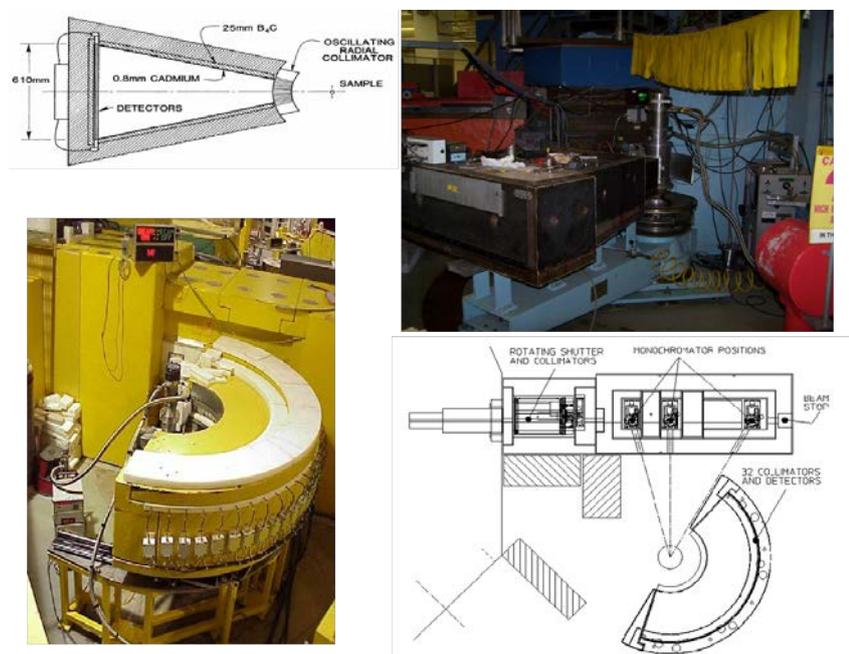
Fig. 6.10 Powder diffractometers are workhorse instruments. The more used they look, the better instruments they tend to be. Shown on the top is the PSD (Position Sensitive Detector) powder diffractometer at MURR. Clearly, this must be a very good instrument. In fact, it is one of the best designs in neutron scattering land. Shown on the bottom is the high resolution BT-1 diffractometer at NIST (www.ncnr.nist.gov).

at a thermal source and both are designed to measure the structure of similar materials. We show photos and schematics of both instruments in Fig. 6.10.

The BT-1 high resolution powder diffractometer looks the most like a 2-axis diffractometer albeit with multiple detectors. The instrument has 32 detectors- spaced 5^o apart- so that the full angular range of 167^o accessible to this instrument can be covered 32 times faster. Collimators are placed between the source and the monochromator, between the monochromator and the sample, and in front of every detector. This instrument achieves its good resolution (Fig. 6.6) by using very tight collimation: the collimators in front of the detectors are only 7'. There is a choice of three monochromators, as shown in the schematic of Fig. 6.10. In all, this diffractometer operates very much like a 2-axis diffractometer discussed in the previous chapter with the same considerations pertaining to instrumental resolution applying in both cases.

The PSD powder diffractometer at MURR achieves its resolution from focusing techniques rather than from collimators. The resolution characteristics of the instrument were designed by Mihai Popovici (Fig. 6.11). The instrument does not use tight

Fig. 6.11 Mihai Popovici (1939-2003) was one of the few people able to think in reciprocal space. One of his achievements is the design of the powder diffractometer at MURR. This diffractometer combines high spatial resolution with high intensity through focusing techniques both in real space and in reciprocal space.

collimators, yet it boasts a resolution not that different from BT-1. The instrument design utilizes focusing both in real space as well as in reciprocal space.

The PSD uses a silicon monochromator to operate at a fixed incident neutron wave length of 1.48 Å, using the (511) reflection of silicon, which is free from second order contamination. The monochromator has been bent in such a way as to focus the wide beam that comes down the beam tube into a narrow beam at the sample position. This is shown in Fig. 6.12.

Normally when one uses focusing in real space, then one loses resolution on q-space. However, in the case of PSD this does not happen because the real space focusing is accompanied by a change in incident wavelength, as sketched in Fig. 6.12. These two effects combine in such a way as to cancel each other's change in q-resolution: increased incoming wave length values are combined with decreased values for $\sin \theta_m$ so that the ratio of the two (determining the momentum transferred from the neutron to the sample) remains almost constant. A clever feat indeed, resulting in a resolution function (shown in Fig. 6.13) comparable to that of BT-1 (shown in Fig. 6.6).

On the PSD, there is no collimation between the monochromator and the sample, but there is an oscillating radial collimator after the sample. Combined with the narrow size of the sample, these two elements ensure that the good q-resolution obtained
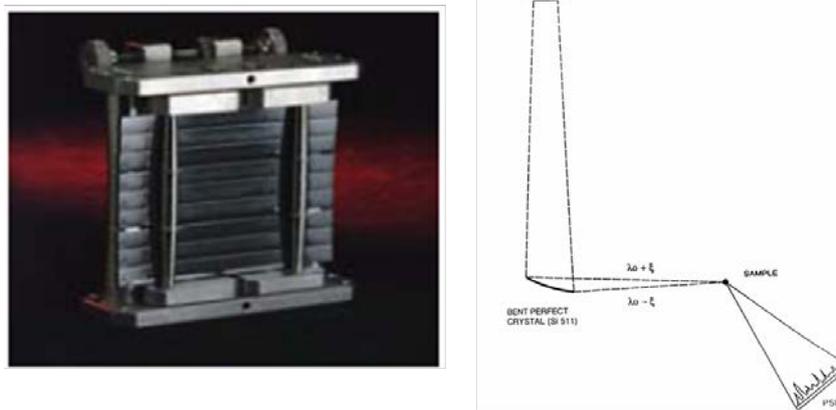
Fig. 6.12 The monochromator for the powder diffractometer at MURR is composed of 9 bent perfect single crystal silicon blades 16.5 x 1.5 x 0.6 cm^3. These are bent horizontally to a 11.5 m radius to provide focusing in q-space and real space as well. Vertical focusing in real space is obtained by mounting the individual blades in a polygonal approximation to a sphere with radius 1.5 m. The monochromator is optimized for the use of small samples - the typical specimen diameter being 3 mm. The focusing that ensues from using this monochromator is shown in the panel on the right.

from the focusing monochromator is not lost. The radial collimator, made of metal blades covered in GdO_2 paint, oscillates so that not a single part of the detector is in the permanent shadow of one of the blades. Behind this collimator is a shielded detector casing that houses five position sensitive detector tubes (Fig. 6.14). Each of the tubes covers an angular range of approximately 20^o, and five tubes are stacked on top of each other so as to intercept a larger fraction of the Debye-Scherrer cone. These tubes can pinpoint the arrival of a neutron to within 1.5 mm. This number implies that the effective detector size only spans 3 arcminutes (1.5 mm at a distance of 1600 mm), ensuring that we do not lose any resolution because of the effective width of the detector. The detector casing moves in steps of approximately 20 degrees to cover the full scattering range of $4.8 < \theta < 105^o$.

Neutron powder diffractometers are very efficient at doing their job. In fact, they are so efficient that one often employs automatic sample changers, such as the ones shown in Fig. 6.15. SANS diffractometers also frequently use these changers. These sample changers work well, as long as one does not require special sample environment such as a furnace or cryostat.

6.3   Crystal Structures and Miller Indices

First, a very good source of information is Wikipedia, Crystal Structures. Also, Ashcroft and Mermin do an excellent job of explaining crystal structures and Miller indices. This
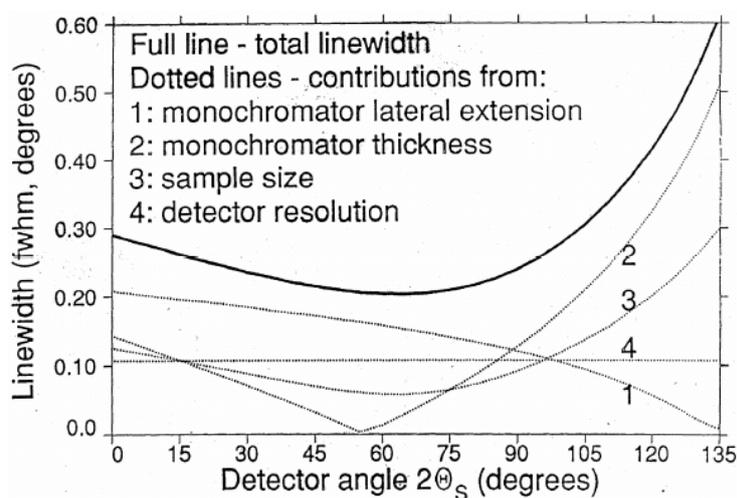
Fig. 6.13 The resolution function for the PSD powder diffractometer at MURR. Despite the lack of collimation, the focusing techniques applied in the design of this diffractometer combine to give a resolution comparable to BT-1 (Fig. 6.6). The fat solid line represents the overall resolution, which is a combination of the four resolution elements shown by the thin dotted lines. The horizontal axis shows $2\theta_S$, which is the full scattering angle we refer to as θ in this booklet.

merely serves as a very short summary, all stolen from various places.

We can view a perfect crystal as being built from many smaller units, called unit cells. These unit cells are small boxes with one or more atoms in them (and with one or more chemical formulae), and they are placed next to each other and stacked on top of each other to fill the entire volume of the crystal. The symmetry of the unit cell, such as being cubic or hexagonal, can sometimes be seen at the macroscopic level in the facets of a crystal. An example is shown in Fig. 6.16.

There are seven distinct possibilities for the symmetry of the unit cell, called the seven lattice systems. Symmetry is related to rotating the unit cell around an axis in such a way that everything looks the same after the rotation is completed. The more ways in which this can be accomplished, the higher the symmetry of the lattice. There are other symmetries possible besides rotating: for instance, we can imagine mirroring the unit cell in some plane that bisects the unit cell. As long as everything looks the same before and after we perform this action, we have a perfect crime and we say that the lattice possesses this symmetry. The notation for these symmetries is not immediately obvious, but letters such as m are used to indicate a mirror symmetry.

When there is more than one atom in a unit cell, we can further distinguish these 7 lattice systems into 14 Bravais lattices. Lattices with the highest symmetry are the

Fig. 6.14 Shown are the position sensitive detector tubes employed on PSD at MURR. Each detector has a 1" diameter and is a 24" long stainless steel tube. The outer wall is the cathode, and the anode is a spring-tensioned nickel-chrome wire. The tube is filled with 8 bar of ^{3}He and 4 bar of argon with 5% CO_2 gas mixed in. The detection efficiency of neutrons with wave length 1.48 Å is 70%.

cubic lattices, the ones with the lowest symmetry are the triclinic lattices. All this is summarized in Fig. 6.17. These 14 possibilities are further subdivided into 230 space groups, depending on the contents of the unit cell. All space group possibilities are described in the crystallographic tables of the society of crystallographers; these are the blue fat books that every university has somewhere, even though they can never

Fig. 6.15 Typically, it does not require a lot of time to measure one powder pattern. A sample changer allows one to measure multiple patterns, and still get enough sleep.

Fig. 6.16 In large perfect crystals, the symmetry of the unit cell shows up in the facets of the gemstone, such as in this piece of rhodochrosite at the Denver Museum of Nature and Science.

be found because someone took them to their office and never returned them.

As an example, we show the three Bravais lattices associated with the cubic symmetry in Fig. 6.17. The 'simple cubic' Bravais lattice, denoted by the letter 'P' for primitive, only has one atom in its unit cell (or 8 times $1/8^{th}$ of an atom in the figure shown). The 'body centered cubic' lattice, denoted by the letter 'I' for innenzentriert (German), has two atoms per unit cell (one in the middle, plus 8 times 1/8 at the corners in the figure). The third possibility, the 'face centered cubic' lattice, has 4 atoms per unit cell (8 x 1/8 + 6 x 1/2).

When doing a powder refinement on a new piece of material, it is often not difficult to figure out which of the 14 Bravais lattices we are dealing with. The space group determination can be more tricky, and it can even be a matter of trial and error to see if one space group (ultimately) results in a better fit than another group when doing a Rietveld refinement. Always bear in mind that there is the possibility of the powder having an impurity phase in it. When this happens, you will know because you cannot find any space group that gives a satisfactory fit.

Neutron diffraction happens because the neutron is scattered by all the atoms in a unit cell. The Bragg condition states that the probing wavelength has to correspond to the separation between the planes that slice up the crystal in such a way that every atom is accounted for. There are many ways (infinite really) of slicing up a crystal in this way. Every which way will correspond to a particular separation d between the Bragg planes and, therefore, to a particular Bragg scattering angle θ_B given by the relation $\lambda = 2d \sin \theta_B$ (with $\theta = 2\theta_B$). Miller indices are a way of denoting these planes by three numbers (hkl), and they provide us with an intuitive way of counting down
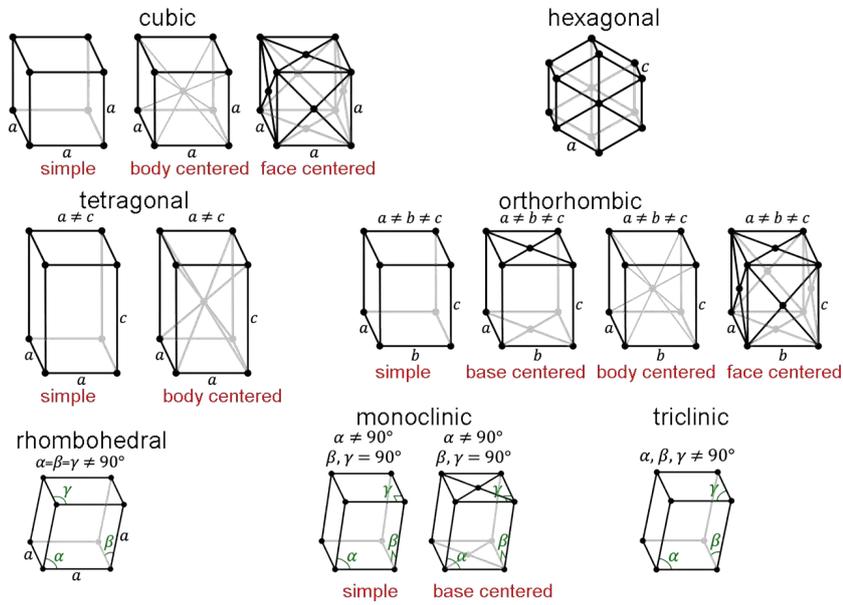
Fig. 6.17 Source individual graphs: Wikimedia Commons, author DrBob. The seven crystal systems (cubic, hexagonal, tetragonal, orthorhombic, rhombohedral, monoclinic and triclinic in order from highest to lowest symmetry) can be divided into 14 Bravais lattices when we allow for more than one atom per unit cell.

over all possibilities.

The Miller indices are determined by where these planes intersect the axes in the unit cell, and they are defined in the following way. If a plane intersect the a-axis at a point halfway along the axis, then the h index of the Miller indices (hkl) is given by the reciprocal of this intersection, $h=2$ in this case. We show some examples in Fig. 6.18 for the cubic case. A Miller index of '0' implies that the intersection happens at infinity or, in other words, the plane would be parallel to that axis. A Miller index with a bar above it, such as $\bar{2}$ implies that the intersection took place on the negative side of the origin.

Miller indices can be used to calculate the amount of momentum transferred by the neutron to the sample: $\vec{q} = h\vec{a}^* + k\vec{b}^* + l\vec{c}^*$ with the reciprocal unit vectors $\vec{a}^*$ etc. given by:

$$\vec{a}^* = \frac{2\pi \vec{b} \text{ x } \vec{c}}{\vec{a}.(\vec{b} \text{ x } \vec{c})}; \vec{b}^* = \frac{2\pi \vec{a} \text{ x } \vec{c}}{\vec{a}.(\vec{b} \text{ x } \vec{c})}; \vec{c}^* = \frac{2\pi \vec{a} \text{ x } \vec{b}}{\vec{a}.(\vec{b} \text{ x } \vec{c})} \tag{6.6}$$

For a lattice where all axes are perpendicular to each other, this results in $|\vec{a}^*| = 2\pi/a$, $|\vec{b}^*| = 2\pi/b$, and $|\vec{c}^*| = 2\pi/c$. Let us look at what this implies for a powder
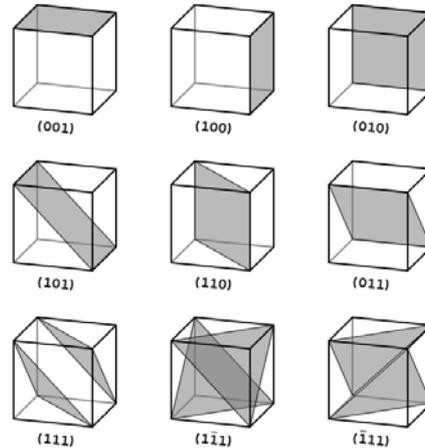
Fig. 6.18 Source: Wikimedia Commons, author Christophe Dang Ngoc Chan. Miller indices (hkl) are both useful for specifying momentum transfer in reciprocal space, as well as for picturing the planes of the lattice that give rise to Bragg reflections.

pattern from a cubic unit cell. The amount of momentum transferred q is given by $q = 2\pi/\lambda_{\mathrm{probe}} = 2\pi/d_{(hkl)}$. We have introduced the notation $d_{(hkl)}$ to indicate the separation between the lattice planes corresponding to the Miller indices (hkl). For a cubic system where the length of the a-axis is a we find

$$d_{(hkl)} = \frac{a}{\sqrt{h^2 + k^2 + l^2}}; q^2_{(hkl)} = \frac{4\pi^2}{a^2}(h^2 + k^2 + l^2). \tag{6.7}$$

What does this mean in practice? When we do powder diffraction, we can calculate q from the scattering angle, independent of what the structure of the sample is. So we can replot our intensity not as a function of scattering angle, but as a function of q^2. The first allowed reflection would correspond to a certain set of Miller indices, let's say for the sake of argument to the (100) reflection. If we do a further replotting, namely instead of plotting versus q^2 we plot versus q^2/q^2_{first}, we would expect all the other Bragg peaks to correspond to integer positions on the horizontal axis. For instance, the (301) reflection would show up at a position of 10, representing $3^2 + 0^2 + 1^2$.

If this is the case (all other reflections showing up at integer values), then we have figured out the symmetry of the unit cell. We have not figured out the particular Bravais lattice yet though. However, if we find that the higher angle Bragg peaks do not occur at integer values when plotted this way, then this is because of two possible reasons. First, perhaps the lowest angle reflection did not correspond to the (100) reflection, perhaps it was the (110) reflection or the (200) one. How to deal with this? We simply try out these possibilities by replotting versus $2q^2/q^2_{\mathrm{first}}$, or $4q^2/q^2_{\mathrm{first}}$ for the case of the (110) reflection or the (200) reflection, respectively.

Second, the above scheme may not work, we may find that some peaks show up at integer values, but others do not. This implies that we do not have cubic symmetry, but perhaps we have a tetragonal or orthorhombic symmetry. To deal with this, we simply look at the peaks that are off, and we look to see if their positions reflect some commonalities. For instance, we might be able to identify the (010) and (020) reflections, giving us the length of the b-axis. It takes some staring at the pattern, and some punching in of numbers on the calculator, but it is the type of input that we need to have *before* we can do a Rietveld refinement.

What is the unit cell is monoclinic, or even lower symmetry? It will still be possible to identify peaks that are located at integer fractions from each other. The best approach here is to plot the data as a function of q^2, and to look for such fractions. Perhaps one can find a peak at $q^2 = 1.4$, and another one at $q^2 = 2.8$. These two peaks have something in common, and we can label them to identify their kinship. Try to find as many of these related reflections as possible, and then try to see if these families can be related. The ones that are related likely belong to $(h00)$ types of reflections, others to $(0k0)$ or $(00l)$ families. Reflections belonging to more complicated lineage, such as $(hk0)$ are likely to remain unidentified at first. However, the ones that have been identified can be used to calculate the lengths of $\vec{a}^*$, $\vec{b}^*$, and $\vec{c}^*$.

The angles between the axes in a monoclinic system have to be found by looking at one particular reflection such as the $(hk0)$ reflection with smallest scattering angle, and the angle between the axes has to be picked so that the reflection identified corresponds to an integer set of Miller indices. If the angle has been identified correctly, then we should be able to identify all possible $(hk0)$ reflections. If not, then we have to try again with another reflection because the one we picked might have corresponded to a more general reflection such as (hkl). It may sound like a time consuming process, but it is not too bad really. The main trick is to make sure to plot the data as a function of q^2 and play around with trying to identify the reflections.

6.4   Rietveld Refinement and Reverse Monte Carlo Analysis

Looking at the powder pattern in Fig. 6.5, one can see that the analysis of the crystal structure cannot be done by hand. In order to determine the structure of the material, we need to determine the intensity of the various reflections, even when they are overlapping. The Rietveld refinement algorithm does this for us, but it is not a black box: it still requires some input.

The main piece of input is the symmetry of the unit cell. This means that we have a choice of 230 space groups, and we have to pick one. In the previous section we discussed how one can go about figuring this out by identifying the crystal structure from the positions of the Bragg peaks, partially by trial and error. In this section we will assume that the correct space group has also been identified, maybe by talking to a chemist or maybe by not being able to obtain a satisfactory fit by choosing the

wrong space group.

What are the next steps in the refinement of the powder pattern? We give an example for the refinement of the nuclear unit cell. Note that there is no particular specified order in what parameters to determine first, just some common sense considerations. As a first step, we have to tell the fit program where the atoms are located in the unit cell, and what type of atoms they are. Most of the time we have a rough idea, and a rough idea is a good starting point. If we really do not know, then ask a chemist what the likely positions are of the atoms given the chemical formula unit. Just put the atoms in, and start with high symmetry positions (identified in the crystallographic tables, the blue book). We also have to tell the program the lengths of the a, b and c axes. We roughly know this number based on the positions of the Bragg peaks. Use these rough numbers.

The way the refinement works is that we tell the program to vary the values of some parameters, while keeping others fixed. We must do this in a step by step process to prevent run-away fits. The first parameter that we will fit is the overall scale factor. We leave all other parameters fixed at their initial guesses while we try to obtain the best fit by varying the overall scale factor. Our next task is to get accurate numbers for the a, b and c axes, for the zero-offset angle and for the incident neutron wave length. We can get these without having to worry too much about where the atoms sit in the unit cell, and what they are up to. As a first refinement, leave the lengths of the a, b and c-axes as free parameters, as well as the zero-offset angle. Leave only 1 free parameters to fit as background, and use isotropic, fixed Debye-Waller factors for all atoms (and fix their values at a fairly large number, we will worry about the details later). Leave all other parameters fixed, including the incident wave length of the neutrons, and run the fit until the goodness of the fit, χ^2, reaches a stable value.

Now repeat the fit, but also allow the incident neutron wave length to be fitted. After this run, we will have fairly accurate values for the a, b and c-axes, for the zero-offset angle, and for the incident neutron wave length. Other than that, the fit will still look awful since we have not looked at the resolution function of the spectrometer, our atoms do not have the correct Debye-Waller factors, and our background is probably underparameterized. And, the atoms might not sit at their correct positions within the unit cell.

As a next step, we can try to refine the resolution parameters u, v and w. These parameters affect the width (in angle) of the peaks, but leaves their integrated intensity unaltered. As a first refinement of these three parameters, fix all other parameters at their current values, and run the fit. Then unfix the parameters that we refined already while leaving u, v and w unfixed, and run another fit. We now have a reasonable value for the resolution function.

Now it is time to play with the contents of the unit cell. Fix all parameters, but unfix the Debye-Waller factors (but leave them isotropic). Run a fit. This should show

a marked improvement in the quality of the fit as higher scattering angles. Unfix u, v and w and run the fit again. Next, add a few more background parameters, perhaps up to 6 in total and run another fit. We can now try to see if we can refine the positions of the atoms in the unit cell that do not sit at high symmetry positions. For instance, if we have atoms at $(0,0,z)$ positions, leave z as a free parameter and run the fit again.

At this point, we are starting to get a reasonable fit, so it might be good to re-refine the basic parameters, such as the length of the axes, the incident neutron wavelength, etc. Do this in groups, that is, unfix certain values while keeping others fixed, before doing a fit with all parameters unrestrained. Now it is time to have a detailed look at the difference plot, the difference between measured data and the fitted curve.

One thing that will likely show up when looking at this difference profile is that the Bragg peaks at the lower scattering angles do not fit very well. This is problematic since these peaks are essentially unaffected by the Debye-Waller factors and, therefore, carry great weight in determining the concentration of the atoms in the unit cell. So we need to fix this as best as we can. The origins of this not-so-great-fit are well known: the curvature of the Debye-Scherrer cone as it intersects with the detector. This causes an asymmetry in the peak shape, and we show an example of this in Fig. 6.19.

As just mentioned, the origin of the asymmetry at low angles is the shape of the Debye-Scherrer cone. When the radius of this cone, where it intersects the detectors, is not much larger than to the height of the detector, then we will see the curvature of the cone. We can exactly calculate this asymmetry effect, but it makes the fitting of the scattering profile more cumbersome as we would have to carry out a numerical convolution at each fitting step rather than simply using a gaussian line shape for our peaks. While computer power is high enough these days to be able to do such a numerical convolution at each fitting step without having to wait for too long, refinement programs instead use asymmetry parameters to compensate for these deviations from gaussian line shapes. We do not approve of this practice, but it is a fact of life.

We have some fitting parameters at our disposal to mitigate the asymmetry effects at low scattering angles. Looking at Fig. 6.19, we can see that a gaussian line shape would do scant justice to what is actually being measured. Adjustable (that is, free parameters in the fit) peak shape parameters can capture most of the deviation from a gaussian line shape, however, one should routinely scrutinize the agreement between fitted peaks and measured intensities at low angles. Note that it actually is not important that the fitted curve follows the data, but it is important that the data and the fitted curve yield the same intensity (surface area) for the peak. After all, it is the intensity that is used for our structural determination, not the peak shape itself. So play around with asymmetry parameters, and then inspect the difference plot to see if we got the intensities correct.

Now that we have most parameters in place, we can try to squeeze the data further to see if we can learn more about our system. Scrutinize the difference pattern,
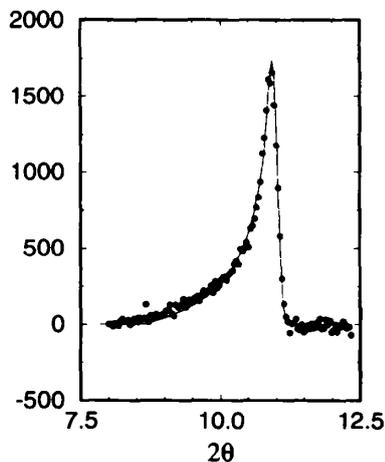
Fig. 6.19 Bragg peaks at low scattering angles clearly show asymmetry effects caused by the curvature of the Debye-Scherrer cone. The data (black circles) are the same as in Fig. 6.5, the solid line shows the exact line shape calculated using geometrical considerations. This exact line shape gives an excellent fit resulting in a close matching of the intensity of the peak, however, using such a line shape requires a numerical convolution during the fitting process. In standard Rietveld refinement, we approximate the asymmetry line shape (by using asymmetry parameters) so that we do not have to carry out this convolution. Figure reproduced from W. Montfrooij *et al.*, Journ. of Appl. Cryst. 29, 285 (1996).

and check if the intensities of the peaks look correct. If not, then we can try leaving the concentrations of the atoms in the unit cell as free parameters; this should only produce modest differences in fit quality to the naked eye, but it still might improve χ^2 noticeably. Another thing to try is to see if anisotropic Debye-Waller factors might give us a better fit, especially in unit cells that have low symmetry.

In a typical refinement we will have reached the point where we can essentially leave all our parameters as free parameters and see what our final fit results looks like. We can try adding more parameters to the background, but make sure that the background (when plotted separately) looks like a smoothly varying function of angle. If it does not, then either the instrument has a problem, or (more likely), we have not learned all there is to learn about our system from the data we took.

If it still does not look right, then we have a few more parameters at our disposal. We can try to see if we potentially have a preferred orientation of our powder grains, or perhaps the grain size shows up as a broadening of the peaks. Perhaps we have some strained grains. These are all parameters that can be used to try to obtain a better fit, but use them with caution. A better fit might be obtained, but the result may not be physical.

Note that there is one correction that we do not have to worry about in powders, but we do in single crystal diffraction: extinction. In powders, the grains are so small that multiple scattering events can be neglected. In contrast, when we do single crystal diffraction to determine the structure of a material as we do using a four circle diffractometer, then we do have to worry about multiple scattering events, and correct for them.

What to do when the result of our powder refinement is simply not good enough, or does not reveal the information that one is interested in? At this point, standard Rietveld refinement is probably not going to help, so it may be time to try a less user-friendly, more hands on method. We briefly discuss the Reverse Monte Carlo (RMC) method.

The Reverse Monte carlo method is a powerful method because it allows for obtaining more information from the data than is possible using standard Rietveld refinement. It is also a very dangerous method, since it is possible to obtain a perfect agreement between the fit and the data at the cost of ending up with an unphysical model. It is time-consuming to implement and, as far as we know, it has only been applied once to the analysis of powder patterns, although it has been used many times in the analysis of diffraction patterns from liquids.

When applying the RMC method to powder patterns, one starts with a collection of unit cells with all atoms at their equilibrium positions. Thus, the input of the RMC method are the results of a previous Rietveld refinement. One chooses the number of unit cells in such a way that the refinement can be completed in one's lifetime, but typically one should choose this number to be as large as possible. One then 'moves' one atom, chosen at random, away from its equilibrium position by a small amount, and one calculates the powder pattern based on a direct calculation of eqn 4.6 using the positions of the atoms. The resulting $S(\vec{q})$ is then averaged over all crystal orientations, the Lorentz factor is put in place, and the pattern is convoluted with the experimental resolution function.

After the 'move', χ^2 is calculated. If the new χ^2 is better than the old one, then the move is accepted. If χ^2 got worse, then the move is accepted with an exponential probability given by $e^{-(\chi^2_{\text{new}}-\chi^2_{\text{old}})}$. This process is repeated until an equilibrium is reached where, on average, one does not find any improvement is the quality of the fit. Thus, RMC is similar to standard Monte carlo, but instead of minimizing the free energy one minimizes the disagreement with the data.

One advantage of applying this RMC method is that the instrumental background that is fitted is no longer angle-dependent. In fact, in the RMC method one actually uses the 'background' to learn more about the sample; the diffuse scattering is an integral part of determining the Debye-Waller factors of the sample. Another advantage is that one can use the exact line shapes for the Bragg peaks, one does not have to resort to asymmetry parameters. We show the result for such a RMC-refinement in Fig. 6.20.
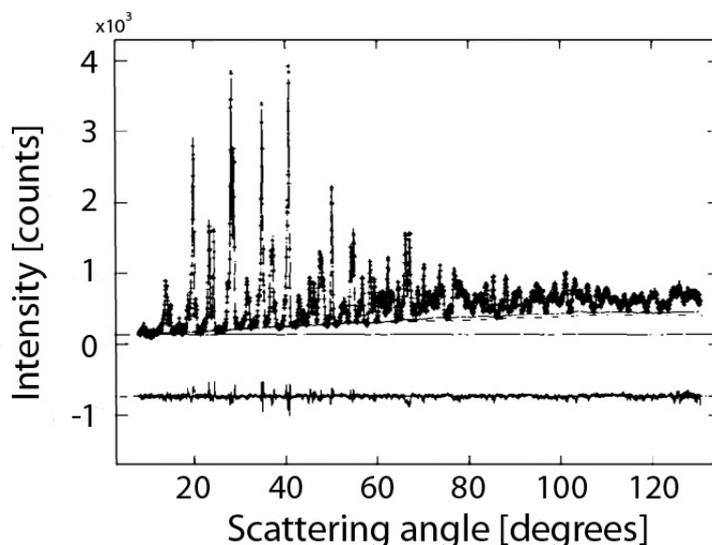
Fig. 6.20 The data shown are the same as in Fig. 6.5, but this time the data have been refined using the Reverse Monte carlo algorithm on 72 unit cells, incorporating the exact line shapes shown in Fig. 6.19. The overall quality of the fit is better than in Fig. 6.5. The flat horizontal line is the angle independent background that was used as a free parameter, the full 'background' (solid sloping curve) was calculated from the displacements of the atoms from their equilibrium positions. This yields a slightly different background from the one found using the standard Rietveld refinement (sloping dashed-dotted curve), with implications for the Debye-Waller factors of the atoms. The agreement of the fit is expressed in two quality factors: $R_w = 6.18$ %, $\chi^2 = 2.31$. Figure reproduced from W. Montfrooij *et al.*, Journ. of Appl. Cryst. 29, 285 (1996).

In RMC, one calculates Debye-Waller factors from averaging over the deviations from equilibrium positions of all atoms in all unit cells. This way, one can immediately see if there are anisotropic Debye-Waller factors, or even split sites. The latter could be a precursor to a (soft-mode driven) structural phase transition. Thus, the wealth of information at one's disposal has increased noticeably, at the cost of a much larger time investment in refining the data.

We should be exceptionally careful when using the RMC method. When we use 100 unit cells with 10 atoms in each of them, we have 3000 free parameters at our disposal when we 'move' these atoms in the three spatial directions. We probably also have about as many data points, so we can obtain a perfect fit, accounting for any random errors in the data, or even for systematic errors in the data. To counteract any unphysical 'movements', we have to build in constraints such as imposing that atoms cannot sit closer to each other than their filled shells allow. When imposing such constraints correctly, then one will end up with a configuration of atoms in one's

computers that is as disordered as possible but still consistent with the data, as well as with the physical reality.

The bottom line about methods such as RMC is that one can use them, but one should only use them when standard Rietveld refinement does not do a good enough job for the problem at hand. One example of this could be in hydrogen storage materials. When trying to determine the positions of hydrogen atoms and molecules inside a crystal matrix, it makes perfect sense to use the information contained in the diffuse background to refine the positions of the hydrogen atoms.

Finally, when doing refinements on magnetic systems, the level of difficulty associated with obtaining a good and physical fit rises greatly. Magnetic unit cells have more possible symmetries than nuclear unit cells, and the orientation of the magnetic moments also affects the scattered intensity. This is good in the sense that one can determine this orientation from the scattering pattern, it is bad in the sense that we are getting to the stage that we have almost too many free parameters at our disposal. We now have Debye-Waller factors, magnetic form factors, background parameters, atomic concentrations, and orientations of magnetic moments that all influence the intensity of the peaks. As such, these fit parameters have a tendency to become cross-correlated, and one is in danger of obtaining a very good fit of the data that does not actually represent what the atoms and their moments are doing in the unit cell.

When doing Rietveld refinement on magnetic systems, one should always collect a powder pattern of the system at a temperature above the magnetic ordering temperature. This pattern can then be used to refine the nuclear unit cell and the instrumental parameters. Then one can use these refined parameters as a starting point to refine the data at the lower temperatures, making sure to leave the instrumental parameters untouched. Some, but not all, Rietveld refinement programs allow for magnetic refinements. An example of such a program is FULLPROF. Magnetic refinements are very much hands-on, and one should always keep a look out for systematic discrepancies between the fit and the data since this likely implies that there is a fundamental problem with our refinement of the magnetic structure. Also, when one finds that the magnetic moments have tilted in some odd way, such as away from a high symmetry axis, then be doubly careful. More likely than not, this is an artifact of the fit.

6.5   Exercises

Exercise 6.1

The graph shows a powder pattern measured on a diffractometer. The aim of this exercise is to figure out as many things as you can about the crystal structure. The diffractometer was operated using an incident neutron wave length of 1.1 Å. The instrumental resolution has smeared out the Bragg peaks into a Gaussian lineshape. Try to answer the following questions, but note that not all questions can be answered necessarily. You will need a calculator and a ruler.
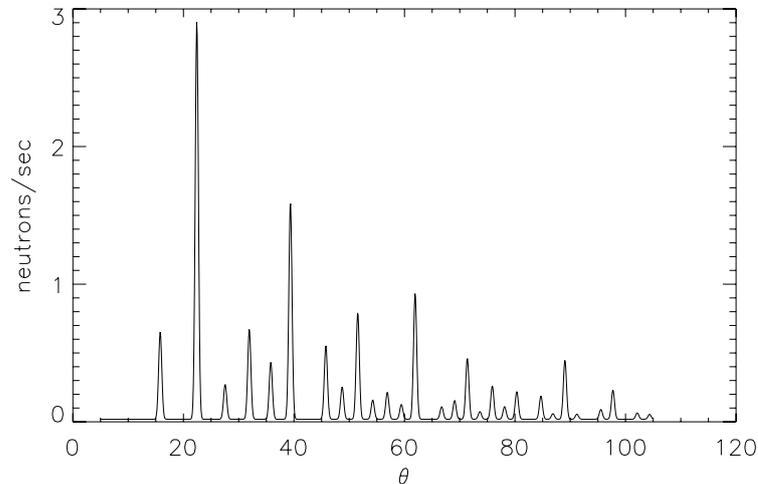
Fig. 6.21 The powder pattern as a function of scattering angle θ for an unknown material measured with incident neutron wave length $\lambda = 1.1$ Å.

a) What is the crystal structure, is it cubic, tetragonal or orthorhombic? (It is not monoclinic, triclinic, rhombohedral or hexagonal).

b) What is the length (in Å) of the a, b and c axes?

c) How many atoms are there in the unit cell? If this cannot be answered, is it possible to say whether it is one or more than one?

d) If there is more than one atom in the unit cell, are these atoms of the same species?

e) The intensity between the Bragg peaks is non-zero. What could be the cause(s) of this?

f) Why do the peaks at large angle become less and less pronounced?

g) Estimate the average displacement of an atom from its equilibrium position.

Exercise 6.2

The PSD diffractometer located at MURR that operates at an incident neutron wave length of $\lambda = 1.48$ Å is very good for determining the positions of atoms in unit cells, but is is not very good at determining the average displacements from equilibrium of those atoms. Why is this?

Exercise 6.3

What are the main consequences of not taking into account the peak shape distortion of the reflections at low scattering angles?

Exercise 6.4

Powders come with a packing fraction; the more one taps the sample holder, the more powder one can fit into one's sample holder and the stronger the scattered signal will be. What can be a possible, unintended consequence of trying to squeeze as much powder into the holder as possible?

Exercise 6.5

What is the crystal symmetry of the specimen shown here (Almandine, $Fe_3Al_2(SiO_4)_3$)?

Fig. 6.22 The Almadine garnet $Fe_3Al_2(SiO_4)_3$ shown in this picture is on display at the Denver Museum of Nature & Science.

7

Small Angle Neutron Scattering Diffractometers

Before we start our discussion on small angle neutron scattering (SANS), we would like to point out an exhaustive reference guide written by Boualem Hammouda of NIST. We will make extensive use of this resource in this chapter. This reference guide, called 'The SANS Toolbox' can be downloaded from:
http://www.ncnr.nist.gov/staff/hammouda/.

SANS is a method by which large structure can be studied. The usage of the word 'large' here reflects that we are talking about structures that are much larger than the size of a unit cell in a solid. SANS is a diffraction technique; hence, we will only obtain information about these structures and their relative positions, we do not learn anything about their movements. Fig. 7.2 displays some of the structures that can be investigated as well as their morphologies. Note that anything 'large' can be measured on a SANS machine, it does not have to be a biological material or polymer. For instance, vortex lattices have been studied extensively using SANS.

The SANS technique is capable of probing large structures through virtue of being able to resolve scattering processes that involve small amounts of momentum transfer $\hbar q$ from the neutron to the sample according to the relationship $\lambda_{\mathrm{probe}} = 2\pi/q$. Using the expression for momentum transfer in an elastic scattering event

Fig. 7.1 Boualem Hammouda is the author of 'The SANS Toolbox'.
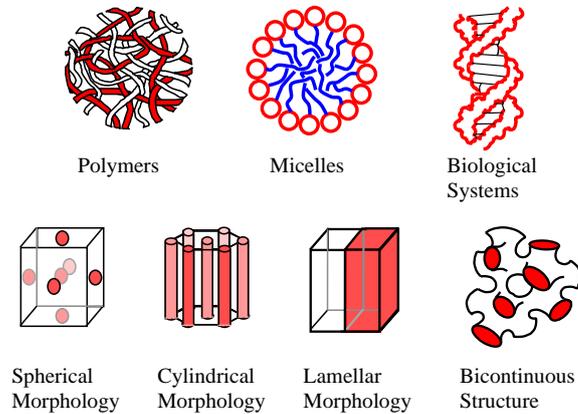
Fig. 7.2 Source: 'The SANS Toolbox'. The top row displays typical structures that are being measured using SANS instruments, the bottom row shows possible morphologies.

$[q = 2k_{\mathrm{initial}} \sin(\theta/2)]$ we find the relationship between probing wave length and scattering angle to be:

$$\lambda_{\mathrm{probe}} = \frac{2\pi}{q} = \frac{\lambda_{\mathrm{initial}}}{2 \sin \theta/2} \approx \frac{\lambda_{\mathrm{initial}}}{\theta}. \tag{7.1}$$

Thus, in order to see large structures, we need to be able to measure scattering events at small scattering angles. It certainly helps if we employ incident neutrons of long wavelength $\lambda_{\mathrm{initial}}$ as available at cold sources, but our main effort is to distinguish between neutrons that have scattered over small angles and neutrons that have not been scattered at all (the straight through beam). This explains why SANS spectrometers tend to be located at cold sources, and why they have their familiar long cylinder shape. We show the two SANS machines located in the HFIR cold source guidehall in Fig. 7.3.

The length of a SANS spectrometer simply reflects that if we have a large distance between the sample and the detector, then we can measure at smaller detector angles. For instance, if we have a detector that represents a 5x5mm^2 target for neutrons to hit, then this detector would measure much smaller scattering angles when placed at 20 m from the sample and 5 cm from the straight through beam, then when placed 2 m from the sample and 5 cm from the straight through beam. In addition, the spread in scattering angles associated with the finite size of this detector would be much smaller at 20 m than at 2 m.

The cylindrical shape of a SANS instruments is because of the disk shape of the SANS detector, which in turn reflects the facts that the samples that tend to be measured on SANS machines do not have a periodic structure (like solids do). Instead,

Fig. 7.3 The cold source guide hall at the HFIR reactor at ORNL houses two SANS instruments, a 30 m and a 50 m SANS. Source: ORNL website.

SANS samples scatter very much like liquids and, therefore, only the magnitude of momentum transfer is important, not its direction. In other words, we only need to know the angle between the straight through beam and the scattered neutron, we do not need to know whether the scattering plane is horizontal, vertical, or anywhere in between. We show a typical SANS scattering setup in Fig. 7.4.
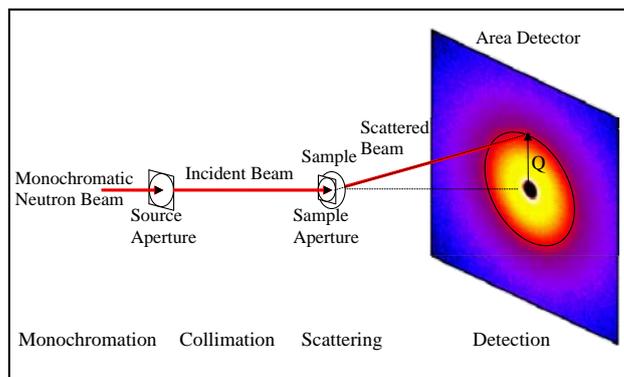
Fig. 7.4 Source: 'The SANS Toolbox'. A SANS experiment is deceptively simple in its setup. Essentially, we have a collimated monochromatic beam that is impingent on a sample. The amount of momentum transferred to the sample is directly related to the polar scattering angle. Note the highly exaggerated vertical scale in this schematic, where vertical distances have been enlarged by a factor of 100.

Finally, SANS instruments are housed inside large tanks. These tanks can be evacuated, or filled with an inert gas. Either way, the air is removed from these tanks so that air scattering of the neutrons that go straight through the sample does not occur, allowing us to measure down to very small scattering angles indeed. A typical SANS machine can measure scattering angles down to $\theta = 0.2^o$, yielding a typical q-range of $0.001 < q < 1\text{Å}^{-1}$ for a 1 m^2 detector. SANS machines and configurations are characterized by their figure of merit, which is equivalent to the minimum q-value that can be achieved for a given configuration. Some highly specialized SANS machines, called ultra SANS or USANS can go down to even smaller scattering angles, however, they achieve these very small scattering angles at the cost of cutting down on the intensity of the beam. A relatively new technique, called SESANS, that might get around this intensity problem is discussed in Chapter 10.

The above paragraphs deal with how we measure our samples from an instrumental point of view, but not with what we actually measure about our samples. Compared to what we have seen in previous chapters about neutron diffraction experiments, we have two important modifications. First, the structures we look at are so large that we cannot hope to be able to elucidate the positions of individual atoms. Instead, we will be looking at scattering that originates from a small volume, so we will be talking about a (three-dimensional) scattering length density $\rho(\vec{r})$, which is the product of the number of particles in a given volume $n(\vec{r})$ and the average scattering length in that volume $b(\vec{r})$. This product $\rho(\vec{r}) = b(\vec{r})n(\vec{r})$ is what we aim to determine about our sample. However, since diffraction experiments take place in reciprocal space, we will only obtain indirect information about this product.

Second, what we can see (measure) depends on the contrast between the particles we are interested in, and the solvent that these particles float around in. Our terminology here is rather loose, and it is borrowed from colloidal suspensions (e.g., smoke, fog, milk, paint). A typical SANS sample consists of the particles we are interested in, and the matrix that surrounds them. For example, we can think of micelles in water, or of a flux lattice penetrating a superconductor. In SANS experiments, the strength of the scattered signal, the differential cross-section, is proportional to $|\rho_{\text{particles}} - \rho_{\text{solvent}}|^2$. We discuss in the next section why this is, but here we already alert the reader to two consequences.

The first consequence is that we can use the technique of contrast matching to our advantage. We can either ensure that we are scattering off of our particles or, if we have multiple types of particles, we can make sure that we only scatter off of one type as shown in Fig. 7.5. The second consequence is that similar to the case of reflectivity experiments (Chapter 8), we have to deal with a 'phase problem'. Since the strength of our scattered signal depends on a number squared, we would measure the same signal independent of whether this number equals- for instance- 2, or minus 2. When we are interested in the sign of this number, then we must perform multiple experiments using various contrasts.
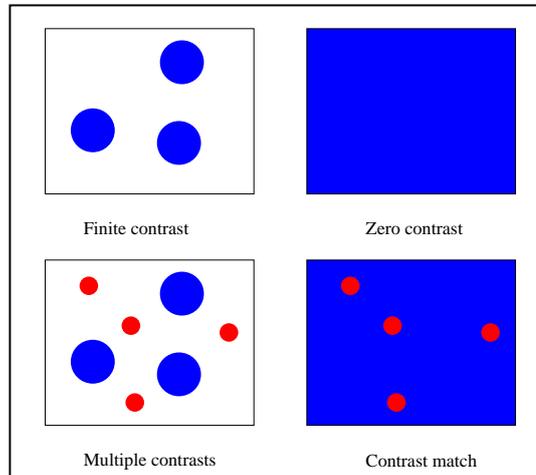
Fig. 7.5 Source: 'The SANS Toolbox'. For the case of particles in a solvent, there can be a distinct difference in scattering length density between the particles and the solvent (top left). When we contrast match the solvent to have the same scattering length density as the particles (top right), then we do not expect any scattering at small angles. For the case where we have multiple particles in our solvent (bottom left), we can selectively match the scattering length density of the solvent to that of one of the particles in order to be able to study the other particles (bottom right).

Contrast matching tells us how likely it is that neutrons will be scattered over small scattering angles by our sample, it does not tell us the shape of the scattering profile when measured as a function of momentum transfer q. The next section will deal with this, followed by a section on the components of a (typical) SANS spectrometer. Our discussion will be mostly at a conceptual level, we refer the reader to Hammouda's write-up for all the details.

7.1   The Skinny

When we discussed elastic scattering in previous chapters, we made a connection between the differential cross-section and the static structure of the system:

$$\frac{d\sigma}{d\Omega} = Nb^2 S(\vec{q}).$$

When we looked (discussion around Fig. 5.5) at magnetic elastic scattering, we essentially kept the same expression, but we put in a magnetic form factor to account for the fact that we were no longer scattering from tiny entities (atomic nuclei), but rather from extended electron clouds (also see eqn 9.1 and Fig. 5.6):

$$\frac{d\sigma}{d\Omega} \sim F(\vec{q})^2 S(\vec{q}).$$

The form factor $F(q)$ is an interference effect. Neutrons scatter from the entire electronic cloud, with the result that, once the probing wavelength becomes comparable to the size of the cloud, we will see reductions in scattered intensity because not all regions of the cloud are being probed with an identical phase.

In small angle neutron scattering experiments we will experience the same phenomena, resulting in a form factor that is determined by the size and shape of our particles. For instance, imagine a large spherical molecule that is probed using a probing wavelength that is comparable to the diameter of this molecule. This implies that if we were to depict this scattering event as an interference diagram such as the one shown in Fig. 2.3, then we would have that half of the molecule would be in a region close to a 'white band', and the other half would be close to a 'dark band'. The scattering from these two regions would be out of phase, and therefore, we would not observe any scattering because of the resulting destructive interference. This is shown in Fig. 7.6.

We can summarize the above reasoning in an equation that tells us the connection between the differential cross-section on the one hand, and between the contrast of the particles and the solvent, the shape of the particles, and the relative positions of the particles on the other hand:

$$\boxed{\frac{d\sigma}{d\Omega} = N_p |\Delta\rho_p|^2 P(q) S(q)}. \tag{7.2}$$

In this equation we have dropped the vector subscript from q since we are dealing with scattering from a non-periodic structure. $|\Delta\rho_p|^2$ is a measure of how much a single particle scatters while submerged in the solvent (details to follow), N_p stands for the number of particles in our system, and $P(q)$ is the form factor of the particles we scatter from (the equivalent of $F(q)^2$ in magnetic scattering). Note the change in definition of the form factor between SANS notation and magnetic scattering notation. We have done this to be consistent with the nomenclature in SANS experiments.

Eqn 7.2 is the master equation for interpreting SANS experiments. It is also a beautiful equation in the sense that everything relevant to the scattering has been nicely factorized. The one caveat about this equation is that it is only valid in the limit of single scattering, something which sometimes can be a problem for hydrogenous samples.

The contrast factor $|\Delta\rho_p|^2$ tells us how much scattering we can expect from a particle. It is formally defined for a particle of volume V_p as

$$\Delta\rho_p = \int_{V_p} d\vec{r}[n(\vec{r})b(\vec{r})_{\text{particle}} - n(\vec{r})b(\vec{r})_{\text{solvent}}].$$

The main function of this equation is for planning what solvent to use. However, we can also look this up in a table as the scattering length densities for most molecules

of interest have been tabulated (see Appendix B). The most interesting aspects of our sample are hidden in the form factor $P(q)$ and, to some lesser extent, in the static structure factor $S(q)$. Before we look at these, a quick, two paragraph intermezzo of why we use $\Delta\rho$ as opposed to ρ. This intermezzo can be skipped and the reader can move on to the discussion about form factors without losing the train of thought.

In the continuity approximation, where we do not deal with individual atoms but rather with scattering length densities $\rho(\vec{r})$, the differential cross-section is given by:

$$\frac{d\sigma}{d\Omega} = |\int_{\text{sample}} d\vec{r}\rho(\vec{r})e^{i\vec{q}\cdot\vec{r}}|^2. \tag{7.3}$$
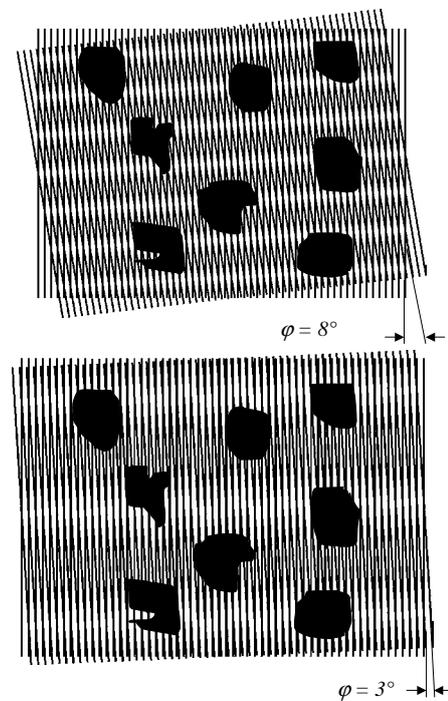
Fig. 7.6 The particles investigated by doing small angle neutron scattering are so large that even at small scattering angles we lose some of our constructive interference. In scattering experiments, this shows up as the particle's form factor $P(q)$. In the bottom panel (scattering angle $\theta = 3^o$) the loss of constructive interference is only moderate since $\lambda_{\text{probe}} >$ particle size. The top panel (scattering angle $\theta = 8^o$), however, represent an almost complete loss of constructive interference since $\lambda_{\text{probe}} \approx$ particle size. Conversely, we can infer the shapes of the particles by measuring this demise of constructive interference with increasing scattering angle. Source: neutron scattering course by Ignatz de Schepper.

This equation is not all that mysterious when we realize that the scattering length density for a system of N particles located at positions $\vec{R}_i$ is formally given by

$$\rho(\vec{r}) = \sum_{i=1}^{N} b_i \delta(\vec{r} - \vec{R}_i). \tag{7.4}$$

Note that in this definition, the summation goes over all atoms, both the atoms that make up the SANS particles we are interested in, and the atoms that make op the solvent. Also note that if one is a physicist, then one can quickly realize that the above equations are nothing more than a rewrite of the expressions for the differential cross-sections that we have already used in previous chapters. If one is not a physicist, then the only important thing to notice is that the differential cross-section depends of $\rho(\vec{r})$, as opposed to $\Delta\rho(\vec{r})$.

We can decompose the scattering length density $\rho(\vec{r})$ into the scattering length densities of the SANS particles $\rho_p(\vec{r})$ and that of the solvent (or matrix) $\rho_m(\vec{r})$: $\rho(\vec{r}) = \rho_m(\vec{r})$ (when in the solvent), and $\rho(\vec{r}) = \rho_p(\vec{r})$ (when we are dealing with a particle). When we substitute this decomposition into our expression for the differential cross-section (eqn 7.3) we find

$$\frac{d\sigma}{d\Omega} = |\int_{\text{sample}} d\vec{r} \rho_m(\vec{r}) e^{i\vec{q}\cdot\vec{r}} + \int_{\text{all particles}} d\vec{r} [\rho_p(\vec{r}) - \rho_m(\vec{r})] e^{i\vec{q}\cdot\vec{r}}|^2. \tag{7.5}$$

The first term does not contribute to the scattering in the SANS regime. We only expect scattering from a uniform solvent when we probe is using wavelengths comparable to the interatomic spacings; when we use SANS probing wavelengths we will not observe any constructive interference patterns given the non-periodic (a.k.a. random) nature of the solvent. We already encountered this in Fig. 4.7 where we did not observe any scattering from a liquid at low q-values. Therefore, for small q-values, only the second term will contribute, yielding the $\Delta\rho$ terms that appears in the cross-section for SANS experiments (eqn 7.2). This ends our little intermezzo.

The form factor tells us about the size of our particles, about their shapes, and about whether they are solid, hollow, or even fractal. In fact, the appearance of the static structure factor can be a nuisance in determining the form factor $P(q)$. We illustrate this in the following example. Picture a collection of particles that have a form factor $P(q)$ and a static structure factor $S(q)$. What we will get out of SANS experiments is the product $P(q)S(q)$. This is shown in Fig. 7.7 for various concentrations of particles.

Fig. 7.7 nicely illustrates how the static structure factor can become a problem in determining the form factor: normally we do not know the form factor before hand, it is what we are after in a scattering experiment. After all, from the form factor we can determine something about the structure of our particles. So we would like to do our experiments in the low concentration limit where we will not find any correlation between neighboring particles, so that $S(q) = 1$, and so that our scattering directly
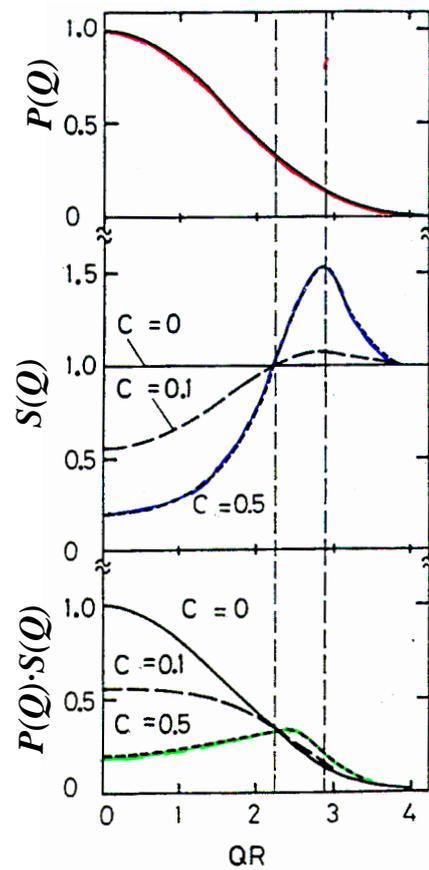
Fig. 7.7 The top panel shows the form factor of 'Gaussian' spherical particles of radius R. The middle panel shows their static structure factor $S(q)$ for three concentrations $c = N_p V_p / V$ of N_p particles of volume V_p in a sample volume V. The bottom panel displays what we can infer from a SANS experiment about the product $P(q)S(q)$. Source: neutron scattering course by Ignatz de Schepper.

yields the form factor $P(q)$. But, in order to get a decent scattering intensity, we have to add more particles to the solution so we cannot work in the very low concentration limit. But when we add more particles, we also have to consider the static structure factor.

We can overcome the above problem through modeling, and by taking proper care when we prepare the sample so that we know what the concentration of particles is. The modeling will involve coming up with a model for the static structure factor of the particles in the solution. These models, based on hard sphere theory, are actually quite accurate (especially at low concentrations), although they are not perfect. Al-

ternatively, the static structure factor can be measured by other means, such as by using light scattering methods. In the end, the experimentalist will have to come up with a consistent model, where the radius R of the particles and their concentration not only describes the form factor, but also the static structure factor.

We now turn our attention to the form factors themselves. The form factors of particles that can be described by an analytic expression can easily be calculated through means of Fourier transform from real space to reciprocal space. We show some examples in Figs. 7.8 through 7.10. In practice, one compares the scattering profile qualitatively to the possible form factors, and one then concludes what type and shape of particle one is looking at.

An important measure characterizing the size of a particle is the Guinier radius R_G (used interchangeably with the term radius of gyration), which is essentially a weighted distribution of the mass of a particle akin to the moment of inertia. For instance, for a hollow sphere of radius R we have that $R_G = R$, while for a solid sphere we obtain $R_G^2 = 3R^3/5$. In neutron scattering experiments, the Guinier radius dictates the fall off of the form factor at small q-values as q: $P(q) = 1 - q^2 R_G^2/3 + ...$. In short, if we are able to determine the form factor of our SANS particles from our SANS experiment, then we can determine the shape and mass distribution from the qualitative behavior of the form factor, and we determine their characteristic size from the fall off (with q) of this form factor.

We can have many different shapes of particles, from needles to disks. The form factors for all these particles can be calculated, and we refer the reader to the SANS
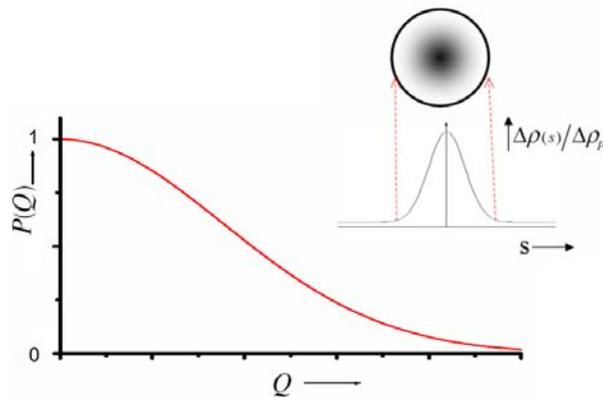
Fig. 7.8 The form factor $P(q)$ for a particle whose scattering length density $\Delta\rho(s)$ falls of according to a Gaussian distribution with distance s to the center of the particle. The 'radius' R_G of this particle is defined by the characteristic width of the Gaussian distribution $[\Delta\rho(s) \sim e^{-\alpha s^2}; R_G^2 = 3/(2\alpha)]$. For this type of density distribution, the form factor is given by $P(q) = e^{-q^2 R_G^2/3}$. Source: neutron scattering course by Ignatz de Schepper.
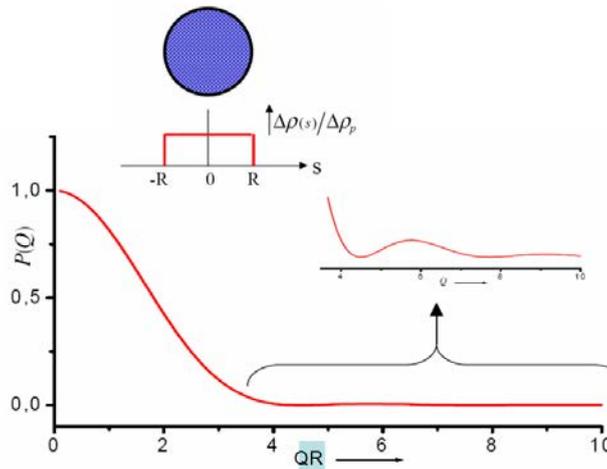
Fig. 7.9 The form factor $P(q)$ for a solid sphere of radius $R = \sqrt{(5/3)}R_G$ is given by $P(q) = 9[j_1(qR)]^2/(qR)^2$. For large q the form factor falls off as $P(q) \sim 1/q^4$. Source: neutron scattering course by Ignatz de Schepper.

toolbox for the details. Every parameter that describes a shape parameter of the particles will show up in the SANS experiment as a characteristic q-value. We show an example for cigar shaped particles in Fig. 7.11. With increasing q-values, the form factor will reflect the decreasing relevant length scales R: $q_{\mathrm{characteristic}} \sim 1/R$. For the largest q-values, where we probe the smallest length scales corresponding to the internal structure of the particle, we find a powerlaw behavior where the scattered intensity is proportional to $1/q^n$. For solid cigar-shaped particles, we find that the drop off in
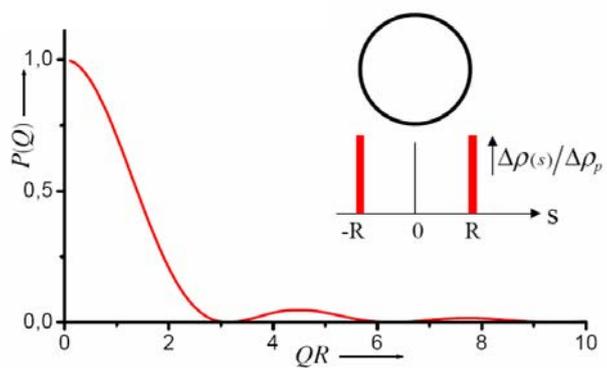
Fig. 7.10 The form factor $P(q)$ for a hollow sphere of radius $R = R_G$ is given by $P(q) = [j_0(qR)]^2$. For large q the form factor falls off as $P(q) \sim 1/q^2$. Source: neutron scattering course by Ignatz de Schepper.

this region is similar to that of a solid sphere ($\sim 1/q^4$), since both objects represent solid, uniform masses when probed on length scales smaller that any characteristic size related to the shape of the particles. In general, when the probing length scale becomes smaller than characteristic length scale, then we refer to this region as the Porod region.

In fact, the high-q behavior of the form factor is very useful when it comes to determining whether a particle is solid, or not. When we plot our form factor on a log-log scale, we can determine this high-q fall off quite easily. We have already seen some examples of this high-q behavior for solid and hollow spheres, but we can have more possibilities in general. For needles, or cigar shaped particles, we will find that $n=1$ as long as our probing wavelength is smaller than the longest dimension of our needles (the length of the needle), while being larger than the shortest dimension (thickness of the needle). As we have seen, a value of $n= 4$ represents a solid object. We can also have non-integer values for n, representing a rough surface characterized by a fractal dimension D. For this latter case we find that $n = 6 - D$. These numbers are summarized in Fig. 7.12.
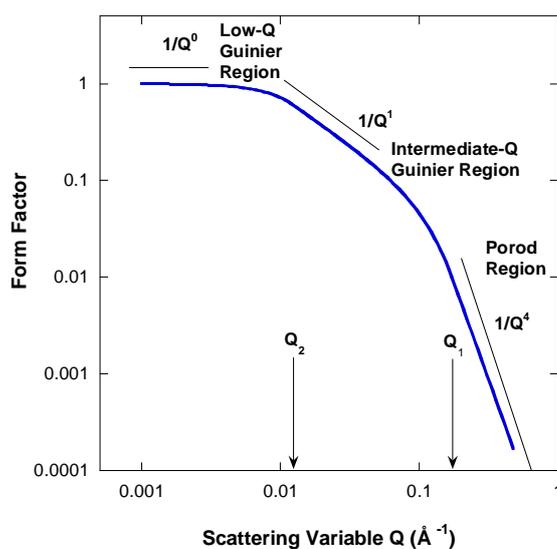
Fig. 7.11 Source: 'The SANS Toolbox'. Shown is the calculated form factor for cigar shaped particles of length 100 Å and of thickness 10 Å. Note that the form factor is shown on a log-log scale. At low momentum transfers we find a drop off of the form factor that is characterized by $P(q) = 1 - q^2 R_G^2/3 + ...$ Once the probing wavelength becomes smaller than the length of the cigars, but still larger than their thickness, then we observe a power law behavior $\sim 1/q$ characteristic of very thin needles. Once the probing wavelength becomes smaller than the thickness of the cigars, then we observe a $1/q^4$ behavior as expected for solid objects.
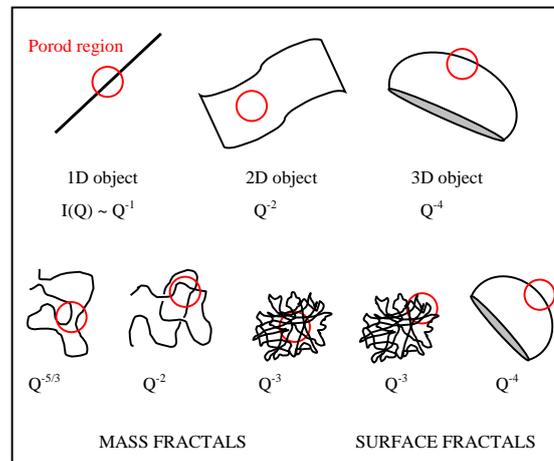
Fig. 7.12 Source: 'The SANS Toolbox'. The Porod region refers to the high-q behavior of the form factor. High q-values imply small(er) probing wavelengths, so that we are probing the structure of the particles themselves. The radius of the red circles is given by the probing wave length. The top row displays smooth objects (1-, 2- and 3-dimensional) and their expected q-signature in the form factor $P(q)$ as measured in a SANS experiment. The bottom row shows the expected q-dependence in the Porod region for fractal objects. The two objects on the left are self-avoiding polymers and spaghetti polymers, the three objects on the right of this lower row are mass fractals.

The above discussion on form factors related to the case where all our particles that are floating around in the solution have the same size and shape. In general, we can expect, at least, some variation in their sizes. This is referred to as polydispersity, and an example of this for hollow spheres is shown in Fig. 7.13. The particle sizes are now characterized by an average radius and by a distribution $D(R)$ around that radius. The average radius will take on the role of uniform radius in the initial fall-off of the form factor, while the distribution $D(R)$ will modify the intermediate fall-off of the form factor. This is shown in Fig. 7.14. For very large q we still find that the form factor falls of according to the (same) power law (in q) that characterizes the shape of the particles (not their size). Thus, the form factor of polydisperse particles is more difficult to interpret, but SANS does provide a very good tool for measuring the degree of polydispersity.

The form factors discussed above are shown as a function of q, whereas the SANS detector is a two-dimensional detector. In order to go from the data as collected by the detector to function of a single variable, we need to perform averaging. For scattering that only depends on the polar angle (and not on the azimuthal angle), this averaging is shown in Fig. 7.15. This is the case where our SANS data will be most accurate in (ultimately) determining the form factor. For the cases where we do have some directional dependence (some examples are shown at the end of this section), we can
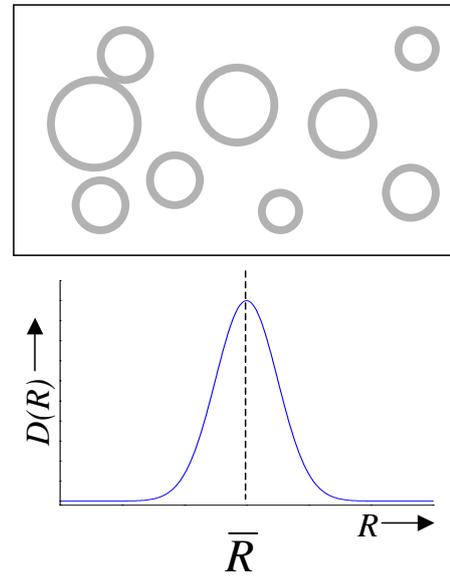
Fig. 7.13 Schematic of a polydisperse collection of hollow spheres. This collection of spheres can be characterized by an average radius $< R >$ and a distribution $D(R)$ around the average radius. Source: neutron scattering course by Ignatz de Schepper.

only average along certain lines of symmetry.

We finish this section by showing some examples of actual systems that have been measured using the SANS technique. As mentioned, the radius of gyration R_G is an
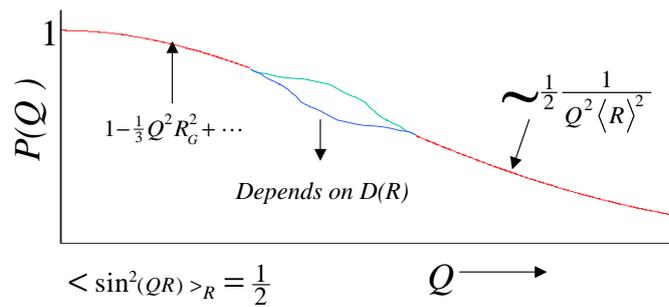

Fig. 7.14 The form factor $P(q)$ for a collection of hollow spheres of average radius $< R >= R_G$ is given by $P(q) =< [j_0(q < R >)]^2 >$. The low and high-q behaviors are qualitatively the same as the case for uniform hollow spheres, however, the intermediate q-region will be modified according to the distribution $D(R)$. Source: neutron scattering course by Ignatz de Schepper.
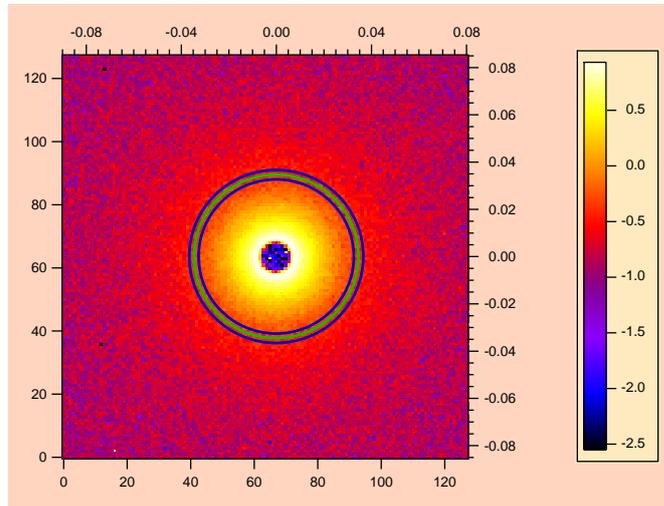
Fig. 7.15 Source: 'The SANS Toolbox'. For scattering by particles that do not show any directional dependence, we perform the averaging over the intensities as measured by the detector along circular paths (one of which is shown). The bottom and left hand scale display the positions on the detector where the neutron hit it in millimeters; the top and right hand scale display the momentum transferred by the neutron to the sample along the horizontal and vertical directions in $Å^{-1}$. The number of neutron counts in each element of the detector is indicated by the color scale in the right panel. Note that the number of counts have been normalized, the numbers on the right no longer directly represent the number of counts.

important measure characterizing the size of the particles, and this radius can be determined from the low-q behavior of the form factor. As long as we are looking at dilute solutions, we can safely assume that $S(q) = 1$ so that we can directly determine R_G from our scattered intensity $I(q)$. An example of this is shown in Fig. 7.16 for various concentrations of dendrimers (large spherical molecules consisting of branched polymers that have potential in drug delivery as they can penetrate the cell's membrane).

In order to determine the radius of gyration, we plot the log of the scattered intensity versus q^2. The reasoning behind this is that the form factor falls off as $P(q) = 1 - q^2 R_G^2/3 + ... \approx e^{-q^2 R_G^2/3}$. Since $I(q) \sim P(q)$ when $S(q) = 1$, we find that $\log [I(q)] \sim q^2 R_G^2/3$. The data in Fig. 7.16 demonstrate that the SANS technique is sensitive enough to observe small changes in the radii of gyration when the concentration of the dendrimers is changed.

The SANS technique is also sensitive enough to observe the difference between a polymer that does not cross itself, and one that does. By plotting the SANS data in the relevant high-q range on a log-log scale, we can determine the power n of the high-q power law $P(q) \sim 1/q^n$. The data for DNA coils in glycol shown in Fig. 7.17 show
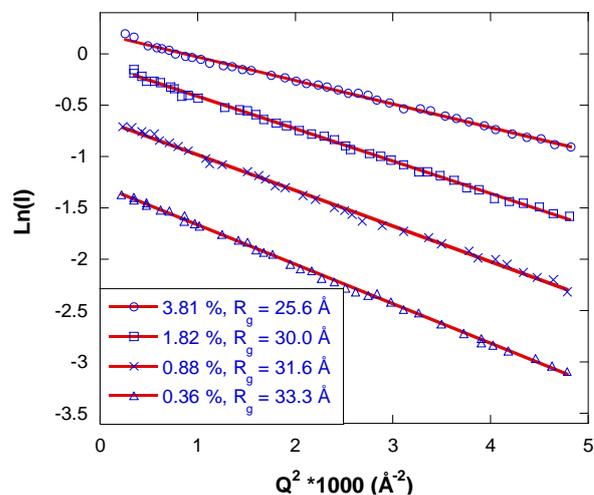
Fig. 7.16 Source: 'The SANS Toolbox'. Guinier plot for SANS data taken from seventh–generation PAMAM dendrimers in D_2O. The dendrimer fraction is varied.

that this molecule represents the equivalent of a self-avoiding random walk, namely a polymer that does not end up looking like spaghetti.

The data shown thus far- and, in fact, the entire discussion in this section- was relevant to samples where the scattering does not show any directional dependence.
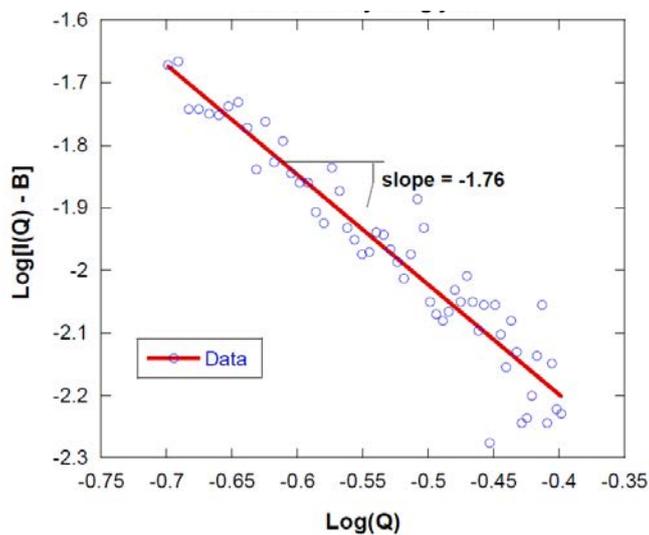

Fig. 7.17 Source: 'The SANS Toolbox'. Porod plot for SANS data taken from 4% (g/g) DNA coils in d-ethylene glycol at 50^oC (above the helix-to-coil transition temperature).
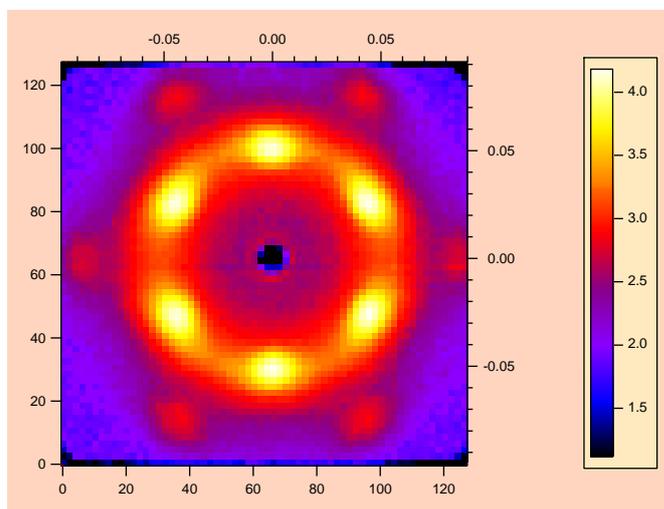
Fig. 7.18 Source: 'The SANS Toolbox', data taken by Slawecki and coworkers. SANS data on 25% P85 Pluronic (PEO-PPO-PEO triblock copolymer) in D_2O under Couette shear (5 Hz frequency) at 40^{o}C. The micelles form a cubic single crystal structure.

However, the SANS technique is by no means restricted to such specimens. One type of *in situ* experiment that can be performed on a SANS diffractometer is to apply a shear to the sample. When doing so, the molecules in the solution can arrange themselves into an ordered structure. An example of this effect is shown in Fig. 7.18 where a collection of Pluronic micelles in a (heavy) water solution are observed to adopt a cubic structure when subject to shear. Vortex lattices in superconductors and superfluids adopt similarly ordered structures.

We conclude this section by showing SANS data on a membrane. This membrane was oriented at 60^o to the neutron beam so that momentum could be transferred both in the plane of the membrane as well as perpendicular to it. The resulting pattern on the SANS detector is shown in Fig. 7.19. The reason we included these data in this booklet is because they look beautiful, and because they reveal the level of detail in structure that can be observed in a SANS experiment.

7.2   Components and Spectrometer Resolution

The components that make up a SANS spectrometer, and the sample equipment that is used on such spectrometers is described in detail in the SANS toolbox. In this section we give the briefest of summaries of the materials presented in the toolbox.

SANS instruments are (mostly) located at cold sources, and preferable at the end of a curved neutron guide. The cold source is essential to obtain neutrons of long wavelengths- required to be able to go down to low momentum transfers- whereas a
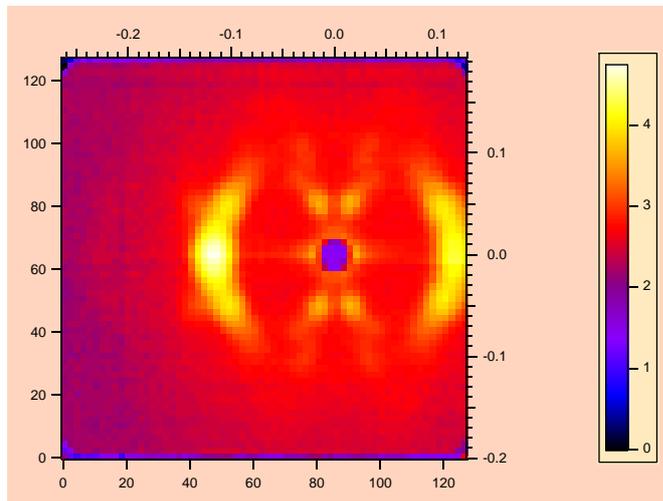
Fig. 7.19 Source: 'The SANS Toolbox', data taken by Yang and coworkers. SANS data from oriented DMPC/DMPG membranes containing magainin peptides in DMPC bilayers and D_2O. The sample was oriented at 60^o to the neutron beam direction in order to observe structures both parallel and perpendicular to the membrane surface.

curved guide is very helpful in eliminating the background of faster neutrons and xrays. If a curved guide is not present, then one tends to use Be-filters that only transmit neutrons with wave lengths longer than 4 Å, and one uses filters such as Bi to deal with the xrays.

In order to select a particular wavelength for the incoming neutrons one uses a velocity selector or a crystal monochromator. Unless the requirements on the wave length resolution are very stringent, one normally uses a velocity selector that selects
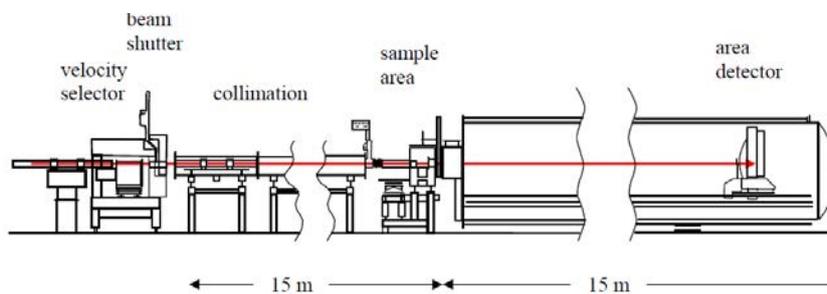

Fig. 7.20 Source: 'The SANS Toolbox'. Shown is a schematic of the SANS spectrometer at NIST.

a wavelength range of about 10% around the average wavelength: $\Delta\lambda/\lambda = 10\%$. This is the first component that is shown in Fig. 7.20.

Before the beam of monochromatic neutrons is directed at the sample, special collimators are used to shape the beam and reduce its divergence. The collimators are in fact sets of circular apertures that are housed in an evacuated casing (see Fig. 7.20). Next, the beam will encounter the sample, followed by another set of circular apertures that define the post-sample collimation. The area detector, which tends to be on tracks so that the distance to the sample can be varied, is housed in the characteristic cylindrical tank.

In between the sample and the detector is a moveable beam stop. The function of this beam stop is to prevent the neutrons that fly through the sample without being scattered from saturating the detector. When a detector has to deal with too many counts, it ceases to function properly, since it takes a certain amount of time for the electronics to be ready to deal with the next counting event. This amount of time is called the dead-time of the detector and, while it is a very short time, it can nonetheless make the detector malfunction when more than (roughly) 50,000 neutrons hit is every second. This dead time effect is not restricted to SANS detectors, all neutron detectors suffer from it.

Even with the beam stop properly positioned, the counting rate can be so high that one actually has to attenuate the incoming beam. SANS spectrometers tend to have sets of calibrated attenuators that can be rotated into the beam before it strikes the sample. The area detector itself has many aspects worth considering, such as the spatial resolution with which it can pinpoint the point of arrival of the neutron when it hits the detector. All aspects are detailed in the SANS toolbox, but as far as the occasional user is concerned, the detector is just a means of detecting the neutrons with roughly 70% efficiency (Fig. 7.21).

The resolution function of a SANS spectrometer describes by how much features related to the scattering by the sample will be smeared out in q-space. This smearing out will be determined by the spread of incident wave lengths, by the angular resolution as determined by the collimators, by the spatial resolution of the detector and by the sample-to-detector distance, and by the effects of gravity. The combined effects of these resolution elements on the q-resolution are shown in Fig. 7.22.

The reader might be surprised to see that gravity plays a role in determining the resolution of the spectrometer, but this aspect can be directly related to the spread in incoming neutron wave lengths. A neutron that travels 10% slower than another neutron, will have 10% more time to fall because of the force of gravity. This would result in a 21% increase in the vertical distance over which it falls. Thus, gravity has a direct effect on the vertical q-resolution, but not on the horizontal q-resolution. A neutron with a wave length of 10 Å takes 40 ms to traverse a 16 m separation between the sample and the detector, during which time it will fall about 0.8 cm. An 11 Å
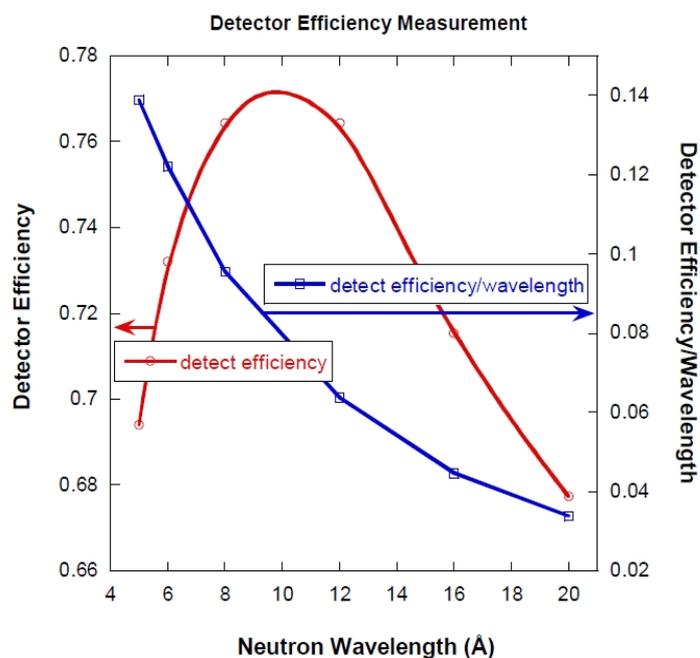
Detector Efficiency Measurement

Fig. 7.21 Source: 'The SANS Toolbox'. The detector efficiency of the SANS detector of the NIST SANS spectrometer displays a slight wave length dependence (red curve). However, given that $\Delta\lambda/\lambda= 10\%$, the user will not need to be concerned with changes in detection efficiency.

neutron will fall about 1 cm.

In fact, the effects of gravity are more severe than what was mentioned in the preceding paragraph. In order for the slower neutrons to actually hit the sample, they need to make it through the sample aperture and, therefore, they had to be traveling in a slightly upwards trajectory in the first place, as shown in Fig. 7.23. This trajectory requirement to cover the pre-sample flight path also affects the vertical q-resolution. Of course, the slower the average speed of the neutron, the more severe the gravity induced resolution smearing will be.

The resolution can be improved, if needed, by tightening up the collimation, by moving the detector to its maximum separation from the sample, by choosing a velocity selector with a more restrictive transmission window such as $\Delta\lambda/\lambda= 5\%$, by replacing the velocity selector with a crystal monochromator, or by performing the experiments on the international space station. It is important to realize though that in general we do not have very restrictive resolution requirements. The reason for this is that we are scattering from individual structures, not from a periodic array of individual structures. The result is that the static structure factor will look much more like that
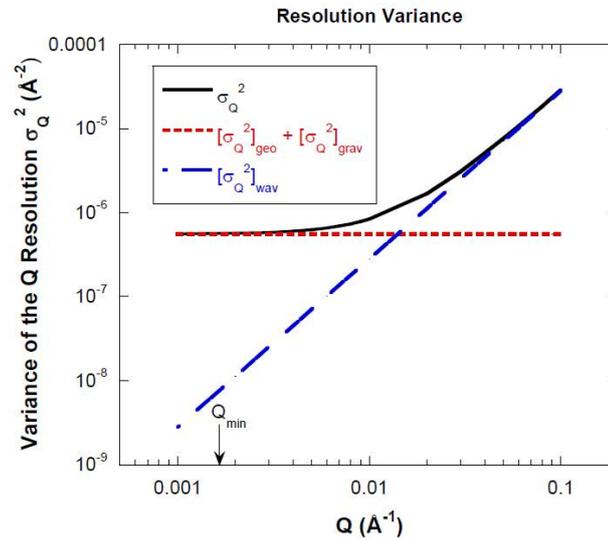
Fig. 7.22 Source: 'The SANS Toolbox'. The overall q-resolution (black line) is determined by the quadratic sum of the component related to the spread in incoming wave lengths (blue line), and by the collimation and effects of gravity (red line). This figure is relevant to the SANS spectrometer located in the cold source guide hall at NIST.

of a liquid than that of a solid: we do not expect any sharp features in q-space and, by extension, there is no need to investigate these broad features employing a very sharp resolution function.

7.3 Exercises

The SANS toolbox comes with many exercises, all of which are strongly recommended. We have borrowed and adapted a few of these problems.

Exercise 7.1

SANS instruments use velocity selector for selecting the incident energy range rather than monochromators. Not only results this in higher flux, but it has an additional advantage. What is it?

Exercise 7.2

Normally we use a standard powder (such as Si) on diffractometers to calibrate the incident neutron wave length. Why can this method not be employed on SANS instruments?
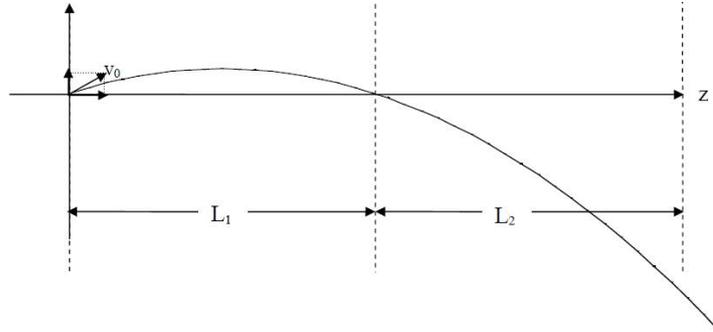
Fig. 7.23 Source: 'The SANS Toolbox'. Gravity affects the vertical q-resolution since the neutrons will have to follow a parabolic trajectory in order to be able to make it through the aperture. The actual trajectory and its vertical fall distance depend on the wave length of the neutron. Note that the vertical distances have been greatly exaggerated this this sketch.

Exercise 7.3

The radius of gyration is determined from the low-q behavior of the $q=0$ peak. How would instrumental resolution affect the determination of the value of the radius of gyrations, if at all?

Exercise 7.4

When doing powder diffraction, we see that the peak shapes at low momentum transfers become asymmetric (Fig. 6.19). Can we expect similar effects when we use a 2-dimensional SANS detector?

Exercise 7.5

Describe the set of experiments that are needed to end up with a dataset free from all instrumental and background effects.

Exercise 7.6

The form factor shown in Fig. 7.11 displays many characteristic q-regions, resulting from characteristic length scales of the cigar-shaped particles. Explain what aspect(s) of the cylinders we are probing, and why we are probing these aspects in the various q-regions. What determines, conceptually speaking, the q-values where we transition from one region to another, such as the Q_1 and Q_2 that are shown in the plot?

8
Reflectometers

In neutron reflectometery, we graze the surface of a material with a beam of neutrons, and we look at the ones that are specularly reflected, that is, the ones that appear to bounce off the surface as if it were a mirror. The intensity of the reflected beam as a function of grazing angle is determined by the structure of the top layers of the material; typically we use this technique to learn about these top layers in materials, such as membranes and semi-conductors.

Neutron reflectometry is still based on the same principles as neutron diffraction; namely, that the neutrons act as waves when interacting with the material. In fact, there is a very close connection between neutron reflectometry and similar phenomena in optics. The main difference with neutron scattering experiments described so far is that we no longer work in the approximation that a neutron is scattered only once. In fact, we have to focus on the opposite extreme to understand what we measure in reflectometry; the extreme being that multiple scattering is the norm and we treat the material as if it is a continuous medium (as opposed to a collection of atoms).

We show a typical reflectometry setup in Fig. 8.1. The angle of incidence of the neutron is identical to the angle of reflection and, as is clear from this figure, the neutron transfers momentum in the direction perpendicular to the surface. Therefore, all we can learn from such experiments is the composition as a function of distance perpendicular to the surface. In essence, we get one-dimensional information from our experiments. This restriction makes it difficult to analyze reflectometry data without ambiguity, the so-called phase problem.

We wish to investigate the layers near the surface of our sample. In general, these layers will be many atoms thick and, therefore, our probing wavelength should be large. This in turn implies that our angle of reflection must be small (Fig. 2.3). And as a consequence, reflectometers tend to be large instruments where a lot of effort has gone into ensuring that the beam is very narrow and very tightly collimated. We show the reflectometer at HFIR in Fig. 8.2 before it was located in the cold source guide hall. And we note that whenever we have to use extreme collimation, we quickly run out of intensity.

For very small reflection angles, we can find 100% reflectivity. This number quickly drops off when we increase the amount of momentum q transferred to the sample. In fact, it drops off as $\sim 1/q^4$. With neutrons, we can measure reflectivities down to
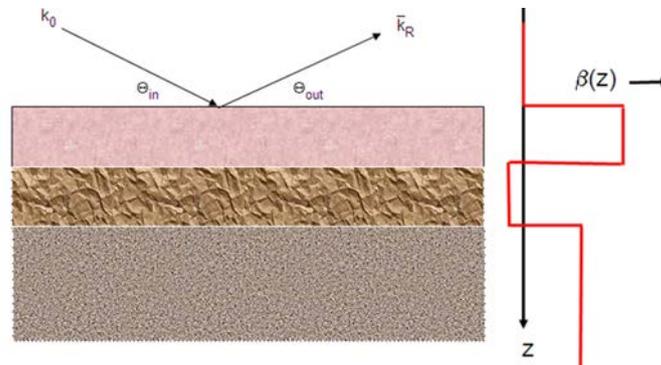
Fig. 8.1 A neutron 'hits' the surface of a material, and either gets reflected (shown), or transmitted (not shown). When the angle of incidence equals the angle of reflection, we call this specular reflection. What fraction of the neutrons is being reflected depends on the angle, the wavelength of the neutrons, and the structure of the material as a function of distance z from the surface. What we can learn from such experiments is the scattering length density profile $\beta(z)$ (to be explained in the next section), shown in red next to the layered sample.

$10^{-5} - 10^{-7}$, so roughly 1 neutron out of every million of neutrons that graze the sample can still be detected as having been reflected. Smaller reflectivities are very difficult to achieve because the incident neutron beam is intensity limited because of the requirements on collimation, and when we push our signal to the limit of around 1 neutron per minute we will be measuring more stray neutrons (background neutrons) than neutrons reflected off of our sample.

Most reflectometers reflect neutrons from a vertical surface, but some can do reflection from horizontal surfaces. The latter types are suited to studying the surface layers of liquids. Reflectometers at reactor sources operate using a constant incident wave length while varying the angle of the reflected neutrons to probe the reflectivity profile as a function of momentum transfer. When the angle of the reflected neutrons is varied, the sample must be rotated at the same time to maintain the specular condition. Reflectometers at pulsed sources have a detector fixed at a specific angle, while allowing neutrons of many wavelengths to hit the sample; this way we can probe the sample as a function of momentum transfer while keeping the detector in place and without having to rotate the sample. Note that varying the incident wavelength is the only way to probe the surface of liquids.

Pulsed neutron sources can normally compensate for their lack of time integrated flux by utilizing more detectors simultaneously. For reflectometry this is not an option since there is only one detector position for which we have that the angle of incidence with respect of the sample surface equals the angle of reflection. All other events are referred to as off-specular reflection, and while these events do carry information about the sample roughness and lateral homogeneity (as in this case momentum is also trans-
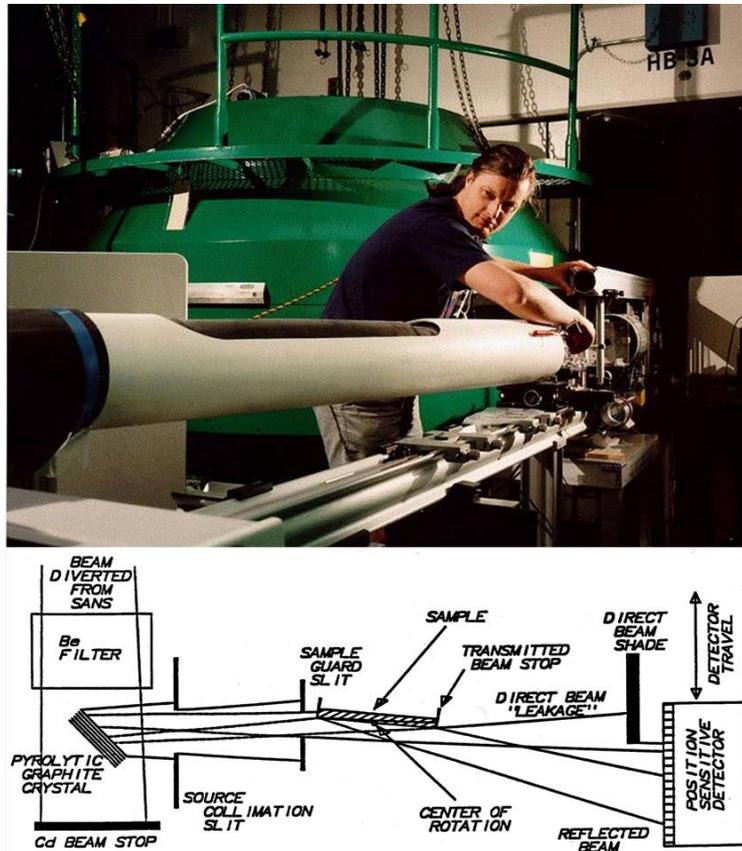
Fig. 8.2 Reflectometers are straightforward instruments that are designed to take data very close (in angle) to the straight through neutron beam. This is achieved by using very narrow slits to shape the beam, and by ensuring that the distance between sample and detector is large. The top panel shows the reflectometer on the beam port floor of the HFIR, the bottom panel shows the schematic of this instrument. Drawing courtesy of Bill Hamilton (pictured).

ferred parallel to the surface of the sample), this type of information tends to be not as useful as that contained in specular reflection. In this chapter, will will only deal with specular reflection except for one figure where we show what off-specular reflection looks like.

The reflectivity depends on the average scattering length density $\beta(z)$ (to be defined in the next section), where z is the distance to the surface. Whenever there is a change in β, we will see a corresponding change in the reflectivity. This is similar to optics where we see reflection and diffraction at interfaces. It is this that makes reflectometry well suited to locating interfaces when hidden under the surface. But there

is more. β is the product of the density of the atoms and their *scattering lengths*. So it serves as a direct measure of the scattering length, not of the scattering length squared. This implies that $\beta(z)$ can be positive or negative. But it also implies that we can do reflection measurements as a function of the magnetic moment direction of magnetic ions within the sample. Whenever we change the angle between magnetic moment and that of the neutron spin, we end up with a different scattering length. This makes reflectometery a very sensitive technique for studying layered structures that contain magnetic ions, such as magnetic semiconductors.

As a final word in this section, since reflectometry depends on the scattering length, it can be used to actually measure the scattering lengths. This is how a large fraction of the scattering lengths listed in Appendix B have been determined. All details about reflectivity can be found in the treatise 'Neutron Optics' by Varley Sears (see our reference in the introductory chapter).

8.1   The Skinny

In order to model the reflectivity profile (the fraction of neutrons that are reflected as a function of momentum transferred to the sample), we use the continuum approximation for the material we are studying. This is an entirely reasonable approximation since the regions we are probing contain thousands of atoms in the planes perpendicular to the direction of momentum transfer. Therefore, we expect to only be sensitive to the average neutron scattering power in the layers. The neutron scattering power, a word we just made up, depends on how many atoms there are in a layer [denoted by the number density $n(\vec{r})$], and what their scattering lengths are. This product is actually referred to as the scattering length density (see also the SANS terminology in Chapter 7) and it is defined for atoms in a region near $\vec{r}$ by

$$\beta(\vec{r}) = 4\pi n(\vec{r}) b_{\text{coh}}(\vec{r}).$$

When we probe a material in a reflectivity experiment, we (indirectly) measure the structure of the material as a function of z, the distance to the surface. This implies that we will only be sensitive to the scattering length density $\beta(z)$, which is obtained from $\beta(\vec{r})$ by averaging over the perpendicular directions. The aim of a reflectivity experiment is to determine $\beta(z)$.

In order to predict what the reflectivity of a material will be we have to solve the Schrödinger equation where the scattering length density will play the role of potential: $V(z) = \hbar^2 \beta(z)/2m_{\text{neutron}}$. We do this for some cases in the exercises at the end of this chapter, but in here we simply will take our cue from optics, and from waving our hands. We can get away with this based on the one-to-one correspondence between the Maxwell equations that govern all optical properties, and the one-dimensional Schrödinger equation. This correspondence can be expressed by introducing a refractive index n for the neutron:

$$n = \sqrt{(E - V)/E}, \tag{8.1}$$

with E the kinetic energy of the incident neutron. Since the potential V can be positive or negative, we can have an index of reflection that is smaller or larger than 1. This

for instance implies that a neutron can speed up when it enters a material.

Bearing this correspondence in mind, we look at three general examples: an interface, a single layer on a substrate, and a multi-layer assembly on a substrate. To understand these examples better, we take another look at Fig. 8.1. Outside of the material, the neutron has a certain energy E [with $E = \hbar^2 k_0^2/2m = \hbar^2/2m(k_{0,x}^2 + k_{0,y}^2 + k_{0,z}^2)$]. For the purpose of understanding reflectivity, we are only interested in the amount of energy associated with the neutron's velocity in the z-direction, perpendicular to the surface of the sample. The other components will not be affected during the reflectivity process, so we will ignore them.

Outside of the material, the component of the neutron's wave vector in the z-direction $k_{0,z}$ is given by $k_{0,z} = k_0 \sin \theta_{\text{in}}$, and $E_z = \hbar^2 k_{0,z}^2/2m$. The neutron's wavelength associated with this $q(z)$ is $\lambda(z) = 2\pi/q(z)$. When the neutron enters the material, its available energy will change to $E_z - V(z)$; thereby either shortening or elongating its wave length according to

$$\lambda(z)^2 = \frac{h^2(E_z - V(z))}{2m}; q(z) = \frac{2\pi}{\lambda(z)}. \tag{8.2}$$

This is shown in Fig. 8.3

When we solve the reflectivity problem mathematically by finding a solution to Schrödinger's equation, what we do graphically is to make sure that the wave length of the neutron matches the wave length given by the available energy

$$E_z - V(z) = \frac{\hbar^2}{2m}[q^2(z) - \beta(z)]. \tag{8.3}$$

We do this matching in every region, and we make sure that the waves link up smoothly between the different regions. With smoothly we mean that both the waves and their slopes match at the interface where they link up. This is physics' speak for saying that the second derivative can be discontinuous, implying that we have a step function in the available energy.

So let's focus our attention on an interface, such as the one shown in Fig. 8.4. If the interface represents a 'step up' in $\beta(z)$, which would be the case for a material with a positive scattering length density, then we can expect 100% of the neutrons to be reflected for very small incident angles. At the air-material interface we have $q_{air}(z) \equiv q_0 = k_0 \sin \theta_{\text{in}}$; this implies that we can end up with a situation where $q_0^2 - \beta(z)$ is smaller than zero (eqn 8.3), in turn implying that neutron waves such as the ones shown in Fig. 8.3 cannot exist in the material (eqn 8.2). The only option neutrons have in that case is to stay out of the material, also known as reflecting off of the surface of it. We can compare this to a ball having a certain amount of kinetic energy, but not enough to roll up a hill. As a result, the ball will roll back down.
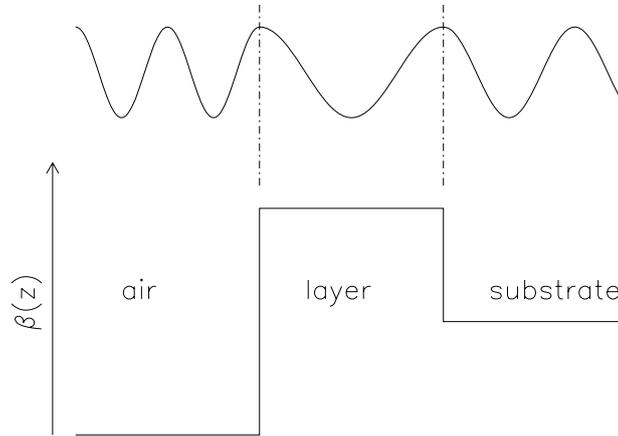
Fig. 8.3 The probing wavelength of the neutron in the material is determined by the kinetic energy of the neutron, the angle of incidence, and by the scattering length density $\beta(z)$ according to eqn 8.2. The probing wavelength in the different types of material that make up the sample is sketched at the top of the figure. Solving the 1-dimensional Schrödinger equation involves making sure that the wave function in the various regions links up with the neighboring regions in a smooth fashion: the wave function has to be continuous, and its derivative has to be continuous. This is sketched for the wave that is drawn here.

We will find that 100% of the neutrons will be reflected off of the material provided that $q_0^2 - \beta(z) < 0$, or up to a critical angle of incidence θ_c given by

$$k_0 \sin \theta_c = \sqrt{\beta(z)}. \tag{8.4}$$

Critical angles are typically very small, but note that they are dependent on the wave length of the neutron: the longer the wavelength, the larger the critical angle. This fact is used in neutron guide tubes, where slow (long wave length) neutrons are redirected over a few degrees, so they can be separated from faster neutrons (which will be transmitted into the material of the guide tube) that would cause a lot of background problems in scattering experiments. For this reason, cold neutron sources have their instruments at the end of guide tubes. Using these optical-like properties of the neutron is a field by itself, referred to as 'neutron-optics' (and written up by Varley Sears).

When the scattering length density of the material is negative, then we will never get total reflection. This follows rightaway from our previous statements since $q_0^2 - \beta(z)$ will always be larger than zero. In the general case, as soon as $q_0^2 - \beta(z) > 0$, then we will have that a large fraction of the neutrons will not be reflected. How many are reflected- the so-called reflectivity R- can be calculated (Exercise 8.1) to be:

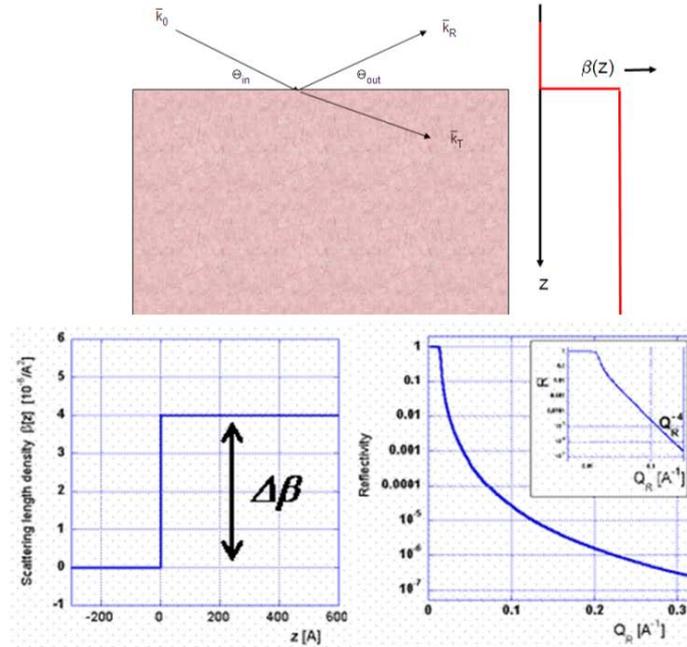$$R = |\frac{q_0 - q_1}{q_0 + q_1}|^2, \tag{8.5}$$

Fig. 8.4 An interface between two materials- such as air and a material with a positive scattering length shown here- is marked by a sudden change in scattering length density (top panel). For small angles of incidence we see that all neutrons are reflected, while the reflectivity rapidly drops as the angle of incidence is increased beyond the critical angle and most neutrons will be transmitted. In this figure $\vec{k}_0$ stands for the incident neutron, $\vec{k}_R$ for the reflected neutron, and $\vec{k}_T$ for the transmitted neutron. The graph on the lower right shows- on a log scale- the expected reflectivity (fraction of the neutrons being reflected under specular conditions) for the scattering length density shown in the lower left panel. Note that we have used the symbols k to indicate the full momentum vector of the neutron, not to be confused with q. Bottom drawings courtesy of Bill Hamilton.

where $q_0 = k_0 \sin\theta$ as defined previously, and q_1 is defined within the material through $q_1^2 = q_0^2 - \beta(z)$. So we use $q(z)$ to specify the general case, and q_0, q_1, q_2 etc. to denote the case within specific layers of the material. This definition ensures that q_1 gives us the neutron wavelength within the material: $\lambda(z) = 2\pi/q(z) = 2\pi/q_1$ (eqn 8.2). The solution of eqn 8.5 is shown in Fig. 8.4. Also, from eqn 8.5 we can see that for large momentum transfers $[q_0^2 \gg \beta(z); q_0 \approx q_1]$ that the reflectivity falls off as $\sim 1/q_0^4$ (see inset Fig. 8.4):

$$R = [\frac{q_0 - q_1}{q_0 + q_1}]^2 [\frac{q_0 + q_1}{q_0 + q_1}]^2 = [\frac{(q_0^2 - q_1^2)}{(q_0 + q_1)^2}]^2 = \frac{\beta(z)^2}{16q_0^4}. \tag{8.6}$$

The curve shown in Fig. 8.4 makes sense. We have a range of incident grazing angles that see all neutrons reflected, followed by an increased transmission coefficient.
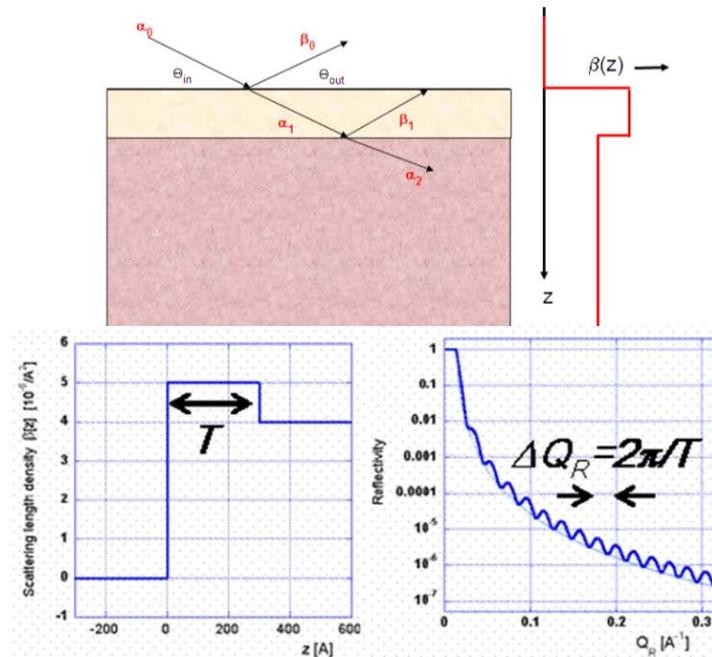
Fig. 8.5 A single layer on top of a substrate provides two changes (interfaces) in the scattering length density profile. In order to mathematically solve this, we have to take into account neutrons that are traveling up and down in every layer, and ensure that the wave function of the neutron (and its derivative) is continuous at each interface. The result is that we see the same reflectivity profile as we have for a single interface, but now on top of it we find interference fringes whose period is given by the reciprocal of the thickness of the single layer, as shown in the panel on the lower right. Bottom drawings courtesy of Bill Hamilton.

Anyone who has ever looked at light reflected from a glass pane will not be surprised by this curve. We now take it one step further, the case where we have a layer on top of a substrate as shown in Fig. 8.5.

When we put a layer on top of the substrate, we can still expect the same behavior we already discussed for the interface: total reflection for small angles and a drop-off in reflectivity proportional to $1/q_0^4$ for large momentum transfers. But, we will also see the effects of this layer. We now have the equivalent situation of a thin film in optics, such as a soap bubble in air or a layer of oil floating on water. When we look at these thin films we see various colors being reflected; the colors that we see are determined by the thickness of the film. Conversely, we can calculate the thickness of the film by just looking at the colors, provided we know the refractive index of the film (that is, what is the film made of).

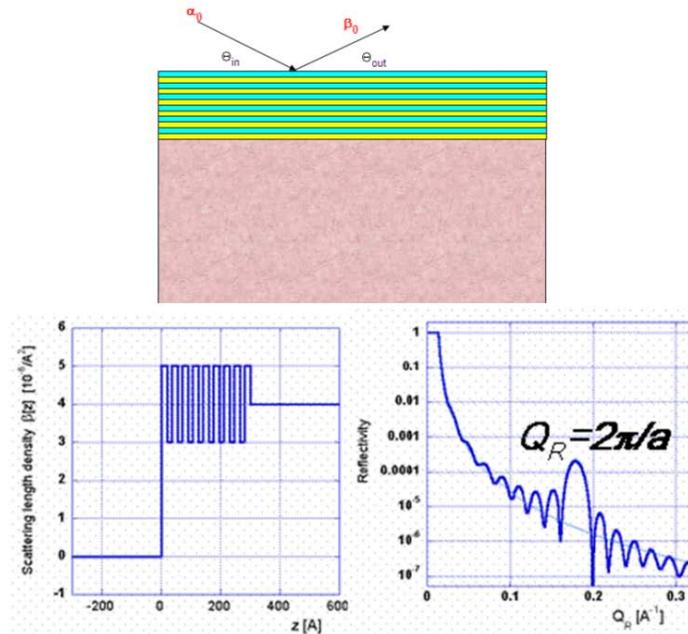What if we neither know what the film is made of, nor how thick it is. Can we still

Fig. 8.6 In case of multi-layers on top of a substrate we can still distinguish the reflectivity profile of the substrate, but now we also see oscillations with the period determined by the overall width of the multi-layer assembly, as well as the onset of a Bragg peak at $q_r = 2\pi/a$ determined by the distance between adjacent multi-layers. Bottom drawings courtesy of Bill Hamilton.

look at the colors that are reflected to determine the film's thickness and make-up? Using neutron reflectometery we can. The thickness of the film (layer) will create oscillations in the reflectivity profile while the make-up of the film (layer) will determine the amplitude of the oscillations, the equivalent of the brightness of the colors. The period of the oscillations is given by $\Delta q_R = 2\pi/T$, with T the thickness of the film. This is shown in Fig. 8.5. These oscillations will be superimposed on top of the reflectivity profile of the substrate (shown in Fig. 8.4).

The substrate plays an important role in how visible the oscillations are. When we do the math, we find that the peaks of the oscillations follow an intensity envelop given by $\beta_1^2/16q_0^4$ while the troughs follow $(\beta_1 - \beta_2/2)^2/16q_0^4$. Thus, when $\beta_1 = \beta_2/2$ we can expect some very pronounced oscillations. This is very helpful when convincing the referee of the trustworthiness of one's data. Conversely, when $\beta_2 = 0$, an unsupported interface, then the oscillations will not be very pronounced.

The next step up in complexity would be a multi-layer assembly on top of a substrate. We now have two characteristic length scales, namely the overall thickness T of the multi-layer, and the spacings a between the individual layers that make up the

multi-layer. A typical reflectivity profile is shown in Fig. 8.6. We can see that the separation between the layers shows up as increased reflectivity at $q_R = 2\pi/a$. These spots of increased reflectivity are sometimes referred to as Bragg peaks. While this nomenclature is not entirely correct, it makes sense as we expect this peak to become increasingly more pronounced when we increase the number of layers. This is entirely equivalent to what happens in liquids and solids: the more ordering (structure) there is, the sharper the features in the resulting interference pattern.

The make-up of the individual layers will be reflected in the amplitude of the oscillations, but this information is now hidden a bit deeper in the reflectivity profile. This is where reflectivity profiles become more and more difficult to interpret. The reason for this can be seen in Fig. 8.6. In order to be able to follow the amplitude of these oscillations, we need to be able to measure very small reflectivities R, preferably down to (and beyond) 1 neutron in a million ($R = 10^{-6}$). Here we quickly run into experimental limitations caused by background scattering. The best we can do is to come up with a model of how the scattering length density varies across the layers, and fit this to the data. Sometimes we find that we cannot distinguish between similar models because we do not have reflectivity data down to small enough reflectivities, or because our surfaces and interfaces are not smooth enough. These are the intrinsic limitations of the reflectivity technique.
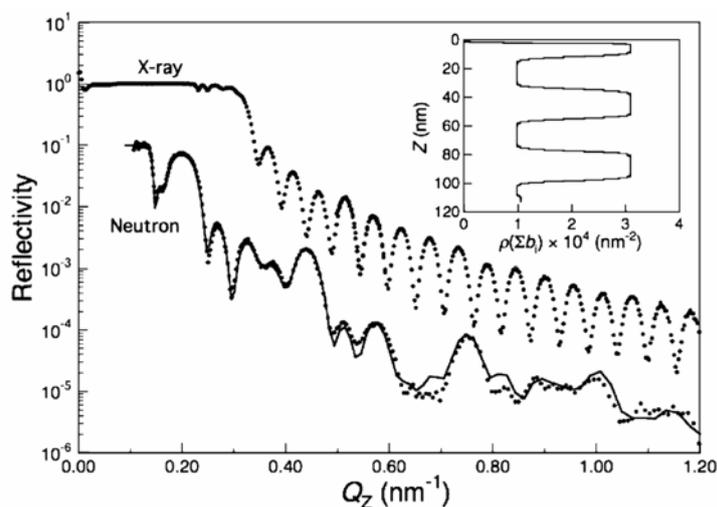

Fig. 8.7 Neutron reflectivity reveals, through modeling, the make-up and thickness of a multi-layered structure. Note the difference in level of sensitivity between neutron and xray curves. The neutron data has been offset by a decade. The layered model shown in the inset yields the solid curve drawn through the neutron scattering data. This figure has been reproduced with permission from Macmillan Publishers Ltd: Polymer Journal (M. Torikai *et al.*, Polymer Journal 39, 1238 (2007)), copyright 2007.
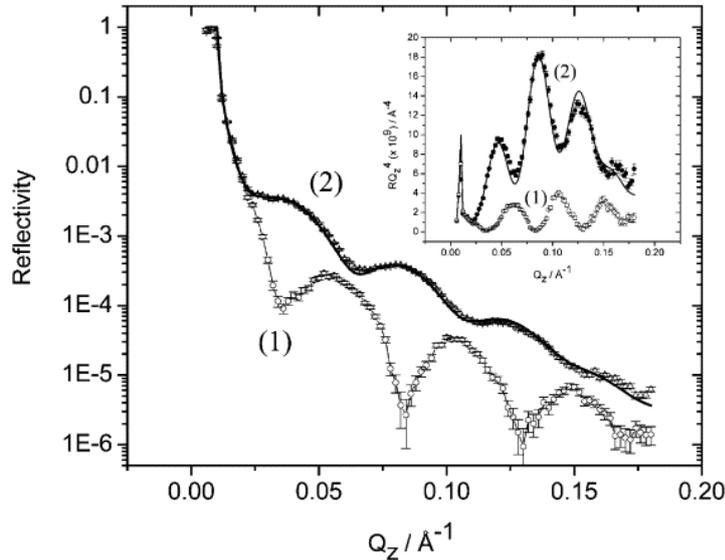
Fig. 8.8 Figure reprinted from I. Burgess *et al.*, Biophysical Journal 86, 1763-1776, copyright 2004, with permission from Elsevier. The curve labeled (1) shows the reflectivity profile for a gold coated block of quartz. The contrast was optimized by performing the reflectivity experiments when the surface was immersed in D_2O. The curve labeled (2) shows the result when the gold surface is covered by a bilayer of the 7:3 mixture of h-DMPC and cholesterol. The inset shows the same results, but now the reflectivity curves have been multiplied by q_z^4 so that the oscillations become better visible. The solid lines show the fitted reflectivity curves.

Before we discuss the components of the reflectometer, we show some examples of the types of experiments that have been performed and what has been learned about the materials that were being studied. We show an example of the reflectivity profile of a multi-layered sample in Fig. 8.7. The agreement between the data and the calculated reflectivity profile based on the model is very good. In practice, to obtain these models one changes the scattering length density over a small range of z values; if this change improves the agreement between the model and the data, then this change is effectuated. If not, another change is tried out. This process is repeated until the fit is as good as the data allows.

Reflectivity experiments are well suited to probing the structure of membranes, making this technique (potentially) very useful for studying biological samples. An example of this is shown in Fig. 8.8 where we can observe the change in reflectivity when a mixture of h-DPMC and cholesterol is deposited on a gold surface. The purpose of this experiment was to study the influence of a substrate charge on the structure and hydration of a bilayer lipid membrane.

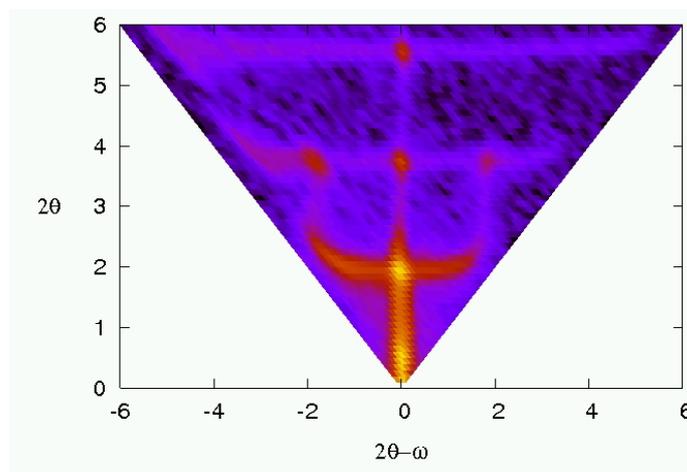We have largely ignored off-specular reflection, that is, reflection events where

Fig. 8.9 Figure reproduced from the PSI website (www.psi.ch). Shown is a contour plot of the reflectivity by a Ni/Ti multilayer. Specular reflection correspond to the vertical line located at $\omega=0$. All scattering not on this line (off-specular) is caused by the roughness of interfaces and inhomogeneities in planes parallel to the surface. Reflectivity points out these 'defects' without mercy.

the neutrons also transfer momentum in the lateral direction. Off-specular reflection happens when either the surface in not exactly perpendicular to the nominal direction of momentum transfer (when we have interface roughness), or when we have a non-uniform make-up of the sample in the planes perpendicular to the z-direction. Therefore, off-specular reflection tells us about surface roughness, or about these lateral variations. As an example we show the full reflectivity profile (specular and off-specular) of a less than perfectly smooth multilayered sample in Fig. 8.9. As can be seen in this figure, interface roughness gives rise to significant (and measurable) off-specular intensity in this Ni/Ti multilayer assembly. Currently, we do not exactly know how to translate the presence of off-specular reflectivity into quantitative statements about our sample, we can only draw qualitative conclusions.

8.2   Components

Reflectometers are straightforward instruments. What is needed, at a reactor source, is a monochromator to select the incident wavelength of the neutrons, a whole bunch of very narrow slits to shape the beam into a very narrow one that can be used to deliver the neutrons onto the sample under grazing angles, a detector, and a means of moving the detector around while rotating the sample. We show the MURR reflectometer in Fig. 8.10.

For a monochromator on a reflectometer, we like to use one with high reflectivity because we would like to have as many neutrons as possible. Monochromators do not have to be horizontally focusing since we will narrow down our beam considerably,
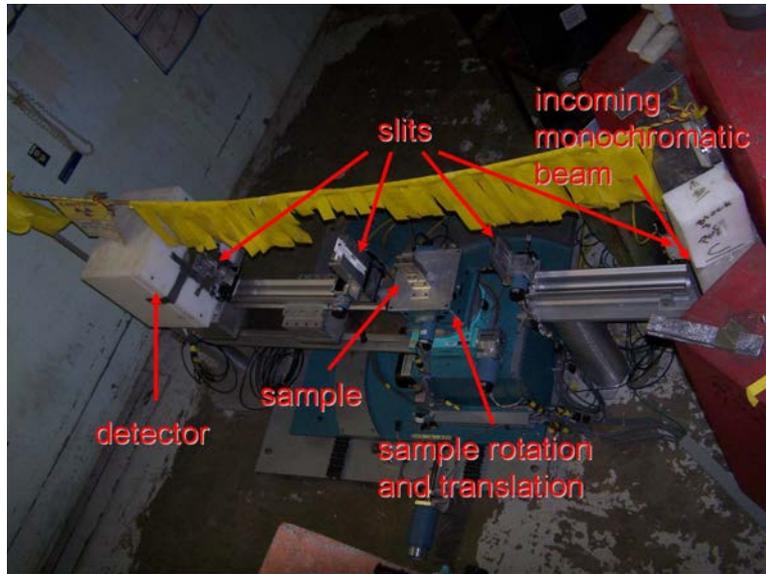
Fig. 8.10 The MURR reflectometer uses a graphite monochromator and slits to deliver a monochromatic beam of neutrons onto a sample that is located at the center of this photo. The detector housing is located on the left side of the photo. The two optical benches (aluminum rails on either side of the sample) are used to position the polarizing mirrors (not shown) and/or additional slits and shielding.

and we are mostly interested in getting a beam with very little angular divergence. As such, a flat graphite monochromator is a natural choice, or a monochromator with slight vertical focussing such as the one shown in Fig. 8.11 that is in use at the MURR reflectometer. There is not much more to be said about the monochromator.

We use very narrow slits to obtain a highly collimated and narrow beam. Typical beam widths at the sample position are of the order of 1 mm wide. A set of slits after the monochromator and another set of slits before the sample position shape this narrow beam, and also cut down on the straight through beam which would interfere with our ability to measure the reflectivity at very small angles. After the sample position we use additional slits, all with the aim of cutting down on background scattering and maintaining our angular resolution so that we know the scattering angle to high accuracy.

Aligning a reflectometer so that the sample is at just the right position, and so that we know exactly where the beam hits the sample and what the scattering angle is, is a time-consuming process. We detail the alignment procedure in Appendix G. Once the spectrometer is aligned, the measurements are straightforward. We have neutrons hit the sample, and we measure the reflectivity profile by rotating the sample over an angle θ_{in} and the detector over an angle $2\theta_{in}$ (with neutron deflection angle θ equal to

Fig. 8.11 The monochromator of a reflectometer is flat or vertically focussing, such as the black looking stack of four graphite crystals shown here. The monochromator is oriented with the aid of several motors shown in this photo.

$2\theta_{\text{in}}$) so that the condition for specular reflection, $\theta_{\text{in}} = \theta_{\text{out}}$ (Fig. 8.1) is always met.

In order to get down to very low reflectivities ($R < 10^{-5}$) we must do accurate background measurements. A typical experiment can consist of measuring the reflectivity of the sample, and repeating this scan without anything on the sample table (to get an idea of the sample independent background caused by the straight through beam and low flying neutrons from elsewhere). If we are interested in layers on a substrate, it should come as no surprise that we will measure the reflectivity of the substrate without layers on it under identical conditions to aid us with the data analysis, in particular with the background correction.

Reducing the number of background counts is essential to our ability to get down to very low reflectivities. For instance, take the case where the background level is 1 neutron per second. Assuming that only 1 neutron in 1 million is reflected (that is, we are interested in $R=10^{-6}$), and we have 1 million neutrons per second incident on the sample after we have narrowed down the beam with our slits, then the signal from our sample and that of the background are of the same order of magnitude, and we are close to reaching the limits of what we can measure on our instrument. Should our

background be 10 counts per second rather than 1 count per second, then we will lose one order in reflectivity sensitivity, which can be detrimental to our ability of measuring the scattering length density profile of our sample. In most cases, the background originates in neutrons being scattered by the sample holder and by other bits of the spectrometer. The only way of eliminating unwanted background is to improve the shielding, a process that normally is done by trial and error.

When we do polarized reflectivity measurements, we also keep track of whether the neutron's intrinsic spin is aligned parallel or anti-parallel to the magnetic field the sample is exposed to, and we keep track of whether the neutron changed its spin orientation in the reflection process, or not. In order to be able to accomplish this, we use polarizing mirrors. A polarizing mirror reflects a neutron of one spin orientation, but transmits neutrons that have the other spin orientation. Thus, we can use these mirrors to polarize the neutron beam: the reflected (off the mirror) beam will only contain neutrons that all have the identical orientation of their intrinsic spin. We place one mirror between the monochromator and the sample to create the polarized beam, and one mirror between the sample and the detector to ensure that we only detect neutrons of one particular spin orientation. Note that we could also reflect out the unwanted spin orientation, and instead use the transmitted beam for doing our experiment.

The setup of a reflectometer with polarizing mirrors is more complicated, because we now have to deal with additional deflection angles at the polarizing mirrors, or we have to deal with transmitted beams that come with more background attached caused by the scattering of the mirror holder. In addition, we have to set up guide fields in order to maintain the neutron's spin orientation. However, this setup process is manageable and it is much better figured out how to do this setup while actually standing at the instrument than to read about it in this paragraph.

There is one additional piece of equipment that comes with the usage of polarizing mirrors, and that is a device called a spin-flipper. This device is aptly named as it does exactly that, it flips the spin orientation of the neutron by 180 degrees. At least, that is the ideal situation. In practice, the device is not perfect and not all of the neutrons get flipped. The effectiveness of the device is denoted by its flipping ratio, which is the ratio of the number of successful flips over the number of failed flips. A flipping ratio of 30 is fabulous, a flipping ratio below 5 is impossible to work with. Optimizing the flipping ratio is time well spent as the signal of interest in an experiment will be based on the difference between two measurements. Insufficient control over the orientation of the neutron spin will significantly affect this difference. We show an example of a polarized experiment in Fig. 8.12.

The final piece of advice in this chapter deals with the footprint correction. In our preceding discussion, we swept one item under the rug that interferes with the data analysis of reflectivity experiments at reactor sources. When we move our detector to a different scattering angle, we must also rotate our sample to ensure that the condi-
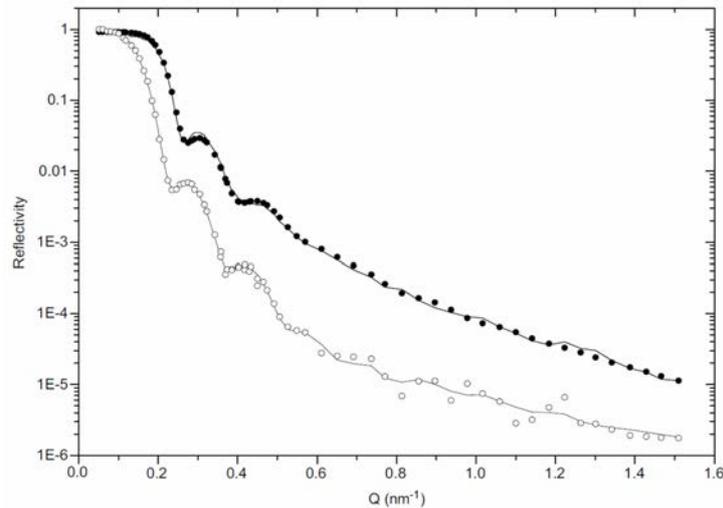
Fig. 8.12 Figure reproduced with permission from R. Stephan *et al.*, Journal of Magnetism and Magnetic Materials 320, 3378 (2008). The difference in reflection between spin-up and spin-down polarized neutrons when reflected off of a magnetized Fe-layer deposited on a Si substrate yields information about the size of the overall iron magnetic moment as a function of substrate and deposition technique.

tion of specular reflection is maintained. This changes the area of the sample that is illuminated by the very narrow neutron beam. This area is referred to as the footprint of the beam, and we need to correct for it so that we get a reflectivity curve that is based on a fixed number of incident neutrons per surface area of the sample.

The surface area over which the neutron beam is spread out is proportional to $1/\sin(\theta_{in})$, which in turn will be proportional to $1/q_{0,z}$. By simply putting this factor back into the reflectivity profile we correct for the changing footprint of the beam. At least, almost. For very small angles of reflection we can have that the beam is wider than the sample, so we have to put in the appropriate cut-off in our correction. For very narrow beams we do not have to worry about this. In fact, for wide beams that go from being wider than the sample to smaller than the sample, we actually have to take into account how the beam falls off in intensity away from the center of the beam. Such footprint corrections tend to involve the error function.

The bottom line is that footprint corrections will have to be carried out before one can observe the flat bit of the reflectivity curve corresponding to total reflection. Once this flat bit has been identified, one takes this level to normalize the reflectivity profile: $R = 1$.

8.3   Exercises

Exercise 8.1

The reflectivity R for a single air-solid interface is given by

$$R = |\frac{q_0 - q_1}{q_0 + q_1}|^2,$$

with q_0 the component of the neutron's q-vector perpendicular to the surface of the interface, while q_1 is given by $q_1^2 = q_0^2 - \beta$.

a) In order to calculate the reflectivity, we have to solve the Schrödinger equation in the direction perpendicular to the surface. We do this by writing down the general solution in air and in the solid, followed by a matching of the wave function and its derivative at the interface. Why do we do this matching? And, why only for the wave function and its derivative, and not for the second derivative?

b) Derive the above equation by finding the solution to the one-dimensional Schrödinger equation. Use $\Psi(z) = \alpha_0 e^{-iq_0 z} + \beta_0 e^{iq_0 z}$ as the general solution for the neutron in air, and $\Psi(z) = \alpha_1 e^{-iq_1 z}$ as the general solution in the solid. The reflectivity is given by $R = |\beta_0/\alpha_0|^2$. Also, why is there no term $\sim \beta_1 e^{iq_1 z}$?

c) Show that the above equation yields a reflectivity $R = 1$ for $q_0^2 - \beta < 0$.

Exercise 8.2

a) Show that the reflectivity R of a single layer on top of a substrate (Fig. 8.5) is given by

$$R = \left| \frac{[1 - \frac{q_0}{q_2}] + i[\frac{q_0}{q_1} - \frac{q_1}{q_2}]\tan(q_1 T)}{[-1 - \frac{q_0}{q_2}] + i[\frac{q_0}{q_1} + \frac{q_1}{q_2}]\tan(q_1 T)} \right|^2 .$$

In this equation, T is the thickness of the layer and the label '0' refers to air, the label '1' to the layer, and the label '2' to the substrate.

b) Show that in the high q-limit the reflectivity oscillates between $R = \beta_1^2/16q_0^4$ and $R = (\beta_1 - \beta_2/2)^2/16q_0^4$.

Exercise 8.3

Take a look at the membrane shown in Fig. 4.13. Discuss whether this membrane can be studied successfully using neutron reflectometry. What can we learn (if anything much)? What would be some of the problems that we might encounter?

Exercise 8.4

For this exercise we make use of the applet that can be found on http://rkt.chem.ox.ac.uk/techniques/nrmain.html. This applet illustrates the principle of contrast matching (this exercise), and that of the 'phase problem' (next exercise). Open the neutron reflectivity applet, which is the third applet on the webpage.
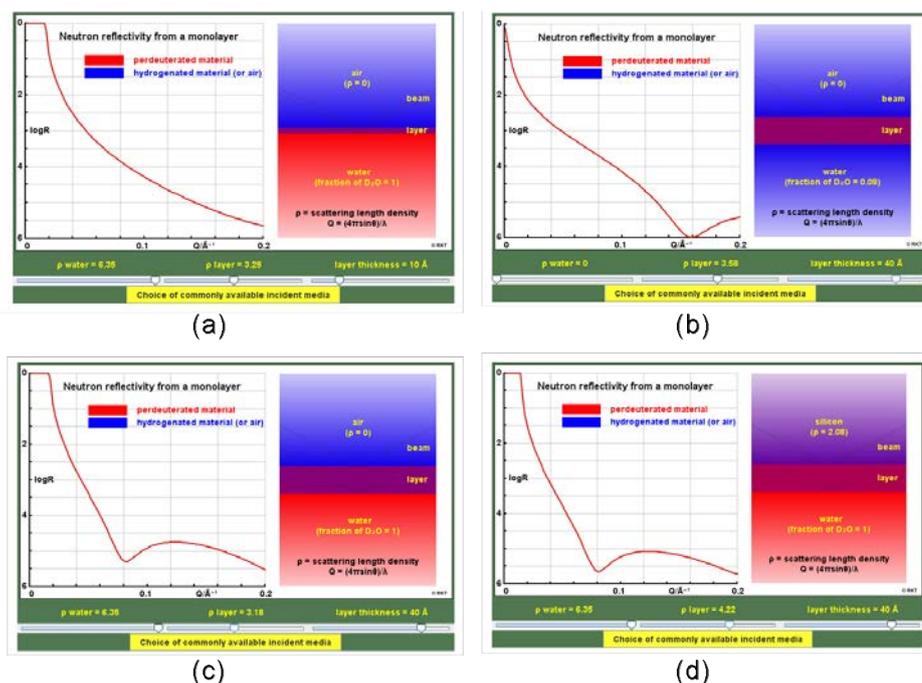

Fig. 8.13 Screen shots from the applet in this exercise. (a) The total reflectivity due to the air-D_2O interface. (b) When the air and the substrate are matched in scattering length density, we get very little overall reflectivity. However, this makes the profile sensitive to what happens in the layer, as can be seen by the minimum in reflectivity given by the thickness of the layer. (c) The position of the minimum in reflectivity occurs at half the q-value shown in panel (b) when we choose a suitable scattering length density of the substrate. (d) Changing air to silicon changes the scattering length density of the layer for which we find a minimum on the reflectivity profile, however, the position of the minimum is determined by the thickness of the layer.

Play around with the applet, and try the exercises that are written up on this web page, points (i) through (v). Note that what we refer to as β, the scattering length density, is referred to as ρ in this applet. Make sure you verify the following:

- When the scattering length density of the air and the substrate match (called non reflecting water NRW on this web page), then all reflectivity originates from the

thin layer. Play around with the thickness of the layer and its scattering length density. For a layer thickness T of 40Å, does the minimum in the reflectivity curve occur at π/T, or at $2\pi/T$? (Fig. 8.13b).

- When we change the substrate to D_2O for a given layer thickness, we see that the minimum in reflectivity shifts to half its original q-value (Fig. 8.13c). Why is this?

- Verify that the minimum is most pronounced when $\beta_1 = \beta_2/2$ (Fig. 8.13c).

- When we change the substrate to silicon (Fig. 8.13d), notice that the position of the minimum does not change. Also notice that the minimum in reflectivity now occurs for a scattering length density of the layer given by $\beta_1 - \beta_{\text{silicon}} = (\beta_2 - \beta_{\text{silicon}})/2$. Why is this?

Exercise 8.5

In this exercise we will use the applet referred to in the previous exercise. In the main text we mentioned that it is not possible to take the reflectivity profile, and determine without ambiguity the make up of the layered structure. For this reason, we normally come up with a model of the structure, and calculate the reflectivity profile of this model and compare it to the data. However, we cannot be sure that we always end up with the right model. This is referred to as the phase problem, and it is easy to demonstrate with the applet.
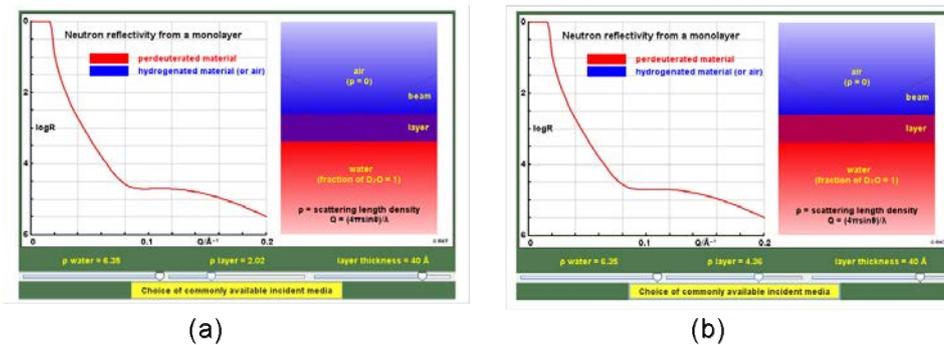

Fig. 8.14 Screen shots from the applet in this exercise. We can obtain identical reflectivity curves for two different values [$\beta_1 = 2.02$ in panel (a) and $\beta_1 = 4.36$ in panel (b)] of the scattering length density of the thin layer.

Set the substrate to D_2O, and pick a layer thickness of 40 Å. This setup will have a minimum in reflectivity at $q = \pi/40\mathring{A} = 0.08\mathring{A}^{-1}$ for $\beta_1 = \beta_2/2 = 3.18$. Play around with the scattering length density of the layer. For instance, set β_1 at 2.02 and 4.36 (Fig. 8.14) and notice that the reflectivity profile is identical. What we are seeing here is, simply put, that $(\beta_1 \pm \Delta\beta - \beta_2/2)^2$ does not depend on the $\pm$-sign when $\beta_1 = \beta_2/2$.

This is the 'phase problem' and it originates from the fact that we have to take the square of the scattering length densities at one point. Why does this hold true for the entire reflectivity curve, not just for the minimum in reflectivity?

Exercise 8.6

We can use the phenomenon of total reflectivity to transport neutrons over large distances in neutron guide tubes. As long as the angle of incidence of the neutron is small enough, we will get total reflection and we will be able to move neutrons away from the reactor into a guide hall. Very good neutron guide tubes consist of smooth surfaces (glass) with coatings that contain ^{58}Ni.

Fig. 8.15 A neutron guide tube at the ILL reactor (www.ill.eu). Notice that, unlike a glass fiber, the tube is almost straight.

a) The curvature of a guide tube is very small (Fig. 8.15), that is, guide tubes look almost like straight tubes. Why is this?

b) Neutron guide tubes are used to direct slow neutrons onto a sample, while faster neutrons (and gamma radiation from the reactor) do not make it down the guide tube. Why do faster neutrons not make it?

c) Some guide tubes use nickel coatings, whereas very good and expensive guide tubes use the isotope ^{58}Ni. What is so special about this isotope, other than that it is expensive?

Exercise 8.7

If we want to lock up a neutron in a tiny box made of smooth walled aluminum (so that we can study the lifetime of a neutron), what is the maximum speed a neutron

can have so that it cannot escape? Aluminum has a (face centered) cubic structure with unit cell length a=4.04 Å, and it has 4 atoms per unit cell.

Exercise 8.8

When we model the reflectivity profile, we change the scattering length density in a layer of thickness Δz in our model. Based on your knowledge of Fourier transforms, what is the minimum thickness for Δz that we can model and why can we not (in principle) adjust the scattering length density in a layer that is only one atom thick?

Exercise 8.9

In reflectivity experiments on membranes, one sometimes uses a bunch of membranes stacked on top of each other, with the entire sandwich placed on a substrate. Compare this method to doing reflectivity measurements on a single membrane placed on a substrate. What are the advantages and disadvantages?

9
Triple-Axis Spectrometers

The triple-axis spectrometer was invented by Bertram Brockhouse at Chalk River, Canada. For this, he was awarded the Nobel prize in 1995.

The triple-axis, or 3-axis spectrometer, allows the user to measure inelastic scattering events at constant momentum transfer. It derives its name from the three axes of rotation (indicated in Fig. 9.2) that can be used during the experiment. The advantage of a 3-axis spectrometer over a time-of-flight machine is that it is very flexible in setup because every angle can be varied, and allows for easy adjustment of sample parameters, such as tilt angles.

One of the disadvantages that the 3-axis spectrometer has over a time-of-flight instrument is that the background can vary because of varying degrees of shielding when the arms of the spectrometer are moved during the course of the experiment. Also, in general, the 3-axis spectrometer has only one detector. This is both an advantage and disadvantage. On the one hand, the resolution can easily be fine-tuned to the desired accuracy. On the other hand, the count rate can be low. As such, inelastic time-of-flight machines and triple-axis spectrometers are often complementary to each other: low dimensional magnets are best measured on a time-of-flight spectrometer,

Fig. 9.1 Bertram Brockhouse at the first 3-axis spectrometer at the E3 beam line at Chalk River Laboratories. The beam line is still operational. Since Brockhouse lived long enough to pick up his Nobel prize, we can safely assume that the beam line shutter was closed during this photo-op.
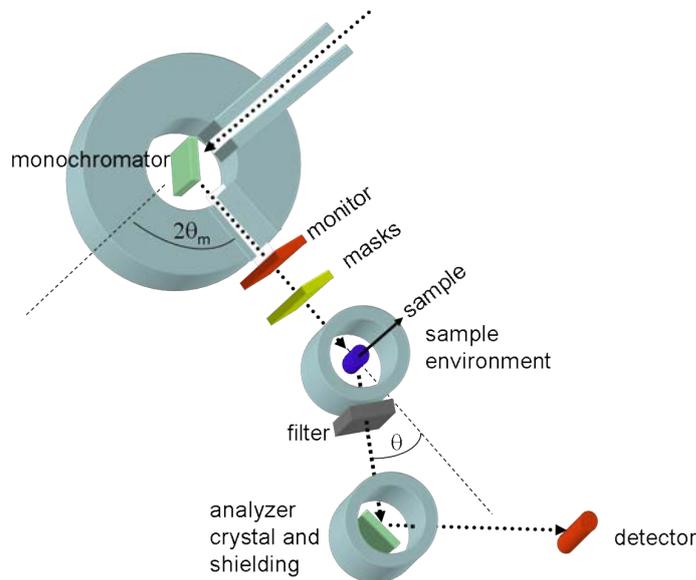
Fig. 9.2 A schematic representation of a 3-axis spectrometer. The monochromator crystal is placed on the first axis, the sample on the second axis, and the analyzer crystal on the third axis. The path of the neutron that is being detected is indicated by the purple line.

three dimensional systems are still best studied on a triple-axis, in most cases.

Actual 3-axis spectrometers are large instruments, as shown in Fig. 9.3. Their size reflects a number of practical considerations. First, neutrons easily penetrate materials and, therefore, one requires a considerable amount of material in order to shield the surroundings from unwanted neutrons and exposure to radiation. After all, neutrons are produced by drilling a hole in the side of a nuclear reactor. Second, the instruments have to be large enough so that they can house large cryostats, pressure cells, and other pieces of sample environments. Third, the neutron beams themselves tend to be fairly large, of the order of 2x4 inches. This is the only way that one can achieve a flux of neutrons that is large enough to complete a scattering experiment within a reasonable time frame.

We first discuss the very basics of how a 3-axis spectrometer does its job, only then will we have a detailed look at the components of the spectrometer.

9.1   The Skinny

In an inelastic scattering experiment, the neutron confers a certain amount of momentum, and a certain amount of energy to the sample. The spectrometer will be set for a certain amount of momentum and energy transfer, and then one measures what fraction of the neutrons that hit the sample will be scattered by the sample. This fraction

is (roughly) proportional to the quantity that one is after in such experiments, namely, the dynamic structure factor $S(\vec{q}, E)$ (with $\hbar\vec{q}$ the amount of momentum transferred to the sample and with $E = \hbar\omega = h\nu$ the amount of energy transferred).

The 3-axis spectrometer utilizes two crystals to keep track of momentum and energy transfer. The first crystal that the neutron beam encounters is called the monochromator crystal, or monochromator for short. Neutrons with a rainbow of energies impinge upon this crystal, but only those neutrons that satisfy the Bragg condition for reflection are reflected out of the beam onto the sample. The remainder of the neutrons are discarded.

The above selection procedure fixes the energy of the neutrons that are sent on towards the sample. This is because the energy of the neutron is given by the wave length of the neutron, and the selected wavelength of the neutron is such that its probing wavelength matches the lattice spacing of the monochromator crystal. This selection procedure does more: it also fixes the momentum of the neutron. The magnitude of the momentum vector is determined by the energy (wave length) of the neutron, and the direction of the neutron's momentum vector points from the monochromator crystal to the sample.

The neutrons that make it to the sample are either scattered by it, or they continue through the sample and will be discarded. Of the neutrons that are being scattered by the sample, a (small) fraction of them will be scattered in the direction of the waiting analyzer crystal. This is the main difference between a 3-axis experiment and an elastic scattering experiment where the neutrons will be allowed to make it directly to a detector; in an inelastic experiment, the neutrons have to undergo one more selection

Fig. 9.3 Two examples of 3-axis spectrometers. The one on the left is the TRIAX spectrometer at the Missouri Research Reactor, the one on the right is HB3 at Oak Ridge National Laboratory. The strapping lad shown in the left photo demonstrates that these spectrometers are in fact quite large, and that things rapidly get quite messy when doing an actual experiment. The HB3 spectrometer is the brightest 3-axis spectrometer in North America, TRIAX is down in brightness by a factor of about 5 compared to HB3.

procedure. During this final selection procedure, only those neutrons that have a certain energy will be Bragg reflected by the analyzer crystal (analyzer for short) in the direction of a detector; all other neutrons are discarded.

The final selection step ensures that only neutrons with the desired final energy make it to the detector to be counted. With the final energy E_{final} of the scattered neutron known, and combining this with the known initial (incident) energy E_{initial} of the neutron, we can now easily calculate the energy transferred from the neutron to the sample as $E = E_{\text{initial}} - E_{\text{final}}$. Moreover, we also know the final momentum vector $\hbar\vec{k}_{\text{final}}$ of the neutron. The direction of this vector points from the sample to the analyzer, and its magnitude can be inferred from the final energy (wavelength) of the neutron. Again, combining this with the initial momentum vector $\hbar\vec{k}_{\text{initial}}$ that was fixed by the monochromator, we calculate the amount of momentum transferred from the neutron to the sample as $\hbar\vec{q} = \hbar\vec{k}_{\text{initial}} - \hbar\vec{k}_{\text{final}}$.

While we may feel sorry for all the discarded neutrons, they do serve a purpose. The measurement consists of tabulating what fraction of the neutrons are actually scattered by the sample while transferring a precise amount of momentum and energy to the sample. If the conditions are just right- for instance, when the amount of energy and momentum transferred are just what is needed to get a sound wave going in the sample- then we will observe that a significant fraction of the neutrons are being scattered. If the conditions are slightly off- for instance, the amount of momentum transferred would be good for creating a sound wave, but the energy transferred is slightly too little or too much- then we will observe that barely any neutrons are scattered.

During the course of an experiment we can probe the scattering power of the sample while a given amount of momentum and energy is being transferred to it. In this way, we will hit the 'sweet spot' that would correspond to a particular excitation. An example of such an excitation is displayed in Fig. 9.4 where we show the scattering power of liquid helium at a particular momentum transfer, while varying how much energy the liquid receives from the neutron. As can be seen in this figure, when the amount of energy transferred is just right, the liquid says, thank you, and an excitation is created. This way we learn about the possible excitations in our sample.

Experiments like these can even be repeated at different momentum transfers. In this way, one can map out the momentum dependence of a certain excitation. An example of this, again for liquid helium, is shown in Fig. 9.5. This figure is an intensity map where a larger fraction of the neutrons being scattered is indicated by a darker color. Many details are visible in this figure. One observes that the energy at which the excitation shown in Fig. 9.4 occurs depends strongly on the amount of momentum transferred. This dependence is referred to as the dispersion curve. Thus, the dispersion curve $E(q)$ is given by how much energy E is required to create an excitation while transferring a certain amount of momentum $\hbar q$. For this particular example of a liquid, only the magnitude of the momentum transferred is important since a liquid
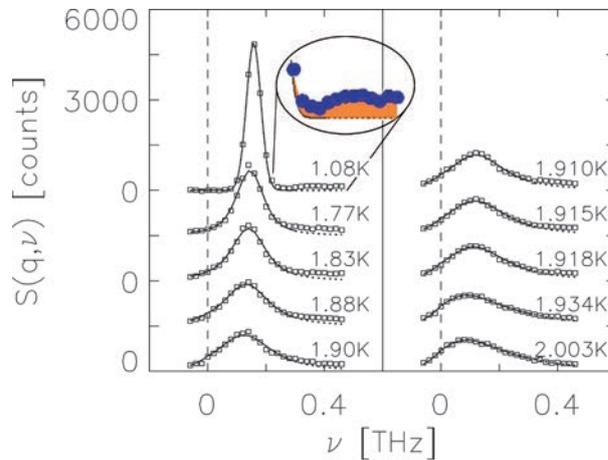
Fig. 9.4 Shown is the dynamic structure factor $S(q, E = h\nu)$ of liquid ^{4}He for $q = 2.0$ Å^{-1} at constant density $\rho = 0.1715$ g/cm^3. The experiments have been repeated at various temperatures so that one can observe the changes that take place in this liquid when the temperature is raised from the superfluid to the normal fluid phase. At the lowest temperature (top left) one can see that neutrons that transfer just the right amount of energy ($\nu = 0.16$ THz) are very successful in creating an excitation in the liquid; a large fraction of the neutrons are being scattered. In contrast, a much smaller fraction of the neutrons are scattered when they transfer a different amount of energy. This results in a well-defined excitation peak. Note, however, (inset) that there are some details visible in the scattering at higher energy transfer; one needs a magnifying glass to see these, but they also represent some excitations in the liquid. For the record, one can also see that the excitation becomes less well characterized (blurry) when the temperature is increased. This reflects an intrinsic property of the liquid that can be studied using neutron scattering techniques.

is isotropic (it looks the same from all directions). This would not be the case for solids, or molecules with a directional dependence. In these cases, the dispersion curve depends also on the direction of the momentum transferred: $E(\vec{q})$.

Before we delve into the details of the components of a 3-axis spectrometer, we show two more examples of excitations that have been studied on 3-axis spectrometers world wide. In Fig. 9.6, we show a high-energy magnetic excitation in the mineral magnetite. This is the mineral that has given magnetism its name, yet not all of the microscopic behavior and arrangement of the magnetic moments (at low temperature) are fully understood. The figure demonstrates the changes that take place in these magnetic excitations when the mineral's temperature is lowered by 10 K, from 125 to 115 K. The underlying change in behavior of the magnetic ions that causes the one excitation to split onto two separate ones is still an object of study.

Another example of magnetic excitations is shown in Fig. 9.7. This figure, showing the dispersion relation of magnons along various crystallographic directions can
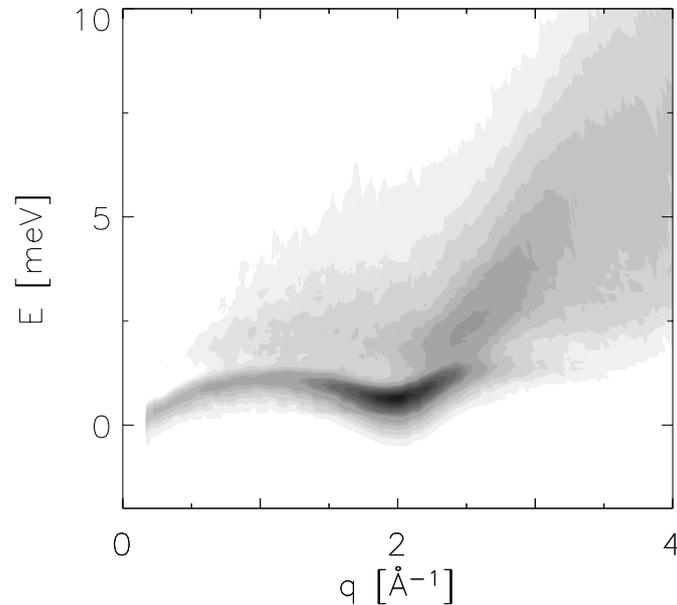
Fig. 9.5 The scattering in superfluid ^{4}He under its own pressure as a function of momentum transfer $\hbar q$ and energy transfer $E = h\nu$, measured by means on inelastic neutron scattering. The strongest scattering corresponds to the darkest bands. A vertical cut through the data at constant $q = 2$ Å^{-1} yields the data shown in Fig. 9.4.

be compared to the phonon dispersion shown in Fig. 4.11. In this particular example on a magnetic material that consists of long chains of strongly interacting magnetic ions with very weak interaction between the chains, the difference in excitation energy between two magnon branches was directly proportional to a magnetic interaction between magnetic ions located on different chains. This interaction strength was needed to figure out whether this system would be capable of supporting an oscillation in its magnetic domains when a constant external magnetic field was being applied. An (predicted) oscillatory response to a constant force is a purely quantum mechanical effect, and the data shown in Fig. 9.7 demonstrated that this particular system was not a good candidate to observe this peculiar effect.

9.2 Components of a 3-Axis Spectrometer

We will review the components of a 3-axis spectrometer, and what to pay attention to when designing an experiment.

The aim of an experiment on a 3-axis machine is to get as accurate and detailed information about one's sample as possible. We outlined the basic workings of a 3-axis spectrometer in the sense that we discussed how the incident and final energies are determined by the monochromator and analyzer crystals, and that these choices in
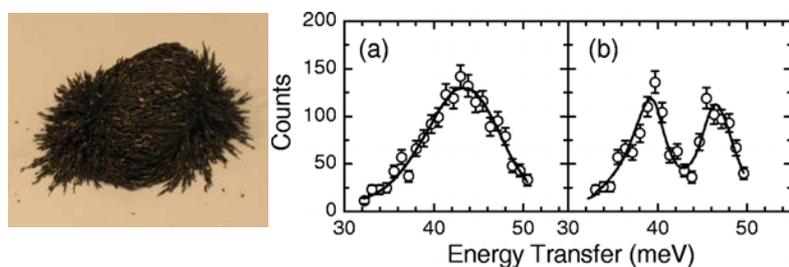
Fig. 9.6 Figure on the right reproduced with permission from R. J. McQueeney *et al.*, Phys. Rev. B73, 174409 (2006). Magnetite is the world's most ancient magnetic material. It has been used as a compass by the vikings, and its behavior is still not fully understood. In particular, when cooled down below 120 K the mineral undergoes a structural phase transition that also affects its magnetic properties. Neutron scattering experiments clearly show the changes that take place in a high energy branch of excitations when crossing from the high temperature structure (middle panel) to the low temperature phase (right panel). The photo on the left shows iron filings becoming magnetized when in contact with a piece of magnetite (buried underneath the iron hairdo).

combination with the angle at which we position our analyzer with respect to the sample also determines the incident and final momentum of the neutron. We can then use the detector to count how many scattering events take place for a given momentum and energy transfer.
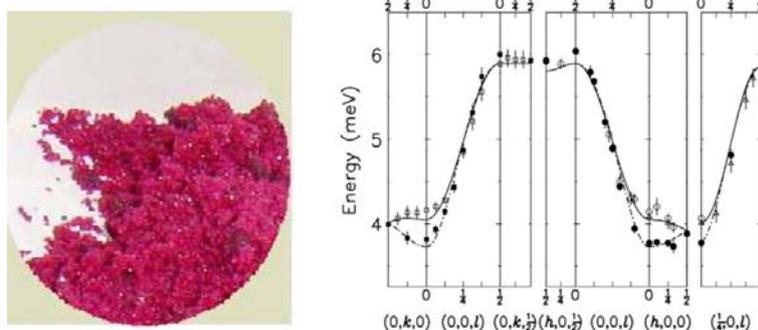

Fig. 9.7 Figure on the right reproduced with permission from W. Montfrooij *et al.*, Phys. Rev. B64, 134426 (2001). Invisible ink is great fun to play with, but it also is a low-dimensional magnetic material. Magnetic chains run through the structure; the coupling between the chains is so weak that a fully ordered structure does not emerge until the ink is cooled down to fairly low temperatures. The strengths with which these magnetic chains couple to each other can be deduced from a measurement of the (magnon) dispersion curve along various directions in the crystal. Photo source: www.dreamtec.or.kr.

In reality, we do have to accept a range of energies, and a range of momenta in order to ensure that a sufficient number of scattering events will actually take place and will be detected. All this is part of choosing the resolution of the spectrometer, which is the most difficult part of any experiment. With resolution we mean the range of energy transfers that will be allowed to constitute a scattering event, and the range of momentum transfers that we are willing to accept.

Relaxing the resolution of a spectrometer results in more scattering events, and therefore in a stronger signal compared to the background scattering. Obviously it also comes with the drawback that our desired signals will get smeared out, even up to the point that they can no longer be identified unambiguously as belonging to a particular excitation. It is up to the experimentalist to walk the fine line between obtaining sufficient counts, and doing so in a way that the referee of the submitted paper is not going to question the validity of the results.

In the following we outline the various spectrometer components and their role in determining the resolution of the spectrometer. Some of the items here have already been discussed before, however, we do not expect the reader to necessarily have read the preceding instrumental chapters.

9.2.1   Shielding

We first start with the shielding, as this is a component that in most cases does not require a lot of thought as it has already been put in place to provide the maximum benefit in both reduction of exposure to radiation on the beam port floor, and in shielding of the detector from unwanted neutrons that come from places other than the sample.

There is one occasion where shielding issues might need to be addressed. One sometimes observes that there is an increase in background counts for a configuration of the spectrometer that does not correspond to an expected increase in background. One expects to see an increase in background when one tries to take measurements at very small scattering angles where the straight through beam (those neutrons that were not scattered by the sample) creates air scattering that can make it all the way to the detector. However, there is an additional possibility of increased background counts corresponding to the detector being in the path of the straight through beam, even though the scattering angle is not small. This situation is depicted in Fig. 9.8.

Whenever one encounters an unexpected increase in background, then simply drive the spectrometer motors to the configuration where this increase occurred, and look for the source of the increased background. More often than not, it will be clear what the source is, and shielding blocks and mats can be used to reduce the source of this heightened background should the level of background scattering interfere with what is being measured.
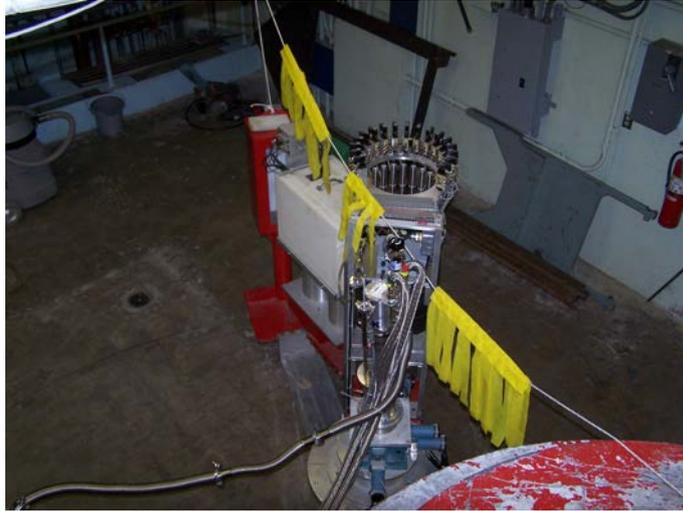
Fig. 9.8 For certain scattering configurations, the detector housing can traverse the straight through beam, leading to an increased level of background counts.

9.2.2   Stuff Upstream from the Monochromator

The monochromator crystal is in a direct line of sight with the beam tube that carries neutrons away from the reactor core, and out onto the beam port floor. When this beam tube represents a direct connection between the reactor core and the outside world, then there will be a great many unwanted neutrons that will make it out of the beam tube. These unwanted, very fast neutrons create a large amount of background scattering, and the aim is to prevent them from ever reaching the monochromator. Note, however, that if the beam tube is looking at a cold source, then this fast neutron problem is greatly reduced.

When fast neutrons need to be removed from the beam, a sapphire (Al_2O_3) filter is placed in the beam tube, or at the very end of it. This filter effectively removes all neutrons with a wavelength shorter than 0.8 Å. It also removes some wanted neutrons, but this is a loss that is well worth the price of cleaning up the beam. At some reactors the filter is cooled to liquid nitrogen temperatures. At these temperatures fewer wanted neutrons are scattered out of the beam by phonons that are present in the sapphire. The only thing to remember here is that, if the filter is cooled, one must make sure that is stays cooled during the course of the entire experiment. When it warms up, it will subtly affect the experiment in the signal to background ratio. If in doubt, use a warm filter. Other than that, either the instrument comes with such a filter, or not, but it cannot be changed during the course of the experiment.

Some 3-axis spectrometers, such as TRIAX at MURR, allow for additional filters to be placed between the beam tube and the monochromator crystal. If one wants to use neutrons with incident energies below, say 25 meV, then it makes sense to remove

as many higher energy neutrons from the beam as soon as possible. Some filters can help with this. Simply look up the characteristics of the filter and decide whether to use it or not depending on whether it removes unwanted neutrons more than wanted neutrons.

The most important component upstream from the monochromator to think about when designing a scattering experiment is the collimation. A collimator is essentially a bunch of parallel channels, as shown in Fig. 9.19. Neutrons that hit the blades that divide the channels will be absorbed. This way, the angular spread, or divergence, of the incoming neutron beam can be controlled. This is important because a neutron that will hit the monochromator crystal under a slightly different angle than the average angle, can still be reflected onto the sample: while the angle is slightly different from the nominal angle, as long as the energy will also be slightly different from the average energy, then reflection is still possible. This results in a spread of the neutron energies that are allowed to make it to the sample.

By using a more tight collimation, implying that a smaller angular range is allowed to make it through the collimator, one can control the range of energies of the neutrons incident on the sample, at least to some extent. The monochromator crystal will also have an intrinsic range of angles that it accepts for Bragg reflection onto the sample (see monochromator mosaicity in the next subsection). In general, the two angular divergences (of the collimator and the monochromator) determine the spread of energies of the neutrons that make it to the sample. Squeezing the collimation down to a divergence that is less than the mosaicity of the monochromator is probably not a good idea; in this case one would throw away quite a few neutrons without gaining any energy resolution because the monochromator crystal will reintroduce angular divergence even in a divergence free beam. But it certainly makes sense to tighten up the collimation such that it is comparable to the monochromator's mosaicity, or a few times that number. Typical ranges for collimation are 40-80 minutes (with 60 minutes in a degree). 20' pre-monochromator collimators are only used when one really has to worry about the energy resolution, and 10' collimators rarely lead to publications.

9.2.3 Monochromator

The monochromator crystal is the most important choice of the experimental setup. Ultimately, this choice is determined entirely by the excitations in the sample that one is interested in. First, the choice of monochromator determines the range of accessible incident energies of the neutron. During an experiment, the angle of reflection Θ_m of the monochromator can be varied from about 12^o to about 50^o. This angular range, in combination with the lattice spacing d_{mono} of the monochromator crystal, determines the wave length of the neutron that is reflected and, thereby, the energy of the neutron. The wave length λ_{in} is given by $\lambda_{\mathrm{in}} = 2d_{\mathrm{mono}} \sin \Theta_m$, and the corresponding energy E_{in} is given (in meV) by (see Appendix A) $E_{\mathrm{in}} = 81.8/\lambda_{\mathrm{in}}^2$.

Given this, the first step is to figure out the maximum amount of energy transfer that is needed in the scattering experiment. If this energy transfer exceeds something
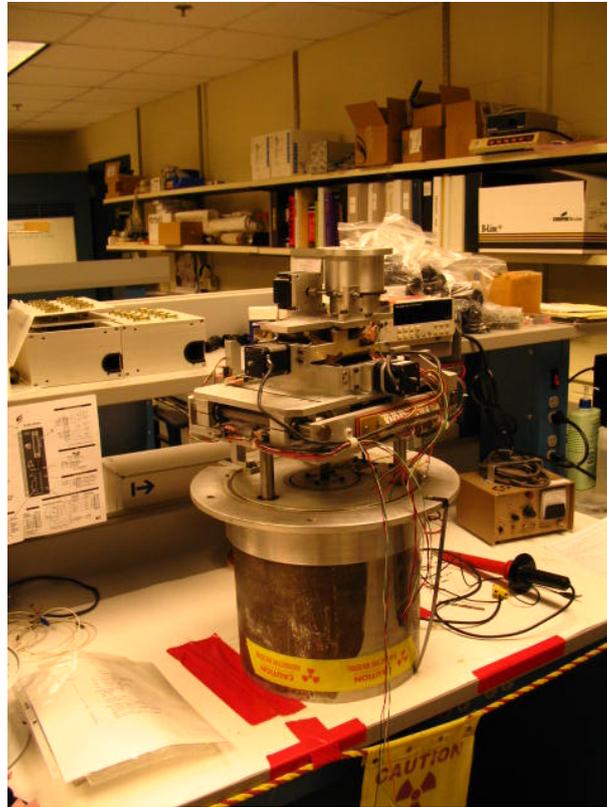
Fig. 9.9 The monochromator crystal is positioned with the aid of multiple motors. The user normally does not have to deal with any of these motors. Shown here are three translational motors, and two tilt motors. The monochromator crystal (not shown) would sit on top of this plug. The plug is then hung upside down, and rotated as a whole when Θ_m is varied.

like 80 meV, then in most cases a triple axis spectrometer is not the right instrument and one would be much better off at a spallation source. In Table 9.1 we show some typical monochromator reflections that are used in neutron scattering, as well as their corresponding energy range.

Within the limitations of these energy ranges, one now has a choice of which monochromator to use. In most cases one selects pyrolytic graphite (PG) as this is a monochromator that reflects a large fraction of the incident neutrons onto the sample position. However, should the spread of energies that is accepted by this monochromator result in too large an energy divergence, then one should use a more selective monochromator, such as silicon.

There is an added benefit to using silicon as a monochromator. The second order reflection- that is, the reflection by the monochromator of unwanted neutrons that

travel twice as fast as the wanted neutrons- is forbidden. This greatly facilitates the analysis of (magnetic) reflections and excitations, and it allows for more accurate line shape analysis of all excitations. Silicon is not used in all cases, however, as the reflectivity- the fraction of desired neutrons that are being reflected out of the main neutron beam onto the sample- tends to be significantly lower than the reflectivity of pyrolytic graphite.

Table 9.1 Commonly used monochromator crystals and reflections. The energy ranges correspond to scattering angles Θ_m of 12^o and 45^o, respectively.

Material	reflection	d-spacing [Å]	E-range [meV]	comments
PG	(002)	3.3539	42-3.6	high reflectivity
PG	(004)	1.6770	168-14.5	”
Cu	(200)	1.8075	145-12.5	used for high energies
Cu	(220)	1.27813	290-25	”
Si	(111)	3.13543	48-4.2	absence of second order
Si	(220)	1.92005	128-11.1	”
Si	(311)	1.63742	176-15.3	”
Si	(511)	1.04514	433-37.4	”
Ge	(111)	3.26651	44-3.8	absence of second order
Ge	(311)	1.70588	163-14.1	”
Be	(002)	1.79035	148-12.8	

Monochromators such as copper are useful because they have a smaller lattice spacing, allowing one to select neutrons with shorter wavelengths and, therefore, higher energies. Copper itself does not have a terribly good reflectivity, so it should only be used when other monochromators do not cover the desired energy range. And of course, when using neutrons with high incident energies, make sure that there are no filters left upstream from the monochromator that would reduce the number of neutrons of the desired energies. Table 9.1 can be helpful in making the monochromator choice.

Quite a few 3-axis spectrometers allow for the monochromator to be bent, resulting in a beam of neutrons that is focused on the sample, as opposed to quite a few neutrons missing the sample. Such focusing techniques increase the angular divergence of the beam and, thereby, it deteriorates the resolution of the spectrometer. Vertical focusing

Fig. 9.10 Images taken from the ILL website (www.ill.eu). The ILL monochromator group has succeeded in growing very large single crystals of copper (left) and a Heusler compound (right). The mosaicity of these crystals can be adjusted post-growth.

does not deteriorate the resolution of the spectrometer nearly as much as horizontal focusing. However, the gains in neutron counts might make even horizontal focusing well worth it. It is up to the user to decide whether to use the focusing options. Focusing tends to be very beneficial when studying weak excitations in single crystals, but biological materials are sometimes better studied under conditions of less angular divergence. Note that not all monochromators can be focused; therefore, when one is pressed for intensity, choose a monochromator that can be bent.

Finally, there are monochromators that only reflect neutrons of a particular spin orientation. These monochromators are used in polarized neutron scattering. An example of a (bad) polarizer is the (220) reflection of magnetite where for one particular neutron spin orientation the magnetic and nuclear scattering mechanism (almost) exactly cancel, resulting in the reflection of only the other spin orientation. Better polarizing monochromators are Heusler crystals. Needless to say, such monochromators are only used when it is important to distinguish between magnetic and nuclear scattering.

The monochromator group at the Institute Laue-Langevin (ILL) can grow very large copper, germanium, and Heusler single crystals, suitable for use as monochromators. This group uses a plastic deformation technique to control and select the mosaic of these crystals. We show two examples of their products in Fig. 9.10.

9.2.4 Monitor

The monitor intercepts a fraction of the neutrons that is being directed from the monochromator onto the sample. The monitor works in the limit that the counting rate is inversely proportional to the speed v of the neutrons: count rate $\sim 1/v$. This has two advantages.

First, only a very small fraction of the neutrons are intercepted so that the rest can be used for learning about the sample. Second, this proportionality cancels a fac-

tor $(k_f/k_i$, see eqn 4.1) that arises in the connection between how many neutrons are scattered in a certain direction within a certain energy range (also known as the double differential cross-section) and the dynamic structure factor (the quantity that describes what is happening in the sample, independent of the probe we use). For the record, this cancelation only works if we vary the incoming energy of the neutron and we keep the final energy (as determined by the analyzer crystal) fixed. However, this is the standard mode of operation of most 3-axis spectrometers.

So why mention the monitor at all? After all, it just tells us how many neutrons are incident on the sample, and it is not likely that we are going to exchange it for a different monitor. One of the reasons has already been alluded to. If we decide to run the spectrometer in a different mode wherein we keep the incident neutron energy fixed but vary the final energy, then we must account for the fact that our measured signal is not directly proportional to what we want to know.

There is however another reason to mention the monitor that has to do with the analysis of the data. In a typical scattering experiment, the monochromator not only reflects the desired neutrons into the beam, but also neutrons that go twice, three times, four times, etc. as fast as the desired ones. This is an inescapable consequence of how the monochromator works. If the probing wave length of the neutron matches the lattice spacing, then half this wavelength will also match the lattice spacing with the result that those neutrons will satisfy the Bragg condition. This is the situation where the wave that is drawn in Fig. 2.4 has only half the wavelength. As can be seen here, all scattering will be in phase and, therefore, Bragg reflection will occur.

These neutrons- referred to as higher order contamination of the beam that hits the sample- are scattered by the sample, but they do not (necessarily) make it to the detector. The reason for the latter is that the detector is shielded from these higher order neutrons by a filter that blocks them. This is further detailed in the subsection on filters.

While it is good that they are not being detected, they nonetheless present a problem because we counted them as neutrons impingent upon the sample. And since our signal is supposed to measure the fraction of good neutrons that are being scattered, we run into a problem when we use the number of counts in the monitor as our normalization. And since this contamination actually depends on what the incident energy of the neutron is, we end up with a distortion of our measured line shapes. This can be a problem when we hope to learn additional information about our sample based upon line shapes, not merely upon whether peaks are present or not.

This is a problem that cannot be avoided other than by employing a velocity selector; in general it can only be corrected for. This correction entails either making a measurement of the level of contamination (which is not a straightforward procedure; it is outlined in Appendix H), or it entails modeling the correction. The important thing to keep in mind is that this problem does exist. Should it have the potential to
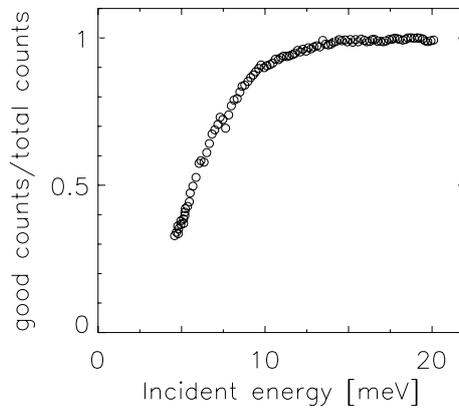
Fig. 9.11 The measured monitor contamination as a function of incident energy for the N5 triple-axis neutron spectrometer at the NRU reactor of the NRC Canadian Neutron Beam Centre. The spectrometer was operated using a Si monochromator with a forbidden second order reflection. The incident beam had a cooled sapphire filter in it to reduce the fast neutron flux. Note that the fact that the monitor correction is essentially absent for experiments using fixed final energies of 14 meV (a common choice in scattering experiments) is largely because of the use of a Si monochromator. Other monochromators such as PG would not eliminate second order contamination resulting in a noticeable effect even at these energies.

be a serious problem, then as a first defense one can switch to the (111) reflection of a silicon or germanium monochromator as these reflections have the bonus property that second order contamination is not possible (third order is however).

When should we expect this monitor contamination to be serious? It is most important when the energy of the incoming neutron is low, such as 5 meV. The second order contamination in this case would correspond to neutrons with 20 meV in energy, third order to neutrons with 45 meV and fourth order to 80 meV neutrons. All such neutrons are abundantly present at thermal reactor sources, and we show an example of monitor contamination in Fig. 9.11. For situations like this, the best remedy is to measure the contamination and correct the data accordingly. It is also helpful to place an additional neutron filter in the incoming beam.

For neutrons with an incoming energy of 14 meV, second order contamination corresponds to neutrons of 56 meV, and third order to 126 meV. The former is definitely present in the beam, the latter tends to be mostly removed by a (standard) sapphire filter.

The bottom line here is to most definitely be on the lookout for this instrumental artifact, and do not get carried away during the analysis stages when uncovering an unexpected line shape.

9.2.5   Masks

In most experiments, the neutron beam tends to be taller and wider than the sample. Masks are used to cut the beam down to the appropriate size. Masks are made of a neutron absorbing material and by removing useless neutrons from the beam they can significantly cut down on the background. The background in this case is caused by the neutrons that miss the sample and scatter off of the sample containment such as the cryostat, or by neutrons that give rise to air scattering.

To position the masks in the case of solid sample one chooses a Bragg reflection of the sample, and squeezes the mask until the count rate goes down. Then, one rotates to another Bragg peak of the sample such that the rotation of the sample corresponds to the rotation range of the sample during the experiment, and one repeats the tightening of the masks. The mask is then set to the loosest opening of the masks (jaws of the mask).

Another method that is very elegant is to actually take a picture of the sample in the beam using a neutron camera. Especially when this camera is coupled to a computer screen one is sure to position the mask correctly. In the case of liquids or biological samples, the masks should be positioned to coincide with the sample container.

Normally there are two sets of masks, one between the monochromator and sample, and one between the sample and the analyzer. There is no preference on whether the masks should be as close as possible to the sample, or far away from it. The take home message is that masks are good, they cut down on unwanted background without reducing the sought after signal. Therefore, use them. Sometimes the jaws of the masks are computer operated, which makes it all much more easy to position them as it does not involve the somewhat tedious process of shutting off the beam, repositioning the masks, turning the beam back on, and so forth.

9.2.6   Sample (Table)

The sample more or less is what it is, we have to setup the spectrometer in such a way as to make the most of it. An important consideration when choosing one's sample (should one have such freedom) is that a typical sample should not scatter more than 10% of the incident neutrons. Samples that scatter more lead to problems with data analysis as the multiple scattering contribution will be too large. Samples that scatter only a wee fraction lead to very weak signals, possibly undetectable within the background signal. This of course holds true for all experiments, not solely for those on 3-axis spectrometers. The preferred sample shape on a 3-axis spectrometer is cylindrical, but strongly absorbing samples will have to be measured in a slab geometry such as shown in Fig. 9.13.

During the course of a scattering experiment, the sample rotation angle can be varied to probe different directions of momentum transfer within the sample. The neutron scattering angle determines the magnitude of the momentum transfer, the

Fig. 9.12 An example of a low-tech set of masks. The (visible) non-adjustable rectangular mask has been chosen to fairly closely match the desired beam profile. The adjustable jaws of the mask behind it (not visible) are controlled by 4 screws for adjustment (the visible screws) .

sample rotation angle determines along what direction within the sample this momentum transfer was applied. Once the sample is properly aligned, the computer will take care of the proper sample and scattering angles.

Fig. 9.13 A strongly absorbing sample should be mounted in slab geometry, such as the one shown in this photo. This sample cell can be filled with a strongly absorbing powder; the thickness of the cell can be varied from a few tenths of a millimeter to a few millimeters. The sample cell can be housed inside a cryostat.

Fig. 9.14 The tilt angles applied to the sample table under the cryostat are easily visible. This is a light weight cryostat, do not attempt these angles with a large cryostat such as a cryomagnet: it is guaranteed to break, or even lift the sample table off the floor.

Aligning a single crystal requires a bit of work, and we outline the procedure for this in Appendix G. When using large pieces of sample environment such as a big fat cryostat, then one is restricted in the tilt angles that can be applied to the sample table before the whole thing falls over (Fig. 9.14). Typically, single crystals need to be mounted within 5 degrees of their intended alignment.

9.2.7   Filters

We already discussed the filters that can prevent very fast neutrons from reaching the monochromator, this subsection deals with removing unwanted neutrons with more modest speeds from the beam. The most commonly used filters are graphite, beryllium and beryllium-oxide. Diamond and plutonium filters are too specialized to be discussed here.

A graphite filter is placed after the sample when running the spectrometer at fixed final energy mode, and it is placed before the sample when running with fixed initial
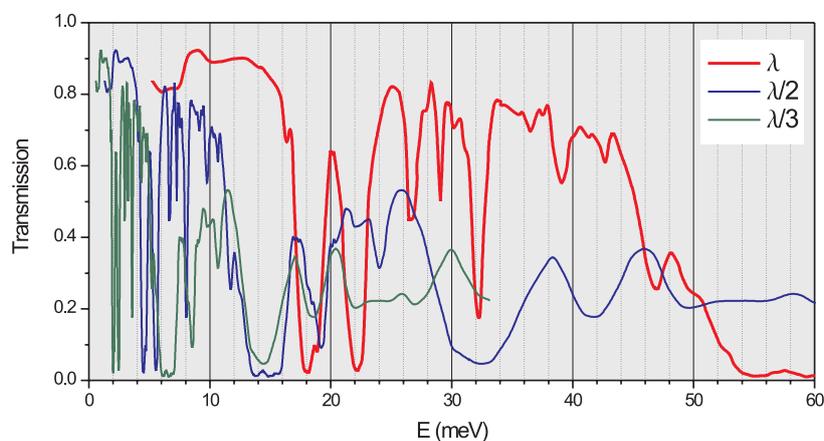
Fig. 9.15 Figure taken from the Neutron Scattering and Magnetism group at the ETH, Zürich (www.neutron.ethz.ch). Shown is the transmission of neutrons through a 2 cm thick graphite filter. The incident neutron energy- corresponding to the wanted neutrons and as shown by the red line- is given on the horizontal axis. Note the region of incident energies around 14 meV where we find 90% transmission for the wanted neutrons, and only a few percent for the unwanted, higher-order neutrons. Increasing the filter thickness to two inches reduces the transmission for the unwanted neutrons to virtually zero, something that is required in a magnetic scattering experiment.

energy. The reason that it needs to be on the part of the spectrometer where the energy is always the same is because graphite only removes specific wave lengths (energies) from the beam. Graphite works best when the fixed energy of the wanted neutrons is around 14 meV. The graphite filter gets rid off neutrons that have energies of 56 meV (second order contamination) and 126 meV (third order). We want to remove these neutrons as the analyzer will reflect those particular wavelengths onto the detector when the analyzer is set to 14 meV. Fig. 9.15 shows a typical transmission spectrum of a graphite filter.

When doing magnetic experiments, one has to make sure that a sufficient quantity of graphite filters is placed in the beam, of the order of 2" (see Fig. 9.16). The reason for this is that weak antiferromagnetic scattering occurs at exactly the same spectrometer configurations as nuclear scattering caused by second order neutrons. This second order nuclear scattering can easily outcount the magnetic signal; however, graphite filters can completely get rid off this unwanted contribution.

Beryllium and beryllium oxide filters get rid off all the neutrons that travel faster than a particular cutoff speed. The cutoff for Be is 4 Å, for Be-O it is 3.7 Å. These filters are used when the spectrometer is set to a low fixed final energy such as 5 meV. This is an energy where a graphite filter would not be useful. Be and Be-O filters need to be cooled to liquid nitrogen temperatures in order to reduce the phonon excitations in
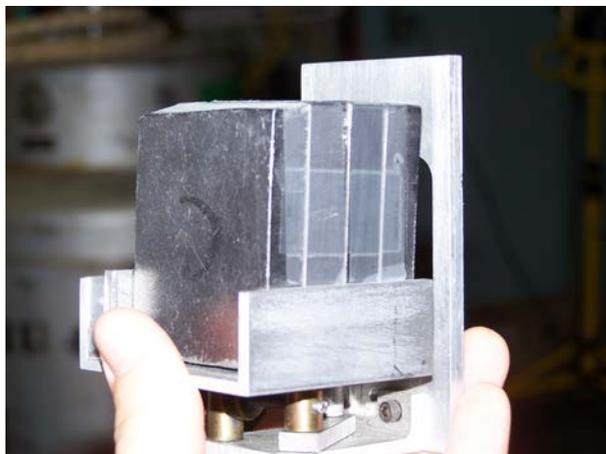
Fig. 9.16 Two one-inch filters of pyrolytic graphite. Once the filters are oriented properly, they are held in place (with screws and scotch tape in this picture) so as to maintain the orientation. These filters are surprisingly expensive, so do not drop them.

these filters that would scatter wanted neutrons out of the beam. Without cooling, the transmission of these filters for wanted neutrons is very low. Therefore, always make sure that the filter is kept at nitrogen temperature during the course of the experiment. For reference, the cutoff energy is determined by the largest lattice spacing available in these materials according to $\lambda_{\mathrm{cutoff}} = 2d_{\mathrm{largest}}$. Longer wavelengths cannot find a lattice plane to scatter out of the beam and, therefore, these slow neutrons will be transmitted by the filter (unless they encounter a phonon that might scatter them out of the beam).

9.2.8 Analyzer

The function of an analyzer is to Bragg reflect neutrons with a particular energy out of the scattered beam and onto the detector. In this sense, everything that holds true for a monochromator crystal also holds true for the analyzer crystal. The acceptance range of energies is determined by both the divergence of the scattered beam, and by the mosaicity of the analyzer crystal. Fig. 9.17 shows the pyrolytic graphite analyzer of the TRIAX spectrometer at MURR.

The choice of analyzer is much easier than the choice of monochromator: when in doubt, choose pyrolytic graphite. Provided the rest of the spectrometer is set up correctly, the only job of the analyzer is to reflect all the desired neutrons onto the detector. The high reflectivity of graphite analyzers make them very well suited for this job. Detectors tend to be so large that even focusing of the analyzer crystal produces only modest effects. There are other analyzer crystals available, but the only crystal that an average user might encounter is the Heusler analyzer crystal used in polarized neutron scattering experiments.
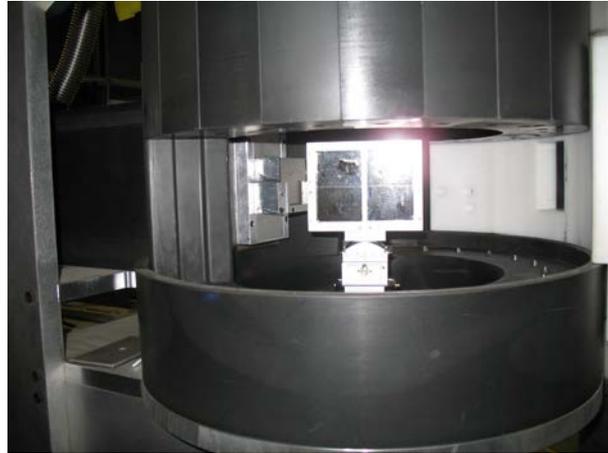
Fig. 9.17 The pyrolytic graphite analyzer crystal as it sits in the shielded housing on the TRIAX spectrometer at MURR. This is an example of a flat (non-focusing) analyzer. A graphite filter can also be seen located just to the left of the analyzer.

9.2.9 Detector

There is little we need to say about detectors. Good detectors (Fig. 9.18) count almost all of the wanted neutrons, and most detectors are good detectors. Only under extreme circumstances, when we are dealing with very high energy neutrons, do we have to worry about whether detectors count all the neutrons that go through it, or preferentially those that travel at slower speeds. In addition, when 3-axis spectrometers are operated using a fixed final energy for the scattered neutron, then all the detected neutrons have identical energy implying that any energy dependent counting efficiency factor would be irrelevant.

On 3-axis spectrometers we only ever need to worry about the detector when we send too many neutrons in its direction. Strong sample Bragg peaks or moving the detector through the unscattered beam will result in the detector being saturated. This manifests itself as the detector maxing out at a high count rate (typically a few ten thousand counts per second). Looking at the counting profile versus detector position immediately makes it clear that the detector has maxed out. If one is interested in the actual peak shape distorted by the saturated detector then one will have to put an attenuator in the beam, such as some perspex glass, a paper copy of an annual report, or this booklet.

9.2.10 Collimators and Resolution

The resolution of the spectrometer is determined by the incident energy, by the choice of monochromator crystal and by the choice of collimators such as the ones shown in Fig. 9.19. Other than the choice of incident energy, the resolution of the spectrometer
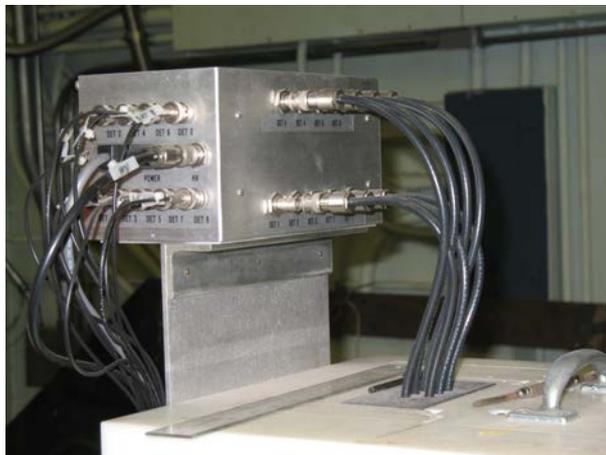
Fig. 9.18 While most detectors on 3-axis spectrometers are rectangular boxes, on TRIAX the detector consists of 7 pencil detector positioned in a bowling-pin-like formation. The user can choose which detectors to add together for the sought after signal, while at the same time getting an indication of the background signal by using the outer pencil detectors. The photo shows the electrical leads going into the shielding, from which the positions of the 7 detectors can be inferred.

is determined by controlling angular divergences of the neutron beam.

As a note up front, the resolution of the spectrometer is a 4-dimensional entity, describing the energy resolution (1 component) and the resolution pertinent to momentum transfer (3 components). Since the energy and momentum transfer of the neutron are both determined by the incident and final momentum vector of the neutron, we have that the energy resolution and the q-resolution are interrelated. Changing one component influences the other three. In general, it is not intuitively clear how a change in collimation, for instance, will affect the energy resolution, but programs exist to (roughly) estimate these effects.

Lowering the incident energy is the most effective way of improving the energy resolution, at the cost of restricting the accessible (q, E)-range for the experiment. This holds true for a 3-axis spectrometer, but note that time-of-flight spectrometers at spallation sources do not abide by this principle. For the latter category, the best energy resolution is found for energy transfers close to that of the incident neutron energy. The second most effective way of improving the resolution is to change the collimators, as already discussed in our chapter on diffractometers.

As mentioned, there are programs available to help the user estimate the resolution function of a 3-axis spectrometer. A very effective and well-documented piece of software is RESLIB, which is written in MATLAB. We reproduce typical output of this program (from the manual by Andrey Zheludev) in Fig. 9.20. This figure represents a

Fig. 9.19 Controlling the angular divergence of the neutron beam between two components of the spectrometer (such as between the sample and the analyzer) fixes the angular spread of the neutrons that can still reach the next component. Since the scattering angles at the monochromator and analyzer determine the incident and final energy, respectively, the choice of collimator affects the range of energies of the accepted neutrons. Shown is a variety of collimators that can be used, with the smallest angular divergence corresponding to the highest energy resolution.

single crystal that has been oriented in such a way that we can transfer momentum in the x and y-directions, which form the scattering plane. We can also transfer energy to and from the sample. This is captured in the plots which show the resolution ellipse in the xy-plane for various values of energy transfer.

The orientation of the ellipse (the direction of the main axes of the ellipse as well as its tilt) depends strongly on energy transfer. Also, with increased energy transfer we see that the ellipse gets larger, implying that the resolution gets worse. The change in orientation of the ellipse in the xy-plane is very hard to quantify intuitively, which is why it is good practice to run programs such as RESLIB routinely when taking data on 3-axis spectrometers.

Scrutinizing Fig. 9.20, we see that the resolution ellipse has a tilt to it when projected onto the xE-plane, or onto the yE-plane (the vertical walls in the figure). This has important implications when we measure the dispersion of phonons or magnons. Depending how well the tilt of the resolution ellipse and the slope of the dispersion curve line up, we will get more or fewer scattering events. When the tilt and the slope match closely, then we will sample the dispersion over a range of momentum transfers and we can expect quite a few scattering events. This is called the focusing condition, or focusing side of the dispersion. When the tilt and the slope are almost perpendicular to each other, then we will get very few scattering events, and we call this the defocusing condition. We show a typical example in Fig. 9.21.
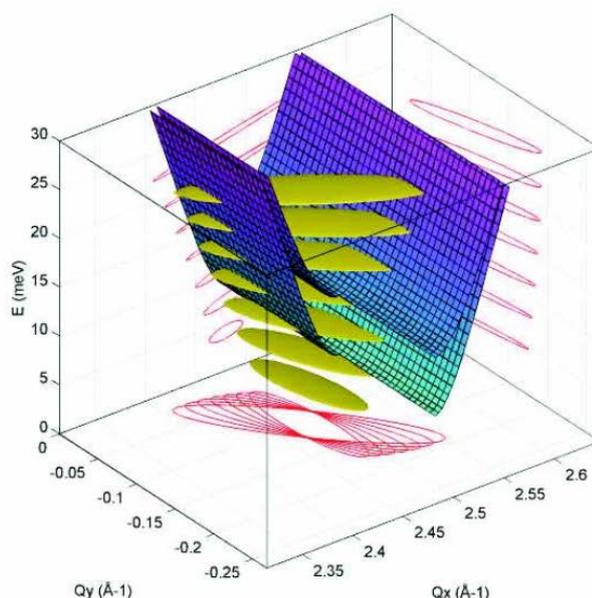
Fig. 9.20 Sketch of the energy resolution's dependence on the amount of energy transferred as a function of direction of momentum transfer. This plot has been reproduced from the RESLIB manual by Andrey Zheludev, and it also dons the cover of 'Neutron Scattering with a Triple-Axis Spectrometer' by Shirane et *al.* The yellow ellipses show the resolution corresponding to the energy transfers plotted on the vertical axis. At the edge of the yellow ellipse, the scattering intensity has dropped to half the maximum intensity that is measured at the center of the ellipse. Projections of the ellipses are shown on the sides of the cube. Note that both the overall size of the ellipse, as well as its orientation, change as a function of energy transfer. This plot has been calculated for neutrons that all have the same fixed final energy.

The tilt of the resolution ellipse is determined by the final energy of the neutrons. In other words, the user has control over this tilt, and the experiment can even be set up to have a perfect match between the slope of the dispersion and the tilt of the resolution ellipse. This would result in a very high scattered intensity. However, use this perfect matching with extreme caution. When the match is perfect for acoustic phonon or magnons, then it implies that our resolution ellipse will also cover the Bragg peak at zero energy transfer. Since the scattering of a Bragg peak is millions of times larger than the inelastic scattering by phonons, we run the risk of actually sampling the tails of the Bragg peak rather than the elementary excitations we are interested in. In practice, it is better to not get too greedy and have a slight misalignment between the tilt and the slope of the dispersion curve.
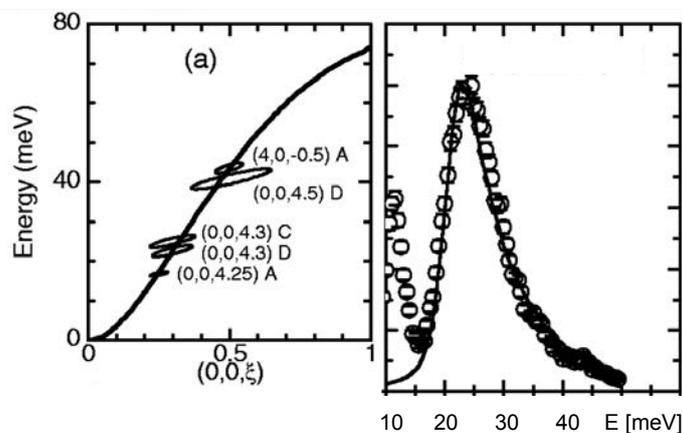
Fig. 9.21 Figure reproduced with permission from R. J. McQueeney *et al.*, Phys. Rev. B73, 174409 (2006). Shown on the left is the dispersion of (one branch of) magnons in magnetite, and the (projection of the) resolution ellipses relevant to various energy and momentum transfers. In this particular experiment, the absolute intensity and the line shape of the scattered signal were needed in order to resolve the behavior of the atomic magnetic moments. As can be seen, these numbers can only be determined accurately from experiment when we take the size and orientation of the resolution ellipse into account. The panel on the right, corresponding to the resolution ellipse label '(0,0,4.3) D' in the left panel, demonstrates the asymmetry that will be introduced in the line shape: the peak displays a high energy tail.

9.2.11   Spurions

Spurions are not actually a component of a 3-axis spectrometer, but they are an unwanted artifact whose appearance depends on the setup and components used on the spectrometer. Spurions look like peaks caused by scattering by the sample and they cause great excitement in the user because they look like a publication in Science or Nature. However, they are spurious signals caused by unfortunate combinations of events. We discuss these spurions in our chapter on data analysis (Chapter 12), this little paragraph merely serves as a warning.

9.3   General Considerations When Doing Experiments

Triple-axis spectrometers are versatile instruments that are used for measuring various excitations. Depending on the type of excitation one is interested in, the spectrometer setup will be slightly different. We will discuss examples for measuring the phonon dispersion and the magnon dispersion in single crystals, and for measuring soft modes.

We start with phonons. Phonons are excitations that can be started by, or absorbed by the neutron. The likelihood of this happening depends on how much momentum the neutron transfers along the direction of movement of the atoms that constitute the phonon wave. This makes sense because atoms will only move (in an oscillatory

pattern or not) in the direction that they are being pushed in. A guy named Newton mentioned something to this effect. This implies that the cross-section for phonon creation or absorbtion (annihilation) is proportional to

$$\frac{d^2\sigma}{d\Omega dE} \sim (\vec{q}.\vec{e})^2,$$

with $\vec{e}$ a vector of length 1 pointing in the direction of movement of the atoms. This vector is also referred to as the polarization vector of the phonon mode.

When the oscillatory movement of the atoms is in the direction of propagation, then we call this phonon excitation a longitudinal mode. When the movement of the atoms is in the direction perpendicular to the propagation direction of the phonon mode, then we call it a transverse phonon. If we want to map out a longitudinal phonon, which can be an acoustic or optical phonon, then we should measure at high momentum transfers, and make sure that most of the momentum transferred is in the longitudinal direction. For example, in a cubic system we would do a good job measuring the longitudinal phonon excitation by imparting a momentum transfer of $(4 + h, 0, 0)$ to the sample. In contrast, if we want to measure a transverse phonon, then we should impart $(4, k, 0)$ in momentum, or $(4, 0, l)$. These choices ensure that the factor $(\vec{q}.\vec{e})^2$ is large.

Why not measure at even higher momentum transfers, such as at $(8 + h, 0, 0)$, in order to further increase the cross-section? To answer this question, we first have to look into why we can measure the same dispersion using different values of momentum transfers. In other words, why do $(h, 0, 0)$ and $(4 + h, 0, 0)$ scans yield the same information about the phonon dispersion? Or, phrased differently, why can we measure in different Brillouin zones? (A Brillouin zone is the region in q-space around a Bragg peak that encompasses those q-values that are closer to this particular Bragg peak than to any other Bragg peak). The answer to this can be seen in Fig. 9.22. Shown is a neutron probing wavelength that matches the relative movements of the atoms in a transverse phonon excitation.

It is probably clear by now that a probing wavelength matching some length scale in the sample would indeed result in constructive interference. For example, in Fig. 9.22 a maximum positive amplitude of an atom partaking in the wave is probed with a phase that is 180 degrees different from the atom that is subject to a maximum negative amplitude. This probing wave (as shown in the figure) would correspond to momentum transfers in the first Brillouin zone: $(0, k, 0)$. However, we can also exactly match the motion of the atoms with a probing wavelength that is much shorter, corresponding to measuring in the second Brillouin zone: $(0, 1 + k, 0)$. Graphically (Fig. 9.22) it is easy to verify this assertion since both the probing wave with shorter wavelength, and the probing wave with longer wavelength, probe the atomic displacements with the same relative phases.
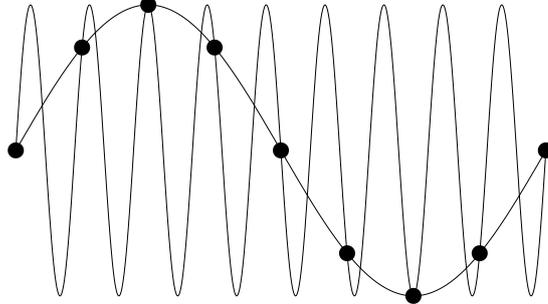
Fig. 9.22 Figure adapted from Squires' "Introduction to the Theory of Thermal Neutron Scattering". Depicted is a transverse phonon whose direction of propagation is from left to right, which is taken to coincide with the crystallographic a-axis. A probing wave with wavelength $\lambda_{\text{probe}} = d/h$ along the a-direction, samples the relative displacements of the atoms with the same phases as a probing wave with a much shorter wavelength $\lambda_{\text{probe}} = d/(1+h)$. In the depicted example, d is the separation between the atoms along the a-direction, and h=1/8.

The same reasoning holds for when we probe the atomic displacements in Fig. 9.22 with even shorter probing wave lengths, such as $(0, 8 + k, 0)$. However, in practice we would not find a large scattering cross-section for the latter case. The reason for this has to do with the Debye-Waller factor. The atoms will not only move in the direction of motion dictated by the phonon, they will also move in other directions because of thermal and zero-point motion. For instance, the atoms in Fig. 9.22 will also move from left to right, albeit in a random way. This randomness will introduce differences in the phase with which the atoms are being probed. Imagine an atom being displaced a fraction to the right by a distance $u = |\vec{u}|$ from its ideal position in the phonon wave. The change in phase angle will be given by $\vec{q}.\vec{u}$. If the probing wave length is large, then this change in phase will be small and almost negligible. However, for short probing wave lengths this change in phase will be large, even to the point that it can easily represent a change in phase of 180 degrees. All this is simply saying in words that phonons are subject to the same Debye-Waller factors as Bragg peaks are:

$$\frac{d^2\sigma}{d\Omega dE} \sim e^{-q^2 u^2} (\vec{q}.\vec{e})^2.$$

The bottom line for phonons is that large momentum transfers are helpful, but momentum transfers that are too large are not so useful because the Debye-Waller factor will quickly kill off the scattered intensity. Combining this consideration with our choice of studying longitudinal or transverse phonons, and combining it with whether we want to be on the focusing side or the defocusing side of the dispersion curve, we automatically find that there is an optimum place to look for a particular phonon. And by place we, of course, mean a particular Brillouin zone.

Propagating magnetic excitations, called spin waves or magnons, can also be studied with neutrons. The cross-section for exciting or absorbing a magnon is proportional to

$$\frac{d^2\sigma}{d\Omega dE} \sim F(q)^2 e^{-q^2 u^2} \sum_{\alpha\beta} (\delta_{\alpha\beta} - q_\alpha q_\beta/q^2) S^{\alpha\beta}(\vec{q}, E). \tag{9.1}$$

In here, $\delta_{\alpha\beta}= 1$ when $\alpha = \beta$ and zero otherwise, with α and β cartesian coordinates. This equation is a bit more formidable than the previous ones. The same Debye-Waller factor shows up because the magnetic moments are located on atoms which are still moving in a random fashion. In addition, we have a magnetic form factor $F(q)$ which is the Fourier transform of the electronic cloud. An example of the q-dependence of such a form factor was shown in Fig. 5.5; essentially, this form factor imposes that the scattered intensity diminishes quite rapidly when the probing wavelength becomes comparable to the size of the electron cloud. This is a straightforward interference effect.

The magnetic dynamic structure factor $S^{\alpha\beta}(\vec{q}, E)$ is a measure of the correlation between the magnetic moments that are pointing in the α-direction at $t'=0$, and in the β-direction at $t' = t$. When we state that it is a measure of this, we mean that we measure the Fourier transform (from time to energy) of this correlation function. How strong the scattered signal is depends on the orientation of the q-vector as given in eqn 9.1. For instance, if one is interested in $S^{zz}(\vec{q}, E)$ and if the direction of momentum transfer is also exactly along z, then the scattered intensity will be zero since in this case we have that $(\delta_{\alpha\beta} - q_\alpha q_\beta/q^2)$= 1-1= 0. Thus, in magnetic scattering the q-vector always needs to have a component perpendicular to the magnetic moment direction that one is interested in. This is a peculiarity of magnetic neutron scattering which makes it a bit harder to carry out, but it also gives the user more power to distinguish between, for instance, $S^{zz}(\vec{q}, E)$ and $S^{xx}(\vec{q}, E)$.

Note that this q-dependence to the cross-section is not restricted to inelastic magnetic neutron scattering; elastic scattering has the same q-dependence. Therefore, if one probes magnetic order where the moments are lined up along the c-axis of a crystal, then one will only observe Bragg peaks when there is momentum transferred perpendicular to the c-direction. This is in fact very handy because it allows the user to determine the orientation of the magnetic moments in an ordered structure.

In summary, magnons are measured at low q-values to counteract the effects of the Debye-Waller factor and the magnetic form factor. The Brillouin zone for measuring the magnon dispersion is chosen such that a large fraction of the momentum transferred is along the direction perpendicular to the magnetic moment direction that one is interested in studying. The considerations for focusing versus defocusing are identical between magnons and phonons.

Collective excitations in liquids and amorphous materials are always measured at low momentum transfers. We can view this case as an extreme case of the Debye-Waller

factor. Without the repeating structure that solids have, only very large probing wave-lengths will generate any constructive interference. In contrast, single atom excitations, such as the motion of an atom locked up in the cage formed by its neighbors, can be measured at very high momentum transfers for the case of atoms with a large coherent cross-section. For the case of atoms with a large incoherent cross-section, these individual-atom excitations can be measured at all momentum transfers.

Lastly, some structural phase transitions are heralded by so-called soft-mode behavior of a phonon. Below the phase transition temperature, the symmetry of the crystal is lower, and what used to be one position for an atom has now been split into two equivalent positions. At high temperatures there is enough thermal energy available to the atom to move from one position to the other, but close to the phase transition temperature it will be easier and easier to 'push' the atom from one position to the other and make it stay there.

When we do this 'pushing' using a neutron we find that we only need to transfer a small amount of energy to the atom to make it move, compared to at high temperatures where we needed to supply a substantial amount of energy. Thus, the energy requirements to excite this phonon have diminished, and the phonon energy drops close to zero. When it drops all the way to zero, then we say that the mode has softened completely, and we observe that a new Bragg peak has formed since we are now looking at elastic scattering where no energy needs to be transferred to the sample. Such a phase transition characterized by a soft-mode is a welcome object of study on a 3-axis spectrometer.

Soft-modes, as well as quasi-elastic scattering associated with a disordered phase in liquids, amorphous systems or paramagnetic solids, are fairly straightforward excitations to study using a 3-axis spectrometer. The energies involved are low and easily achievable. The main consideration in performing these experiments are related to the resolution of the spectrometer. There will always be some unwanted elastic scattering events (such as neutrons being scattered through the incoherent cross-section) that will interfere with our low-lying (in energy) excitations. The resolution has to be tight enough so that these unwanted scattering events do not mess up our sought-after signal. Other than that, the main errors that one is likely to make are in the analysis of the data, in particular in deciding whether a soft-mode is just very soft, or whether it has softened completely. We will detail how to deal with this in Chapter 12.

9.4   Exercises

Exercise 9.1

Make a sketch using the vectors $\vec{k}_{\text{initial}}$, $\vec{k}_{\text{final}}$ and $\vec{q}$ laid out in a triangle that shows the difference in measuring inelastic scans at constant scattering angle, versus doing inelastic scans at constant momentum transfer.

Exercise 9.2

What is the main problem with the following experimental procedure? We want to study the low-temperature excitations in a single crystal. We put the crystal in a cryostat, and align it using two Bragg peaks in such a way that the sample table tilts under the cryostat do not exceed a few degrees. Then, before we go home we tell the instrument computer to cool the sample down, and once it has reached the desired temperature to wait for an hour to make sure the sample is at the right temperature, and to start measuring the excitations by probing the correct Brillouin zone.

Exercise 9.3

The energy transfer from the neutron to the sample on most 3-axis spectrometers is done by varying the incident energy, while keeping the final energy of the neutron fixed. Give two reasons why this is the preferred method (in most cases) over the method where we vary the final energy while keeping the incident energy fixed.

Exercise 9.4

In order to measure the dispersion of an acoustic sound wave, whose dispersion relationship is given by $E = \hbar c q$ with c the adiabatic speed of sound, we have the requirement that neutrons need to have an incident speed that is at least as large as the speed of sound c. Use the laws of energy and momentum conversation to verify this statement.

Exercise 9.5

Use the instrumental fact sheet for HB3 (Table 11.2) to determine whether the magnetic excitations in magnetite, shown in Fig. 9.6, can be measured on HB3. Discuss the choice of monochromator, instrumental setup, anything relevant. The excitations in this cubic system occur at momentum transfers of (0,0,1/2) and (4,0,1/2), with the reciprocal lattice vector $a^* = 2\pi/a = 0.75$ Å^{-1} so that a momentum transfer of (4,0,0) corresponds to $q = 3$ Å^{-1}.

10

Other Noteworthy Instruments and Methods

In this booklet we have focused on the scattering of the neutron by the nuclei of the atoms, be it coherent or incoherent scattering. We cursorily mentioned magnetic scattering, but we have not paid much attention to the intrinsic magnetic moment of the neutron, most commonly referred to as the spin of the neutron. In this chapter we look at techniques that utilize the spin of the neutron to learn more about the sample, either through encoding information into the neutron's spin (spin echo technique), by using the spin to split the neutron in two (yes, we know) as is done in the SESANS technique, or by utilizing the orientation dependent scattering cross-section of a neutron and an atomic magnetic moment (semi-polarized neutron scattering). All three methods can only be used for very specific problems, but when they can be used they reveal information about the sample that cannot be obtained by other means.

10.1   Spin Echo Spectrometers

Spin echo spectrometers, originally developed by Ferenc Mezei and John Hayter, are used for inelastic neutron scattering experiments during which very small amounts of energy are transferred to the sample. Typical energy transfers are in the range of neV-μeV, amounts much smaller than are accessible through the scattering experiments discussed so far.

Spin echo spectrometers do not actually yield the dynamic structure factor, rather they measure the demise of correlations as a function of time. Thus, the user gets direct information about the intermediate scattering function $F(q, t)$ (See Table 4.1), the same function that is being accessed in molecular dynamics computer simulations. The time scale of the motion that is being probed is quite long (compared to the picosecond time scales one tends to measure on a triple-axis spectrometer). One typically can follow the decay of correlations from a few picoseconds all the way up to tens of nanoseconds. The ability to probe such slow decays make spin echo spectrometers well suited to study the dynamics of polymers and biological systems, but also the slow dynamics in the glass state of amorphous materials. The dynamic range of spin echo spectrometers is shown in Fig. 10.1.

The spin echo technique is based on encoding the change of speed of the neutron during a scattering event into the spin orientation of the neutron. This encoding is
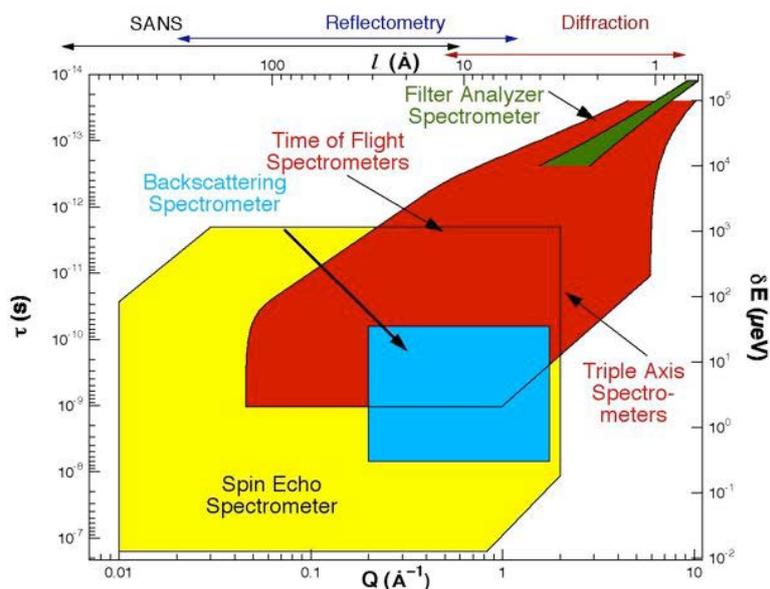
Fig. 10.1 Source: 'NCNR Summer School on Methods and Applications of neutron Spectroscopy', 2007. Shown is the dynamic range (q,E) that can be probed by various neutron scattering spectrometers. Momentum transfer is shown at the bottom, the corresponding probing length scale at the top. The energy transfer is shown on the right hand scale, and the corresponding characteristic time scale of relaxation is shown on the left hand scale. As is clear, spin echo spectrometers are instruments that probe the dynamics of systems on much longer timescales and larger length scales that more commonly used neutron scattering instruments.

done by using magic, and it is sketched in Fig. 10.2. One starts by only selecting neutrons with a particular spin orientation (called either spin-up or spin-down neutrons) using a polarizing device. We do not need to look into the details of the polarizing device, we will just assume that it works very well (as it does), and that only half of the neutrons that were part of the beam that came out of the reactor make it through this device, yielding a fully polarized beam of neutrons.

Next, the spin of the neutrons is made to precess by guiding the neutrons through a region where a permanent magnetic field is pointing perpendicular to the spin direction of the neutron. The angle over which the neutron's spin orientation precesses depends on how long the neutron takes to traverse this region of space where there is this magnetic field. This in turn depends on the speed of the neutron and, therefore, the angle of precession that the neutron endured and the speed of the neutron are related to each other in a reproducible manner. Thus, the speed of the neutron has been encoded into the spin of the neutron.
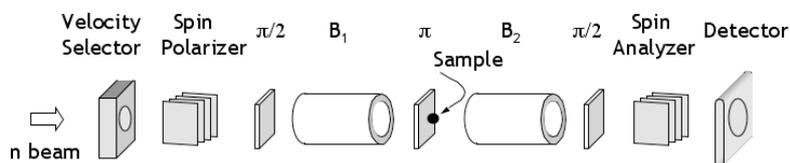
Fig. 10.2 Source: 'NCNR Summer School on Methods and Applications of neutron Spectroscopy', 2007. Sketched is a schematic diagram of the components that constitute a spin echo spectrometer. A velocity selector monochromizes the beam, which is then sent through a polarizer to select one of the two possibilities for the orientation of the neutron spin. The neutrons are then sent through a magnetic field where the orientation of the spin is forced to precess. The neutrons then interact with the sample, and the precession of the spin is reversed in the second region of identical magnetic fields. A final polarizer only allows neutrons to pass with the same spin orientation as the ones that were selected in the first polarizer. How many neutrons make it through this second polarizer is a measure at how effective the second magnetic field was in undoing the spin precession. Less than perfect undoing of the precession implies that the neutron has gained or lost some of its speed during its interaction with the sample. The labels π and $\pi/2$ denote over what angle the spin orientation of the neutron is flipped at some stages, but this flipping is instrument dependent and not essential to understanding how changes in speed are encoded in the neutron's spin orientation.

After emerging the magnetic field region, the neutron will interact with a sample. During this interaction, the neutron will transfer momentum to the sample, and it will gain or lose some speed. As always, we will gauge the amount of momentum that has been transferred by placing our detector at a certain scattering angle. The change in speed, however, is not measured by using an analyzer crystal, rather is is determined by forcing the spin orientation of the neutron to undergo the opposite precession.

After exiting the sample, the neutron is guided through a region where there is once again a permanent magnetic field, in all aspects identical to what was used before. Actually, after exiting the sample, but before entering this second region of magnetic fields, the spin of the neutron is flipped by a device called a spin-flipper (again, no details in this write-up). A spin with the opposite orientation will precess in the opposite direction and, therefore, we can expect the spin precession to be completely undone, provided the neutron is traveling just as fast as before interacting with the sample. Should this be the case, then we regain the exact same orientation of the spin as it had when it emerged from the polarizing device. Sending this neutron through a second polarizing device that only transmits (or reflects) neutrons of a particular spin orientation, we would obtain a 100% signal in the detector (that is, all neutrons that emerged from the initial polarizer and that were scattered over a particular angle, will be counted in the detector). All this is correct provided the neutron did not change its speed. This 100% recovery of the initial spin orientation is called the spin echo.

Fig. 10.3 Source: 'NCNR Summer School on Methods and Applications of neutron Spectroscopy', 2007. The spin echo spectrometer NSE located in the cold source guide hall at NIST. In the magnetic field region (green cylinder marked '1') the spins of the polarized neutrons are forced to precess, after which they interact with the sample which is located between the two magnets. The magnet on the left (labeled '2') undoes the spin precession as far as the interaction with the sample allows, and the resulting degree of polarization is measured in the black box on the far left. The scattering angle can be varied by moving magnet '2' across the floor using air pads.

Note that this spin echo effect does not depend on the speed of the neutron, it should work for neutrons of all wave lengths. This is very helpful as this implies that we can utilize neutrons with a pretty large spread in incident speeds (energies), greatly enhancing the data collection rate of spin-echo instruments. For this reason, one tends to use a velocity selector to set the (average) incident wave length of the neutron, rather than using a much more restrictive monochromator.

If the neutron did suffer a change in speed, then the precession will not be undone exactly, since the neutron will now spend more or less time in the second region that has a magnetic field. As a result, not 100% of the neutrons will make it past the second polarizer device, and our detected signal will be weaker. Exactly how much weaker depends on the change in the angle of the spin orientation, which in turn depends on the change in speed. Thus, the strength of the signal is a measure of the change of speed of the neutron. There is, of course, some math involved in figuring out exactly what the relationship between the strength of the signal and the change in speed of the neutron looks like, but the important thing to note here is that the two are related to each other and, therefore, by measuring one (the loss of polarization), we glean information
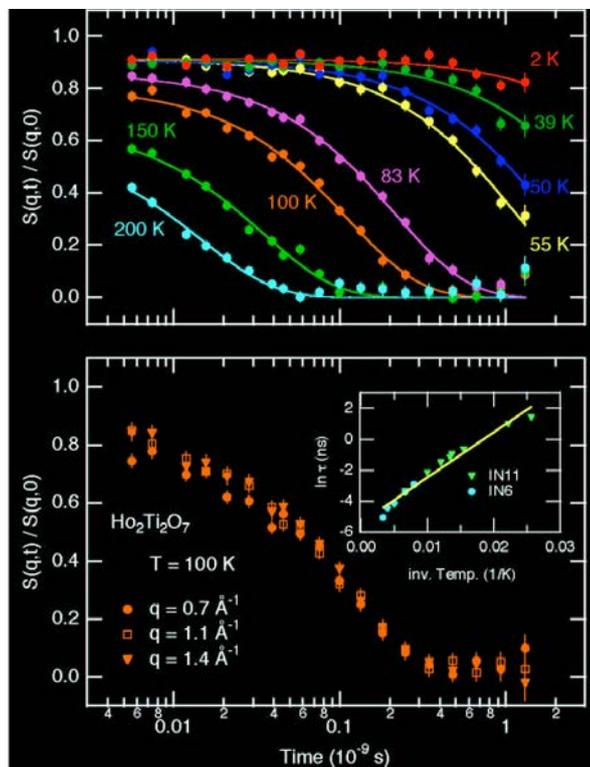
Fig. 10.4 Source: 'NCNR Summer School on Methods and Applications of neutron Spectroscopy', 2007. Data taken by G. Ehlers *et al.* on IN11 at the ILL. Spin echo data on the spin-ice system $Ho_2Ti_2O_7$ show that the relaxation time of excitations increases with decreasing temperature by two orders of magnitude (upper panel). The bottom panel shows that the relaxation time is independent of momentum transfer, implying that the processes that are being studied must have a local, single-ion origin.

about the other (the change of speed).

The above description is about the operating principles of a generic spin-echo spectrometer such as the neutron spin-echo spectrometer NSE located at NIST (Fig. 10.3), the actual details vary from instrument to instrument. As mentioned, the strength of the signal will be a measure of the intermediate scattering function $F(q,t)$. In here, we will ignore the fine details of how measuring the change of speed of the neutron through the overall loss of polarization leads to a signal that yields $F(q,t)$. Obviously, q will be directly related to the scattering angle through the Moireé interference patterns that we have seen in previous chapters.

The appearance of the time variable t is not straightforward to see. In spin-echo language it is referred to as the Fourier time, reflecting that the spin-echo instrument

effectively does a weighted average over all inelastic scattering events that produce scattering at a certain momentum transfer. Unlike the case of a detector in a diffraction experiment which simply counts all scattering events and thereby integrates over the dynamic structure factor $S(q, E)$ to produce $S(q)$, the weighing over all inelastic scattering events in a spin-echo experiment is achieved in the polarizer through a transmission rate that reflects the spin precession of the neutron, thereby introducing a term $\sim \cos(\omega t)$ (with $E = \hbar\omega$) in the weighing procedure. The averaging over all possibilities weighed by $\cos(\omega t)$ now produces the Fourier transform of $S(q, E)$, the intermediate scattering function $F(q, t)$.

The (Fourier)-time t itself can be manipulated by the user by varying the strengths of the magnetic fields. Therefore, in a spin-echo experiment, one positions the detector at a certain scattering angle, and then measures $F(q, t)$ for all accessible t's by scanning the magnetic field strength. An example of what can be measured is shown in Fig. 10.4

10.2   SESANS

SESANS stands for spin echo small angle neutron scattering and it is a technique to perform small angle neutron scattering without using restrictive collimators. It does this by using the spin of the neutron to encode information about the scattering angle. The technique itself is still under development, with the one instrument that is currently operational located at the Delft research reactor (Fig. 10.5).

The SESANS technique is used to probe very large structures, on length scales ranging from tens of nanometers to tens of micrometers. The signal that one measures is the fraction of the initial polarization of the neutron that survives, similar to the spin echo technique described in the preceding section. However, this loss of polarization is not measured as a function of (Fourier)-time, rather it is measured as a function of distance. This distance, called the spin-echo length, is a direct measure of distances in the sample. Thus, SESANS measures correlations in real space, rather than in reciprocal space as is the case for all neutron scattering techniques described thus far. Also, SESANS is an elastic scattering technique, it does not measure a signal as a function of energy transferred to the sample.

The SESANS technique was largely developed by Roger Pynn and Theo Rekveldt, and it is based on having a single neutron follow different paths through a sample depending on the spin orientation of the neutron. When the setup is just right, we can have that a neutron of known polarization 'splits' up into two. With 'splitting up' we mean that this neutron would be just as likely to follow one trajectory as it is to follow a second trajectory. As long as we do not look, the neutron will follow both trajectories simultaneously according to the main law of quantum mechanics that states that a particle behaves like a wave when one does not look.

The way this splitting is realized in practice is sketched in figure 10.6. A neutron whose spin orientation has been chosen to point in a particular direction, such as out
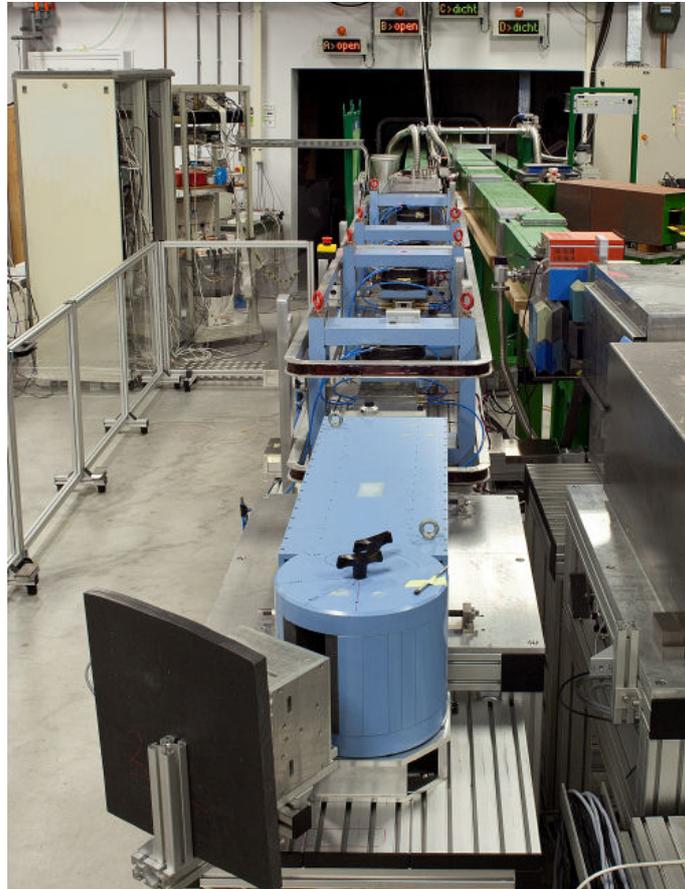
Fig. 10.5 Source: website TUDelft (www.tnw.tudelft.nl, keyword sesans). The spin echo small angle neutron scattering spectrometer located at the research reactor of the Technical University Delft (the blue instrument on the left). The silver box in the foreground houses the detector, the blue cylindrical shape has the polarizer crystal in it. Further upstream we can see the magnetic field coils.

of the page in Fig. 10.6, can be regarded as being a superposition of two states (that is, possibilities): one of a neutron whose spin is pointing to the top of the page, and one of a neutron whose spin is pointing to the bottom of the page. This neutron is then guided through a region of magnetic fields where the equivalent of diffraction will occur. In fact, it is not the equivalent of diffraction, it is identical: diffraction governed by the wave-like properties of the neutron.

The path along which the neutron is diffracted depends on its spin polarization, resulting in a real physical separation between the two possible paths. And since our neutron was equally likely to have either spin orientation, both paths will be traveled
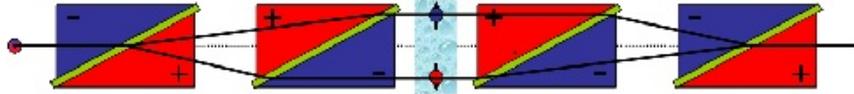
Fig. 10.6 Source: website TUDelft (www.tnw.tudelft.nl, keyword sesans). The way the SESANS technique works can be visualized by showing the two possible paths that a neutron can take. The real neutron will exist as a superposition of two neutrons taking these two paths. The two trajectories are determined (and controlled) by the magnetic fields that exist in the blue and red colored regions. The fields are constructed in such a way that the two paths merge at the exit point, unless the neutron suffered some interaction with the sample (blue area) that originated in inhomogeneities within the sample that exist on the length scale corresponding to the separation between the two paths.

with equal probability. After the neutron (both halves) has traveled through the region where we can place the sample (Fig. 10.6), the spin orientation of the neutron is flipped, and the neutron travels through identical magnetic fields, undoing any separation between the two trajectories. When we place a polarizer- that only transmits neutrons with the same polarization as the initial polarization- at the exit point of the magnetic field, then we will find that 100% of the neutrons make it through the polarizer. This is the spin echo signal similar to the signal in standard spin echo measurements. This 100% signal will only be measured, of course, if there was no sample placed at the center of the magnets.

When we place a sample in the center of the magnets, then this sample can scatter neutrons. If this sample is not homogenous on length scales corresponding to the separation between the two trajectories, then the two trajectories no longer correspond to equal likelihood of being scattered. Should this be the case, then the two trajectories can no longer merge into one identical trajectory before exiting the magnetic field region. As a result, we we lose our 100 % polarization, and not all neutrons will make it through our polarizer and be detected. Thus, we observe a loss of polarization directly attributable to sample inhomogeneities on length scales corresponding to the separation between the two trajectories (the so-called spin echo length).

By varying the strengths of the magnetic fields, the separation between the two trajectories can be fine-tuned, and we can measure the loss of polarization caused by our sample as a function of spin echo length. Spin echo lengths can range from 20 nm to 20 μm. The geometry of the magnetic fields is such that a neutron that comes in not exactly along the horizontal line sketched in Fig. 10.6 will still undergo the identical 'splitting' and 'recombining'. Thus, there is no need for tight and restrictive collimation in front of the magnetic field region. This greatly enhances the intensity of our signal compared to SANS and ultra SANS techniques.

As mentioned, this is a new technique and not all about what can be measured- and what is being measured- have been worked out. In Figs. 10.7 and 10.8, we show two
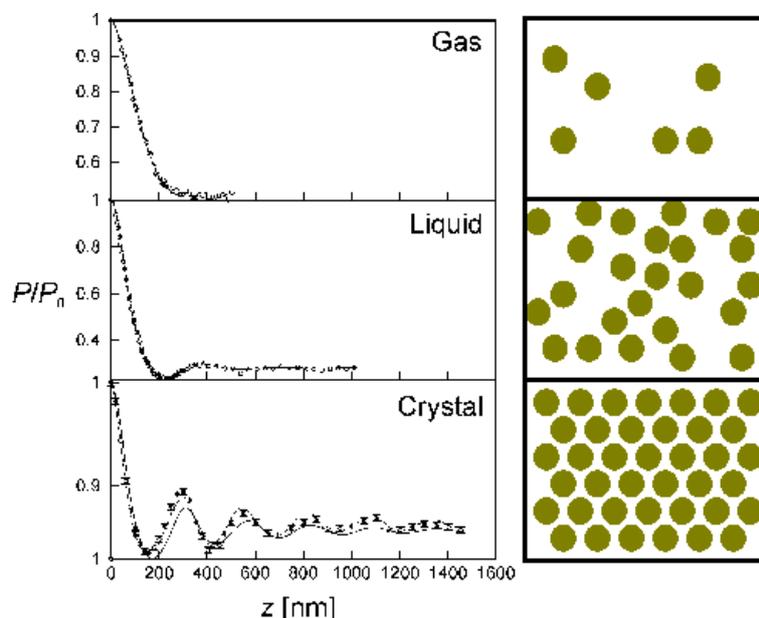
Fig. 10.7 Source: website TUDelft (www.tnw.tudelft.nl, keyword sesans). Suspensions of (hard-sphere like) silica spheres provide a good example of the direct measurement of the real space correlation function. This function measures the likelihood of finding another sphere at a distance z away from a sphere located at the origin. For a dilute suspension (top panel), the correlation function simply displays a drop when we move away from the origin, and the polarization level drops to 50 % once we go beyond the diameter of the sphere (300 nm for the silica spheres). Increasing the density of the suspension (middle panel) to liquid like densities makes a dip in the correlation function visible which can be attributed to the excluded volume effect of the hard spheres. When the density reaches close packing densities, we can clearly see the peaks corresponding to nearest and next nearest neighbor distances (bottom panel).

examples of some of the experiments that have been performed at the Delft spectrometer. As can be seen in these figures, this technique is sensitive to measuring correlations over large length scales (large compared to what is measured in SANS experiments). It is not clear at the moment whether this technique will find as widespread application as the (now) standard spin-echo technique. As an ending note in this section we mention that this technique is now being developed for use in spin-echo neutron reflectometry so that advantage can be taken of not having to deal with restrictive collimators.

10.3   Semi-Polarized Scattering

Semi-polarized neutron scattering is used when we wish to study weak ferromagnetism. A neutron can be scattered by the nucleus of an atom, or by the unpaired electrons orbiting the nucleus. The latter is referred to as magnetic scattering, and it gives us information about the magnetic order and magnetic excitations within a system. One
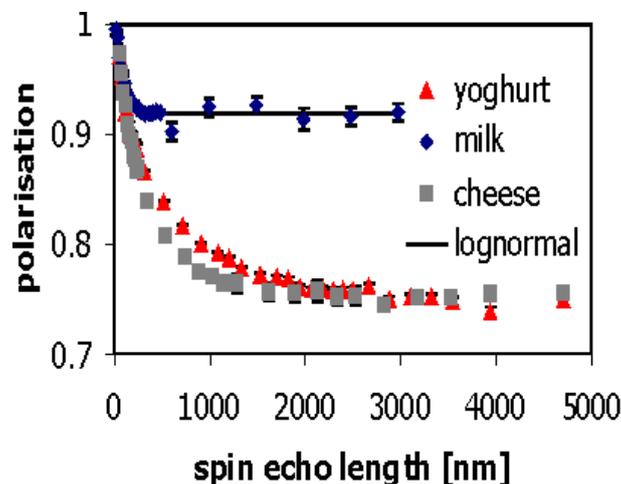
Fig. 10.8 Source: website TUDelft (www.tnw.tudelft.nl, keyword sesans). Milk tastes good because of the caseine micelles that can reach diameters of several hundreds of nanometers. Once the SESANS probing length has reached the maximum size of these micelles, then we do not see any further change in the signal (blue data points). When the milk is turned into yogurt (red points), the aggregation of the micelles into larger entities is easily observed since now the loss of polarization easily extends up to 1 μm. Turning the milk into cheese instead (grey points) reveals subtle differences between cheese and yogurt. The faster decay of the cheese correlation function compared to the yogurt one can be interpreted as being caused by the more open aggregate of the micelles in cheese.

of the problems with magnetic scattering is that sometimes it is difficult to distinguish this scattering from scattering by the nucleus. In most cases, we can employ some tricks such as mapping out the q-dependence of the scattering, since this q-dependence is different for the nuclear and magnetic scattering. We can employ more sophisticated techniques, such as measuring whether the spin of the neutron is flipped during the scattering process or not. This keeping track of the orientation of the neutron's spin, and its flipping, is referred to as polarized neutron scattering.

When we do polarized neutron scattering experiments, we first polarize the beam of incoming neutrons by only selecting those neutrons that have the desired spin orientation. When we detect the neutrons, we once again select neutrons of a particular spin orientation, corresponding to neutrons that maintained their orientation during the scattering process, or to neutrons that flipped their orientation. Polarized neutron scattering tends to be used only when the nuclear scattering interferes so much with our sought-after magnetic signal that we have to separate the nuclear scattering from the magnetic scattering.

The way polarized neutron scattering works is that some magnetic excitations involve a flipping of the neutron spin during the scattering process, whereas coherent

nuclear scattering never induces a change in the orientation of the neutron spin. However, there are some situations where we are interested in the magnetic structure of our sample for which polarized neutron scattering does not help very much. This is the case for weak ferromagnetic order.

When we do elastic scattering by the unpaired electrons, then we will not observe a flipping of the neutron spin. Thus, we cannot use polarized neutron scattering to distinguish between nuclear and magnetic scattering. This is not a problem for anti-ferromagnetic ordering since this will give rise to Bragg peaks that are forbidden for nuclear reflections. Of course, when we are interested in this antiferromagnetic ordering we do not need to resort to polarized scattering in the first place, since we will not have any problems in separating magnetic peaks from forbidden nuclear peaks (provided we are not troubled by higher order contamination).

When we are interested in a ferromagnetic material, then we will find magnetic Bragg peaks at the same positions (in reciprocal space) as the nuclear Bragg peaks. If the magnetic peaks are strong enough, then we can measure the change in Bragg peak intensities at nuclear positions to determine the strength of the magnetic signal. However, when the ferromagnetic Bragg peaks are weak, then we cannot do this separation with enough accuracy. For this situation, we use semi-polarized neutron scattering.

The method of semi-polarized neutron scattering is based on the fact that the magnetic scattering length, p, depends on the relative orientation of the magnetic moment (spin) of the neutron, and that of the atomic magnetic moment. The overall strength of the magnetic scattering is determined by the magnitude of the atomic moment and by the orbitals of the unpaired electron. The sign of the interaction, being $+p$ or $-p$, is determined by whether the neutron spin and direction of magnetic moment are parallel or anti-parallel. Of course, as is the case for nuclear scattering where the cross-section of individual nuclei is proportional to b^2, the magnetic cross-section is proportional to p^2 and therefore, the sign of p does not matter for isolated moments.

When we have nuclear and magnetic scattering contributing to the same Bragg peak, then the strength of a ferromagnetic Bragg peak will be proportional to $(b+p)^2$, or to $(b-p)^2$, depending on the relative orientation of the spin of the neutron and that of the atomic magnetic moment. In a semi-polarized neutron scattering experiments, we perform two separate measurements. In the first experiment we polarize the incoming neutron in a certain way (either up or down), and in the second experiment we choose the opposite orientation. We do not perform any polarization analysis after the scattering event, all scattered neutrons will be counted independent of their final polarization. This is why the technique is called semi-polarized neutron scattering. Next, we subtract the two experiments from each other to yield:

$$\text{difference counts} \sim (b+p)^2 - (b-p)^2 = 4bp.$$

Thus, the difference between the two experiments will be proportional to the product of the nuclear and magnetic scattering strengths. For example, we will see a strong

difference signal whenever the strength of the nuclear Bragg peak is strong, even when the magnetic scattering is only weak. This allows us to study weak ferromagnetism in the presence of strong nuclear peaks, and it is particularly useful to study the onset of ferromagnetism when cooling the sample down through the Curie temperature. As with the other techniques described in this chapter, these techniques are used only sporadically and the average neutron scatterer can lead a perfectly happy life without ever encountering the need for any of these specialized techniques.

10.4   Exercises

Exercise 10.1

What would be the reason for having a 180 degree flipper near the sample on a spin-echo spectrometer, followed by an identical second region of magnetic field, as opposed to doing away with the flipper, and having a region of opposite magnetic fields?

Exercise 10.2

In SESANS experiments, one observes a rapid, initial loss of polarization as a function of spin-echo length. An example of this is shown in Fig. 10.8. What information can be obtained from the steepness of the slope? What information can be otained from the saturation polarization level for large spin-echo lengths?

Exercise 10.3

In our discussion on semi-polarized scattering, we simply use expressions such as $(b + p)$ and $(b - p)$, whereas we are looking at an interference problem so that we should include phase factors such as $e^{i\vec{q}\cdot\vec{R}_i}$ for atom i. What is the reason that we did not have to worry about those phase factors in our discussion, even though we were looking at a magnetic system?

Part III

Pre- and post-experiment

This (short) part deals with how to prepare for an experiment, how to get beam time, how to complete a successful experiment, and how to analyze the data.

11

Planning a Neutron Scattering Experiment

As in any experiment, the best laid plans of mice and men can be thwarted by circumstances out of our control. However, it still pays off to think an experiment through beforehand as opposed to just showing up with a sample to measure. The more experiments one does, the more routine these experiments become. This is pretty much obvious. This chapter can be used as a checklist that should help in avoiding some common mistakes. Especially when one has (finally) been granted beam time at a national user facility, it would be nice if the data were to lead to a publication rather than to a return to take the missing data.

This chapter is written as if we are planning for an experiment at a national user facility. We will first give some general considerations, and then we will go over two actual examples, one in hard condensed matter, and one in soft condensed matter.

Once one has identified what one wants to learn from studying one's sample, then the first step is to identify the best spectrometer for performing this study. Globally speaking, if one is interested in the structure of the material, then the choice of spectrometer is not as stringent as when one is interested in the dynamics of the sample. Choosing a spectrometer at a reactor source tends to be slightly more straightforward than choosing one at a spallation source, but in general similar considerations apply to all cases.

11.1   Choice of Spectrometer

When one is interested in the structure of a material, the only thing to make sure of is that the spectrometer has the desired range that allows one to determine all the characteristics of the structure. For instance, if the structural units are large, then the spectrometer has to be able to detect neutrons that are being scattered at a sufficiently small angle so that the probing wavelength determined by the interference pattern (Fig. 2.3) matches the size of the structural units. Typically, this range is listed in the spectrometer characteristics as the d-space range of the instrument. By the same token, in order to be able to accurately determine the positions of the atoms within the structure, one needs to be able to access small d-spacings as well. As a rule of thumb, large d-spacings are used to figure out the overall structure, such as the size of the unit cell and the space group of the crystal, whereas smaller d-spacings are

needed to figure out which atom sits where in the unit cell and how far they wander from their time-averaged positions. Smaller d-spacings are probed using smaller probing wave lengths (Fig. 2.3).

It is also important to note that spectrometers that determine the structure of materials (these spectrometers are called diffractometers) are, as a rule of thumb, the most versatile and the most powerful at spallation sources. The main difference is that spallation sources have diffractometers where the detectors cover a much larger solid angle than they do at reactor sources and, therefore, experiments can be done faster at spallation sources. In addition, at spallation sources it is much more easy to separate out Bragg reflections that occur at angles that are close together. This is important when one has a great many different atoms in the unit cell, such as is the case in protein crystals. Having said that, the (standard) powder diffractometer that can be found at most reactor sources still does the job it was designed for, namely, determining the structure of materials that are not too complex or changes in that structure when the system undergoes a phase transition.

When one is interested in the motion of the atoms, then the choice of spectrometer plays a much more important role and there are many more considerations to take into account. First of all, the dynamic range of the spectrometer has to be able to allow one to follow the movement of the atoms corresponding to the excitations one is interested in. In particular, it will take a certain amount of energy to start a sound wave, and a certain amount of momentum has to be transferred before the sample can be excited. The range of energy transfers that one is interested in largely determines the choice of spectrometer.

Suppose one is interested in an excitation that requires 40 meV to ignite, and compare this to studying an excitation that only requires 0.1 meV to get going. These two problems will require two different spectrometers and spectrometer settings. Typically, the spread in energies of the incident and scattered neutrons is of the order of 1%. We somewhat discussed these numbers when we discussed the resolution function of a spectrometer in previous chapters; in here, we will leave those details aside for a moment.

In order to be able to transfer at least 40 meV in energy from the neutron to the sample, the neutron needs to have at least this amount of incident energy. Therefore, the spread in incident energies will be of the order of 0.4 meV. This is okay for studying an excitation that requires 40 meV, but what for the 0.1 meV excitation? Clearly, the settings of the spectrometer would not work since a spread of 0.4 meV in incident energies would completely wash out an excitation that requires on 0.1 meV to get started.

In addition, samples and the stuff that keeps the sample at the right temperature (sample environment) always scatter neutrons without changing their energies. This is called elastic scattering and it presents a problem when studying low lying (in energy) excitations such as an 0.1 meV excitation. The unwanted elastic scattering tends to
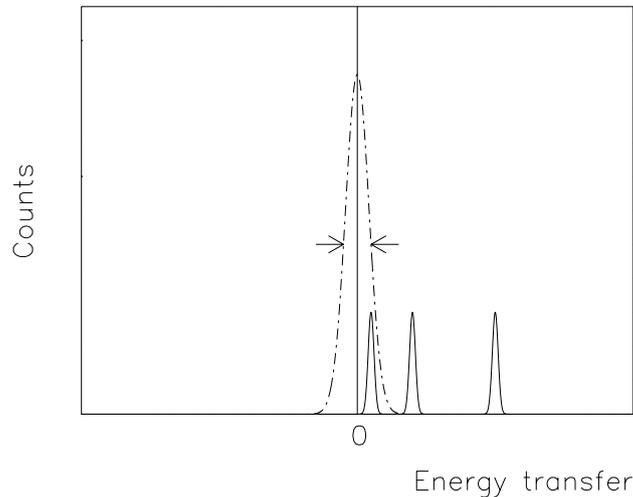
Fig. 11.1 Because of the finite resolution function of the spectrometer, the elastic scattering (dashed-dotted curve) caused by sample holders and incoherent scattering, will be spread out over a range of energy transfers. The characteristic energy resolution width is indicated by the arrows. For a given amount of elastic scattering combined with a given resolution width, inelastic excitations can either be drowned out (solid curve within the elastic scattering contribution), or they can be well separated (the two other solid curves) from the elastic intensity.

be quite strong, and given the spread in energies of the incident neutron, the elastic scattering events will sometimes show up at positions that do not strictly correspond to zero energy transfer. This is shown in Fig. 11.1 and it is called the energy resolution function of the spectrometer. Obviously, if one wants to study excitations at 0.1 meV, then the elastic scattering should no longer be a problem by leaking into this energy range.

For the latter case of the low energy excitation, one would select a spectrometer that uses incident neutrons of low energy, perhaps something like 1 meV. In this case, the spread in incoming energies would be of the order of 0.01 meV, and they would not interfere with an excitation that is located at 0.1 meV. (See Fig. 11.1). Neutrons with such low incident energies (See Table 3.1) require a cold source to produce, so one would be looking for a cold source spectrometer.

Sometimes one is not only interested in studying whether a particular excitation exists and at what energy, but also in how long the excitation lasts for, i.e., how long it takes to decay back to equilibrium. When studying the lifetimes of excitations, one has to think very carefully about the spread in energies of the incident and scattered neutron. In other words, one has to think very carefully about the energy resolution

of the spectrometer. There is a direct relationship between the energy resolution of a spectrometer, and the lifetime of the excitations that one can study. The better the energy resolution, the longer (in real time) one can follow the demise of an excitation that the neutron started in the material. We give some numbers in Table 11.1.

Table 11.1 Correspondence between the energy resolution of a spectrometer and the time an excitation can be followed for.

Energy resolution [meV]	characteristic time
1	0.7 ps
0.1	7 ps
0.01	70 ps
0.001	700 ps

The way to read this table is the following. If an excitation might persist for 10 ps, then one should not use an energy resolution of the order of 1 meV; with such a resolution function one can only determine what happens during the first 1 ps of the existence of the excitation, but is is not possible to follow the demise of the excitation over the full 10 ps. To study (the lifetime of) such an excitation one would need an energy resolution of 0.1 meV or better. And since typical typical energy resolutions are of the order of 1% of the neutron energies used in a scattering experiment, one should use incident neutrons that have energies less than about 10 meV. This would be an experiment that could benefit from a cold source spectrometer, or from a spectrometer setup where one uses 5 meV neutrons from a thermal source with the approximate beryllium filters in place.

What if one wants to measure a high energy excitation (such as a 50 meV excitation) with a very good energy resolution? Here one runs into fundamental limitations of a thermal source spectrometer: in order to be able to transfer 50 meV, one will likely end of with an energy resolution of at least 0.5 meV. One can try to improve upon this energy resolution by using various collimators, but improvements will be very difficult to achieve. For situations like this, one is much better off at a spallation source: at spallation sources the energy resolution actually improves when the neutron gives up more and more of its energy in a scattering event. An example of this is shown in Fig. 11.2. At thermal reactor sources the opposite tends to be the case: the more of a neutron's energy is transferred, the worse the energy resolution becomes. This is yet another consideration to take into account when looking for the best spectrometer to do one's experiment.

The bottom line is that when choosing a spectrometer one has to take into account the dynamic range of a spectrometer (how much energy can be transferred) as well as
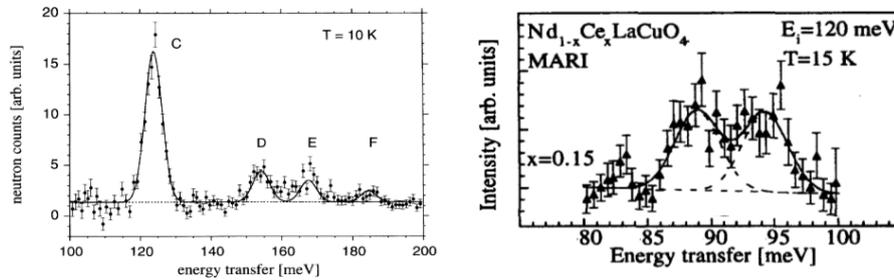
Fig. 11.2 Figure reproduced with permission from M. Guillaume *et al.*, Phys. Rev. Lett. 74, 3423 (1995) (left panel) and from M. Gutmann *et al.*, Physica B 234-236, 812 (1997) (right panel) with permission from Elsevier. High energy (crystal-field) excitations in $SmBa_2Cu_3O_7$ (left) in Ce doped $NdLaCuO_4$ (right) measured on the MARI spectrometer at the ISIS spallation source in the United Kingdom. Note that the energy resolution at such high energy transfers is so good that two excitations that are closely spaced together (in energy) can still be observed as being separate excitations; they have not merged into each other because of the (resolution function) allowed spread in transferred energies.

the energy resolution of a spectrometer (will the resolution wash out the signal, will it allow for following of the demise of the excitation). These ranges and resolutions are available on the websites of the instruments at national user facilities. We reproduce such a table below for the HB3 spectrometer at the HFIR reactor at Oak Ridge national Laboratory.

Table 11.2 The instrument specifications for the HB3 spectrometer at the HFIR reactor at Oak Ridge National Laboratory. This table has been adapted from the HB3 web pages.

component	options
Monochromators	Variable vertical focusing PG(002), Be(002), Si(111)
Analyzer	Fixed vertical focusing PG(002), Be(101), Si(111)
Mono take-off angle	$18^o < 2\theta_m < 75^o$
Scattering angle	$-90^o < \theta < 120^o$
Resolution (elastic)	5-10% of E_{initial}
Collimator (pre-mono)	15', 30', 48'
Collimator (mono-sample)	15', 30', 48'
Collimator (sample-anal)	20', 40', 60', 80'
Collimator (anal-detector)	30', 70', 90', 120', 210', 240'

The way to read Table 11.2 is the following. One's first reaction probably is: yikes, time to call the instrument scientist. However, this table was part of the spectrometer's fact sheet, and it stated on the fact sheet that the spectrometer covered an incident neutron energy range of 2-100 meV. This is a nice shortcut to having to look up the monochromator's lattice spacings, and combining it with the monochromator take-off angles listed in the table to calculate the range of incident energies. Having said that, these calculations are not that difficult, just a little time consuming.

Having either found the range of incident neutron energies in the write up, or calculated from the table, then one can figure out the energy resolution. For HB3 (Table 11.2), it states that it is 5-10% of the incident neutron energy. Thus, if one choose an incident neutron energy of 25 meV, then the energy resolution will be of the order of 1-2 meV. This is a large number, and we say that the energy resolution is not very tight, definitely not of the order of 1% that we talked about before. This is an example of how different spectrometers are designed differently. HB3 is designed to measure the positions (in energy) of excitations, but not to study their lifetimes. For a lot of systems this is good enough. For instance, in superconductors one is interested in the excitation energies which would allow one to say something about the glue that holds Cooper pairs together, not in whether a particular excitation takes 1 ps or 0.1 ps to decay.

For the interested reader, HB3 uses a focusing monochromator. This greatly enhances the number of neutrons that hit the sample, but it increases the spread in neutron energies that are allowed to be reflected out of the main beam onto the sample by allowing for a larger spread in angles to be Bragg reflected at the monochromator. In short, it looks like this spectrometer (HB3) is very well suited for small samples and/or for studying weak excitations over a large dynamic range, but it would not be one's first or second choice for studying the separation (in energy) between two crystal field levels (Fig. 11.2). The majority of experiments that tend to be performed on HB3 are hard condensed matter experiments such as experiments on high-temperature superconductors.

Let's compare HB3 to the SPINS spectrometer (NG5) at NIST. According to the instrument website on the NIST web pages, the incident neutron energy can be varied between 2.4 and 14 meV, with an energy resolution ranging from 0.02 to 1 meV. These numbers state that when choosing an incident energy of 2.4 meV, the best resolution possible is 0.02 meV. It should not be read as being able to have an energy resolution of 0.02 meV using an incident energy of 14 meV (when in doubt, contact the instrument scientist).

Poking around on the NIST website more, one can look up the details of the High Flux Neutron backscattering spectrometer (NG2, details under the instrument performance menu option). This instrument is designed for energy transfers of up to 50 μeV (0.05 meV) with an energy resolution of the order of 1 μeV. Clearly this is an instrument designed to follow very slow motions that take place on timescales of up to

1 nanosecond. The website also states that the instrument maintains a high neutron flux despite its highly stringent energy resolution requirements. This high energy resolution in combination with a high neutron flux was achieved at the cost of sacrificing the angular resolution of the instrument.

This brings us to considerations related to the angular resolution of a spectrometer. So far we have been talking mostly about the energy resolution of a spectrometer, but the angular resolution and angular ranges are important too. However, these considerations are of secondary importance as a spectrometer that does not have the required dynamic range will not be useful. It is only when we have narrowed down our choice of spectrometers that we should be looking into the angular range and resolution.

The amount of momentum transferred from the neutron to the sample depends on the scattering angle, as can be seen in Fig. 2.3. If we do not know the scattering angle with great accuracy, we also do not know the momentum transfer with great accuracy. This implies that during the course of an experiment we are averaging over scattering events that span a range of momentum transfers. This is similar to the averaging that happens when the experiment spans a range of energy transfers (as dictated by the energy resolution function of the spectrometer). So the question is, do we have to worry about this averaging over a range of momentum transfers?

Most materials have excitation energies that depend on how much momentum is transferred. This is called the dispersion relation of the material, and two examples of dispersion relations are shown in Figs. 4.11 and 9.7. Looking at such a relation (such as the sketch in Fig. 11.3), imagine that we average over a small range momentum transfers because of the less than perfect angular resolution of the spectrometer. This would imply that the energy ranges of the excitation would vary accordingly as dictated by the slope of the dispersion curve. If the spectrometer's energy resolution should exceed this range, then we would not have to worry about this lack of angular resolution; however, if we had picked an energy resolution that is much better than this range, then we would have negated any positive effect we obtained from using such a tight resolution with the result that we would have thrown away neutrons for no good reason.

In other words, it is the range of excitation energies within a range of momentum transfers that matters. This range should be less than the energy resolution. If not, we are throwing useful neutrons away. These considerations also tell us whether we have to worry about the angular resolution in the first place. If the dispersion does not vary very much with momentum transfer (we call this a flat dispersion), then we do not have to worry. Conversely, if the dispersion is very steep, then we probably should not be using instruments such as NG2. For a typical spectrometer, the angular resolution tends to be such that we do not have to worry; it is just for atypical spectrometers such as NG2 that angular resolution considerations tend to become important.
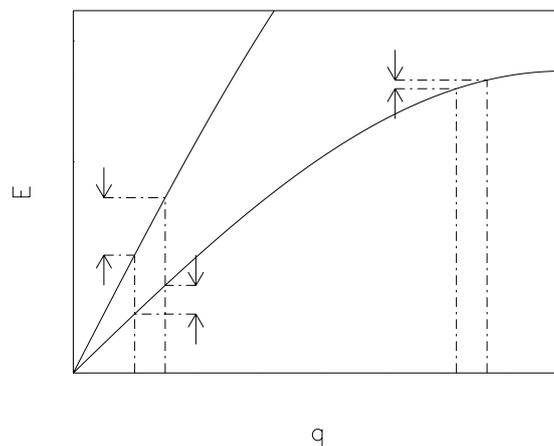
Fig. 11.3 The angular resolution (and the accompanying q-resolution) leads to a spread in energies depending on the slope of the dispersion, and on which part of the dispersion is being measured. In this sketch, the q-resolution is given by the distance between the vertical dashed lines; the resulting spread in energies is indicated by the arrows in the figure.

However, certain dispersions are very flat as they represent the excitations associated with individual molecules, such as the rotational states of a molecule confined to a cage formed by its neighbors. Such local excitations are not collective excitations, and as such they do not show any dependence on scattering angle. Dependence on scattering angle only occurs when one neighbor influences the next. Such localized excitations are ideal for instruments like NG2. An example of such a local excitation is shown in Fig. 11.4.

11.2   Ancillary Equipment

After the selection of spectrometer, we want to make sure that the required sample environment is available, and that it can be used on the spectrometer we have selected. In most cases, sample environment includes such items as cryostats, cryomagnets, furnaces, pressure cells, and hydration chambers. Most instrument websites do not have specific details on what piece of equipment fits the spectrometer, and which piece of equipment is being refurbished (a.k.a. broken). Since it is always good practice to contact the instrument scientists anyways (especially as this tends to improve the chances of being awarded beam time), this would be a good time to discuss with the local contact whether the desired equipment is available for use on the instrument. Other than these generic notions, beyond checking what equipment is available in general at the user facility, there is not much more that can be learned without actually contacting someone.
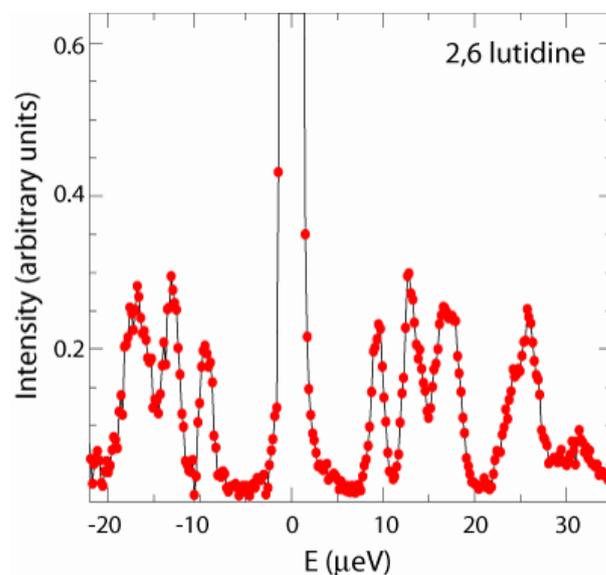
Fig. 11.4 Shown are the excitations resulting from a molecule tunneling from one state to another state. These excitations do not depend very much at all on the amount of momentum transferred between the neutron and the sample and, as such, they are ideal candidates to be studied on specialized spectrometers such as NG2 at NIST. The data on lutidine shown here were summed over a range of momentum transfers between 0.6 and 1.6 Å^{-1}. Note the scale on the horizontal axis as being micro-eV, not milli-eV. Source: NIST NG2 instrumental website.

11.3   Performing the Experiment

The success of an experiment depends in a large part on how much the user does during an experiment. Experiments can be hard work, and not necessarily a lot of sleep is possible. Therefore, always make sure that there is enough manpower available, do not count on the instrument scientist to do (all of) the hard work for you.

When doing an experiment, try to analyze the data as you take them. Not only does this make sense because the instrument scientist will be there to help you and make you more familiar with the analysis software, it gives you a chance to fine tune the experiment. This implies that the experiment may well have to start with an empty sample cell run, or some other background scan.

Be prepared that you will not be able to take your sample home with you. It might be stuck in a cryostat, but more likely, the sample will have become radioactive and will need to be stored for a while at the user facility before it can be shipped. You can estimate beforehand how long it will take your sample to cool off, but these are estimates at best. Elements in particular to be on the look out for that get radioactive

are the ones that have a sizeable neutron absorbtion cross-section. The NIST website has a feature that allows the user to input the chemical formula of the sample and obtain an estimate of the activation of the sample. This feature can be found at www.ncnr.nist.gov/cgi-bin/neutcalc.

One thing that tends to pay off is to do the following: ask the instrument scientist to take a neutron picture of the sample while it is in the beam. This way, you will learn how the sample is actually situated in the beam (is it in the middle, etc.), and it will give you piece of mind when something odd pops up during the analysis stages. Also, try to make sure that you have transmission data of the empty cryostat and of the cryostat with the sample in it. There are occasions that you will need this during the analysis, and it only takes a minute to gather these data.

Finally, sometimes stuff breaks or does not work out the way it was originally planned. While it seems to make sense to have a back-up sample on hand, in most cases you will not be allowed to substitute one experiment for another. This actually makes sense as beam time for a similar experiment might already have been granted to another user. So if something does break, do not dwell on it.

11.4   Writing a Beam Time Proposal and Examples

The stage where a beam time proposal is being written also tends to be the stage where some issues come to light; therefore, do not postpone writing the beam time proposal until the last second. Beam time proposals are reviewed by quite a few scientists. Not all of them will be familiar with the nuances of the particular field, but they will be familiar with the ins and outs of neutron scattering experiments. The beam time proposal should be written so as to make the life of the reviewers easy.

Start the proposal with a one or two sentence summary of what you want to measure. Follow this by a brief background of the field, by what the open questions are that you hope to resolve, and why resolving them is important. Then explain the experimental procedures. This step is quite frequently added as an afterthought, but it tells the reviewer that you have given the experiment the necessary amount of thought and it makes it easier for the reviewer to judge whether the requested amount of beam time is correct and adequate.

The experimental procedure should make clear that time has been reserved for completing not just the data taking part, but also for completing background scans (such as empty sample holder, reference data at a temperature above that of a certain phase transition, etc.) and for cooling down samples, changing samples, and so forth. Ideally, the requested time also deals with measuring the resolution function of the spectrometer, when needed and when not known. In short, make sure the reviewer can follow the plan of action and that it leaves no doubt that all the data for a successful publication will be taken.

The following are two examples of beam time proposals. The first one was put in to Oak Ridge National Laboratory for a hard condensed matter experiment; the second one was put in to NIST for a soft condensed matter experiment. We have blocked out the names of the applicants. Both proposals were granted beam time, but their ranking by the reviewers is unknown.

HFIR
Center for Neutron Scattering

NEUTRON SCATTERING RESEARCH PROPOSAL FORM

Submit all proposals to: Oak Ridge National Laboratory, HFIR Center for Neutron Scattering, User Services Office, Oak Ridge, TN 37831-6393. Phone: 865-574-5231   Fax: 865-576-7747   Email: ns_user@ornl.gov

Title of Proposal: Identity theft at the quantum critical point	Date Submitted: June 2nd, 2004

Status: ☒ New   ☐ Continuation of Proposal No. _____*   ☐ Resubmission   ☐ Rapid Access
* If this is a continuation, a report on previous experiments must be included

Principal Investigator: ▉▉▉▉▉▉▉▉   Phone: 865-576-6069

Institution & Mail Address: HFIR

FAX:

Email:

Citizenship: USA

Co-Proposer (attach additional sheets if necessary)	Institution and Address	Citizenship	Email	Please Check If Attending	Please Check If First-time User
▉▉▉▉▉	▉▉▉▉▉▉▉▉	▉▉▉	▉▉▉▉▉	☒	☐
▉▉▉▉	▉▉▉▉▉▉▉	▉▉▉		☒	☐
▉▉▉▉	▉▉▉▉	▉▉▉▉		☐	☐

Scheduling Information			
Desired Dates ASAP	Impossible Dates	Est. Beam Time (days) 8	Instrument(s) HB3

Have You Contacted an ORNL Staff Member to Discuss The Feasibility of Your Experiment?

☐ No   ☒ Yes   Contact Name: M. Yethiraj

Please Categorize Your Proposal (For DOE Reporting Purposes)

Research Area		Organization (check all that apply)	Funding Agency (check all that apply)
☐ Materials Sciences	☐ Polymers	☒ University	☒ DOE
☒ Physics	☐ Engineering	☐ Industry	☒ NSF
☐ Chemistry	☐ Geoscience	☒ Government	☐ DOD
☐ Biology	☐ Other:	☐ Foreign	☐ NIH
		☐ Other:	☐ Other:

Proposal Review Committee: ☐ Large Scale Structures   ☐ Diffraction   ☒ Inelastic   ☐ Other

Sample/Instrument Information

Instrument Configuration: ☒ Standard   ☐ Nonstandard

Description of Sample (Check Applicable Fields):

☐ Powder   ☒ Single Crystal   ☐ Polycrystal   ☐ Polymer   ☐ Liquid   ☐ Thin Film
☐ Other (please describe):

Crystallographic Information - Space Group (if applicable):	Range of Momentum Transfer (Å^{-1})

Crystallographic Information - Space Group (if applicable):

a = <u>4.1</u> (Å)   alpha = <u>90</u> (°)
b = <u>4.1</u> (Å)   beta   = <u>90</u> (°)
c = <u>10.4</u> (Å)   gamma = <u>90</u> (°)

Range of Momentum Transfer (Å^{-1})
0.2 - 4
Range of Energy Transfer (meV)
-3 to 75
Chemical Formula
$Ce(Ru_{1-x}Fe_x)_2Ge_2$ with x= 0.76

Size: 6x0.5x0.7 cm   Weight: Mosaic (if applicable):   No. of Samples: 1

How will sample be transported?

☒ Hand carry   ☐ Mail, UPS, Fed Ex   ☐ Other (please describe):

Resources Needed/Special Requirements (please specify)

☐ Sample Preparation   ☐ Mechanical Preparation
☐ Chemical Preparation   ☐ Sample Storage
☐ Computer Requirements   ☐ Sample Environment
☐ Health Physics   ☐ Other:

Sample Safety Issues (*attach MSDS for each material)

☒ No major safety issues   ☐ Corrosive Material*   ☐ Toxic Material*   ☐ Cryogenic
☐ Flammable Material*   ☐ Radioactive Material* ☐ Explosive Material   ☐ High Pressure
☐ Carcinogenic*   ☐ Biohazardous*   ☐ Electrical   ☐ Other:

By signing or by electronic submission, I certify that the above information is correct to the best of my knowledge.

Signature: **Printed Name: Mohana Yethiraj**   **Date: June 2nd, 2004**

Safety and Feasibility Reviews

☐ No further safety review required   ☐ To be reviewed by Safety Officer

Recommended number of days: | Recommended instrument:

Comments:

Instrument Scientist signature: | Date:

Facility Safety Officer signature: | Date:

DETAILED STATEMENT OF THE EXPERIMENT OR PROPOSED RESEARCH

(Include scientific context; relevance of proposed experiment; why neutrons as opposed to other techniques; preliminary work performed using neutron scattering and other techniques; details of experimental approach. ***Proposals are limited to two pages.***

IMPORTANT! List or attach a list of publications resulting from HFIR neutron scattering experiments for last three years. If this proposal is for a continuing experiment, attach appropriate experiment report or the new proposal will not be submitted for review.

During the past decade, a multitude of strongly-correlated metallic systems [1] have been uncovered in which the low-temperature response departs from that predicted by the standard Landau Fermi-liquid (Fl) theory for metals. This non-Fermi-liquid (nFl) behavior is observed in systems that have been prepared to be on the verge of ordering magnetically at T= 0 K, the quantum critical point (QCP).

Recent IN6 experiments [2] have shown that the intermetallic compound $Ce(Ru_{1-x}Fe_x)_2Ge_2$ is unique among quantum critical point systems [1]. Like other compounds, $Ce(Ru_{1-x}Fe_x)_2Ge_2$ exhibits nFl behavior when prepared to be at the QCP (x= 0.76, see Fig. 1). nFl behavior manifests itself in the specific heat, susceptibility and most strikingly, in the scaling behavior of the dynamic response (see Fig. 2). However, in contrast to QCP-systems studied before [3], $Ce(Ru_{1-x}Fe_x)_2Ge_2$ evolves towards an ordered phase in a manner that cannot be accounted for using either of the two scenarios proposed in the literature [4,5] for the physics near a QCP: the response is somewhat similar to the predictions for both scenarios, yet it is in qualitative disagreement with both of them.

However, the IN6 experiments were performed on polycrystalline samples over a limited (q,E) range, and as such, left many questions unanswered. We have now grown a large single crystal (~2 cm^3) of the quantum critical composition which, combined with the flexibility of a triple axis spectrometer, should allow us to complete the investigation of this unique system. By probing the directional information in this magnetically anisotropic crystal over a larger q and E-range, we will answer the following fundamental questions about this unique system, and about quantum phase transitions systems in general.

- What is the exact form of the scaling function for large E/T values?
- What is the ordering wavevector (we only know the magnitude, not the direction), how different is the response along the easy axis compared to the hard direction? Do both directions show scaling?
- Does the sample also exhibit H/T scaling? (follow up experiment)
- Is the lineshape in q-space near the ordering wavevector Lorentzian, or does it involve an unusual exponent as is the case for the lineshape as a function of energy [2]. If so, is this the same exponent?
- Can S(q,E) be separated in a localized moment and a itinerant electron contribution, i.e., do the two subsystems evolve independently?

Experimental details:
We would like to perform the experiment for q-values in the q-range of 0.2- 4 Å^{-1}, for energy transfers up to 75 meV to ensure maximum E/T range. For the same reason, we require a minimum temperature of 2K. Most scans will be centered (in q and E space) around the ordering wavevector of the parent compound, measuring departures from the quantum critical point along the q-axes (two independent directions), E, and T directions. These measurements will enable us to answer the $64,000 dollar question: **whether an overall scaling function exists that shows that all variables can be treated on the same footing** [E and T are all just measures of distance to the QCP, as well as the generalized Weiss temperature Θ_q]. In other words, have all characteristic time and length scales been lost at the QCP, is it a sample without identity? Indications are that this is indeed the case in $Ce(Ru_{1-x}Fe_x)_2Ge_2$

3

We ask for an initial allotment of 8 days of beam time on HB3. This includes searching for the incommensurate ordering wave vector, whose magnitude is known from polycrystalline experiments. We prefer HB3 over HB1 since we are investigating quasi-elastic magnetic scattering which is somewhat intensity limited, while at the same time the energy signature of the critical fluctuations extends up to very high energy transfers. In a future experiment, we would also like to investigate the magnetic field direction to identify the H/T scaling. However, for the current experiment, any helium flow cryostat will work.

Fig. 1: Ce$(Ru_{1-x}Fe_x)_2Ge_2$ phase diagram. The QCP is at x= 0.76, the arrow denotes our single crystal sample. nFl stands for non-Fermi-liquid, Fl for Fermi-liquid, AF for the anti-ferromagnetically ordered phase, Fe for ferromagnetic order.	Fig. 2: Scaling behavior of the neutron scattering data at the QCP. The data are plotted as a function of the scaling variable E/T, for four q-values denoted in the figure. Every curve represents 3000 data points, with temperatures ranging from 1.8 (darkest symbols to the right) to 200 K (lightest symbols to the left). The curves are offset by a factor 5.

[1] G. R. Stewart, Rev. of Mod. Physics 73, 797 (2001), and references therein.
[2] ████████████████████████
[3] e.g., A. Schroeder et al., Nature 407, 351 (2000); M.C. Aronson et al., Phys. Rev. Lett. 87, 197205 (2001).
[4] A. J. Millis, PRB 48, 7183 (1993).
[5] Qimiao Si et al., Nature 413, 804 (2001).

Previous HFIR publications:

████████████████████████████████████

████████████████████████████████████

NIST Center for Neutron Research

Proposal for Neutron Beam Experiment

Submission ID:17273   Proposal Number: null

Experiment Title
Title: Studies of Model Membranes
Proposal Type: New Proposal
Time Received: 04-OCT-11 23:46

Scheduling
Desired Dates: 06/22/2012 to 07/31/2012
Impossible Dates: 03/01/2012 to 06/21/2012
Estimated Duration: 8 days

Participants

	Name	Address	Country	Telephone/e-mail
�	▬	▬	▬	▬
▬	▬	▬	▬	▬
▬	▬	▬	▬	▬
▬	▬	▬	▬	▬
▬	▬	▬	▬	▬
▬	▬	▬	▬	▬

Instrument

Instrument Requested:	NG-2 -- HFBS, High-flux backscattering spectrometer (CHRNS)
Suggested Local Contact:	Madhu Sudan Tyagi
Instrument Resolution:	1 microeV
Instrument Configuration:	Default instrument configuration

Sample Description

	Sample 1
Name	DLPG: 1,2-dilauroyl-sn-glycero-3-phospho-(1'-rac-glycerol) (sodium salt)
Chemical Formula	$C_{30}H_{58}O_{10}PNa$ + H_2O (low hydration)
Mass (grams)	2 mg + 117 g Si
Form	Amorphous Solid
Temperature Measurement Range (K)	250 - 400 K
Number of Runs	8
Total Collection Time (hrs)	96
Sample Availability	2012-03-01 00:00:00.0

	Sample 2
Name	DLPG: 1,2-dilauroyl-sn-glycero-3-phospho-(1'-rac-glycerol) (sodium salt)
Chemical Formula	$C_{30}H_{58}O_{10}PNa$ + D_2O (high hydration)
Mass (grams)	2 mg + 117 g Si
Form	Amorphous Solid
Temperature Measurement Range (K)	250 - 400 K
Number of Runs	8
Total Collection Time (hrs)	96
Sample Availability	2012-03-01 00:00:00.0

Sample Environment

Sample Environment Equipment: CCR; 30-600K

Special Requirements

Please describe any non-routine needs for sample temperature, magnetic field, etc., or other ancillary equipment. Specify any equipment needed at NIST for sample loading, treatment, storage, etc. (inert atmosphere, refrigeration, dry box, etc.). Also describe any equipment you plan to bring to NIST.

Safety

Check *at least* one box that describes your sample
[X] No Hazards
[] Toxic [] Corrosive [] Radioactive [] Explosive [] Flammable
If there are any hazards associated with your proposed experiment, please indicate how any risks are to be handled.

Categorization

For reporting purposes, please categorize your proposal:

Research Area:	Biomolecular Science
Funding Agency:	NSF National Science Foundation

Publications

(If you have been allocated beam time through proposals to the NCNR during the past three years, please list the instrument and resulting publications)
HFBS:
"Study of the Dynamics of on Single-supported Bilayer Lipid Membranes by Quasielastic Neutron Scattering," ▮▮
▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮▮) submitted to Europhys. Lett.

Description of Proposed Research

(Please include scientific context; relevance of proposed experiment; preliminary work performed using neutron scattering and other techniques; details of proposed experimental approach; appropriate references.)

This work is a critical part of the graduate thesis of: ▮▮▮▮▮▮▮▮▮▮▮

Studies of Model Membranes

Scientific background:

Membrane proteins represent a significant frontier in structural biology – they are ubiquitous in nature and are notoriously difficult to study with traditional techniques [1]. Therefore, the structure and dynamics of nearby water molecules and their impact on the function of membrane-embedded proteins remains one of the most fundamental problems in biological physics today.

Fig. 1. Sketch of a hydrated single-supported bilayer lipid membrane as suggested by Ref. 4.

The experiments described in this proposal will investigate the diffusion of water in proximity to a single-supported DLPG ($C_{30}H_{58}O_{10}PNa$) membrane. In contrast to DMPC that we have studied previously, DLPG has a net negative charge that facilitates protein insertion. Our principal aim is to compare the dynamics of water on this charged membrane with that which we have found for water associated with the electrically neutral DMPC membrane.

In a second experiment (see Proposal #17423), we will investigate the effect on the diffusion of the membrane-associated water caused by a low concentration of the protein Cytochrome *c* residing on the DLPG membrane. Cytochrome *c* (Cyt *c*) is an electron shuttling protein that accepts and donates electrons from and to the bioenergetic integral membrane proteins found in the inner membrane of mitochondria and bacteria. On the surface of charged lipid membranes, it undergoes an ordered-to-fluid globule structural transition, thereby influencing its native role in soluble electron transport chains. Thus, understanding how the solvation of a membrane's surface is influenced by or influences the protein-lipid head group interactions is an important biological question.

Aims of the proposed experiment:

Recently, we have demonstrated the use of high-energy-resolution quasielastic neutron scattering to investigate the diffusive motion of water and of the lipid molecules in single DMPC
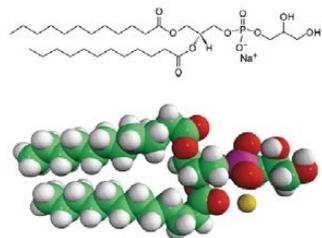
Fig. 2. **(a)** Schematic diagram of the DLPG; and (b) 3D model.

membranes supported on a solid substrate [2,3] as depicted in **Fig. 1** [4]. We propose now to extend these experiments by investigating the interaction of water with single-supported DLPG membranes. As shown in **Fig. 2**, the DLPG molecule contains two aliphatic chains of 12 carbons compared with the 18-carbon chains in DMPC. The head group differs from that in DMPC membranes by not having an amino group attached to the phosphate group. By protonation, the phosphate group can be made electrically neutral as in the choline head group of DMPC; and, by deprotonation, it can be electrically charged. Thus, we have the interesting possibility of studying the effect of the head group's electric charge on the water dynamics in the membranes.

Preliminary work:

During May and October 2010, we performed measurements with the backscattering spectrometer BASIS at the Spallation Neutron Source on a single-supported protonated DMPC membrane ($C_{36}H_{72}NO_8P$) hydrated with H_2O in order to be sensitive to the diffusive motion of both the water and the lipid molecules. The membrane was deposited on a SiO_2-coated Si(100) substrate with an amount of water sufficient to hydrate the membrane fully. A total of 100 wafers of diameter 50 mm filled the neutron beam and were oriented with the wave-vector transfer **Q** parallel to the wafer surface (see **Fig**. **1**) [5]. One-hour spectra were obtained at 0.5 K increments in the temperature range from 250 K to 315 K. In order to verify sample quality and locate possible water freezing transitions, we plotted the elastic intensity summed over all Q vs.

Proposal #17273

temperature as shown in **Fig. 3(a)**. The sharp vertical step at 267 K indicates a homogeneous sample and that the equilibration of the wafer stack.

We fit the BASIS spectra by folding the instrumental resolution function with a scattering law composed of three terms: a delta function representing the elastic scattering plus two Lorentzians [5]. Utilizing the combination of energy resolution and dynamic range of BASIS, we could resolve two diffusive processes occurring at different rates: a "fast" motion (time scale < ~40 ps) that can be fit to a broad Lorentzian and a "slow" motion (time scale ~ 0.5 ns) described by a narrower and weaker Lorentzian. The full-width-at-half-maximum (FWHM) of the broad Lorentzian has a Q^2 dependence at low Q (not shown) characteristic of translational diffusion of the H_2O molecules. We tentatively identified the slower motion corresponding to a FWHM ~ 5 μeV that is nearly Q-independent (not shown) with conformational changes of the alkyl tails of the lipid molecules. We have identified similar motion in solid bulk alkane particles [5] and in monolayer alkane films [6,7].

The temperature dependence of the broad Lorentzian intensity is shown in **Fig**. **3(b)**. We see that it has a step-like decrease on cooling that coincides with the step-like increase in elastic intensity at 267 K, which is consistent with about 2/3 of the H_2O sample becoming immobile on a nanosecond time scale. In contrast, the intensity of the narrow Lorentzian in **Fig. 3(c)**, attributed to conformational changes of the lipid tails, has only a weak temperature dependence. The solid and open data points correspond to samples hydrated with H_2O and D_2O, respectively.

Proposed measurements:

We request 8 days of beam time on the HFBS to measure elastic scans on two samples of single-supported DLPG membranes: one at a high level of hydration as in **Fig. 3** and one at a factor of 10 lower level of hydration [3]. From our previous measurements on single-supported DMPC membranes [5], we have a good estimate of the measurement time for each sample; i.e., 4 days to perform 3 to 4 heating/cooling cycles at a slow ramp rate of 0.04 K/min.

Fig. 3. Normalized intensity of the three components of the quasielastic spectra from a well-hydrated single-supported DMPC membrane: (a) elastic; (b) broad Lorentzian; and (c) narrow Lorentzian.

Of particular interest will be to determine whether there is a sharp, step-like freezing transition for the high-hydration sample on the charged DLPG membrane as is the case for DMPC as shown in **Fig. 3(a)**. Due to stronger water/head group interactions for DLPG, we might expect a further suppression of the water freezing transition below 267 K. In the case of the DMPC, the membrane of lower hydration did not show a step-like decrease in elastic intensity on cooling. Instead, beginning at 267 K, there was a continuous and monotonic decrease in the elastic intensity. Perhaps with the stronger water/head group interactions for DLPG, there will now be a step-like freezing transition as in the case of DMPC at high hydration.

These studies of the water freezing transition for pure single-supported DLPG membranes conducted on the HFBS should establish the feasibility of further measurements of the full quasielastic spectra on BASIS as we have done for DMPC. In addition, these experiments with membrane-associated water on the bare DPLG membranes are an important prerequisite for studies of the effect of the protein Cyt *c* on the water dynamics as discussed in the introduction and in Proposal #17423.

Proposal #17273

References:
1. M. Luckey, *Membrane Structural Biology* (Cambridge University Press, Cambridge, 2008).
2. C. L. Armstrong, M. D. Kaye, M. Zamponi, E. Mamontov, M. Tyagi, T. Jenkins, and M. C. Rheinstädter, Soft Matter **6**, 5864 (2010);
3. M. Bai, A. Miskowiec, S.-K. Wang, H. Taub, T. Jenkins, M. Tyagi, D. A. Neumann, and F. Y. Hansen, Bull. Am. Phys. Soc. **56**, B40.3 and B40.4 (2011); M. Bai, A. Miskowiec, F. Y. Hansen, H. Taub, T. Jenkins, M. Tyagi, D. A. Neumann, S. O. Diallo, E. Mamontov, K. W. Herwig, S.-K. Wang, submitted to Europhysics Letters.
4. E. T. Castellana and P. S. Cremer, Surf. Sci. Rep. **61**, 429 (2006)..
5. S.-K. Wang, E. Mamontov, M. Bai, F. Y. Hansen, H. Taub, J. R. D. Copley, V. García Sakai, G. Gasparovic, T. Jenkins, M. Tyagi, K. W. Herwig, D. Neumann, W. Montfrooij, and U. G. Volkmann, Europhys. Lett. **91**, 66007 (2010).
6. F. Y. Hansen, L. Criswell, D. Fuhrmann, K. W. Herwig, A. Diama, R. M. Dimeo, D. A. Neumann, U. G. Volkmann, and H. Taub, Phys. Rev. Lett. **92**, 046103 (2004).
7. D. Enevoldsen, F. Y. Hansen, A. Diama, H. Taub, R. M. Dimeo, D. A. Neumann, and J. R. D. Copley, J. Chem. Phys. **126**, 104704 (2007).

12
Data Correction and Analysis

The data analysis and interpretation of an experiment can be involved, or it can be straightforward. If we are looking for the existence of some effect, the analysis will likely be easy. If we are looking for something subtle that involves line shape analysis, then we are in for it. This chapter serves as a general guide on what to do to get the most out of the data, but it cannot cover every possible experiment.

There are some standard steps in correcting the data of an experiment:
- remove all the scattering that does not come from the sample,
- remove all the scattering from the sample that is unwanted (e.g., incoherent scattering, multiple scattering),
- correct the scattering for effects that changed the intensity of the signal (e.g., attenuation, monitor contamination).

.

At the completion of the data correction procedures we have the signal that we are after, ready for us to fit our model to and interpret what we see. There are some potential pitfalls in this step that largely fall into two categories:
- fitting the model to the wrong (scattering) function,
- restricting the model too much (or too little) with the result that we get a good fit to an unphysical model.

Both these mistakes are surprisingly common so, be warned, they might sneak up on you.

12.1   Data Correction

There are quite a few steps in the data correction procedure, depending on what type of experiment one is doing. It is good practice to do the correction steps one by one, and to look at the ensuing results before undertaking the next correction step. In particular, start by ignoring all correction steps, and look at the raw data as if they are the final product. This already gives a good indication of how accurately one has to perform the data correction steps, or whether some of them can be ignored altogether.

The first step in our data correction procedure is to get rid off the background scattering. With background scattering we mean neutrons that end up in the detector and that did not come from our sample. They could have come from the spectrometer next door, from the cryostat, from the sample holder, from air scattering, or from the
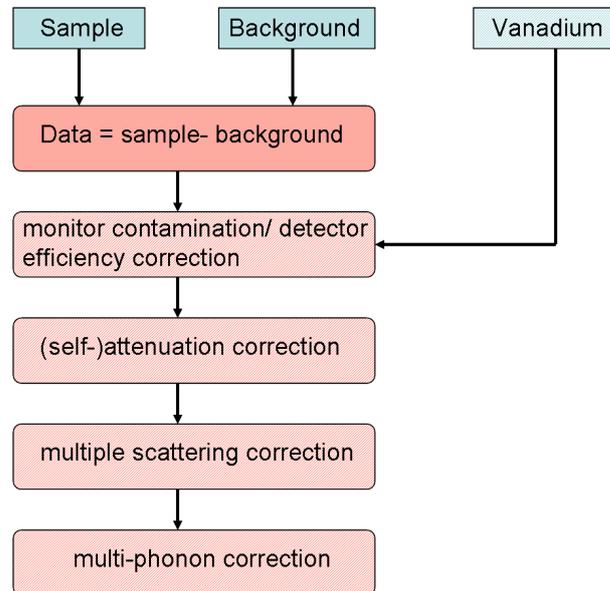
Fig. 12.1 The basic steps in correcting neutron scattering data. The blue boxes denote the experiments, the red boxes the steps in data reduction. When a box is more lightly colored, then these steps are not necessarily required, depending on the type of experiment.

straight through beam. A pretty good way of correcting for this scattering is by doing a separate run (experiment) where everything is measured in the same way, but the sample has been removed from the sample holder. It is pretty good, not perfect.

It is pretty good because we can think of the scattering- when the sample is in place- as being made up of:

$$\text{total scattering} = \text{sample scattering} + \text{sample environment scattering} \\ + \text{time-independent scattering.}$$

With time-independent scattering we mean the neutrons that come from next door and air scattering. Hence, when we remove the sample we can expect the scattering to be given by:

$$\text{total scattering} = \text{sample environment scattering} + \text{time-independent scattering.}$$

Subtract the two measurements from each other, and what remains is the scattering by the sample. However, this is not entirely correct.

It is not entirely correct because the sample changes how many neutrons will hit what parts of the cryostat, and even how many neutrons are left in the straight through beam that will give rise to air scattering. With the sample in place, the back of the

sample environment will be in the shadow of the sample. Since the sample scatterers some 10 % of the neutrons out of the beam, the back of the cryostat will generate fewer scattering events with the sample in place than without the sample in place. In addition, the sample will scatter neutrons that in turn can scatter off of the cryostat and sample holder in ways that are not possible without the sample in place. For these reasons this straightforward correction procedure is not perfect, it over subtracts.

Under what circumstances is this subtraction procedure good enough? Mostly when we have a sample that only scatters very weakly, that does not absorb many neutrons, and when the signal we are interested in is pretty strong. Conversely, the procedure does not work well when we have weak signals or strongly scattering and absorbing samples. It is up to the user to be aware of this, to not simply do this straightforward subtraction and no longer think about it. For instance, when ending up with negative intensities after the correction procedure it is a sure sign of having subtracted too much background scattering.

What to do when this procedure is not good enough? We can use sophisticated computer programs for the case of liquids, and we can do even better for the case of magnetic systems. We will elaborate upon this a bit because it is probably time for a figure. Shown in Fig. 12.2 are the raw data that led to the magnetic peaks that we displayed in Fig. 5.5.

In this experiment, the aim was to measure some very weak peaks in the presence of a lot of unwanted scattering, namely cryostat scattering, air scattering, nuclear incoherent, and nuclear coherent scattering. A sample subtraction procedure between the two experiments with sample in the cryostat and with the sample removed would not have worked in view of the weakness of the signal. In order to measure this weak signal in this case the following experiments were performed. Instead of carrying out a background experiment with the sample removed, the sample was left in but its temperature was raised from 2 K to 56 K. The latter temperature is above the magnetic ordering temperature so that the scattering will contain all nuclear scattering (coherent and incoherent), cryostat scattering and air scattering, but no magnetic scattering. Subtracting the higher temperature from the low temperature data leaves purely the magnetic scattering.

The above procedure is almost perfect in getting rid off unwanted scattering. All the cryostat and air scattering is exactly subtracted since we left the sample in place, and all the nuclear coherent and incoherent scattering is subtracted almost perfectly. "Almost" in this case, refers to the fact that the crystal will have contracted a little bit between 56 K and 2 K so that the nuclear Bragg peaks will have shifted a little bit. This effect caused the gaps in scattering at (112) and (114) shown in Fig. 5.5. In addition, the incoherent scattering might show a very small amount of temperature dependence, but that amount was negligible here. Finally, we can expect some multiphonon scattering at 56 K caused by phonons present in the crystal scattering neutrons, phonons that are no longer present at 2 K because of lack of thermal energy.
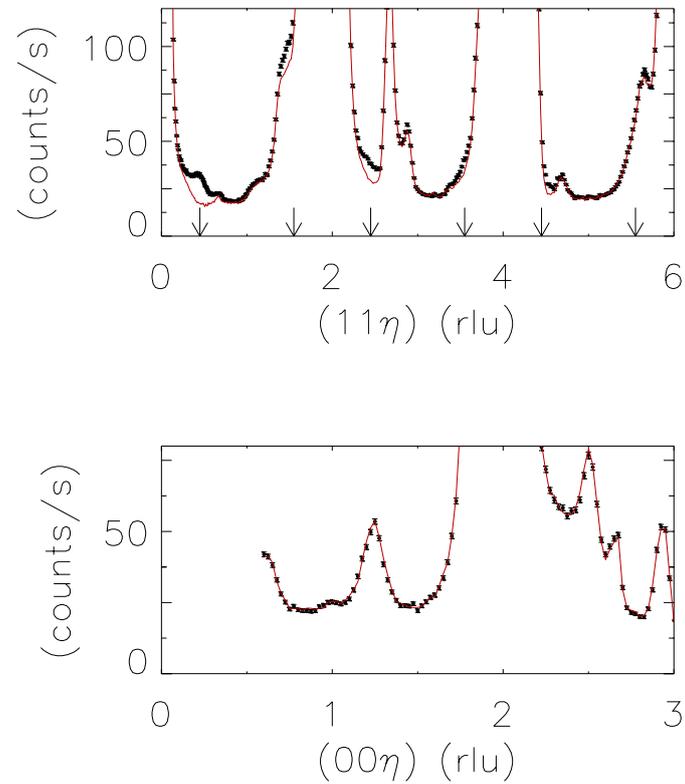
Fig. 12.2 The notation is the same as in Fig. 5.5. The top panel shows the change in scattering between 56 K (solid line) and 2 K (points). This change is entirely due to the onset of magnetic correlations. Subtracting the two data sets yields the peaks shown in Fig. 5.5. The vertical arrows are the positions of maximum difference between the two temperatures. The bottom panel displays the same two temperatures, but now for momentum transfers that are purely along the crystallographic c-axis. We observe no difference, and hence, we must conclude that the magnetic moments are pointing along the c-axis.

But over all, the subtraction is essentially perfect, as demonstrated in Fig. 5.5.

The procedure outlined above can be applied whenever we want to measure a signal that is absent at a different temperature, or different magnetic field, etc. Or, we can apply it to biological samples using contrast matching techniques to set up a background experiment. Or, we can rotate a single crystal sample away from a Bragg reflection, leaving the same amount of incoherent scattering. The bottom line is that if it is possible at all to measure the background (unwanted) scattering with the sample in place, then this will result in much more accurate data.

When it is not possible to set up a background measurement with the sample in place, and we are after highly accurate scattering data, then only models will help. An example of this can be an experiment on a supercritical liquid, such as argon at high pressure. What needs to be done in this case in order to obtain an accurate background subtraction is to measure the sample in its pressure cell, to measure the empty pressure cell, followed by using a computer program that has as input the way argon scatters. The computer program can then figure out, to some extent, how much the presence of the sample attenuated the scattering of the empty pressure cell. People have even improved upon this by filling the cell (during the background measurement) with an amount of helium-3 gas that absorbs just as many neutrons as the argon would scatter out of the beam.

In summary, the correction for unwanted scattering requires thinking before performing the experiment. One needs to have a pretty decent idea of how accurate the background correction needs to be done in order to be able to see and analyze the sought-after signal. At this point, look at the results of the subtraction procedure. If the signal that one is interested in is not recognizable, then skip all following correction steps and redesign and redo the experiment. All following steps ensure that the data are corrected for, but they will not create data where there are none.

The next step- now that we have the scattering intensity that comes from our sample and not from the cryostat, air scattering, or (in some cases) from scattering events within the sample that we are not interested in- is to make sure that our counts are what they should be. In this step we have to correct for monitor contamination (if present), we have to make sure that all of our detectors (if we have multiple) are counting just as efficiently, and we have to make sure that the detectors count the slow and the faster neutrons just as efficiently.

We can expect to have to perform a monitor correction if we are doing an inelastic 3-axis experiment with fixed final neutron energy. This was detailed in section 9.2.4. If a monitor correction is not available (that means, if it has not already been measured), then measure it by doing a set of tin-foil experiments, by measuring the scattering of a liquid for both energy gain and energy loss so that the monitor correction can be inferred form the detailed balance condition, or use a model based upon the incident neutron flux. Unless the correction is already known, none of these options sounds very attractive and the first course of action can be to ignore this correction. However, if we are interested in the line shape (as a function of energy transfer) of a particular peak, then at one point we will have to do this correction. In figure 12.3, we show the difference in line shape before and after the correction on TRIAX.

When using a 3-axis machine operating with fixed incident energy (so we do not have to worry about the monitor correction), then this would be a good time to divide the data with the $k_{\text{final}}/k_{\text{initial}}$ factor to get rid off this factor by hand (eqn 4.1), since our incident beam monitor does not take care of it when we use a fixed incident neutron energy.
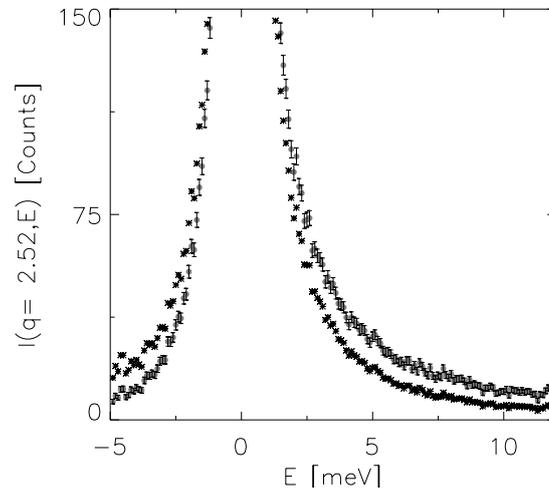
Fig. 12.3 The scattered intensity (circles plus error bars) for supercooled Ga at 293 K for $q=$ 2.52 Å^{-1} as measured at a thermal source triple-axis spectrometer (TRIAX). This particular q-value corresponds to the peak in the static structure factor. Applying the measured monitor correction leads to a substantial change in line shape (stars), illustrating the need to account for this correction before modeling the line shapes.

Should our instrument have multiple detectors, then it is unlikely that all detectors have an identical detection efficiency. We can actually measure this relative detector efficiency by sticking a purely incoherent scatterer in the beam, such as a rod of vanadium. All detectors should have an identical count rate since the scattering power of vanadium does not depend on scattering angle. We simply measure the count rate for each detector, and then we divide our measured signal (sample minus background) by this count rate. This is an easy correction, as long as we made sure that we actually performed this very quick vanadium measurement during our experiments.

Detectors are more efficient at counting slow neutrons than they are at fast neutrons. Therefore, a faster neutron is more likely to make it through the detector and not be counted than a slow neutron. If this is a sizeable effect, then we have to correct for it. We can do this correction either by measuring it, or by modeling it. However, in most cases there is no need to do this correction.

We will only need to perform this correction when we are doing inelastic experiments and the final energy of the neutron is not fixed. On top of that, we only need to do this correction when we are looking at very high energy transfers, something which we do not frequently encounter at thermal reactor sources. The correction itself can easily be measured by performing a vanadium experiment where we look at neutrons that are scattered elastically while we vary the energy of the neutrons we use for this elastic experiment.
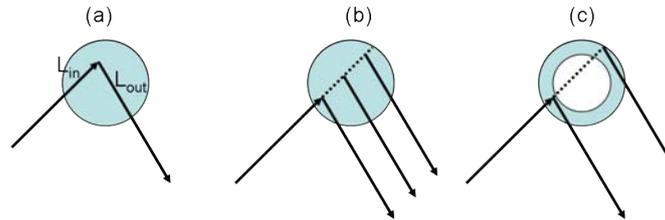
Fig. 12.4 Panel (a) shows a scattering event where the neutron travels a distance L_{in} before scattering, and a distance L_{out} after scattering. The scattered signal will be attenuated because the sample (shown in blue) will absorb the neutron, or will scatter the neutron before it reaches the scattering point, or before it exists the sample. Panel (b) shows multiple possibilities for the source of the scattering event. When correcting, we have to sum over all possibilities along the dotted line, and then translate the dotted line to the left and right of the incoming neutron until the entire sample has been modeled. Panel (c) shows how we can do similar ray tracing for samples of all shapes, such as the annulus depicted here.

Modeling this correction is not too terribly hard to do. We can simply do a calculation based on the amount of helium-3 gas in the detector, the shape of the detector and the energy of the neutron to be detected. Note that this correction might be more important when using smaller detectors, such as the 1/2" pencil detectors used on TRIAX.

The next step of our correction procedure involves correcting for (self-)attenuation. This is the effect that the front of the sample receives more neutrons than the back of the sample by virtue of absorbing some neutrons and scattering some away. In addition, neutrons that try to make it out of the sample on their way to the detector can be absorbed by the sample, or scattered once more, before they get to the detector. The importance of this correction depends on how good a scatterer the sample is, on its shape, and on how many neutrons the sample absorbs.

The good news about this correction is that it can be calculated exactly (Appendix I), so there is no reason not to perform this correction. The correction will be of least importance for cylindrical, weakly absorbing samples, and it will be most important for slab-geometry samples that are strongly absorbing. In order to calculate this correction, we only need to know the average scattering and absorbtion power of the sample, and its dimensions.

The details of how to calculate this correction were worked out by Varley Sears, but the essence is that we need a computer program to do ray tracing. We show the basics in Fig. 12.4. We assume a scattered neutron has been detected having been scattered over a certain scattering angle. We want to figure out how many neutrons were *not* detected at this scattering angle because something happened to them on the way to the scattering event, or after it. In the scattering event depicted in Fig. 12.4, the

neutron traveled a distance L_{in} to the scattering event, and a distance L_{out} after the scattering event before it made it out of the sample. Denoting the average scattering and absorbtion cross-section by σ (details to follow), and the average number density of scattering centers by n, we find (see also Chapter 3) that the number of neutrons that were scattered along this path are down by a factor W:

$$W = e^{-L_{\text{in}} n \sigma} e^{-L_{\text{out}} n \sigma}.$$

Thus, had the sample not had any scattering or absorbing power, we would have counted more neutrons being scattered along this path. To be precise, we would have measured $1/W$ more. Other points of scattering in the sample will change the respective values for L_{in} and L_{out} and, therefore, the value of W. Then, in order to perform this scattering, we simply calculate this factor W for all possible points of scattering anywhere in the sample, and we divide our measured signal by this factor, properly averaged over all points of scattering. Appendix I shows some examples that were calculated by doing numerical integrations using a computer.

When we are using a computer, we can also take into account the fact that the likelihood of absorbtion depends on the speed of the neutron. This is a straightforward modification. We simply make the value σ dependent on the energy of the neutron: $\sigma(E) = \sigma_{\text{coh}} + \sigma_{\text{inc}} + \sigma_{\text{abs}}(E)$. Thus, we use the following energy dependent correction factor

$$W(E_{\text{in}}, E_{\text{out}}) = e^{-L_{\text{in}} n \sigma(E_{\text{in}})} e^{-L_{\text{out}} n \sigma(E_{\text{out}})}.$$

Once we have these attenuation programs, we can actually use them for more purposes. We can use them to correct (after the facts) our empty sample holder experiment to pretend there is a sample in it. Rather than dividing by the factor W, we would multiply our empty sample holder experiment by it. We can also use such programs to calculate the transmission of a sample, and compare it to our measured transmission. The latter we do when we try to put the scattering of our sample onto an absolute scale, a topic not detailed in this booklet.

The last two corrections we need to consider are multiple scattering and multi-phonon corrections. Despite their name, the two are very different entities. We try to sketch this in Fig. 12.5. Multiple scattering events are events where the neutron scatters more than once within the sample, whereas multi-phonon events are events where the neutron scatters only once, but it creates multiple excitations when it does scatter. Or, it absorbs an already existing excitations and creates a new one. Hence, multi-phonon events are a true part of the double differential cross-section and of the dynamic structure factor, whereas multiple scattering events are not part of the dynamic structure factor; they merely show up in the double differential cross-section as unwanted scattering events.

Multiple scattering events are nasty. In liquids and amorphous materials, they show up as a broad (in energy) background that does not display a lot of angular dependence. In single crystals multiple scattering events can show up as sharp peaks mimicking magnetic Bragg peaks, or they can manifest themselves as a reduction in
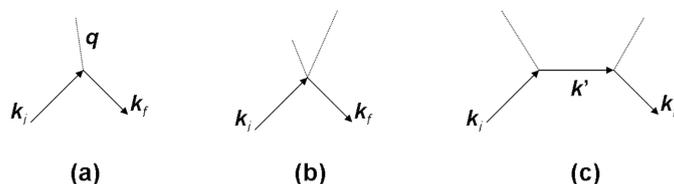
Fig. 12.5 When a particle with incident wave vector $\vec{k}_i$ (arrow) gets scattered by a sample it can create an excitation of momentum $\vec{q}$ (dotted line, part a), or it can excite multiple disturbances (part b). The latter is referred to as the multi-phonon component. A particle can also undergo multiple scattering events before exiting the sample (part c), a process referred to as multiple scattering.

scattered intensity of a Bragg peak. In liquids we can do a very good job of correcting for multiple scattering, provided the level of multiple scattering is not too high. In solids we cannot correct, but as a rule of thumb, we do not have to anyways.

The way to correct for multiple scattering in liquids is by number crunching. The procedure is similar to how we correct for attenuation effects, but now at every point of scattering we allow the neutron to scatter in any direction, transferring any amount of energy. We weigh each of these many possibilities with the dynamic structure factor, which gives us the likelihood that a neutron would actually transfer that much momentum and energy. Then we follow the neutron on its way out of the sample, but we force the neutron to scatter a second time. This second scattering event is such that the neutron will make it to the detector, and it will transfer just the right amount of energy to the sample so that when we detect the neutron, it looks just like a neutron who scattered only once. Then we tell the computer to sum over all possibilities. This is sketched in Fig. 12.6.

This multiple scattering procedure for liquids is actually very easy to do on computers, the only drawback is that we already need to know what the dynamic structure factor looks like, the very entity that we are trying to determine. Luckily, since these types of calculations do not require exact matches for the dynamic structure factor, we can get away with approximate models. By and large, multiple scattering corrections in liquids can be performed; they are just a bit time consuming in the sense that the computer needs very specific instructions (sample shape and dynamic structure factor) for how to calculate them.

Multiple scattering events in solids cannot be corrected for. This is because they will involve Bragg scattering, and Bragg scattering is simply too strong. The only thing that can be done is to avoid having to deal with it in the first place. Should multiple scattering occur and it stands in the way of obtaining the data, then the best bet is to repeat the experiment, but now with a different incident wave length of the neutron. The biggest danger in multiple scattering in solids is that these events do have a strong angular dependence, and they can even look like a genuine magnetic
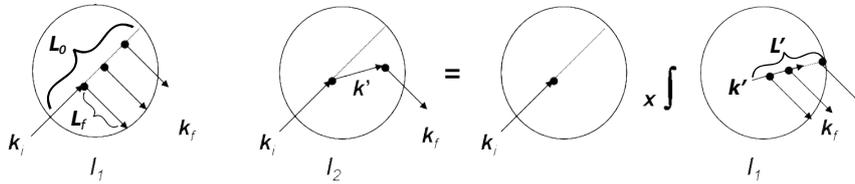
Fig. 12.6 Some of the potential paths that a particle can take when it is scattered by the sample. At every scattering event we have conservation of momentum and energy. The left panel depicts a neutron with incoming momentum $\hbar\vec{k}_i$ that is scattered exactly once by the sample, somewhere along its path L_0. The length of the path L_f that the scattered particle has to travel before exiting the sample with momentum $\hbar\vec{k}_f$ depends on where it scattered initially. The other panels on the right sketch how we can decompose events where the particle is scattered twice. In these double scattering events, the neutron exits the sample with the same momentum vector as in single scattering. The computer now has to sum (integrate) over all possible intermediate scattering events, all weighed by the dynamic structure factor.

Bragg peak. More on this is in the section on spurions in this chapter.

The last correction we look at in this section is the multi-phonon correction. This is an effect that we only need to worry about under special circumstances, namely when we are trying to measure close to a Bragg peak, or when we are interested in the density of states. Multi-phonon corrections are only relevant to crystals. Multi-phonon scattering takes place when the neutron creates two excitations simultaneously, or when it absorbs an existing excitation while creating another. The cross-section for the latter is highly temperature dependent, since the number of existing excitations present in a material is given by the temperature of the material.

When we are measuring the density of states, then we can correct for multi-phonon effects according to the equations given by Lovesey. We will not give any details here other than the mention that standard software can deal with it rather easily. When we are dealing with multi-phonon effects around Bragg peaks, then the take-home message is less helpful: it is very difficult to correct for it. This is not as bad as it sounds though; as long as we are aware of its existence, we can work around it. The most difficult part is identifying its presence, something we do by varying the temperature of our sample. Once we have isolated, or identified, the multi-phonon component, then we can (somewhat) model its temperature dependence. Sometimes we can work around it when we are interested in magnetic scattering. We can take an iso-structural, non-magnetic sample and use that as part of our background experiments. But again, this is a procedure far too detailed for this booklet where we merely preach awareness of this component.

12.2 Spurions

Before we look into how to analyze our (fully corrected) data, let's make sure that what we call our data actually are what we want, and not some spurious effect masquerading

as a new discovery. We already showed the source of a spurious peak in Fig. 9.8. We will look at some more of these spurious effects in no particular order. We recommend the chapter on spurions in the book by Shirane, Shapiro and Tranquada for a more detailed discussion.

12.2.1   A New Excitation Below 4 K

When doing temperature dependent experiments at low temperature, it is possible to see an excitation appear below 4 K. The excitation looks like it is part of the sample scattering because it disappears as soon as we raise the temperature of the sample. It even displays a dependence on scattering angle, and it becomes better defined the more we lower the temperature. The strange thing is that its angular dependence does not seem to match a natural length scale in our sample.

More likely than not is that the cryostat is malfunctioning. Especially orange cryostats are susceptible to this. To cool down the sample, cold helium gas is pumped out through a needle valve. Sometimes the rate at which helium goes through this valve is so high that liquid helium will condense at the bottom of the cryostat, and slowly fill the space around the sample. This helium liquid will be around 4 K (the boiling point of helium is 4.2 K), and it can scatter neutrons quite well. It disappears as soon as we raise the temperature above 4.2 K and we let the condensed helium boil off. This helium scattering will be strongest for momentum transfers around $q = 1.9 - 2.0 Å^{-1}$. To get rid off this effect, do not open the needle valve as much.

The user should also be on the lookout for air scattering, or scattering by any gas that is used to increase the thermal conductivity of one's sample. In particular, there can be a sharp reduction in background when the sample is cooled below 65 K and N_2 freezes out. A nice example is shown by Shirane, Shapiro and Tranquada in Fig. 6.10 of their book 'Neutron Scattering with a Triple-Axis Spectrometer'.

12.2.2   A New Excitation At 41 meV

Picture the following experimental setup on a 3-axis spectrometer. The spectrometer is run with variable incident energy, and the final energy is set to E_{final}= 13.7 meV. With this choice of final energy, we can use a graphite filter to get rid off unwanted higher order contamination (see Fig. 9.15), and we have put a 1" PG filter in the scattered beam. We are looking for weak excitations in a high-temperature superconductor where a (pseudo)-gap opens up around 41 meV. After counting for long times, we do indeed find this excitation at 41 meV but, strangely enough, we even see its presence above the superconducting transition temperature in the form of increased scattered intensity at 41 meV.

The problem with the above setup is that the graphite filters used were not sufficient. Even when the filter cuts out 99% of unwanted higher order neutrons, our inelastic peak is likely to be contaminated with elastic scattering. This is what is going on. In order to get to an energy transfer of 41 meV, we use an incoming neutron energy of 54.7 meV. Some of these neutrons will scatter elastically at the sample

through the incoherent cross section of the sample, or perhaps in the wall of the sample holder. The analyzer happily accepts those neutrons, since 54.7 meV corresponds to 4 x 13.7 meV. This is, of course, the energy corresponding $\lambda/2$ neutrons, the second order contamination that the analyzer allows. So we are relying entirely on the filter to remove this bunch of elastically scattered neutrons. Even 1% of these elastically scattered neutrons can still represent a strong signal, especially when compared to some weak high energy excitation. So we need to add more filters to reduce this 1% seep-through rate.

In general, we can expect this type of spurions to occur whenever the second (or third) order contamination allowed by the analyzer corresponds to the energy of the incoming neutrons, provided we are measuring at fixed final energy. When we measure using fixed incident energy, similar reasoning applies. Obviously, this type of spurion is restricted to inelastic scattering experiments.

Another type of 3-axis spurion has been discussed by Currat and Axe, and it is referred to as accidental Bragg scattering. We refer to the tome on 3-axis spectrometry by Shirane, Shapiro and Tranquada for all the details. Accidental Bragg scattering happens when, during the course of an inelastic scattering experiment, the sample has been rotated in such a way that elastic Bragg reflection is allowed to occur, and an intense beam of neutrons is sent down to the analyzer. The graphite filter (or any other filter that is in place) might stop most of them, but this beam is so intense that quite a few make it to the analyzer. The analyzer will not Bragg reflect these neutrons, since their wave length is not the right wavelength to be Bragg reflected into the detector. However, the analyzer is at room temperature, and has quite a few phonons present in it. Additionally, the analyzer can have an incoherent cross-section. The result is that a small fraction of these neutrons impingent on the analyzer will be scattered in all directions, including the detector. It will give rise to only a weak signal in the detector, but this signal can still be substantial compared to the excitation we are studying.

This accidental Bragg scattering peak is a nasty spurion, and the spurions are recognizable by a dispersion that does not follow the symmetry of the sample crystal. They are not easy to recognize as spurions, but some spectrometers have software running on them that may help the user identify such spurions, or at least warn the user to be on the lookout for them.

12.2.3 The Magnetic Moments Are Canted Away From a High Symmetry Direction

Most magnetic systems- that is, systems where the magnetic moments line up parallel or anti-parallel with their neighbors below a certain temperature- harbor magnetic moments that line up along a particular axis (the easy axis), or they lie within a plane (the easy plane). Neutron scattering is very well suited at figuring out the direction of the magnetic moments, because the magnetic cross-section is present only when the magnetic moment direction and the direction of momentum transfer are not parallel

to each other.

An example of this sensitivity of the cross-section to the direction of momentum transfer is shown in the bottom panel of Fig. 12.2. In this panel, we observe a total lack of magnetic scattering because the direction of momentum transfer is exactly parallel to the orientation of the magnetic moments, in contrast to the top panel of this figure where the momentum transfer vector has a component perpendicular to the orientation of the magnetic moments.

Suppose the magnetic moments would have been tilted away slightly from the easy direction. In this case, we would have gotten a small amount of magnetic scattering in the bottom panel of Fig. 12.2. If we are in a situation where we have sharp Bragg peaks (unlike the example in Fig. 12.2 where we only had short-range order), then slight tilting (or canting) would give rise to very weak Bragg peaks along directions that would otherwise be forbidden. When comparing the intensity of such weak and forbidden Bragg peaks to other Bragg peaks, we can figure out the canting angle. Typically, we find small numbers, such as around 1-4^o.

The problem with weak Bragg peaks is that they can also originate from multiple scattering events. The resolution function of the spectrometer can be so accommodating that we can pick up the tail of a nuclear reflection and combine it with the tail of a non-forbidden magnetic reflection in such a way that it looks like the neutrons are scattered as if there exists a weak magnetic peak. We show an example of such a multiple scattering event in Fig. 12.7.

These spurious peaks can fool, and have fooled, the best of us since they appear and disappear exactly at the magnetic ordering transition. Whether a particular spurious magnetic peak will show up, or not, depends strongly on the wave length of the neutron. When in doubt, which should be always, simply repeat the experiment with a different neutron energy, or look for other forbidden Bragg peaks using the same settings. Only report on canting of the magnetic moments if the peaks survive these tests.

12.3   Pitfalls In Data Analysis

When our signal is easy to identify, such as a nice sharp excitation in an inelastic scattering experiment, or perhaps the appearance of a new Bragg reflection below some ordering temperature, then there is no need to read this section. It is just when our signals get weak, or not so sharp, that we have to be extra vigilant. This even holds when we already have been extra careful in correcting our data, and when we have avoided all spurions.

Excitations that only persist for a short amount of time will show up as broad features in inelastic scattering experiments. When these broad excitations have low excitation energies, or when we are looking at quasi-elastic features, then we have to be very careful in our line shape analysis. The data shown in Fig. 9.4 are an excellent
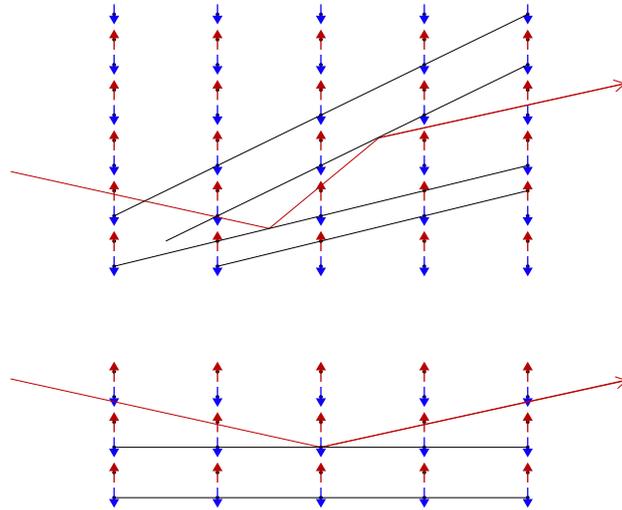
Fig. 12.7 The bottom panel shows an antiferromagnetically ordered system, with the moments in the vertical direction. Shown is a forbidden reflection as the direction of transferred momentum is parallel to the orientation of the magnetic moments. Therefore, such a reflection can only be observed if the moments are canted (slightly) away from the vertical. The top panel shows how double Bragg reflection can mimic the presence of a forbidden reflection. The neutron first reflects off of a set of nuclear Bragg planes, followed by a reflection off of an allowed magnetic reflection. The result is a neutron exiting with an angle very close to that of a forbidden reflection, giving the appearance of canted moments. It is only through bad luck that such double reflection might occur, but they do occur. The situation depicted corresponds to a lattice spacing of 11.7 Å in the horizontal direction, to 2.72 Å in the vertical direction, and an incident neutron wave length of $\lambda = 2.35$ Å. The numbers were chosen to represent an actual experiment of $YBa_2Cu_3O_6$, studying the forbidden reflection $(1/2, 1/2, 0)$. See exercise 12.2 for a calculation of the angles involved.

example of how we can get it wrong.

The data in this figure (Fig. 9.4) show the transition from a superfluid to a normal fluid. The transition temperature is at $T = T_\lambda = 1.9203$ K. What we observe in this figure- where the dynamic structure factor is plotted- is the following. We see how a very sharp excitation at the lowest temperature broadens in energy. This broadening appears to be very similar just below, and just above the transition temperature. The position of the peak (that is, the excitation energy) does not seem to shift very much at all over the entire temperature range, from deep within the superfluid phase to well into the normal fluid phase. It certainly does not look like we are following one of the most spectacular transitions in nature.

The reason why nothing much seems to happen upon going through the superfluid

to normal fluid transition is because we are inspecting the dynamic structure factor. This function is a measure of the spontaneous fluctuations present in the fluid, and it is what the double differential cross-section is proportional to. However, the shape of this function is strongly influenced by the temperature in the sense that when we raise the temperature, more excitations will be present in the liquid. This alters the shape of the function, even when the amount of energy that is required to create an excitation does not change. This is the detailed balance effect, and it tends to mask subtle changes in excitation energies.

In order to study how much energy it takes to create an excitation, we should scrutinize the dynamic susceptibility. The dynamic susceptibility is a complex function whose poles tell us the excitation energies of particular excitations, as well as their lifetimes. The dynamic structure factor is closely related to the imaginary part of the dynamic susceptibility. The exact relationship is given in table 4.1, but in essence the relationship between the two is given by the thermal population factor, a measure of how many excitations are present in the system at a given temperature.

Plotting the (imaginary part of) the dynamic susceptibility provides a better visual representation of what the excitation energies are. We can even do one step better, and plot a symmetrized dynamic structure factor $S^{\mathrm{sym}}(q, E)$ that accounts for the thermal population factor as well as for the asymmetry implicit in the susceptibility. This new function can be plotted straight from the neutron scattering data as

$$S^{\mathrm{sym}}(q, E) = \frac{1 - e^{-\beta E}}{\beta E} S(q, E), \qquad (12.1)$$

with $\beta = 1/k_B T$ the reciprocal thermal energy.

The visual peaks in this function correspond very closely to the excitation energies of the system, since they occur at the pole positions of the dynamic susceptibility. This can be shown mathematically, but is is much easier to show using our actual helium example. In Fig. 12.8, we simply replot the data of Fig. 9.4 by putting in the sample independent prefactor of eqn 12.1. The change in appearance of the data is easy to see. Not only do we see this change in appearance, we can now also directly follow the fate of the excitations when going through the superfluid to normal fluid transition. We still witness a broadening of the excitation, but now this broadening is combined with a change in excitation energy. In particular, we now see how the excitation energy changes from a non-zero value in the superfluid phase to a zero value in the normal fluid phase. Notice that there is nothing in the replotting that would have forced this change from non-zero (propagating) to zero (overdamped) to happen exactly at the superfluid transition temperature.

The take home message of this replotting is that it is something that can be done in the absence of any data analysis, and that it is a procedure that reveals, to the naked eye, what the changes are in fundamental excitation energies of the system. Therefore, when doing an experiment on excitations that are not terribly sharp (that
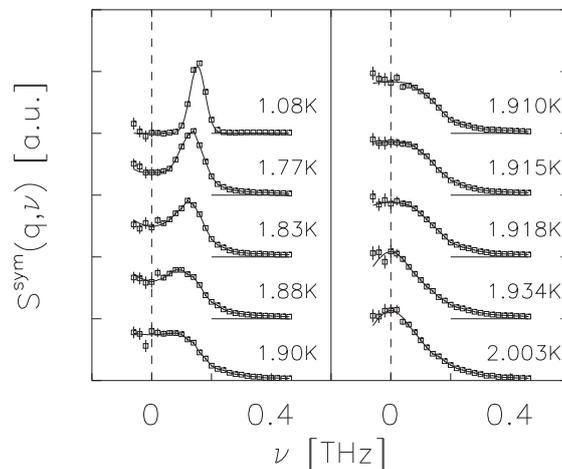
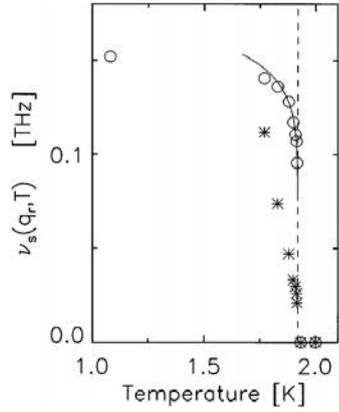**Fig. 12.8** Figure reproduced with permission from W. Montfrooij *et al.*, Phys. Rev. Lett. 77, 4398 (1996). Shown is the symmetrized dynamic structure factor $S^{sym}(q, E = h\nu)$ for liquid ^{4}He for $q = 2.0$ Å^{-1} at constant density $\rho = 0.1715$ g/cm^3 as a function of temperature (given in the figure). The data in this figure are obtained through a simple replotting of the data in Fig. 9.4 (see text). The solid lines are guides to the eye. The data have not been corrected for the resolution of the spectrometer. It is now possible to follow by eye the complete softening of the elementary excitation at $T = T_\lambda = 1.9203$ K.

is, not terribly long lived), always do a replotting of the data while performing the experiment to ensure that all the relevant temperatures are measured. Since this is an important point, we will hammer it home in the next paragraph.

We can fit the data shown in Fig. 12.8 to a model such as the harmonic oscillator function. We show the results of such a fitting procedure in Fig. 12.9. What this graph demonstrates are the very rapid changes extremely close to the phase transition. These changes are visible in Fig. 12.8, but they are totally invisible in Fig. 9.4. The consequence of this is that if we simply rely on Fig. 9.4, then we are likely to not take any measurements at any temperatures between 1.90 K and the phase transition, thereby missing out on most of the changes.

The wrong conclusions that would be drawn from (only) visually inspecting Fig. 9.4 are not restricted to liquid helium. We can expect similar misinterpretations in any broad excitation close to zero energy transfer. This can be magnetic scattering close to an ordering transition, or it can pertain to the scattering by a biological membrane. When excitations are very sharp in energy, then there is little likelihood of misinterpretations since the prefactor in eqn 12.1 would only need to be evaluated at one particular energy transfer, namely the excitation energy.

The above description dealt with visual inspection of the data and how easy it is

Fig. 12.9 Figure reproduced with permission from W. Montfrooij *et al.*, Phys. Rev. Lett. 77, 4398 (1996). The open circles are the excitation energies of the roton (sharp peak in Fig. 12.8), as determined through a fit to a harmonic oscillator model. The excitation energies can be seen to soften upon raising the temperature until they soften completely exactly at the superfluid to normal fluid transition at $T = T_\lambda = 1.9203$ K. We leave it up to the reader to guess what the solid line and the stars stand for.

to conclude that nothing much is going on when in fact some very rapid changes are occurring. It did not deal with the details of modeling the data in order to obtain excitation energies and life times of these excitations. During the modeling stages we use a model that has excitation energies and life times as free parameters, and we fit such a model to the data taking proper care to broaden our model according to the experimental resolution function of the spectrometer. While the folding of the model with the spectrometer's resolution function might appear to be the place where most pitfalls can be found, it is actually in the model itself where most mistakes are being made.

We take the harmonic oscillator function as an example. A harmonic oscillator is described by an oscillation frequency ω_p, and by a damping Γ, provided the oscillator is not critically damped, or overdamped. It is straightforward to arrive at the expression that describes the neutron scattering data by doing a Fourier transform from the time domain to energy transfer (with $E = \hbar\omega$):

$$
\begin{aligned}
S^{sym}(q,\omega) &= \frac{1 - e^{-\beta\hbar\omega}}{\beta\hbar\omega} S(q,\omega) \\
&= \frac{S_{sym}(q)}{\pi} \left[\frac{\Gamma}{\Gamma^2 + (\omega - \omega_p)^2} + \frac{\Gamma}{\Gamma^2 + (\omega + \omega_p)^2} \right].
\end{aligned}
\tag{12.2}
$$

This equation has a nice, straightforward interpretation. The neutron can create an excitation with energy $\hbar\omega_p$ (the first term), or it can absorb an excitation of the same

energy (second term). The prefactor $S_{sym}(q)$ in front of the excitations is the equivalent of the static structure factor. Note that it is not actually the static structure factor, rather it is directly related to the static susceptibility. But these are details, let us instead focus on where one can easily go wrong in determining the excitation frequencies from a beautifully executed and analyzed neutron scattering experiment.

When we use eqn 12.2, we assume *a priori* that the excitations are not critically damped, or overdamped. Bluntly put, it is like sticking a pendulum clock under water and assuming it will run. One might think that not running is a possibility that can be accommodated by using eqn 12.2, since a fitting procedure to the data would presumably reveal a value of $\omega_p = 0$. This is not the case since an harmonic oscillator that is overdamped has two damping rates, Γ_1 and Γ_2, whereas eqn 12.2 only has one damping rate when $\omega_p = 0$.

The way to fit the data is to not assume that the excitations are going to be propagating ($\omega_p \neq 0$), but instead fit the data to a more general expression that encompasses eqn 12.2. We do this by using the oscillation frequency in the absence of damping, Ω_0, and the damping parameter Z that embodies the mechanism by which oscillations slow down. However, Z is not identical to the damping rate Γ used in eqn 12.2. The following expression is the correct expression for when using a damped harmonic oscillator form to model the excitations in one's sample, without presuming that $\omega_p \neq 0$:

$$
\begin{aligned}
S(q,\omega) &= \frac{\beta\hbar\omega}{1 - e^{-\beta\hbar\omega}} S^{sym}(q,\omega) \\
&= \frac{\beta\hbar\omega}{1 - e^{-\beta\hbar\omega}} \frac{S_{sym}(q)}{\pi} \frac{\Omega_0^2 Z}{(\Omega_0^2 - \omega^2)^2 + (\omega Z)^2}.
\end{aligned}
\tag{12.3}
$$

The reader can verify that eqn 12.3 and eqn 12.2 are equivalent for $\omega_p \neq 0$ by making the substitution $\omega_p = \sqrt{\Omega_0^2 - Z^2/4}$ and $\Gamma = Z/2$. When the excitations are overdamped however (that is, when $Z/2 > \Omega_0$), then one finds $\Gamma_1 = Z/2 + \sqrt{Z^2/4 - \Omega_0^2}$ and $\Gamma_2 = Z/2 - \sqrt{Z^2/4 - \Omega_0^2}$. In short, eqn 12.3 can always be used to model a harmonic oscillator, whereas eqn 12.2 or its equivalent expression below can only be used for the case when the excitations are not overdamped:

$$
S^{sym}(q,\omega) = \frac{S_{sym}(q)}{\pi} \frac{2\Omega_0^2 \Gamma}{(\omega_p^2 + \Gamma^2 - \omega^2)^2 + (2\omega\Gamma)^2}.
\tag{12.4}
$$

For the sake of completeness (after all, this is a physics book), the equivalent expression for the case of overdamping would read:

$$
S^{sym}(q,\omega) = \frac{S_{sym}(q)}{\pi} \frac{\Gamma_1 \Gamma_2 (\Gamma_1 + \Gamma_2)}{(\Gamma_1 \Gamma_2 - \omega^2)^2 + (\omega(\Gamma_1 + \Gamma_2))^2}.
$$

There are other mistakes one can make during the data analysis stages, such as using a model that has too many free parameters in it. For instance, when neutrons

are scattered by a liquid, there are certain sum rules to the scattering that one has to build into one's model. We will not detail this here, we refer the reader to the textbook on liquids ('Excitations in simple liquids, liquid metals and superfluids') in case one needs to be made aware of such sum rules and how to build them into one's model.

The bottom line of this chapter is to be aware of how subtle changes in excitation energies can be masked because our signal in experiments is proportional to the dynamic structure factor, not to the dynamic susceptibility. And sometimes we can mess up by using a model that is no longer applicable to our data. The latter is especially sad since we might end up with a very good fit between model and data, just not the best fit, or even a physical fit.

12.4 Exercises

Exercise 12.1

Derive an analytical expression for the attenuation factor for a slab of finite thickness τ and infinitely large in all other directions, oriented at an angle ϕ_0 with respect to the incoming neutrons beam (with $\phi_0 = 0$ the situation where the normal of the slab coincides with the incident neutron direction). Assume the sample does not absorb neutrons, and that the scattering power of the sample is given by $n\sigma = \Sigma$.

Exercise 12.2

Calculate all relevant angles involved in the double magnetic scattering events shown in Fig. 12.7, using the numbers listed in the figure caption. Verify that as long as the resolution function is of the order of 0.5-1^o, that the tails of the resolution function allow for the depicted scattering events to occur.

Exercise 12.3

Suppose eqn 12.2 describes the observed scattering in a liquid system. With 'describes' we mean that we observe very good agreement between the model and the data. Show that in this case we have the relationship:

$$\Gamma^2 + \omega_p^2 = \frac{k_B T q^2}{m S_{\text{sym}}(q)} = \frac{q^2}{2\pi m \chi(q)}.$$

Hint, you will have to use the sum rules valid for $S(q, E)$.

Part IV

Appendices

Appendix A
Conversions

Depending on whether one is talking to chemists, physicists or biologists, and depending where they come from, different units are being used to quantify the amount of energy transferred from the neutron to the sample. For instance, one can encounter various units such as meV, THz and ps^{-1}, as well as some less common ones such as cm^{-1}. The table below lists the conversions between the various common units, and the not so common ones. For example, if one wants to evaluate the exponential $e^{-E/k_B T}$ with E expressed in meV and T in Kelvin, then this exponential would become $e^{-E*11.605/T}$.

Table A.1 The conversions between the energy units employed in scattering experiments. The conversions are carried out using : $E = h\nu = \hbar\omega = \hbar c k = k_B T = \mu_B H$.

		E [meV]	ν [THz]	ω [ps^{-1}]	k [cm^{-1}]	T [K]	H [Tesla]
1 meV	=	1	0.24180	1.5193	8.0655	11.605	17.326
1 THz	=	4.1357	1	6.2832	33.356	47.994	71.655
1 ps^{-1}	=	0.65821	0.15912	1	5.3088	7.6384	11.404
1 cm^{-1}	=	0.12398	0.029979	0.18837	1	1.43883	2.1481
1 K	=	0.086170	0.020836	0.13092	0.69500	1	1.4930
1 Tesla	=	0.057717	0.013956	0.087689	0.46551	0.66980	1

The energy of a neutron is measured in meV, but it can also be characterized by its speed v, its wave length λ or its wave number k. If we use the units [meV] for energy, [km/s] for speed, [Å] for wave length and [Å^{-1}] for wave number, then the numerical conversions read:

$$E = 5.2267v^2 = \frac{81.799}{\lambda^2} = 2.072178k^2. \tag{A.1}$$

For example, a 4 Å neutron has an energy of about 5 meV and travels roughly at 1000 m/s.

Appendix B
Scattering Length Tables and Sample Thickness Calculations

In this appendix we list commonly used information that is useful in planning and analyzing neutron scattering experiments.

The tables are for the scattering lengths for the atoms, and for some isotopes. All absorbtion cross-sections are listed for neutrons of 25.3 meV. At a glance, the tables will tell the experimenter whether a particular element absorbs so strongly that it can be a problem in neutron scattering experiments. But the tables will also help in calculating how much material one needs to obtain a sample that scatters neutrons by the desired amount. We give an example of the latter.

Suppose we wish to investigate a sample of $CeFe_2Ge_2$. The unit cell of this sample contains 2 formula units, and the tetragonal unit cell measures 4 x 4 x 10 $Å^3$. We wish to calculate what thickness we would need for about a 10 % scatterer. In other words, we are looking for the thickness L that yields a 90 % transmission as in $T = 0.9e^{-n\sigma L}$, or $n\sigma L \approx 0.1$. n is the number density, which in this case would be 2 formula units per 160 $Å^3$, or 1.25 x 10^{22}/cm^3. We use units of cm, so that we get our thickness in cm.

Next we need to calculate the scattering cross section per formula unit, which is a combination of the coherent and incoherent scattering cross-sections. We use the table below to find

$$\sigma_{\text{scat,formula unit}} = \sigma_{\text{Ce,coh}} + \sigma_{\text{Ce,inc}} + 2\sigma_{\text{Fe,coh}} + 2\sigma_{\text{Fe,inc}} + 2\sigma_{\text{Ge,coh}} + 2\sigma_{\text{Ge,inc}}.$$

Putting in the numbers we find $\sigma_{\text{formula unit}} = 43.78$ barn = 43.78 x 10^{-24} cm^2. With these numbers we find $L = 0.18$ cm, so not very thick at all.

Let us also check on the absorbtion cross-section, to see if this would be a problem. We evaluate this for neutrons that have an energy of 25.3 meV, thus,

$$\sigma_{\text{abs,formula unit}} = \sigma_{\text{Ce,abs}} + 2\sigma_{\text{Fe,abs}} + 2\sigma_{\text{Ge,abs}}.$$

Using the tables below we find $\sigma_{\text{abs,formula unit}} = 10.15$ barn. For a sample of 0.18 cm thickness, this would lead an absorbtion of about 2.2 % of the neutrons. This is a very manageable number, so we do not have to worry about the feasibility of the experiment, but we probably should correct for it during the analysis stages.

These tables have the incoherent cross-sections listed per element. This incoherent cross-section can come from different isotopes that make up the element as found in nature, and it can come from the incoherent cross-section of individual atoms. We can have another source of incoherent scattering in samples, namely when we have more than one atom that can sit on identical positions in the unit cell. For instance, going back to our example of $CeFe_2Ge_2$, we can decide to substitute some ruthenium in place of the iron atoms. This will give rise to incoherent scattering since now the scattered waves originating from nominally identical positions within the crystal lattice, will have a different scattering strength. We would have to recalculate the coherent and incoherent cross section in this case.

This calculation follows the following recipe. We define an average scattering length $\bar{b}$ and an average of the scattering length squared $(\bar{b^2})$. With these numbers we define the coherent and incoherent cross-sections by

$$\sigma_{\text{coh}} = 4\pi\bar{b}^2$$

$$\sigma_{\text{inc}} = 4\pi[(\bar{b^2}) - \bar{b}^2].$$

Looking at the total scattering, we see that we do not add or subtract from the total level of scattering which still is $4\pi(\bar{b^2})$, but rather some of the coherent scattering now becomes incoherent scattering.

Suppose we substitute 25 % Ru on the Fe sites in $CeFe_2Ge_2$. We then find that the average scattering length on the Fe/Ru site is 8.96 fm, and the average scattering length squared is 81.25 fm^2. This gives us $\sigma_{\text{coh, Fe/Ru}}$= 10.09 barn and $\sigma_{\text{inc, Fe/Ru}}$= 0.122 barn. To this latter number for the incoherent cross-section, we would still have to add the incoherent cross-section per atom, that is, 0.75 x $\sigma_{\text{inc, Fe}}$ + 0.25 x $\sigma_{\text{inc, Ru}}$= 0.4 barn. The reason for this is that we still have the incoherent cross-section that each individual atom contributes, but we have added to the loss of coherence by randomly substituting Ru atoms on Fe sites. We would then use those numbers ($\sigma_{\text{coh, Fe/Ru}}$= 10.09 barn and $\sigma_{\text{inc, Fe/Ru}}$= 0.52 barn) in our further calculations to calculate the scattering cross-section per formula unit.

Table B.1 The element and isotope dependent scattering lengths and cross-sections. '*' refers to the naturally occurring isotope mixture for an element. 1 barn $= 10^{-24}$ cm^2.

Z	Name	A	mass [au]	b_{coh} [fm]	σ_{coh} [barn]	σ_{inc} [barn]	σ_{abs} [barn]
1	H	*	1.008	-3.739	1.757	80.263	0.333
1	D	2	2.016	6.671	5.592	2.048	0.001
1	T	3	3.024	4.940	3.067	0.000	0.000
2	He	*	4.003	3.260	1.336	0.004	0.007
2	He	3	3.003	5.740	4.140	1.460	5333
2	He	4	4.003	3.260	1.336	0.004	0.000
3	Li	*	6.941	-1.900	0.454	0.916	70.500
3	Li	6	6.015	1.870	0.439	0.541	940
3	Li	7	7.016	-2.200	0.608	0.832	0.045
4	Be	*	9.012	7.790	7.626	0.004	0.008
5	B	*	10.810	5.300	3.530	1.710	767.000
5	B	10	10.013	0.000	0.000	0.980	3837
5	B	11	11.009	6.660	5.574	0.226	0.001
6	C	*	12.011	6.646	5.550	0.001	0.004
7	N	*	14.007	9.360	11.009	0.501	1.900
8	O	*	15.999	5.803	4.232	0.000	0.000
9	F	*	18.998	5.654	4.017	0.001	0.010
10	Ne	*	20.179	4.566	2.620	0.008	0.039
11	Na	*	22.990	3.580	1.611	1.674	0.530
12	Mg	*	24.305	5.375	3.631	0.079	0.063
13	Al	*	26.982	3.449	1.495	0.008	0.231
14	Si	*	28.086	4.153	2.168	0.003	0.171
15	P	*	30.974	5.130	3.307	0.005	0.172
16	S	*	32.060	2.847	1.019	0.007	0.530
17	Cl	*	35.453	9.577	11.526	5.274	33.500
18	Ar	*	39.948	1.909	0.458	0.225	0.675
19	K	*	39.098	3.670	1.693	0.267	2.100
20	Ca	*	40.080	4.760	2.847	0.043	0.430
21	Sc	*	44.956	12.290	18.981	4.519	27.500
22	Ti	*	47.900	-3.438	1.485	2.865	6.090
23	V	*	50.942	-0.382	0.018	5.082	5.080
24	Cr	*	51.996	3.635	1.660	1.830	3.050

Table B.2 The element and isotope dependent scattering lengths and cross-sections, continued

Z	Name	A	mass [au]	b_{coh} [fm]	σ_{coh} [barn]	σ_{inc} [barn]	σ_{abs} [barn]
25	Mn	*	54.938	-3.730	1.748	0.402	13.300
26	Fe	*	55.847	9.540	11.437	0.383	2.560
27	Co	*	58.933	2.780	0.971	4.829	37.180
28	Ni	*	58.700	10.300	13.332	5.168	4.490
29	Cu	*	63.546	7.718	7.485	0.545	3.780
30	Zn	*	65.380	5.680	4.054	0.077	1.110
31	Ga	*	69.720	7.288	6.675	0.155	2.750
32	Ge	*	72.590	8.185	8.419	0.181	2.200
33	As	*	74.922	6.580	5.441	0.059	4.500
34	Se	*	78.960	7.970	7.982	0.318	11.700
35	Br	*	79.904	6.795	5.802	0.098	6.900
36	Kr	*	83.800	7.810	7.665	0.015	25.000
37	Rb	*	85.468	7.090	6.317	0.483	0.380
38	Sr	*	87.620	7.020	6.193	0.057	1.280
39	Y	*	88.906	7.750	7.548	0.152	1.280
40	Zr	*	91.220	7.160	6.442	0.018	0.185
41	Nb	*	92.906	7.054	6.253	0.002	1.150
42	Mo	*	95.940	6.715	5.666	0.044	2.480
43	Tc	*	97.000	6.800	5.811	0.489	20.000
44	Ru	*	101.070	7.210	6.533	0.067	2.560
45	Rh	*	102.906	5.880	4.345	0.255	144.800
46	Pd	*	106.400	5.910	4.389	0.091	6.900
47	Ag	*	107.868	5.922	4.407	0.583	63.300
48	Cd	*	112.410	5.100	3.269	2.431	2520
49	In	*	114.820	4.065	2.076	0.544	193.800
50	Sn	*	118.690	6.225	4.870	0.022	0.626
51	Sb	*	121.750	5.570	3.899	0.001	4.910
52	Te	*	127.600	5.800	4.227	0.093	4.700
53	I	*	126.905	5.280	3.503	0.307	6.150
55	Cs	*	132.905	5.420	3.692	0.208	29.000
56	Ba	*	137.330	5.070	3.230	0.150	1.100
57	La	*	138.906	8.240	8.532	1.128	8.970

Table B.3 The element and isotope dependent scattering lengths and cross-sections, continued

Z	Name	A	mass [au]	b_{coh} [fm]	σ_{coh} [barn]	σ_{inc} [barn]	σ_{abs} [barn]
58	Ce	*	140.120	4.840	2.944	0.000	0.630
59	Pr	*	140.907	4.450	2.488	0.042	11.500
60	Nd	*	144.240	7.690	7.431	9.169	50.500
61	Pm	*	145.000	12.600	19.950	1.350	168.400
62	Sm	*	150.400	0.800	0.080	38.920	5922
63	Eu	*	151.960	7.220	6.551	2.649	4530
64	Gd	*	157.250	6.500	5.309	174.691	49700
65	Tb	*	158.925	7.380	6.844	0.000	23.400
66	Dy	*	162.500	16.900	35.891	54.409	994
67	Ho	*	164.930	8.010	8.063	0.357	64.700
68	Er	*	167.260	8.160	8.367	0.833	159
69	Tm	*	168.934	7.070	6.281	0.099	100
70	Yb	*	173.040	12.430	19.416	3.984	34.800
71	Lu	*	174.967	7.210	6.533	0.667	74.000
72	Hf	*	178.490	7.770	7.587	2.613	104.100
73	Ta	*	180.948	6.910	6.000	0.010	20.600
74	W	*	183.850	4.860	2.968	1.632	18.300
75	Re	*	186.200	9.200	10.636	0.864	89.700
76	Os	*	190.200	10.700	14.387	0.313	16.00
77	Ir	*	192.220	10.600	14.120	0.000	425
78	Pt	*	195.090	9.600	11.581	0.129	10.300
79	Au	*	196.967	7.630	7.316	0.434	98.650
80	Hg	*	200.590	12.692	20.243	6.557	372.300
81	Tl	*	204.370	8.776	9.678	0.212	3.430
82	Pb	*	207.200	9.405	11.115	0.003	0.171
83	Bi	*	208.980	8.532	9.148	0.008	0.034
88	Ra	*	226.025	10.000	12.566	0.434	12.800
90	Th	*	232.038	10.520	13.907	0.003	7.370
91	Pa	*	231.036	9.100	10.406	0.094	200.600
92	U	*	238.029	8.417	8.903	0.005	7.570
93	Np	*	237.048	10.550	13.987	0.513	175.900

Appendix C
Some Physics Background on Scattering Lengths

In this appendix we give some background about the scattering process of the neutron by the nucleus. Physicists who read this part will feel good about themselves, non-physicists who read this appendix probably would have done quite well as a physicist. We will discuss three points relevant to the scattering process:

- how can a plane wave that is incident on a nucleus give rise to a spherically symmetric scattered wave, or in other words, why is the scattering likelihood by a single atom not a function of the angle over which a neutron is scattered?
- why are the scattering lengths roughly the same as the radius of the nucleus?
- why are some scattering lengths negative, and why are some very large?

It is always possible to represent a function of any shape by a summation over a set of standard functions, as long as this set of standard functions satisfies certain requirements. An example of such a set would be the set of sine of cosine functions. A plane wave, which is a good representation of the incident neutron wave, can be represented by a summation over a set of partial waves. We show this in Fig. C.1. For the physicists amongst us, we decompose the plane wave into a summation over partial waves as follows:

$$e^{i\vec{q}\cdot r} = \sum_{\ell=0}^{\infty} (2\ell + 1) i^\ell j_\ell(qr) P_\ell(cos\theta). \tag{C.1}$$

In here, the $j_\ell(qr)$ are Bessel functions, and the $P_\ell(cos\theta)$ are the Legendre polynominals.

In the scattering process, the partial waves that make up the plane wave get scattered by the nucleus. After the scattering process, we would add up all the partial waves again to see what the resultant neutron wave would look like after the interaction of the neutron with the nucleus. In order for a partial wave to be modified in the scattering process, this partial wave will have to have an appreciable amplitude in the region of the nucleus. Only the partial wave of order 'zero' has an appreciable amplitude. All other partial waves, the ones that have some angular dependence to them, fall off very rapidly in amplitude near the origin (which was chosen to be at the center of the nucleus). In fact, the partial waves of order ℓ fall off as $\sim (qr)^\ell$ with
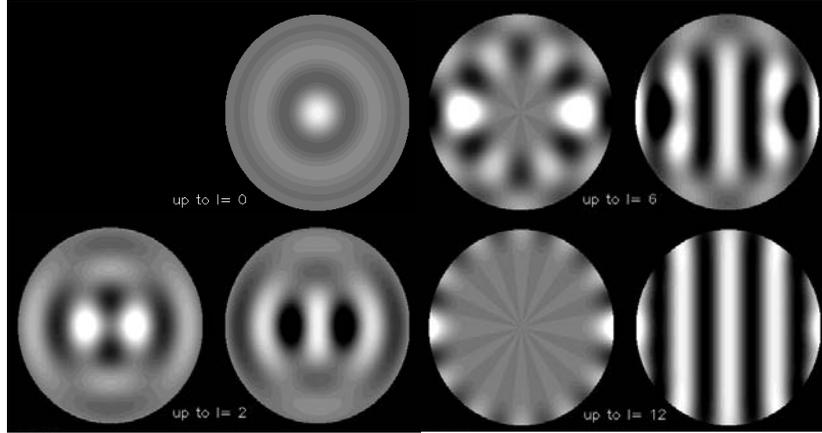
Fig. C.1 By adding up successive partial waves we can end up with any functional shape that we desire. The figure shows the results after adding up terms according to eqn C.1. The $\ell=0$ result is shown top left, all terms up to $\ell=2$ are shown at the bottom left (the leftmost figure is the partial wave, the figure just to the right of it is the summation over all partial waves up to $\ell=2$), all terms up to $\ell=6$ are shown at the top right, and all terms up to $\ell=12$ are shown at the bottom right. The characteristic plane wave pattern is almost complete for this latter summation.

r the distance to the center of the nucleus since $j_\ell(x) \sim x^\ell$. Given typical q-values and nuclear sizes, we have that qr is of the order of 10^{-4}-10^{-5} inside the nucleus, so that only the $\ell = 0$ wave can be changed in the scattering process. By extension, the scattered wave will not show any angular dependence, it will be spherically symmetric. This is called s-wave scattering.

Since the scattering by a single nucleus will not have angular dependence, we can write the wave function Ψ of the neutron (the solution to the Schrödinger equation) as a combination of the incoming plane wave and the outgoing scattered wave:

$$\Psi = e^{i\vec{q}.r} - be^{iqr}/r. \tag{C.2}$$

In here, b is a coefficient of the partial wave of order zero, a number we have identified as the scattering length of a nucleus. The $1/r$ dependence of the scattered wave ensures that when we are far away from the nucleus, that the intensity of the scattered wave fall of as $1/r^2$ since it will be given by the amplitude squared of the scattered wave. In particular, we would find that

$$|\Psi| = b\frac{e^{iqr}}{r} b\frac{e^{-iqr}}{r} = \frac{b^2}{r^2}$$

Also note that if we integrate over all 4π angles across the surface of a sphere a distance r away from the scattering center, we get the overall scattering amplitude of

$4\pi b^2$, a number we have referred to as σ, the scattering cross-section.

Now that we have established that the scattered waves are spherically symmetric because of the smallness of the argument qr, we can estimate the order of magnitude of the scattering length b. We imagine the nucleus to be an impenetrable hard sphere of radius R_N. The wave function inside this hard sphere will have zero amplitude, and because of the condition of continuity on wave functions, we must also have that $\Psi(r_N)$ equals zero. Evaluating eqn C.2 at $r = R_N$ and making use of the approximation that $e^{i\vec{q}\vec{r}} \approx e^{iqr} \approx 1$ given the smallness of qr, we immediately find that $b \approx r_N$. This order of magnitude estimate was shown in Fig. 3.3 and it can be seen to be a good overall approximation for b.

We can do a a better job at calculating the scattering length b. While treating the neutron-nucleus interaction as a collision between hard spheres gives us the correct order of magnitude, it fails to capture the details of the scattering process, and the fact that the interaction is attractive. A more realistic approximation is the square well potential:

$$V(r) = -V_0 \text{ if } r < r_N \text{ and zero elsewhere.}$$

In order to find a solution to the Schrödinger wave equation for a neutron of energy E that is continuous at $r = R_N$, we now have to match the two solutions inside and outside the square well at $r = R_N$. In fact, we have to match the wave functions, and their derivatives for a square well-potential that is not infinitely deep. The solution inside a square well can be looked up in standard textbooks, or the (physics) reader can verify it by plugging it into the Schrödinger equation. Thus, we need to match:

$$\Psi_{\text{inside}}(r) = A sin(Qr)/r \quad \text{with } \hbar Q = [2mE + V_0]^{0.5}$$
$$\Psi_{\text{outside}}(r) = e^{i\vec{q}.r} - be^{iqr}/r \quad \text{with } \hbar q = [2mE]^{0.5}.$$

Note that Qr can be very large when V_0 is very large, which is the case inside a nucleus. qr on the other hand will remain small so that we can make the approximation $\Psi_{\text{outside}}(r) = 1 - b/r$. Evaluating and equating the wave functions and their derivatives as $r = R_N$ we find

$$A\sin(Qr) = R_N - b$$
$$AQ\cos(Qr) = 1.$$

Dividing the two equations we find:

$$\frac{b}{R_N} = 1 - \frac{\tan(QR_N)}{QR_N}. \tag{C.3}$$

We show the solution of eqn C.3 in Fig. C.2 for illustrative values of V_0=-50 MeV and E= 25 meV. While picking just one number to describe all atoms, the figure does show why some scattering lengths can be negative, and also why most scattering lengths are of the order $b \approx R_N$: the latter is because of the shape of the tangent curve.
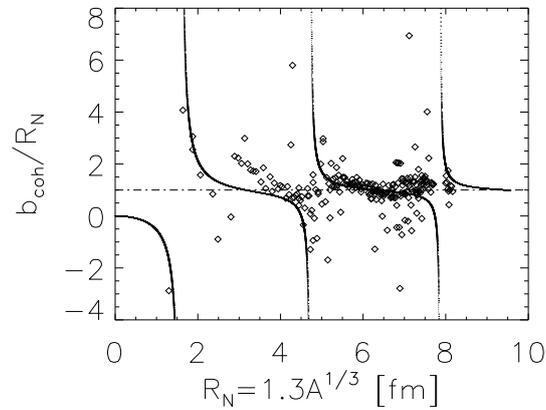
Fig. C.2 The solution (solid line) to eqn C.3 evaluated for $V_0=$-50 MeV and $E=$ 25 meV and compared to the measured scattering lengths of various atoms. R_N is approximated by $R_N = 1.3A^{1/3}$, with A the number of nucleides (neutrons and protons).

We can also use the above solution to understand, conceptually, why some scattering lengths can be negative, and why some are very large. We show this graphically in Figs. C.3 through C.5. What we have done is that we picked one value for the nuclear radius, and we varied the depth of the square-well potential. We matched the solutions inside and outside the square well by ensuring that they are at the same level, and that they have the same slope. When plotting the wave functions multiplied by r, then we can read of the scattering length as the intercept of the horizontal axis. The reader
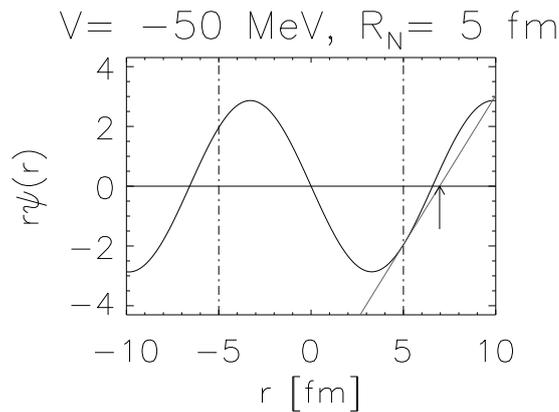
Fig. C.3 Plotted is $r\Psi_{\text{inside}}(r)$ and $r\Psi_{\text{outside}}(r)$ in(r). The nuclear radius has been fixed at 5 fm, and the depth of the square well potential has been chosen to be 50 MeV. Both wave functions, and their derivatives, are made to match at $r = R_N$. The arrow points to the point where $r\Psi_{\text{outside}}(r) = r - b = 0$. This type of solution is characteristic for most atoms.

can verify from these figures that this intercept can be both at positive or negative values, and that this intercept can occur at very large values indeed. This discussion is meant to illustrate the variety of scattering lengths, as well as the fact that most scattering lengths are of the same order as the nuclear radius.
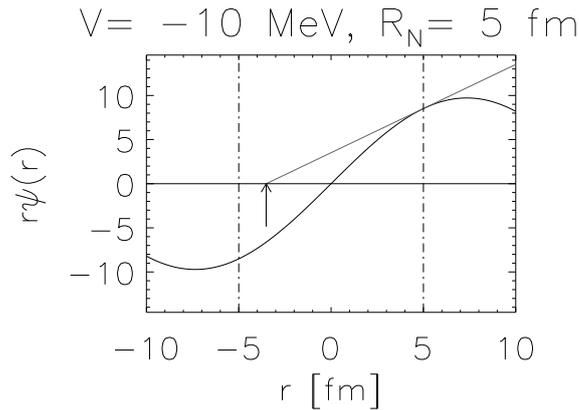
Fig. C.4 Plotted is $r\Psi_{\text{inside}}(r)$ and $r\Psi_{\text{outside}}(r)$ in(r). The nuclear radius has been fixed at 5 fm, and the depth of the square well potential has been chosen to be 10 MeV. Both wave functions, and their derivatives, are made to match at $r = R_N$. The arrow points to the point where $r\Psi_{\text{outside}}(r) = r - b = 0$.
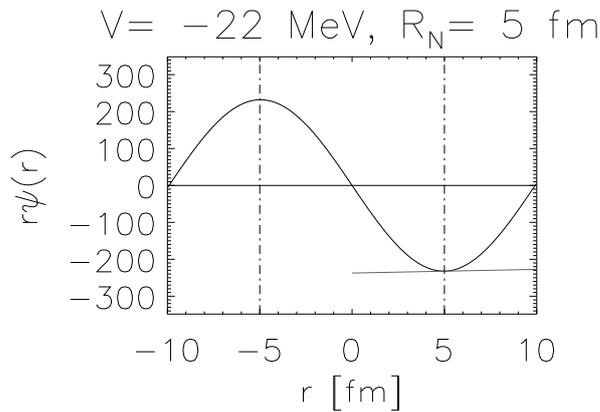
Fig. C.5 Plotted is $r\Psi_{\text{inside}}(r)$ and $r\Psi_{\text{outside}}(r)$ in(r). The nuclear radius has been fixed at 5 fm, and the depth of the square well potential has been chosen to be 22 MeV. Both wave functions, and their derivatives, are made to match at $r = R_N$. The invisible arrow points to the point where $r\Psi_{\text{outside}}(r) = r - b = 0$.

Appendix D
Fourier and Laplace Transforms

Fourier and Laplace transforms connect the scattering function between reciprocal space and real space, and between energy transfer and time. Because of quantum mechanics, which states that not all operators commute, the scattering functions are defined (deep down) as commutators, and the more familiar functions such as the dynamic structure factor are related to these commutator based entities. In this appendix we give the formal definition for the dynamic susceptibility, followed by some examples for more familiar functions.

In real space and in time, we define the imaginary part of the response function between two variables A and B as follows:

$$\chi"_{AB}(\vec{r} - \vec{r}', t - t') \equiv \frac{1}{2\hbar} < [A(\vec{r}, t), B(\vec{r}', t')] >_{\text{eq}}.$$ (D.1)

The square brackets denote the commutator, while $< ... >_{\text{eq}}$ stands for the ensemble average. The variable most relevant to neutron scattering is the number density n, and moreover, we are mostly interested in those response functions between the density and the density: $A = B = n$. Thus (classically speaking), we perturb the density away from its equilibrium density near some location $\vec{r}$ at time t, and we study how the density near point $\vec{r}'$ at time t' responds to this disturbance. We use a Fourier transform to go from real space to reciprocal space, and from time to energy transfer $E = \hbar\omega$:

$$\chi"(\vec{q}, \omega) = \int_{-\infty}^{\infty} dt" e^{i\omega t"} \int_V d\vec{r}" e^{-i\vec{q}.\vec{r}"} \chi"(\vec{r}", t").$$ (D.2)

In the above equation we dropped the label $A = B = n$. This imaginary part of the dynamic susceptibility is directly related to the dynamic structure factor, the one that we measure in inelastic neutron scattering experiments:

$$\chi"(\vec{q}, \omega) = \frac{1 - e^{-\beta\hbar\omega}}{2\hbar} S(\vec{q}, \omega).$$ (D.3)

The above equations are correct, but note that we did not directly carry out a Fourier transform of the intermediate scattering function $F(\vec{q}, t)$ to obtain $S(\vec{q}, E)$. We did this because the interpretation of the intermediate scattering function is not very straightforward. For starters, it is a function that takes on complex values, while we think of it as a real function. In fact, the intermediate scattering function that is

determined through computer simulations is a real function, but it is not the Fourier transform of the dynamic structure factor. We wished to leave those difficulties behind and instead we hid behind the formal definitions in terms of commutators.

Instead of a Fourier transform, one also tends to use a Laplace transform to go from a time-dependence to an energy dependence. For the dynamic susceptibility, this Laplace transform is given by

$$\chi(\vec{q}, z) = 2i \int_0^\infty dt'' e^{izt''} \int_V d\vec{r}'' e^{-i\vec{q}.\vec{r}''} \chi''(\vec{r}'', t''), \qquad (D.4)$$

for Im$[z]>0$ (corresponding to 'physical', positive times $t''>0$), with t''=t-t' and $\vec{r}''=\vec{r}$-$\vec{r}'$. A similar definition holds for Im$[z]<0$ (corresponding to 'unphysical', negative times $t''<0$). Substituting eqn D.2 into eqn D.4 one finds, for Im[z]$\neq 0$,

$$\chi(\vec{q}, z) = \frac{1}{\pi} \int_{-\infty}^\infty d\omega \chi''(\vec{q}, \omega) \frac{1}{\omega - z}. \qquad (D.5)$$

Ideally, one would like to determine the poles of the dynamic susceptibility; these poles give us both the energies as well as the life times of the excitations of the system that we are studying. We only have access to the imaginary part of the susceptibility, and performing the transformation embodied in eqn D.5 is a tour de force. But as mentioned in Chapter 12 (eqn 12.1), by using the symmetrized dynamic structure factor we can circumvent most of the difficulties associated with only having access to the imaginary part of the susceptibility.

Now that we have given the formal and correct definitions (so that we will not get expelled from the APS for being too handwavy), we will move on to some more useful, and slightly incorrect examples of frequently encountered functions.

D.1 Examples

In this section we pretend quantum mechanics does not really exist in the sense that functions such as the intermediate scattering function are real valued. The purpose of this section is to give the reader an idea of what kind of line shapes to expect for certain processes within the system.

When we perform quasi-elastic scattering experiments, where we probe the scattered signal close to zero energy transfers, we often find a line shape in energy that looks Lorentzian. Classically speaking, this is the line shape that can be expected for a diffusion process. We will discuss how diffusion manifests itself in the scattering functions, and their relationship to real space functions.

The microscopic number density $n(\vec{r}, t)$ for a system of N particles is defined in real space as a function of time as

$$n(\vec{r}, t) = \sum_i^N \delta(\vec{r} - \vec{r}_i(t))/\sqrt{N}, \qquad (D.6)$$

with $\vec{r}_i(t)$ the position of particle i at time t. The density-density correlation function $G_{nn}(\vec{r} - \vec{r}', t - t')$ measures what effect a density disturbance at position $\vec{r}$ and time t has on the density of the fluid at position $\vec{r}'$ at a later time t':

$$G_{nn}(\vec{r} - \vec{r}', t - t') \equiv < n(\vec{r}, t)n(\vec{r}', t') >_{eq} . \tag{D.7}$$

The ensuing density disturbance $n(\vec{r}', t')$ can take on various guises. For instance, we could observe that at time t' the atoms alternately compress and rarefy, until they return back to their average equilibrium state. This of course would have happened because a sound wave that originated at $\vec{r}$ passed through the region around $\vec{r}'$ with a propagation speed of $|\vec{r} - \vec{r}'|/(t' - t)$. This can happen both in solids and in liquids.

Alternatively, in liquids, we could observe that some new atoms streamed into the region around $\vec{r}'$, pushing out some of the atoms already there. This would correspond to the density disturbance at $\vec{r}$ relaxing back to equilibrium through the collective diffusion mechanism with the diffusion constant for this process given by $D = |\vec{r} - \vec{r}'|^2/(t' - t)$.

What line shape would be associated with a diffusion process? Imagine increasing the density locally at some point in a liquid with the result that a deviation $n(\vec{r}, t)$ from the average equilibrium density n_{eq} has been created. For this excess density to disappear, it must be that more particles are leaving the region than are entering it, so we must have that $\partial/\partial t\, n(\vec{r}, t) \sim -\vec{\nabla}.\vec{u}(\vec{r}, t)$. Here $\vec{u}(\vec{r}, t)$ is the velocity of the fluid near $\vec{r}$ at time t. Furthermore, the velocity itself must have arisen in the first place because of a gradient in the density; after all, particles move from places of high density to places of low density. Thus, $\vec{u}(\vec{r}, t) \sim -\vec{\nabla}n(\vec{r}, t)$. After eliminating the velocity, we get

$$\partial/\partial t n(\vec{r}, t) \sim \nabla^2 n(\vec{r}, t) \equiv D\nabla^2 n(\vec{r}, t).$$

This is Fick's law for diffusion. This equation can be solved directly, or it can be solved by Fourier or Laplace transforming. It is easily verified that the solution is

$$G(r, t) = 1/\sqrt{4\pi D_s t}e^{-r^2/4Dt},$$

$$F(q, t) = e^{-Dq^2 t} \equiv e^{-t/\tau},$$

and the associated dynamic structure factor is obtained through a Laplace transform (valid for our classical approximation) to read:

$$S(q, \omega) = \frac{1}{\pi}\frac{Dq^2}{\omega^2 + (Dq^2)^2}. \tag{D.8}$$

In order to obtain this result we used the following Laplace transformation which works very nicely for our classical approximation:

$$S(q, \omega) = \text{Re}[\int_0^\infty dt e^{i\omega t}F(q, t)].$$

In experiments, one frequently uses this line shape to describe quasi-elastic scattering data. In practice, one convolutes this line shape with the resolution function of the

spectrometer in order to obtain a good agreement between the diffusion model and the measured line shape.

Eqn D.8 shows us that what we would call the relaxation rate Γ is given by $\Gamma = Dq^2$, or equivalently, a characteristic decay time τ of $\tau = 1/Dq^2$. This $\sim q^2$ dependence is characteristic of all damping processes involving diffusion.

When we have excitations that are propagating, while decaying at the same time, we also will get Lorentzian lineshapes, or something close to it. Suppose we have an excitation that takes an amount of energy $\hbar\omega_p$ to create and that persists for a characteristic time τ. The intermediate scattering function of such an excitation would be given by:

$$F(q,t) = e^{i\omega_p t - t/\tau},$$

yielding a dynamic structure factor

$$S(q,\omega) = \frac{1}{\pi}\frac{1/\tau}{(\omega - \omega_p)^2 + (1/\tau)^2}.$$

This equation immediately tells us that we need to transfer $\hbar\omega = \hbar\omega_p$ in energy to get a nice peak in our dynamic structure factor, and that the sharpness of the peak is determined by the decay time. The shorter the decay time, the broader the peak. As mentioned in Chapter 12, this equation is not complete. Excitations can already spontaneously exist in the system, and the neutron can absorb them. hence there would be an additional term in the dynamic structure factor given by

$$S(q,\omega) \sim \frac{1}{\pi}\frac{1/\tau}{(\omega + \omega_p)^2 + (1/\tau)^2}.$$

The weight of the two terms is determined by the temperature of the system since the temperature determines how much energy is available for excitations to spontaneously exist. When we add up the two parts of the structure factor, in the correct way as explained in the chapter on pitfalls in data analysis, then we obtain expressions such as eqn 12.4. The main point here is that when we have disturbances that take a certain amount of time to decay, then we expect (almost) Lorentzian lines in energy in the dynamic structure factor. They are just not exactly Lorentzian because real systems do not behave according to classical laws, so that we get slightly more complicated expressions such as those in eqn 12.4.

In Chapter 4 and exercise 4.3 we encountered a Gaussian line shape in energy for very large momentum transfers in liquids. This is one of the few instances that we actually have a Guassian line shape that is caused by an intrinsic process of our system. Very frequently we will encounter line shapes that are Gaussian because the width in energy of the resolution function of the spectrometer is so much larger than $1/\tau$ that the underlying Lorentzian line shape has been rendered unrecognizable.

The above discussion deals with line shapes in energy, but what does the scattering look like as a function of momentum transfer, especially in the case of elastic

scattering where we look at correlations between neighboring positions or orientations of magnetic moments? We can expect a variety of shapes under these circumstances, depending on the dimensionality of our system.

Suppose we have correlations between neighboring atoms that persist over a characteristic distance ξ, the correlation length. The easiest way to visualize this is by picturing magnetic moments that have a tendency to align with their neighbors, but do not exhibit long range order from one side of the sample to the other because thermal fluctuations disturb this ordering tendency. As a result, moments will typically be correlated over a distance ξ.

When we do a neutron scattering experiment, we will find that the correlation between neighboring moments decays as

$$g(r) \sim e^{-r/\xi}. \tag{D.9}$$

We did not include any time dependence in the above correlation function since in this example we are talking about the instantaneous correlation between moments. Instantaneous implies that we average over all energies of the scattered neutron, which implies that we are doing a diffraction experiment. Which is what one typically does when one is interested in the degree of correlation between neighboring moments (or positions between neighboring atoms in a liquid).

In neutron scattering we would measure the Fourier transform, from real space to reciprocal space:

$$S(\vec{q}) = \int_V d\vec{r} e^{-i\vec{q}\cdot\vec{r}} g(r).$$

The line shape (in q) that we can expect depends on whether the moments are correlated with each other in a three dimensional sense, or whether the moments first line up in a plane, or even along a line. As an example, we calculate $S(q)$ for a system of magnetic moments that order with respect to their neighbors, but in such a way that the moment vectors are oriented in a plane, with no preferred overall direction in the plane. What we mean by this is that all the moments are lining up with their neighbors, but the net moments of this lined-up conglomerate can point in any direction of the plane. Thus, we will not have any preferred direction of momentum transfer $\vec{q}$, so we will have to average over all directions in the plane. We have a 2-dimensional problem on our hand.

The integral is straightforward to carry out, and we actually give some details of the steps involved so that the reader can perform similar integral in other dimensions:

$$S(q) \sim \int_0^\infty dr \int_0^{2\pi} d\theta r \sin\theta e^{-iqr\cos\theta} e^{-r/\xi}$$

$$= 2 \int_0^\infty dr r \int_{-1}^1 d\cos\theta e^{-iqr\cos\theta} e^{-r/\xi}$$

$$= \frac{2}{-iq} \int_0^\infty dr [e^{-iqr} e^{-r/\xi} - e^{iqr} e^{-r/\xi}]$$

$$= \frac{2}{-iq} [\frac{-1}{-iq - 1/\xi} - \frac{-1}{iq - 1/\xi}]$$

$$= \frac{4}{q^2 + 1/\xi^2}.$$

Thus, $S(q)$ is described by a Lorentzian line shape in q, and $1/\xi$ takes on the role of the width of the peak in q-space. Note that we did not keep track of normalization factors since we merely stated that the decay of correlations follows an exponential dependence on the distance. Should the orientation of the moments not be restricted to lie in a plane, but instead be allowed to point in any direction in space, then we would have to carry out a 3-dimensional integral. For this case we would find the Lorentzian-squared function:

$$S(q) \sim \frac{1}{(q^2 + 1/\xi^2)^2}$$

The bottom line of this discussion is that there is important information hidden in the line shape of the scattering data. Not merely what the characteristic width is, in energy or in q, but also in the details of the line shape, be it Lorentzian or Lorentzian-squared.

As a last mention, in liquids where there is no preferred direction, one frequently uses the pair correlation function to interpret the scattering data. This pair correlation function $g(r)$ is given by a Fourier transform of the static structure factor as

$$S(q) - 1 = 4\pi n \int_0^\infty dr r^2 [g(r) - 1] \frac{sin(qr)}{qr}.$$

Appendix E
Sumrules, Limiting Behaviors and Absolute Normalization

We have sumrules on the cross-sections such as the double differential cross-section, on the scattering functions such as the dynamic structure factor, and we have limiting behaviors of some scattering functions. Sumrules are obtained by summing (integrating) over all possibilities, such as integrating over all possible energy transfers in an inelastic scattering experiment. Sumrules can serve as a check on one's data, or they can be used to put one's data on an absolute scale. The same holds for the limiting behavior of some functions, these limiting values can be used to absolutely normalize one's data.

An example of (a not particularly useful) set of sumrules are those that pertain to the neutron scattering cross-sections directly:

$$\int \frac{d^2\sigma}{d\Omega dE'} dE' = \frac{d\sigma}{d\Omega},$$

$$\int \frac{d\sigma}{d\Omega} d\Omega = \sigma.$$

A much more useful sumrule is the one related to the above set, namely the relationship between the dynamic structure factor and the static structure factor:

$$\int_{-\infty}^{+\infty} S(\vec{q}, E) dE = S(\vec{q}). \tag{E.1}$$

This relationship can be used to put ones scattering data onto an absolute scale, in some cases. When doing an inelastic scattering experiment, it is possible to obtain $S(\vec{q}, E)$ apart from some normalizing prefactor that allows us to convert from counts to the units of $S(\vec{q}, E)$. Frequently, especially in the case of experiments on liquids, one has already measured the static structure factor in some other experiment so that by performing the integration in eqn E.1 one can determine the normalizing prefactor by comparing to the known static structure factor.
 Even in the cases where the static structure factor is not known, eqn E.1 can still be used to put one's data onto an absolute scale. The first way that this can be done is when $S(\vec{q}, E)$ has been determined for very high momentum transfers. In this limit, we have that $S(\vec{q})$ will no longer oscillate as a function of q but instead it will reach its limiting value of 1.

The second way that absolute normalization can be achieved is for the case of liquids where we can use the so-called f-sum rule:

$$\int_{-\infty}^{+\infty} S(q,E)EdE = \frac{\hbar^2 q^2}{2m}, \tag{E.2}$$

with m the mass of the atoms in the liquid. This sumrule can be used to normalize one's data on liquids at low momentum transfer. The energy on the right hand side is also known as the recoil energy.

Note that we also have the following sumrule, which relates the static susceptibility $\chi(q)$ to the dynamic structure factor:

$$\int_{-\infty}^{+\infty} \frac{S(q,E)}{E}dE = \pi\chi(q) = \frac{S_{\text{sym}}(q)}{2k_B T}. \tag{E.3}$$

As mentioned, the static structure factor has the limiting behavior:

$$\lim_{q\to\infty} S(\vec{q}) = 1.$$

This relationship is a direct consequence of the fact that we cannot expect any constructive interference to originate from neighboring atoms when we probe the system using very short probing wave lengths. When our scattering is properly normalized, $S(\vec{q})$ should reach the value of 1. Therefore, we simply obtain the limiting value from our data in terms of counts, and we have obtained our normalizing factor.

In liquids, we also have the low-q limiting behavior of S(q) which is given by

$$\lim_{q\to 0} S(\vec{q}) = nk_B T\chi_T = \frac{c_p/c_v}{nmc^2}, \tag{E.4}$$

with χ_T the isothermal compressibility of a liquid with number density n, c the adiabatic speed of sound, and c_p and c_v the specific heats at constant pressure and volume, respectively. In principle, eqn E.4 can be used to put one's data onto an absolute scale, in practice such a normalization will not be very accurate, although it may still be accurate enough for one's purposes.

Another very useful relationship is the detailed balance condition:

$$e^{-E/k_B T}S(\vec{q},E) = S(\vec{q},-E). \tag{E.5}$$

This relationship is used to check the quality of one's data, such as whether there is a potential monitor correction lurking in the data correction procedure. Of course, this relation cannot be used to put one's data onto an absolute scale, but one should always try to use it to double check one's data correction procedures.

The above sumrules are very useful in scattering by liquids and amorphous materials, but in solids they are not as useful since one has to reach much higher q-values

before the limiting behavior will be reached. However, we still have some tools at our disposal to put our data onto an absolute scale. We can use the fact that neutrons are scattered coherently, incoherently, and (sometimes) through the magnetic cross-section, to our advantage. For example, suppose that we want to put the coherent or magnetic cross-section onto an absolute scale for a system that also scatters incoherently. We employ the following scheme.

We first measure the incoherent cross-section $(d^2\sigma/d\Omega dE')_{\text{inc}}$ for some q-value where we do not expect any coherent scattering, such as in between Bragg peaks. Incoherent scattering tends not to extend very far in energy, so that we can easily perform the following integration:

$$\int (\frac{d^2\sigma}{d\Omega dE'})_{\text{inc}} dE' = \frac{\sigma_{\text{inc}}}{4\pi}.$$

We simply look up the expected incoherent cross section per chemical formula unit σ_{inc} in a table, and we have our sought after absolute normalization factor since this factor will be the same for the incoherent, coherent and magnetic cross-section.

We could use the same trick for the magnetic cross-section at temperatures where the system behaves as a paramagnet, that is, a non-ordered collection of magnetic moments. For paramagnetic scattering on a system of N magnetic moments of spin S we have the following sum-rule:

$$(\frac{d\sigma}{d\Omega})_{\text{paramagnetic}} = N\frac{2}{3}(\gamma r_0)^2 [\frac{1}{2}gF(q)]^2 e^{-2W(q)} S(S+1).$$

In here, $(\gamma r_0)^2 = 0.291$, g is the Landé g-factor (which is 2 for most transition metals), $F(q)$ is the magnetic form factor and $e^{-2W(q)}$ is the Debye-Waller factor. For instance, when this relationship is employed at low-q values for a system where the magnetic moments are associated with transition metals we would find:

$$(\frac{d\sigma}{d\Omega})_{\text{paramagnetic}} = 0.194NS(S+1).$$

Conversely, when we have normalized our system using the incoherent cross-section, we can use this normalization to fathom the size of the magnetic moments. In the magnetically ordered phase we can use the nuclear scattering to normalize our magnetic scattering. We can do this by direct comparison of a single nuclear Bragg peak and a magnetic Bragg peak, provided we know what atoms contribute to the nuclear Bragg peak. Most of the times, the latter will be known.

In order to perform the above absolute normalization of a magnetic Brag peak, we do the following. We first measure a longitudinal and transverse scan of a nuclear Bragg peak that occurs at roughly the same scattering angle as our magnetic peak of interest. We then integrate the scattering over these two scans to get an overall count rate for the nuclear reflection, given our spectrometer setup. Of course, the count rate will be influenced by the spectrometer resolution function, such as the out-of-plane-resolution and by the size of our detector, and by the masks we employed.

Next, we measure a longitudinal and transverse scan of our magnetic Bragg peak, and we integrate to obtain the countrate. Since both the magnetic and nuclear Bragg peak will have been measured under virtually identical spectrometer resolution functions, we can take the ratio of the two peaks as an accurate measure of the relative strength of the magnetic and nuclear scattering. Since we know the strength of the nuclear peak (a prerequisite for employing this method) in terms of the scattering lengths b that contributed to this peak, we can now infer the magnetic scattering lengths from the ratio of the two peaks.

 As a last mention in this appendix, we can verify our absolute normalizations by comparing our results to that of a transmission measurement. When we measure what fraction of the neutrons are not scattered by the sample, then in principle we have the same information as we have coming from our absolute normalization efforts. We must be a little careful in measuring the transmitted fraction, and we must be sure to have corrected our scattering for attenuation effects. The best way to measure the transmitted fraction is to measure the detector countrate when positioned in the straight through beam (with perhaps an attenuator in place), then to remove the sample, and measure the count rate again. The ratio of these two measurements yields the transmission factor.

The ratio of transmitted neutrons depends on the orientation of the sample in the case of solids. Therefore, when we do a transmission experiment on a solid, we should think carefully about the orientation of our sample and choose an orientation that corresponds to some scattering, but not to a Bragg peak. In liquids of course, the orientation of the sample does not matter.

Appendix F
Alignment of a Reflectometer

In order to obtain reliable data on an instrument, we must make sure that the instrument is aligned. This appendix details the alignment procedures for a reflectometer.

F.1 Reflectometer Alignment

The principle in aligning a reflectometer is in one aspect much like any other scattering instrument and in another quite different. As with any instrument, one will always work from upstream (the source end) to downstream (the detector end). The major difference in the case of the reflectometer, however, is that- because of the necessity of shaping the beam with narrow slits the actual beam position is not fixed and it is therefore necessary in general to center the reflectometer in the beam once shaped. It is the narrow beam shape that lends the reflectometer its requisite angular accuracy at very small angles.

Before detailing a step by step procedure (7 steps), we first give a general description in order to, hopefully, add a little insight and clarity to the process. We use the relfectometer at MURR as our example. The neutron reflectometer at MURR is similar to those at other continuous source facilities. Neutrons are extracted from the reactor through an open beam tube that illuminates a monochromator with a wide spectrum of neutron wavelengths. Ours uses as pyrolytic graphite (PG) monochromator to select out a narrow wavelength beam centered at lambda=2.35 Å. This is achieved using the (002) reflection of graphite which has a d-spacing of 3.35 Å. The positioning of the monochromator is set and this procedure will begin assuming that the monochromator is in good alignment.

From here, the idea is to collimate the beam using two slits located between the monochromator and the sample. The smallest beam divergence can be obtained when these two slits are separated from each other as far as possible. For this reason, the first slit (monochromator slit, Sm) is positioned within the biological shielding, upstream from the secondary shutter. The second slit (sample slit, Ss) is positioned on the optical bench as near to the sample as possible.

The alignment begins by extracting a beam through the monochromator slit while positioning the sample slit so as to maximize beam intensity. Now, it is important that the center of rotation for the detector arm (and for the sample rotation stage which are one and the same) is positioned directly beneath the now extracted beam. In general,

this positioning will be slightly different each time there is a change upstream from the sample, though in practice the sample slit position should be fairly reproducible, e.g., when changing sample slits.

The beam onto the sample is defined at this point for the experiment, but in order to control the experimental background we use two additional slits downstream of the sample. These slits will allow the detector to accept neutrons reflected from the sample while rejecting those scattered by the air. One (guard slit, Sg) is positioned immediately after the sample and the other (detector slit, Sd) is positioned immediately in front of the detector. The purpose of the detector slit is, essentially, to define the detector size since we want this to be identical to or very slightly larger than the beam size in order to eliminate the need for detector convolution and to reduce the overall noise. The guard slit minimizes the amount of stray neutrons that scatter off of the cryostat (for example) and reach the detector. Again, this slit is selected to be just bigger than the size of the main beam.

The final piece to adjust before moving on to the sample is the borated poly (B-poly) background shield. Because the slits and shielding are imperfect at blocking neutrons, it is useful to add an additional neutron absorber around the beam as it exits the biological shielding. We have fabricated an assembly for this purpose consisting of two blocks of B-poly held together with spacers in between. The assembly is then attached to a positioner, which helps with getting the slot between the two B-poly blocks lined up with the beam. Having reached this point of the alignment procedure, one is ready to mount and align one's sample. Before doing any of this, it is important that one makes good choices pertinent to a particular experiment, i.e., choices for slit widths.

Before beginning the alignment, one needs to remove slits from any position that are intended to be changed, as well as all the slits downstream from there. We now detail the spectrometer alignment procedure step by step.

Reflectometer alignment procedure:

1. Setting the monochromator slit:
If the slit that is previously put in place at the Sm position is suitable for the experiment, then this step can be skipped; it requires some steps that we prefer to avoid having to do very often. If not, however, here is what needs to be done. The Sm slit is mounted on the upstream end of a (lead) collimator, which is housed within the biological shielding just upstream from the secondary (pneumatic) shutter. The first thing to do is to close the main beam shutter: this primary (rotary) shutter needs to be closed by the on-duty operations crew in the control room. The shutter has to be opened (by someone from operations) on a Monday, so plan ahead.

With the primary shutter closed, one will have to remove the secondary shutter as the collimators (there are two of them) do not fit through the opening in the shutter.

This means unbolting it and lifting it out. Then, one will need to pull out the first (Masonite) collimator. Be careful, though, because the monitor detector is inserted into the Masonite housing of the collimator and will not pull out directly with the cabling attached to the monitor. Therefore, one must pull the monitor out; the easiest way is to turn off the power supply to the monitor, detach the cable from the monitor detector, and pull the Masonite collimator out with the monitor detector inside.

Next, one will pull out the lead collimator. This requires three tools that are specific to the job. One is a plastic base that fits where the Masonite collimator used to be before extraction. Its function is to prevent the lead collimator from dropping off of a lip inside the shielding since this could easily damage the lead given its softness. Another tool is a long handle that bolts to the collimator; this requires the third tool which is a long handle (~ 30") rod with an Allen key head on the end. Once the handle is attached to the collimator, it can be carefully pulled out. Then, the easy part is to bolt the (appropriately chosen) slit onto the upstream end of the lead collimator. There is a unique set of slits for this slit position, whereas the other three slit positions are interchangeable. Also, note that the internal dimension of the lead collimator itself is 2mm so that any slit wider than this will not be of use unless alterations are made to the collimator itself. There is an adjustable (rotary) slit that is available and operated through the control software, but it presents problems at small slit widths because of parallax. It is, in general, best to use fixed slits whenever possible. Finally, the extraction process is reversed in order to replace the two collimators in position, keeping in mind the Monday rule for having the shutter opened by the operations group.

2. Finding the beam:
If one has not changed the Sm slit then it may not be entirely necessary to follow this step, although it is a good exercise. The detector needs to be pulled out of the detector box, together with its B4C cylindrical housing. To do this one must turn off the power on its power supply before the cable can be detached. After re-attaching the cable once its out of the detector, place the detector on the in the middle of the sample table with the detector as close to the middle of the transverse direction as possible. One will need to insert some attenuators in the beam to prevent the detector from being saturated. Now scan the detector (i.e. the sample translation; str) in order to find the center of beam. Drive the detector to this center position.

3. Positioning the sample slit:
Next, the selected sample slit is placed at the sample slit position. With the detector at the position determined in step 2, the sample slit is translated through the beam while noting where the intensity sets in and where it drops off. This procedure should be repeated, after which the sample slit is moved to the midpoint. At this midpoint, the extracted beam coincides with the position of the sample, to be inserted later.

4. Centering the reflectometer:
Next, the detector and B4C housing are placed back in the detector box, making sure

not to detach the power cable with the power supply turned on. Now, the detector arm (tth) is scanned and moved to the peak position. One needs to identify the assembly for aligning the reflectometer; it consists of a plate that mates with the sample table and it has a slit holder (and most likely already a slit in it) that is offset from the center. The plate has to be mounted on the sample table so that the slit is in line with the beam direction. A narrow slit needs to be inserted should there not be one already.

The procedure form this point is an iterative one. The idea is that if the axis of rotation of the detector arm (and the axis of rotation of the sample) coincides with the center of the beam, then one should be able to find the beam with the alignment slit at some forward facing angle and at an angle exactly 180 degrees from there. Of course, barring some lightning strike, this is not where one will start after having made alterations upstream. Hence, one will have to translate the baseplate until this condition is met. By how much? The way to figure this out is to scan theta to find the beam, to note the peak position, to rotate the sample stage (manually) by 180 degrees, and to find the beam again. By using a little geometry one should be able to predict how far and in which direction one will have to move the baseplate in order to achieve the coincidence condition of the beam and center of rotation. In practice, this will take a few iterations to get as closely as needs to be.

A word of warning: it is easy to make mistakes. When rotating the sample stage, one really must do it by hand. Because the reflectometer requires such high precision at small angles, there is a gear reducer that makes moving large distances exceedingly slow. The way to do it manually is to disengage the worm shaft from the brass gear in the goniometer by simply loosening the knurled knob on the goniometer body. Then, pressure needs to be applied to the motor body while pushing it to the right. One should now be able to feel the rotation stage free to move. Note where the red mark is and rotate the stage to that same position relative to the mark 180 degrees opposite. Now re-engage the worm shaft to the gear. Because the gear is soft and the worm shaft is hard it is critically important to do this in a way that will not damage the soft brass gear. One will have to apply gentle pressure to the motor while pushing it to the left and feeling for the worm shaft re-engaging with the gear. The rotation stage might shift slightly as they re-engage. One needs to convince oneself that these are engaged properly before re-tightening the knurled knob. It is important to keep track of where one is at. If the worm shaft and the motor or engaged in the wrong place one will just add iterations to the process. Once the coincidence condition has been achieved, move on to step 5.

5. Positioning the detector slit:

One should use a spreadsheet to calculate the predicted width of the beam both at the sample and at the detector. This will guide one in what slit size to expect and to ultimately use for the detector slit and for the guard slit. Now, we will also have to measure the beam width at the detector. To do this, a small slit is selected (~ 0.3 mm) and placed in the detector slit holder. In this way, one has limited the effect of convolving the detector size (which effectively has been made negligibly small) with

the beam width (which will generally be significantly larger).

Next, the detector arm is scanned and moved to the peak position. From this, the beam width is determined. The next step is to replace the small (~ 0.3 mm) slit with one that is just wider than the beam width that has just been measured. The last step (of step 5) is to make sure that this slit is well centered in the beam. In order to do this one will have to translate it and note the rise and drop positions just as how it was done with the sample slit; following this, the slit needs to be positioned at the midpoint. If the beam is very wide, one may need to create one's own slit out of pieces of cadmium (this is also true for the guard slit, or the sample slit for that matter).

6. Positioning the guard slit:
One can, in principle, use another fixed-width slit at the guard slit position but it is often useful to use the barndoor slit here, because it is more effective at its job of reducing stray scatter that leads to higher background counts. One also has continuous control over the height and the width. This procedure is very similar to the positioning of the detector slit and the sample slit only now one also has the extra degree of freedom of adjusting the width. One really should not need to do anything with the height, in general.

7. Positioning the B-poly shield:
The last step before getting to actually mount one's sample is to mount and position the borated poly shielding block. The background levels particularly in the low to medium q-range require this shielding in order to achieve an acceptable dynamic range. The B-poly assembly simply bolts to the linear translation stage located above the secondary shutter opening. Through some trial and error, one will need to rotate and translate the assembly in such a way that the beam passes through the spacing between the two B-poly blocks, preferably with the west-most block (this should be the larger block and it should be on the left when facing the beam opening) not quite cutting the beam. One should pay attention to the detector rate before mounting the B-poly; one should get all of the counts back unless the spacers are too small, in which case one will have to adjust them.

F.2   Sample Alignment

Aligning the sample is pretty straightforward compared to aligning the instrument. The procedure outlined below is assuming a simple open air experiment. The essential features remain the same but, of course, mounting the sample will certainly vary depending on what sample environment, e.g. cryostat or humidity chamber, one might need for the experiment.

Sample alignment procedure:

1. Zeroing the axes:
In order to give oneself a good starting point, it pays dividend to line up the sample rotation (th) and detector arm (tth) either at their last known zero positions (i.e. if

one just completed the reflectometer alignment, these should already be known), or by eye, or to the hardware zero. Next, the detector arm is scanned through the beam and set to the position of the center of the peak. This center is then encoded to be the new tth zero ("set th 0").

2. Mounting the sample:

There are two different options for mounting the sample. Select the one that allows usage of a clamp, if possible. Use the smaller clamp if possible. If it is not possible to clamp, then one can use the very smallest amount of vacuum grease that will hold the sample onto the mount. Next, the sample and holder are placed on the goniometer stage and bolted in place.

3. Locating the surface of the sample:

The sample is scanned through the beam with the detector in the straight through position (i.e. the detector is at the zero position that was set in step 1). One will have to move pretty far, such as ± 10 mm in the sample translation (str). Both sample holders have a sheet of cadmium mounted to the back of them to help identify features in this profile as it will be the best at attenuating the beam. One may be able to identify steps in the translation profile that correspond to the aluminum holder and the aluminum backing plate (with the cadmium as the meat of the sandwich). If the starting theta is too far off, then these features may all be too rounded to identify without ambiguity. Once the best assessment of the sample surface position has been made, one can move to this position.

4. Cutting the beam:

Now that the surface of the sample has been identified, we want to get the surface as parallel to the beam direction as can be. Because this process affects where exactly the best str position is, we will ultimately once again use an iterative process to hone in on the best str, th, and ch. In this step, one will rock the sample theta about the beam. One may have to play around a bit with the scan range that one is using. Once this rock has been completed, one rotates the sample to the apex position.

5. Reflection peak:

To further hone the theta positioning, we move to a 2theta-theta position away from the main beam. There is a tradeoff here. Going to higher 2theta-theta gives one better accuracy, but the counting times quickly become unacceptably long. A 2theta angle on the order of 0.5-1^o is a typical choice. Next, theta is rocked about half of the 2theta position. At the end of this rock, one moves the motor to the center position. This position is encoded to be exactly at half the 2theta position while noting this adjustment in one's log book. This scan and the tilt scan in the next step will require some longer counting times than the main beam scans for obvious reasons (fewer counts).

6. Sample tilt:

Next the tilt angle (ch) is scanned and the tilt angle is set to the center position and encoded as such. It is not at all uncommon that the tilt direction is very broad but

in some cases it can be important. It is always good practice to check it at the very least.

7. Rinse and repeat:

Well, okay, not rinse, but repeat steps 3 through 6 until one is satisfied that all axes are reproducible within the instrumental resolution. Once they are, then one is ready to perform the experiment.

Appendix G
Alignment of a 3-Axis Spectrometer

In order to obtain reliable data on an instrument, we must make sure that the instrument is aligned. This appendix details the alignment procedures for a 3-axis spectrometer. The alignment of a 2-axis spectrometer is the same as that of a 3-axis one, with some steps missing.

G.1 3-Axis Spectrometer Alignment

The following are the basic steps for aligning a 3-axis spectrometer. The most important thing to remember when aligning a spectrometer is to never trust the alignment of the previous user.

1 Monochromator crystal angle:
 First, drive the monochromator to the energy you want to align at using the software controls, then put the desired monochromator-sample collimation in place. Next, rock the monochromator angle θ_m. For this, use the monitor as the signal (this might require physically switching cables on some spectrometers and plugging the monitor into the detector socket) and count for a fixed number of seconds. Because the monitor spans a fairly large area, it is often a good idea to put a small mask in front of the monitor, centered on where the beam center should be. If possible, move the monitor out towards the sample for better definition (especially when using a focusing monochromator). Record the center of mass and peak position of the rocking curve. Drive θ_m to the center of mass. Replace any switched wires, and switch the way counts are collected to a fixed number of count on the monitor for the rest of the alignment.

2 Offsets in $2\theta_m$ and scattering angle θ, part 1:
To determine the zero offsets in $2\theta_m$ and scattering angle θ (referred to as $2\theta_{m,0}$ and θ_0) a powder pattern has to be measured. For this, the analyzer has to be removed, and $2\theta_a$ should be driven to $2\theta_{a,0}$. If the analyzer arm zero-offset angle $2\theta_{m,0}$ is not known, drive $2\theta_a$ to the approximate zero angle position 'by eye'. Remove the collimator between the analyzer and the detector. Next, place an attenuator in the beam to avoid saturating the detector, put a tight collimator in between the sample position and analyzer position, and rock θ through the straight through beam. Drive θ to the so-determined straight through beam position, and place a fairly tight collimator between the analyzer position and detector. Now rock $2\theta_a$. Drive $2\theta_a$ to the center of mass, and replace the sample-analyzer and analyzer-detector collimators with the

Fig. G.1 A standard powder is a powder for which the lattice spacings are known to high accuracy, such as the silicon powder shown here. Standard powders tend to be small in diameter and tall. This reflects that we care about the angular divergence of the beam in the horizontal plane, but that the vertical divergence is not very important at all.

ones to be used in the experiment.

3 Offsets in $2\theta_m$ and scattering angle θ, part 2:

Remove the attenuator and place a standard powder on the sample table (Fig. G.1). If the spectrometer does not have a ready-to-go callibration routine such as available in SPICE, then do the following steps. Do a quick calculation where (roughly) in θ the powder peaks are expected, and measure about 4 such peaks. In principle it can be done with 2 peaks, but the accuracy of the alignment increases with increased number of powder reflections. Set up a set of θ-scans to cover the expected powder peaks and go for coffee. Determine the peak positions for each reflection and use a least squares method to determine θ_0 and $2\theta_{m,0}$ (see below on one way on how to do this). Encode $2\theta_{m,0}$ and θ_0 into the software and calculate the incoming neutron wavelength. Repeat one scan on a powder peak to ensure that the zeroes were put in with the correct signs.

4 Offset in θ_a:

Replace the powder by an incoherent scatterer, such as a vanadium rod shown in Fig. G.2. Some people use a nice big single crystal Bragg peak instead of vanadium, but this leads to results that are slightly off. However, it is not always possible to use vanadium, in which case one can use the incoherent scattering by the sample (if any), or the incoherent scattering of a Ni-powder or an ice cube. If all this fails, then a Bragg peak provides the next best thing. Put the analyzer in place. Drive θ to a fairly large angle (to avoid the increased background at small angles), bearing the W-configuration in mind for a focused setup. Drive $2\theta_a$ to the angle corresponding to the incoming neutron wavelength [$\lambda_{in} = 2d_{\mathrm{analyzer}}\sin(2\theta_a/2)$], and rock θ_a to find the center of mass. This determines $\theta_{a,0}$. Drive θ_a to the center of mass and encode $\theta_{a,0}$.

Fig. G.2 A thin vanadium rod is placed at the center of the sample table to measure the energy resolution. Since vanadium scatters equal amounts in all directions, the mounting of the vanadium can be done by eye in an ad hoc manner as amply demonstrated in the mounting used here.

5 Offset in $2\theta_a$:

For the next step, make sure that $2\theta_a$ and θ_a are operating in half angle mode (i.e., if θ_a moves by five degrees, $2\theta_a$ should move by ten). Do a θ_a-$2\theta_a$ scan to find the maximum (center of mass) of the vanadium scattering. This tends to be a rather long scan. Once the scan is completed, drive to the center of mass. This point corresponds to $\lambda_{in} = 2d_{anal}\sin([2\theta_a - 2\theta_{a,0}]/2)$. Encode $2\theta_{a,0}$.

6 Energy resolution function:

It might be useful to measure the energy resolution for the fixed initial or final wavelength to be used in the experiment. Use the vanadium sample for this (Fig. G.2), and stick the filters on the spectrometer (if using any). Do an energy scan around the elastic channel, either in fixed initial or fixed final energy mode.

7 Make notes of the physical scales:

Most triple axis spectrometers have scales on the motors. If so, read off the scale values corresponding to the values encoded in the computer. By checking the offsets between the scale values and the computer values about once a day, it is possible to spot a misalignment quickly. Misalignments during the course of an experiment occur quite frequently (at least more frequently than wanted). A motor may have lost some steps, an arm may have run into a dewar etc. Stuff happens.

The following is a recipe for determining $2\theta_{m,0}$ and θ_0 from the powder peaks. The quickest way is to do a linear least squares fit. For this, a guess for θ_0 has to be made, but this can be obtained from the rock of θ through the straight through beam. If
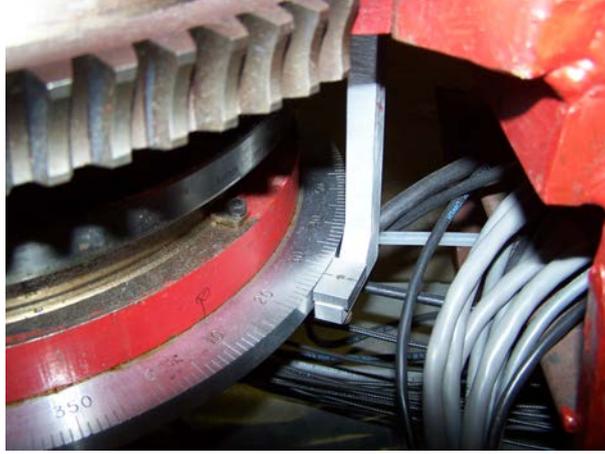
Fig. G.3 Shown is a close up of a hardware scale on a triple axis spectrometer. To utilize the hardware scale readings, simply keep track of the difference between the scale reading and the angle that the software claims the motor is at. This difference should remain constant throughout the experiment, independent of the actual value of the angle. If the difference changes outside of the accuracy of the hardware scale, then a problem has occurred.

this value is larger than 5 degrees, correct the observed peak positions for this value. We label the peak positions, accurate to within 5 degrees, as θ_i^P where the index i runs over all the peaks obtained in the standard powder measurement (typically about 5 peaks are measured). The remainder of the correction to the estimated θ_0 is now small enough to do a linear least squares fit procedure. The equation that describes the observed peak positions reads

$$\lambda_{true} = \lambda + d\lambda = 2d_i \sin([\theta_i^P - \theta_0]/2) = 2d_i \sin(\theta_i^P/2) - 2d_i \cos(\theta_i^P/2) \sin(\theta_0/2). \quad \text{(G.1)}$$

In here, d_i is the lattice spacing of the i^{th} powder reflection. Note that in the above equation we have made use of the fact that $\cos\theta_0=1$ given that our initial guess of ϕ_0 is accurate to within a few degrees. This can be rewritten in matrix form for N peaks as:

$$\begin{pmatrix} 1 & 2d_1\cos(\theta_1^P/2) \\ 1 & 2d_2\cos(\theta_2^P/2) \\ ... & ... \\ 1 & 2d_N\cos(\theta_N^P/2) \end{pmatrix} \begin{pmatrix} d\lambda \\ \sin(\theta_0/2) \end{pmatrix} = \begin{pmatrix} 2d_1\sin(\theta_1^P/2) - \lambda \\ 2d_2\sin(\theta_2^P/2) - \lambda \\ ... \\ 2d_N\sin(\theta_N^P/2) - \lambda \end{pmatrix} \quad \text{(G.2)}$$

This equation is of the form $Ax = u$, with A the 2x N matrix, x the column vector with 2 elements we are trying to solve for, and u the N-element column vector of measured outcomes. This equation can be solved by using the transpose matrix A^T as follows:

$$A^T.Ax = A^T.u; x = [A^T.A]^{-1}A^T u. \quad \text{(G.3)}$$

Now that x is known, one can directly determine θ_0; $2\theta_{m,0}$ follows from $\lambda_{true} = \lambda + d\lambda = 2d_{mono}\sin([2\theta_m - 2\theta_{m,0}]/2)$.

G.2   Sample Alignment

When doing inelastic scattering experiments on a single crystal, the crystal needs to be aligned. This means that we need to know the orientation of the crystallographic axes with respect to the scattering plane, and we need to orient our crystal in such a way that the desired axes are in the scattering plane. This can sometimes be problematic since, unlike the beryl crystal shown in Fig. G.4, not all crystals have facets that reveal their crystal symmetry.

Fig. G.4 The unit cell symmetry is easily visible in this 6" chunck of beryl on display at the Denver Museum of Nature & Science.

In order to align a single crystal that has unhelpful facets, we need to know the length of the crystallographic axes and the symmetry of the unit cell. This information will invariably be available to the user because powder experiments will have been performed before someone decides to grow a single crystal. One very helpful piece of equipment that greatly facilitates the alignment of a crystal is a neutron camera with a computer display. Such a camera can directly display neutrons hitting the camera, with the display not only showing whether a neutron has hit the camera, but also whether a great many neutrons hit it (as in a Bragg peak), or just very few.

The advantage of neutron cameras is that the span a very large solid angle, so that any Bragg peak anywhere within a large angular range will be registered. Typically, one positions the camera close the the sample, and then plays around with the orientations of the sample until a peak appears on the camera. Then the orientation of the sample is fine tuned, and the camera pulled back from the sample, until we have a Bragg peak that is roughly within the scattering plane. Next, we remove the camera, and drive our standard neutron detector to the position where the camera was, and

we fine tune the scattering angle and the sample orientation.

Once we have done this, knowing that the orientation of this particular Bragg peak is not perfect yet, we will have to adjust the tilt under the sample to make sure that the incident and the scattered beam lay within the plane scattering plane of the spectrometer. We would like the Bragg planes to coincide with one of the tilt axes of the sample table, so we will (likely) have to physically rotate our sample with respect to the surface of the sample table. To do so, turn off the beam, and look carefully at the sample. The Bragg planes that reflect the neutrons in a mirror-like geometry are oriented such that the neutrons are reflected as if bouncing off a mirror. We need to pick up the sample and reposition it on the sample table such that these planes are parallel with one of the tilt axes of the table. So do this, pick up the sample, roughly guess by how many degrees the sample was reoriented when put back on the table, and go back to the controls, open the beam, and rotate the sample the opposite direction and wait for the peak to reappear.

Again, do an alignment scan of the scattering angle and the sample rotation angle, and drive the motors to the point of maximum scattered intensity. Then adjust the tilt angle under the Bragg planes (think about which of the two tilt angles to change) until maximum intensity has been found. Then repeat the scattering angle and sample rotation alignment. When all this is done, we have found one Bragg peak and the sample is oriented in such a way that the neutrons are deflected into the detector, while sticking to the scattering plane. From the scattering angle we should be able to determine exactly what Bragg peak we are looking at, such as the (200)-reflection.

What if we do not have a neutron camera at our disposal? In this case, we will have to use the standard neutron detector instead. To make our setup as receptive as possible, remove all collimation before and after the sample, remove the analyzer and position the detector in a 2-axis configuration. Set the scattering angle so that it corresponds to a strong Bragg peak. Once all this is in place, then change the orientation of the sample (only 4π angles to cover) until a peak appears in the detector. Typically one rotates the sample angle over 180 degrees, then- if no peak appears- one repositions the sample with a change in tilt of 5-10 degrees, and rotates the sample angle over another 180 degrees until a signal appears in the detector. Once a peak has been found this way, then one aligns on this peak in exactly the same way as the procedure we employed when we used a neutron camera.

Interestingly enough, the first peak is the easiest to find, even for a shapeless crystal. Next, we have to find a second, independent Bragg peak so that we know the orientation of the crystallographic axes. Select a peak of choice, set the scattering angle to correspond to this Bragg peak, and rotate the sample angle and keep one's fingers crossed. Likely, no peak will appear. When this happens, one will have to change the tilt under the sample. Do this by changing the 'other' tilt angle by 5 degrees, and start rotating the sample again. Of course, if one has a neutron camera, then first try to find this second Bragg peak on the neutron camera by placing this camera between

the detector and the sample.

If one is lucky, then the second Bragg peak will appear within the range of tilt angles that the spectrometer allows. If not, then one will have to change the mounting of the crystal, and change the orientation as best as one can in order to induce a larger tilt angle. After this, put the sample back on the sample table, realign on the first Bragg peak, and then repeat the procedure hoping to find the second Bragg peak. All this can be a time consuming and a somewhat frustrating experience, but there is no other way around it. Eventually, the second Bragg peak will be found. It is a particularly good exercise for developing one's inner strength trying to orient a shapeless sample that has monoclinic or triclinic symmetry.

Now that two Bragg peaks have been found, one should try to position the sample on the sample table in such a way that the tilts under the two Bragg peaks are minimal, say less than 5 degrees. Once this is done, then put all the collimation back in and find the points of maximum intensity for both Bragg peaks, and tell the computer what scattering angles, sample rotation angles and tilt angles these two peaks correspond to. Once this has been punched into the computer, then the sample is oriented.

We should check on the orientation by instructing the computer to drive to another Bragg peak, not one of the two we used for alignment, and make sure that our detector actually detects a peak where it is supposed to find the peak. If the expected peak is not where it is supposed to be, then you have our sympathy.

Appendix H
Higher Order Contamination

When contamination from the monochromator onto the monitor cannot be avoided by either using a monochromator reflection with forbidden second order contamination, or by employing a velocity selector, then one should simply to measure the contamination. By placing well characterized absorbers of varying thickness in front of the monitor, one can determine the contamination. We detail this procedure for a neutron scattering experiment.

In a neutron scattering setup for a given monochromator reflection, one measures the monitor count rate $R(\lambda_i)$ as a function of incident energy E_i, corresponding to an incident wave length λ_i (see eqn A.1 in Appendix A). Because of higher order neutrons also reaching the monitor, this count rate $R(\lambda_i)$ is given by

$$R(\lambda_i) = \sum_{j=1}^{N} \phi(\lambda_i/j)(1 - e^{-C\lambda_i/j}) = C \sum_{j=1}^{N} \phi(\lambda_i/j)(\lambda_i/j). \tag{H.1}$$

In here, C is a measure of the efficiency of the monitor, which is a small number by design. $\phi(\lambda)$ is the neutron flux for neutrons of wave length λ that will make it to the monitor. The constant C will be eliminated since only the ratio $\phi(\lambda_i)/\sum_{j=1}^{N} \phi(\lambda_i/j)$ is needed to perform the monitor correction.

Next one places a thin piece of neutron absorbing material before the monitor, such as a tin foil of known thickness, and one repeats the measurement of the monitor count rate. Subsequently, one doubles the thickness of the foil, and so on, until one has at least as many foil measurements (m) as there are higher orders of contamination present in the incident beam ($m = N$). For m layers of foil of thickness d, the monitor count rate is given by

$$R(\lambda_i, m) = C \sum_{j=1}^{N} \phi(\lambda_i/j)(\lambda_i/j)e^{-mb/j} \tag{H.2}$$

with the foil parameter b is given by the number density of the foil n_f, the absorbtion cross-section of an atom in the foil σ_{abs} at neutron energy E_i, and its thickness: $b = dn_f\sigma_{abs}(E_i)$.[1]

---

[1] The absorbtion cross-section depends inversely on the wave length of the neutron, or equivalently, the cross-section is proportional to $1/v$. An almost correct way to think about this is that the longer

All measurements can be written in matrix form as $M\phi(\lambda_i/j) = R(\lambda_i, m)/C$, with the matrix elements given by

$$M_{j,m} = (\lambda_i/j)e^{-mb/j}. \tag{H.3}$$

Next one finds the sought-after ratios for each λ_i by applying the least squares formalism:

$$\phi(\lambda_i/j)/\phi(\lambda_i) = (M^T.M)^{-1}M^T[R(\lambda_i, m)/R(\lambda_i, 0)], \tag{H.4}$$

with M^T the transpose of the matrix M. The results of such a set of foil measurements are shown in Fig. 9.11. A final word of warning: the foils will get highly activated because of the neutron capture events.

a neutron spends in the nucleus of the scattering atom, the more chance it has of being absorbed. Thus, when we halve the wave length of a neutron, we also halve the absorbtion cross-section. This process is behind the $\sim 1/j$ dependence in the exponents of eqns H.1 and H.2.

Appendix I
Attenuation Corrections

We use computers to carry out attenuation corrections numerically. The principle is easy, the implementation is a little more difficult, especially for unusual geometries such as a hollow annulus. We recommend Varley Sears' original publication, such as V.F. Sears, Adv. Phys. 24, 1-45 (1975).

The fundamental quantity for calculating the attenuation factors is referred to as $H_1(\vec{k}_{\text{initial}}, \vec{k}_{\text{final}})$ by Sears, and as the transmission factor by xray scattering people. This transmission factor can be cast in a form suitable for numeric computation:

$$H_1(\vec{k}_{\text{initial}}, \vec{k}_{\text{final}}) = \frac{1}{V} \int_V d\vec{r} e^{-n\sigma(k_{\text{initial}})L(\vec{r}, \hat{k}_{\text{initial}}) - n\sigma(k_{\text{final}})L(\vec{r}, \hat{k}_{\text{final}})}. \qquad \text{(I.1)}$$

This equation is actually very straightforward. It states that scattering can originate from any point within the sample, and hence, we have to integrate over all points within the sample. The paths that the neutron will have to traverse depend on the point of scattering, and on the initial $\hat{k}_{\text{initial}}$ and final directions $\hat{k}_{\text{final}}$. These paths are weighted by the chance of scattering and absorbtion, as contained in the product $n\sigma$, with n the number density and σ the sum of the scattering cross-section, and the energy-dependent absorbtion cross-section.

When we correct for sample attenuation, we divide the measured counts by the calculated factor $H_1(\vec{k}_{\text{initial}}, \vec{k}_{\text{final}})$ for every scattering angle and energy transfer. When we wish to subtract an empty container contribution from our scattering data, we multiply the empty container spectra by this factor so that we can pretend it was filled with a sample.

Calculating the attenuation factor H_1 is simply a matter of geometry. For each point within the sample, and for each scattering direction, we have to calculate the length of the path from the point of entry to the point of scattering, and we have to calculate the length of the path from the scattering point to the point of exit. If the sample absorbs, then we also have to remember to make sure we put the proper energy dependent absorbtion cross-sections in. Then we integrate over all possible points. Then we are done.
 This correction can be calculated exactly for a slab geometry (that is, it can be calculated analytically), but it has to be calculated by computer for other geometries. For single scattering (which is what the '1' in H_1 signifies), we only have to calculate the scattering in a cross-section of the sample, we do not have to integrate over the sample
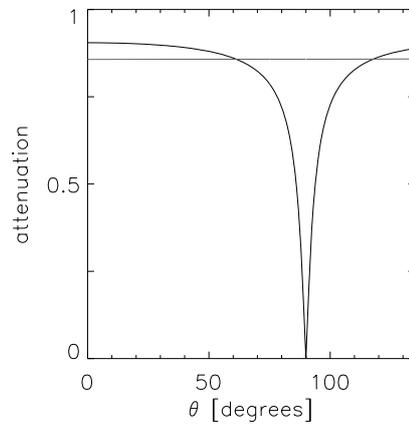
Fig. I.1 Shown is the calculated attenuation for a sample that does not absorb, but that scatters the neutrons with a scattering length of 0.2 cm^{-1} ($n\sigma = 0.2$ cm^{-1}). When this sample is put in a slab geometry with 0.5 cm thickness, and when the slab is oriented perpendicular to the beam, then the attenuation factor will drop from 0.9 at small scattering angles to 0 at 90 degrees (black line). When the same sample is put in an annular geometry with inner radius 1cm and outer radius 1.25 cm, then the attenuation factor (horizontal line) becomes angle independent.

dimension perpendicular to the scattering plane. On the other hand, when we calculate multiple scattering, we do have to take this third dimension into account. Note that multiple scattering calculations are somewhat similar. We would still average over all possible scattering points, and perhaps even force more than one scattering event to take place in the sample, and we would still weigh all scattering events by the attenuation based on the path lengths. The only real difference is that in multiple scattering we would also weigh all paths by the likelihood of scattering from, e.g., $\vec{k}_{\text{initial}} \rightarrow \vec{k}_{\text{final}}$, whereas in calculating the attenuation correction we weigh all events equally.

We give some examples for slab geometry and cylindrical geometry. Suppose we have a sample that does not absorb, but merely scatters. We can place this sample inside a slab, or we can put it inside of a cylinder, or we can even put it in the beam as an annulus (a hollow cylinder). In Fig. I.1 we show the calculated attenuation for a sample that scatters roughly 10% of the incoming neutrons, both for a slab and for a cylindrical (annulus) geometry. This figure demonstrates that a cylindrical geometry would be preferential in this case, unless we are only interested in very small scattering angles. It also demonstrates that if we were to use an empty container measurement as our background data, that we would over subtract. For this case, we should first multiply the measured container data by the calculated attenuation factor. Overall, this figure demonstrates that when our sample is not absorbing, we should have no problems in correcting our data.

Next, we look at the case for a sample that not only scatters, but also absorbs
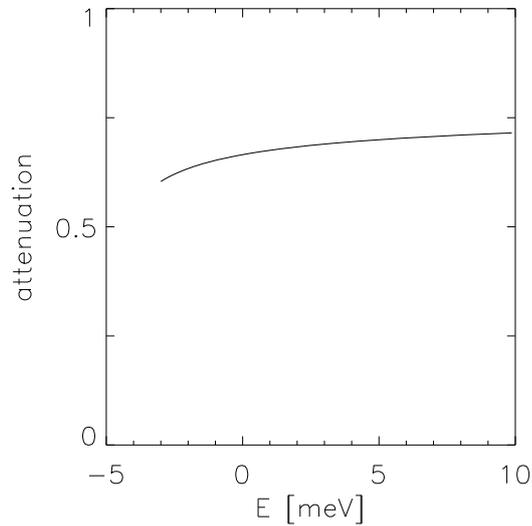
Fig. I.2 Shown is the calculated attenuation for a sample that absorbs and scatters the neutrons. For this example we have set the final neutron energy at 5 meV while varying the incoming neutron energy. The curve has been calculated for a scattering angle of 25 degrees when the sample is in a slab geometry of 0.5 cm thickness. The scattering length was set at 0.1 cm^{-1}, and the absorbtion length at 0.3 cm^{-1} for 25 meV neutrons. Note the increased attenuation caused by the absorbtion of neutrons for neutrons that enter the sample with low energies. Also note that this sample does not scatter a lot, so even though the attenuation factor is always significant, we do not expect insurmountable problems when it comes to calculating multiple scattering corrections.

the neutrons. We show the calculated attenuation as a function of energy transfer for one particular scattering angle in Fig. I.2. It should be clear from this figure that we absolutely will have to correct our scattering for the attenuation factor if we are to have any hope of obtaining line shapes out of our scattering experiment. Given how straightforward it is to punch eqn I.1 into a computer, there is no reason to not carry out this attenuation correction. The correction is just as easy to calculate for time of flight instruments as it is for reactor source instruments. The attenuation correction has to be calculated for every scattering angle, and for every energy transfer, but that is what computers are good at.

Appendix J
Solutions to Selected Exercises

Exercise 2.2

 The maximum amount of momentum that can be transferred is when the neutron bounces back, or scattering over a 180 degree angle. The amount of momentum transferred is twice the momentum of the neutron: $\hbar q = \hbar k_{\text{initial}}$.

Exercise 2.3

When there is music playing in one room, and the door is open, it is easy to hear the music everywhere in an adjacent room. The sound waves diffract through the door opening, which makes sense since a typical wave length of a sound wave is 1 m, comparable to the width of the door. Examples for light waves are the diffuse scattering that occurs when light goes through a small opening, such as a pinhole in a piece of paper, or the bright diffraction spots when light is incident on a diffraction grating.

Exercise 2.4

 The equivalent would be inelastic scattering. For instance, when we wish to excite a sound wave in a crystal, then not only has the period of our interference pattern match the wave length of the sound wave, the oscillation in the amplitude of the interference pattern has to match the natural oscillation given by the frequency of the sound wave. Thus, the scattered wave has to have a slightly different wave length (frequency) than the incoming wave so that we can set up this beat frequency.

Exercise 3.2

 a) see table.

 b) Water and heavy water have the same density, so we can work with the scattering lengths per molecule rather than the scattering length density. The scattering length for a H_2O molecule is $b_{H_2O} = 2B_H + b_O$ = 2 x (-0.3739) + 0.5803 = -0.1675 10^{-12}cm. That od a heavy water molecule H_2O equals 1.9143 10^{-12}cm. In order to get a mixture with a fraction x of normal water with zero scattering lengths, we have that 0 = -0.1675 x + 1.9143 (1-x), and hence, x= 0.9195.

 c) When this mixture is probed on long wave lengths, it would not scatter since all the scattered waves would add up with zero total amplitude. This is very useful when doing SANS experiments. However, when the mixture is probed on much shorter wave lengths, such as those comparable to the atomic spacings within a molecule, we would still see scattering since at these short probing wavelengths (the incoherent limit) we would see individual molecules.

 d) The relevant measure is the overall scattering power, which necessarily must be a

Table J.1 Scattering lengths and cross-sections for biological relevant elements. Note that 1 barn= 10^{-24}cm^2. Source: The SANS toolbox.

Element	b	σ_{coh}	σ_{inc}
	10^{-12}cm	10^{-24}cm^2	10^{-24}cm^2
H	-0.3739	1.757	80.30
D	0.6671	5.592	2.05
C	0.9946	5.550	0.001
N	0.936	11.01	0.50
O	0.5803	4.232	0
F	0.5654	4.232	0.001
Na	0.363	1.66	1.62
P	0.513	3.307	0.005
S	0.2847	1.017	0.007
Cl	0.9577	11.526	5.3

product of how well the individual units (molecules) scatter, times how many molecules there are in the sample volume.

e) This can be advantageous in two ways. In some experiments we might not be interested in what the proteins are doing, so it is nice if they do not contribute to the scattering (when using long probing wave lengths). Alternatively, we might be interested in what the proteins are doing, but the scattering attributable to the proteins might be hidden in the total scattering of all the components in the system. We can do two experiments with different levels of deuteration, and we can use one experiment as our background measurement. So we would have one measurements where the scattering contains the scattering by the (proteins + water + other stuff) and one experiment that would effectively be the scattering by (water + other stuff). Subtract the two, and you get the scattering by the proteins.

Exercise 4.1

Hydrogen has a very large incoherent cross-section, and therefore, any material that contains water tends to scatterer a lot. Therefore, we would like to use very thin samples such as those contained in a slab geometry. This way, we can use a slab that is as big as the dimensions of the neutron beam, but we do not have to worry about multiple scattering too much because of the limited thickness of the slab.

Exercise 4.2

Since the molecules will have random orientations, we should average over all possible

directions of $\vec{q}$. A typical term will involve $< e^{iqa\cos\theta} >$ where the brackets denote averaging over all 4π directions. Carrying out this averaging we find:

$$
\begin{aligned}
< e^{iqa\cos\theta} > &= (1/4\pi) \int_0^\pi \sin\theta d\theta \int_0^{2\pi} d\phi e^{iqa\cos\theta} \\
&= (1/2) \int_{-1}^1 d\cos\theta e^{iqa\cos\theta} \\
&= (1/2) \frac{e^{iqa} - e^{-iqa}}{iqa} \\
&= \frac{\sin(qa)}{qa}.
\end{aligned}
$$

Adding up all 4 terms in the summation yields $2 + 2\sin(qa)/(qa)$.

Exercise 4.4
The first panel for $q = 0.05$ Å^{-1} displays a peak located roughly at 0.2 ps^{-1}. Thus the slope of the dispersion curve is given by 0.2 ps^{-1}/0.05 Å^{-1}= 4 Å/ps = 400 m/s.

Exercise 4.5
A typical bilayer separation is about 40 Å. Doing neutron scattering on randomly oriented bilayers we can expect to see this separation as an oscillation in the static structure factor, very similar to that seen in oxygen and nitrogen molecules. The only difference would be that the characteristic distance (40 Å) is much larger than in those molecules (2 Å), and therefore, the period of oscillation will be much smaller, roughly by a factor of 40Å/2Å= 20. Thus, we can expect a similar pattern as shown in Fig. 4.6, but with the horizontal axis shrunk by a factor of 20.

Exercise 4.6
We would have to do an inelastic experiment at high momentum transfers. We would choose high momentum transfers since then we are sensitive to individual atoms and molecules (the so-called incoherent limit), and we would have to do an inelastic experiment since sometimes we will absorb some of the kinetic energy of the atom in its oscillation, and sometimes we would give the atom some energy. So we expect to see a broad (in energy) feature when we do these inelastic experiments, and the width in energy as well as the average position of the peak would give us information about the characteristic energy scales that determine the motion of the atom or molecule.

Exercise 5.2
Very large lattice spacings need to be probed with large probing wave lengths, which can most easily be done using neutrons of fairly long wave length. At the same time, we need to reduce higher-order contamination from reaching our sample, and even creating unwanted background counts in our sample through multiple scattering. The best way is to clean the incident beam from higher-order neutrons, which can be done by placing a Be or a BeO filter in the incident beam, preferably before the incident beam monitor (why?)

Exercise 5.3

We use the equation

$$\frac{\Delta\lambda}{\lambda} = \frac{\Delta d}{d} + \frac{\Delta\theta}{2\tan(\theta/2)},$$

with $\Delta\lambda = 0$ for a reactor based diffractometer. Detectors are typically located at $\theta=$ 90 degrees, so by plugging in these numbers into the above equation we find $\Delta\theta = 0.002$ radians, or about 0.1^o. Thus we require a very tight angular resolution. Since the monochromator crystal is perfect, the angular resolution in controlled entirely by the acceptance angles of the collimators. So we would like to use collimators of about 0.1^o. Likely, the best we can find is 10' collimators, so we will have to make do with those. We should still be able to see the peak shift using this set of collimators, but we would have to measure the entire peak in order to get a reliable estimate of the average peak position.

Exercise 6.1

a) In order to figure out the symmetry of the unit cell, we need to figure out the lattice spacing or q-value that corresponds to each peak. To do so, we have to read off each peak position, and convert to lattice spacing or q-value using the relation $\lambda = 2d\sin(\theta/2) = (4\pi/q)\sin(\theta/2)$. Once we have the q-values, we can divide all the q-values by that of the first peak, and square the ratio. Any peak that is related to the first peak will have an integer ratio, as explained in the main text. If there are more than one family, then we would get non-integer ratios. When we do it for the given pattern in this problem, we find integer ratios for all peaks, and hence, the system must be cubic.

b) Now that it is known that the system is cubic, we can determine d from any of the reflections using $\lambda = 2d\sin(\theta/2)$.

c) There is more than one atom. This follows from the oscillation in peak intensity. If there is only one atom in the unit cell, then all peaks should be equally strong, and we would only see a gradual overall decrease in intensity because of the Debye-Waller factor.

d) The atoms must be of different species. If the atoms were of the same species, then we would observe some forbidden reflections such as the (100)-reflection. Since all reflections are present, we never get complete cancelation of say, the atoms located in the center of the cube and the atoms located on the sides of the cube. This can only be if the atoms are different, scattering the neutron with a different strength (scattering length).

e) This could be caused by background scattering, scattering by the sample holder, or by scattering through the incoherent cross-sections of the atoms that make up the system.

f) When the probing wave lengths get shorter, more and more of the constructive interference pattern will get lost because the atoms have wandered ever so slightly from their equilibrium positions because of their thermal and zero-point motion.

g) In order to do this, we must identify two reflections that are formally equivalent in intensity, such as the (400) and (800)-reflections. The difference in intensity between those two reflections can be attributed to the Debye-Waller factor. We can express the

ratio of these two reflections as $e^{-(q_{800}^2 - q_{400}^2)<u^2>}$ from which we can calculate $<u^2>$.

Exercise 6.4

 Grains can have facets that reflect the underling symmetry of the unit cell. When we tap the holder repeatedly, facets of neighboring grains can align, and as a consequence, we lose the random orientation of the grains. As a result, we will find that our powder profile suffers from preferred orientation.

Exercise 6.5

It is cubic. All visible edges are of equal length, only cubic symmetry can accomplish this. However, as the model structure shows, figuring out the exact space group was much harder than figuring out the overall symmetry.
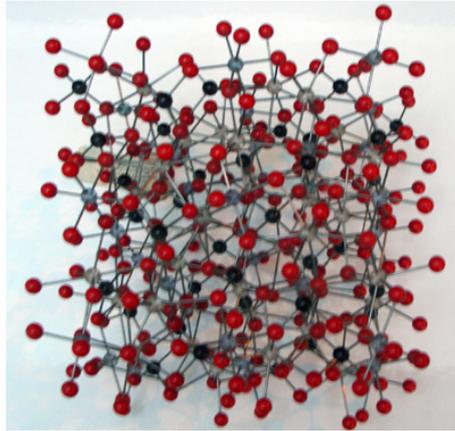
Fig. J.1 The structure of garnet might be cubic, but it is still complex. Source: Wikimedia Commons, author Peter Murray-Rust.

Exercise 7.1

A monochromator reflects neutrons onto the sample. If we were to choose a different incident wave length, then we would have to rotate the entire spectrometer along with the change in monochromator take-off angle.

Exercise 7.2

We are only measuring at very small momentum transfers, and hence, we need a material with a very large lattice spacing. Typical samples that are being used on SANS spectrometers are Silver Behenate ($d= 58.38$ Å) and kangaroo tail tendon ($d= 667$ Å).

Exercise 7.3

 Instrumental resolution functions smear out peaks, including the peak at $q= 0$. Since the radius of gyration is related to the slope of the tail of this peak, a broadened peak

implies a lower slope and therefore, a lower radius of gyration.

Exercise 7.4

Yes. SANS detectors are flat, so what we measure is not the differential cross-section, but rather the projection of this cross-section onto a flat surface. For instance, for a disc-shaped element $d\Omega$, we would measure an ellipsoidal shaped element in the detector.

Exercise 7.5

In order to obtain a complete set of SANS measurements the required runs are: transmission and scattering from the sample, transmission and scattering from the empty cell, scattering run from the blocked beam, transmission run from the empty beam and the detector sensitivity run from plexiglass (SANS toolbox, Chapter 21).

Exercise 8.1

a) The second derivative of the wave function is a measure of a particle's energy. Since there is a jump in energy level between two materials, the second derivative should reflect this jump. Therefore, the first derivative should be continuous, but the slope of the first derivative should be discontinuous. By extension, the wave function should be continuous, and its slope should be continuous.

b) We have freedom of choosing our origin for $z = 0$, so we do this right at the interface. We then find two equations that follow from matching the two halves of the wave function at the interface $(z = 0)$, and their derivatives:

$$\alpha_0 + \beta_0 \qquad = \alpha_1$$
$$-q_0\alpha_0 + q_0\beta_0 = -q_1\alpha_1.$$

We use the top equation to eliminate , and we divide the bottom equation by α_0 to find

$$-q_0 + q_0\frac{\beta_0}{\alpha_0} = -q_1[1 + \frac{\beta_0}{\alpha_0}].$$

Solving for β_0/α_0 we find

$$\frac{\beta_0}{\alpha_0} = \frac{q_0 - q_1}{q_1 + q_0}.$$

There is no term $\sim \beta_1$ since neutron do not travel back to the surface inside the medium since they cannot be reflected from the back of the material since the material is chosen to be infinitely thick.

c) If $q_0^2 - \beta < 0$, then q_1 will be a purely imaginary number. Writing this number as $q_1 = ia$ we get for $|\beta_0/\alpha_0|^2$:

$$|\frac{\beta_0}{\alpha_0}|^2 = \frac{[q_0 - ia]}{[ia + q_0]}\frac{[q_0 + ia]}{[-ia + q_0]} = 1.$$

Exercise 8.6

a) Critical angles are given by $\sin\theta_c \approx \theta_c = \sqrt{\beta}/k_0$ for an air-material interface.

These angles are very small, and hence, the neutrons have to be incident on the wall of the tube under very shallow angles. This can only be achieved if the tube is almost straight.

b) The critical angle is proportional to $1/l_0$ and hence, it is proportional to λ. Thus, long wave length (slower) neutrons can be 100% reflected, but faster neutrons will be transmitted through the wall of the beam tube, disappearing form the beam.

c) ^{58}Ni has two advantages: it has a very large scattering length b and therefore, a relatively large critical angle, and this isotope does not have any incoherent cross-section. The latter is an advantage because fast neutrons cannot make it down the beam tube by means of incoherent scattering.

Exercise 8.8

The smallest resolution Δz we can model in real space is reciprocally related to the extent of our spectrum in reciprocal space. Thus, $(\Delta z)_{min} = 2\pi/q_{max}$. For example, in an experiment where we probe out to $q_{max} = 0.2$ Å^{-1}, we have $(\Delta z)_{min} \approx 30$ Å.

Exercise 8.9

The advantage is that we introduce more structure into our reflectivity profile, so it will be easier to identify various features. The disadvantage is that the way a single, unsupported membrane behaves in nature might be very different from a bunch of membranes slapped on top of each other.

Exercise 9.1

The left sketch shows the situation on the left where we scatter at constant scattering angle. When we vary the incident or final wave length of the neutron (black arrows), we also vary both the length and direction of momentum transfer (red arrow). We can rotate the sample to ensure that the momentum transferred will be along the same crystallographic direction, but we cannot do anything about the change in magnitude of the momentum transfer. Even for liquids, where direction of momentum transfer does not matter, we would have this problem.

The sketch on the right shows how we can keep a constant magnitude of momentum transfer for differing incident or final neutron wave lengths. As long as we choose our scattering angle correctly, we can keep a constant amount of momentum being transferred (indicated by the circle). For a liquid we are done, for a crystal we would need to rotate the crystal to ensure that the direction of momentum transfer with respect to the crystallographic axes does not change.
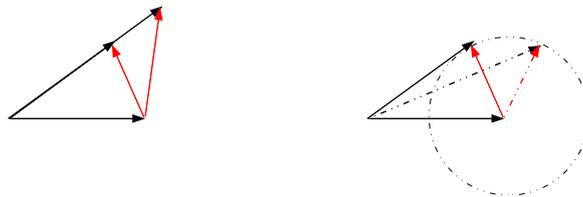
Exercise 9.2

When we cool the crystal down, the lattice will (most likely) contract, with the result that the Bragg peaks will occur at higher scattering angles. Therefore, we need to wait until the crystal is cold, and align the crystal.

Exercise 9.3

When we vary the incident energy, the monitor count rate efficiency will also vary; the monitor will be more efficient for slower neutrons than for faster neutrons. This variation in efficiency exactly compensates for the pre-factor $k_{\text{final}}/k_{\text{initial}}$ that shows up in the connection between neutron counts and dynamic structure factor. Therefore, our counts will be directly proportional to the dynamic structure factor when we use fixed final energy, and variable incident energy.

The second reason is that we would like to get rid off higher order contamination. If we keep the incident energy fixed, then we do not have a lot of choices for fixed incident energy that allow us to get rid off higher order contamination. In fact, the best we can probably do is to use $E_{\text{incident}} = 14$ meV, greatly restricting how much energy the neutron can transfer to the sample since the neutron cannot give up more energy than it has.

Exercise 10.3

The discussion about how to make weak ferromagnetic peaks visible applied to Bragg peaks, and therefore, the exponent would equal one. Also, the ordering is ferromagnetic (not antiferromagnetic) so that the magnetic lattice coincides with the nuclear lattice. In other words, why bother with the lattice if it is the same for both magnetic and nuclear scattering?

Exercise 12.1

We will do the easier to visualize case for which $\phi_0 = 0$. We use eqn I.1 to calculate the attenuation factor, the factor by which the non-attenuated scattering has to be multiplied in order to get the attenuated scattering. The problem is a 1-dimensional problem, where the volume integral can be replaced by a 1-dimensional integral over the thickness of the slab. Denoting the scattering angle by θ, and the point of scattering by x- with x varying between 0 and τ- we find for the attenuation factor (with x the path length through the sample before scattering and $(\tau-x)/\cos\theta$ the path length through the sample after scattering provided that $|\phi| < 90^o$)

$$H_1(\theta) = \frac{1}{\tau} \int_0^\tau dx\, e^{-\Sigma x - \Sigma(\tau-x)/\cos\theta}.$$

Carrying out the integration we find for $|\theta| < 90^o$:

$$H_1(\theta) = \frac{e^{-\Sigma\tau} - e^{-\Sigma\tau/\cos\theta}}{\Sigma\tau/\cos\theta - \Sigma\tau}.$$

Doing the calculation for $|\theta| > 90^o$ we find:

$$H_1(\theta) = \frac{e^{\Sigma\tau/\cos\theta}[e^{-\Sigma\tau} - e^{-\Sigma\tau/\cos\theta}]}{\Sigma\tau/\cos\theta - \Sigma\tau}.$$

When we extend the calculation to $\phi_0 \neq 0$ we find for $|\theta - \phi_0| < 90^o$:

$$H_1(\theta) = \frac{e^{-\Sigma\tau/cos\phi_0} - e^{-\Sigma\tau/\cos(\theta-\phi_0)}}{\Sigma\tau/\cos(\theta-\phi_0) - \Sigma\tau/cos\phi_0},$$

and for $|\theta - \phi_0| > 90^o$ we find

$$H_1(\theta) = \frac{e^{\Sigma\tau/\cos(\theta-\phi_0)}\left[e^{-\Sigma\tau/cos\phi_0} - e^{-\Sigma\tau/\cos(\theta-\phi_0)}\right]}{\Sigma\tau/\cos(\theta-\phi_0) - \Sigma\tau/cos\phi_0}.$$

Part V

Modules

This part lists 6 modules and a final project that constituted the lab component of the IGERT neutron scattering course. This part should be skipped, it was merely included so that we could keep track of what modules were given during the course.

Appendix K
Basic Diffraction Module

Aim: measure microscopic spacings using an interference technique.

Take a look at the figure shown below. In this sketch we see a water wave arriving at two openings, after which each opening will act as a source of a new water wave. Given the geometry of this situation, the two newly formed waves will oscillate in phase since the wavefronts arrive at the two openings at the same time.

1) Come up with an expression that shows the relationship between regions of maximum constructive interference, how much farther one wave had to travel compared to another (path difference PD) and the wave length of the original wave that created the two new waves.
In particular, we are looking for an expression that reads PD= f(λ) where f(λ) is some function of the wave length λ of the original wave.

2) Repeat this exercise, but now for regions of perfect destructive interference.
In particular, we are looking for an expression that reads PD= g(λ) where g(λ) is some function of the wave length λ.

3) What is the relationship between the radial angle at which we observe regions of constructive interference and the other parameters of the
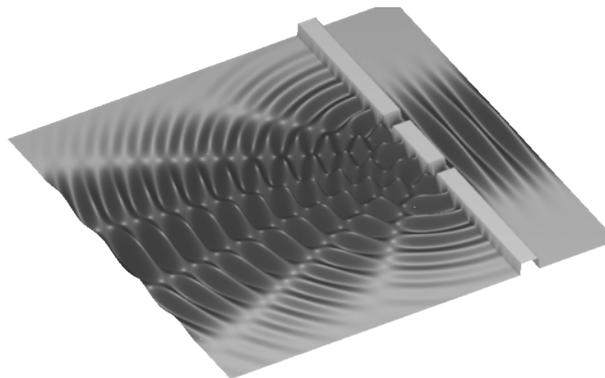
Fig. K.1 Water waves making it through two openings will create an interference pattern.

problem?

We can define the angle Θ as the angle between the direction of the incoming wave, and the direction of the resultant waves. We measure Θ with respect to the center of the two openings. Find an expression that expresses PD in terms of Θ and any other parameter that defines the situation. In particular, find an expression of the sort $h(\Theta,...)= f(\lambda)$. Assume that the distance between the two openings is much larger than the width of the individual openings.

4) Repeat exercise number 3) for the case of destructive interference.

5) If instead of water, we had used sound waves or electromagnetic waves, how would your expressions under parts 3) and 4) have changed?

6) Imagine a wall made up of many, identical openings all spaced the same distance apart. We would still find an interference pattern, and the relationships derived under parts 3) and 4) would still hold. Why is this?

7) You have been given a laser pointer that emits light of known wavelength. You have also been given a grating where the distance between the openings is unknown. See the figure below for a close up of a diffraction grating. Shine the laser of the grating, and observe what happens to the light that passes through the grating. Make a sketch.
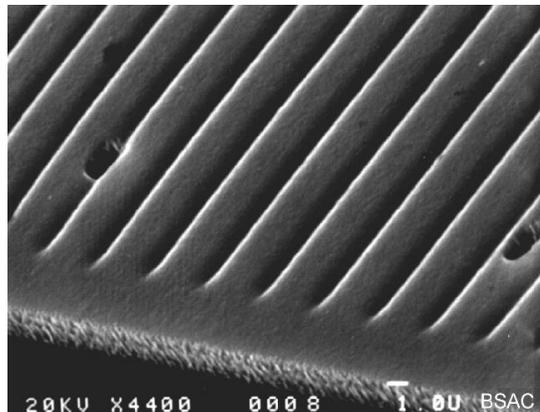
Fig. K.2 An electron micrograph close-up of a diffraction grating. Photo by Jocelyn Nee, www.eecs.berkeley.edu/IPRO/.

8) Observe what happens to the interference pattern when you change the angle between the incoming light direction and the normal of the grating. Make a sketch and explain.

9) Calculate the distance between the scratches based upon the interference pattern that you are observing.

10) What do you expect to happen to the interference pattern when we half the distance between the openings? What if we double the distance?

Appendix L
Counting Chain Module

In the course booklet, the operation of a typical ^{3}He neutron detector has been described. That is, a ^{3}He nucleus absorbs a neutron, splits into two charged particles (a tritium nucleus and a proton); these charged particles ionize the gas so that an electronic pulse is picked up across the two electrodes of the detector. In this module, you will get your hands on a detector and tune the electronic settings as you would for an experiment at an actual neutron scattering instrument when you have discovered that the counting chain part of the spectrometer does not perform as expected.

The following components are utilized in the counting chain:
- a ^{3}He filled neutron detector tube,
- a high voltage power supply,
- a single channel analyzer,
- a ratemeter,
- a preamplifier,
- cabling (high voltage and coaxial).

In this module you will also be making use of an oscilloscope and a multichannel analyzer (which is contained within a pc).

1) Neutron detection events reach the instrument control software via a series of electronic components called a counting chain. **Rationalize the relative order in which the above components should be hooked up in order to deliver the pulse of the detector to the computer (or equivalently, a ratemeter, for the purposes of this module).** (Note: Complete this part one before leaving the classroom for the lab.)

2) To save time, the counting chain is set up for you. **Inspect the order of linkage and compare this to your answer from part 1). If the detector is not connected, do so before powering up.**

There is a power switch (red circle in Fig. L.1) to the NIM bin (the rack that houses all the modules such as the preamplifier) that powers everything except the high voltage power supply, which has its own separate switch (blue circle). The ratemeter provides a quick read out your detected signal. The objective is now to optimize the settings to get the most neutron counts while limiting any spurious (gamma ray, etc.) counts. We will do this in the following sections in turn.

Fig. L.1 The basic setup of the counting chain of a single detector. Note that the order of the components in the rack does not have to coincide with the logical order of the components within the counting chain.

3) What are some guiding principles for how one should adjust the high voltage supplied to the detector?

Fig. L.2 As a neutron source we use a chunk of californium. Needless to say, do not remove the source from its bucket.

4) Place the detector near to the Californium neutron source, i.e on top of the paraffin shielded bucket shown in Fig. L.2. As you do so, observe the signal on the oscilloscope; this will require setting the trigger so that you can see the predominant pulses coming in. Adjust the high voltage bias of the oscilloscope accordingly and mark down your optimized value.

5) Now that we have some signal, we want to eliminate any unwanted counts that result from things other than the neutrons from our source. To do this we must adjust the band pass filter settings of the single channel analyzer (SCA). Any pulses that are too weak or too strong will get tossed out. **To get a feel for it, play around with the lower level discrimination and the window size.**

Fig. L.3 The magic icon to press to get the multichannel analyzer to run.

6) Not all of the counts visible on the ratemeter are coming from neutrons. In order to eliminate the unwanted signal, we will record a spectrum using the multichannel analyzer (MCA). Make sure that the (BNC) cable from the SCA is plugged in to the back of the PC and open the Maestro software. **Acquire a spectrum on the multichannel analyzer and use this spectrum to (fine-)tune the lower level discrimination and window settings on the SCA to their optimal values; record these values. Does your MCA spectrum look similar to the one shown in the figure?**
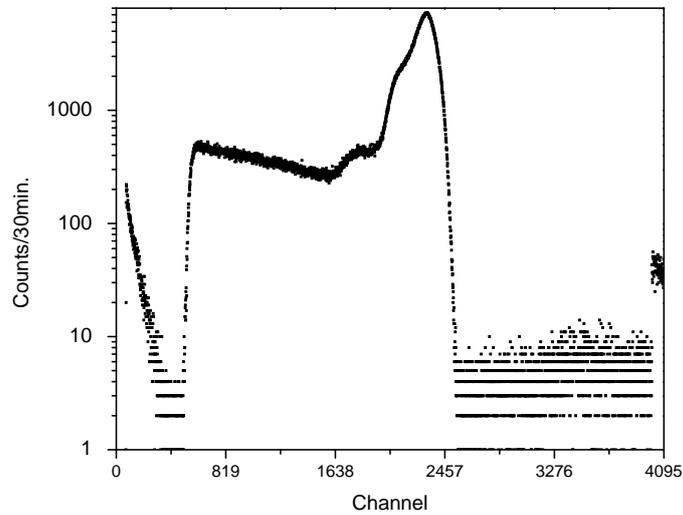
Fig. L.4 An example output of the MCA, with the counts shown on a log scale.

7) To be completed after the lab: describe the features that you observe in the spectrum and what their origin might be, i.e. neutrons, gamma rays, double counted neutron events, etc.

8) Why didn't we remove the neutron source from the paraffin bucket?

9) With your now optimized counting chain, make a note the reading from the ratemeter and do the same with the source removed from the bucket (taking appropriate precautions as the source is hot). Is this what you expected? Explain.

10) With the source placed back in the bucket, make a comparison of the count rate when using the various materials provided for shielding (Cd, boral, aluminum, borated poly, plastic). Mark down the resulting ratemeter readings and describe in what capacity each might be used on a neutron spectrometer.

Appendix M
Diffractometer Module

We will use the spectrometer as a 2-axis single detector diffractometer for this module. There is a sheet with basic spectrometer control commands, such as how to drive the θ-motor, available to you. The main shutter of the instrument is open. These shutters are in the wall of the reactor, and they cannot be operated by us. If we want them closed, we have to ask the control room to do it for us. We have control over the secondary shutter(s), which are the shutters that are after the monochromator and before the sample position.

1). Open the secondary shutter, and use the xray and neutron handheld monitors to assess the level of radiation that is present near the various parts of the instrument. Close the secondary shutter, and repeat with both monitors. What do you conclude about how effective the secondary shutter is at stopping xray radiation?

The next exercise is to look at the straight through beam, and get an idea of its intensity and the amount of air scattering it generates.
2). Open the secondary shutter, and drive the detector to the straight through position, that is, to θ=0. Look at the count rate on the rate meter. Move the detector to θ=0.5 ^{0}. Again, look at the rate meter. Do you notice anything unusual?

3) Stick an attenuator in the beam, and repeat part 2. What you actually observed in part 2) are the limitations of the detector electronics. It takes a certain amount of time for the detector to reset, so that it is ready for the next neutron to be counted. If neutrons come in faster than this reset time, they are not counted. Be aware of this not only for the straight through beam, but also when you notice an unusual counting profile when you are looking at a strong Bragg peak of a big sample. Most likely, if you notice a peak shape that is flat, then your detector is being saturated.

4) With the attenuator in, place a mask before the sample position, and repeat part 2. What do you conclude?

You have been given a sample of $LiMn_2O_4$ in powder form. The unit cell (see figure) is cubic with a lattice constant of about 8 Å(we are being vague here on purpose). The wave length of the neutrons that are coming off the monochromator is listed on the command sheet at the spectrometer.

5) calculate at what angle you would find neutrons Bragg reflected by your sample. You can pick any Bragg reflection that you like.

6) Remove the attenuator and drive the detector to your calculated position. Look at the rate meter to see if you are actually seeing any scattered neutrons. If you do not see any scattered neutrons, move the detector to slightly different angles. If you still do not see a lot of neutrons being scattered, sit back and think.

7) Take a scan of your Bragg peak. Make sure the width of your angular window that you are scanning over is large enough so that you see the peak intensity drop off to back ground level. What is the characteristic width (in angle) of your Bragg peak and how do you know you reached background level?

8) Place a collimator in front of the sample, and repeat part **7)**. Write down the acceptance angle of the collimator that you used.

9) Drive the detector to zero angle. Then drive the detector to its highest allowed angle (see command sheet) in steps of five degrees while looking at the rate meter; take note of the angular positions where the rate meter shows a large number of scattered neutrons. You may have to backtrack a bit using the drive command to get the position of these new peaks reasonably accurate (say to within 0.2^o).

10) To be completed after leaving the beam port floor. **Identify the peak positions of part 9) with their Bragg reflection indices.**

11) To be completed after leaving the beam port floor. **Estimate the q-resolution of the spectrometer. How does this compare to the angular resolution you measured in parts 7) and 8)? You may wish to refer to Chapter 5, eqns 5.3 and 5.4.**

Appendix N
Powder Diffractometer Module

In this module we will take a diffraction profile on a powder (the easy part) and we will refine the structure of our material using the FULLPROF Rietveld refinement program (the hard part).

At your disposal you have an unknown powder sample, and a spectrometer. The spectrometer is already set up to operate at fixed incident neutron wave length. The approximate value of this wave length is written on the instrument cheat sheet. The instrument has also been aligned, but not very well. This implies that the nominal (as per computer readout) zero scattering angle $\theta=0$ does not exactly correspond to the direction of the straight through beam. So be aware of this when analyzing the data.

1) Place the powder sample on the spectrometer, and take a quick pattern that covers all angles of interest in about half an hour. Use the instrument cheat sheet to figure out how to do this.

2) Inspect the obtained powder diffraction pattern. Look at how many counts there are in the most intense peak, and compare this to the background level. Calculate how long you will have to count for so that the maximum intensity in the peak is known to at least 1 % accuracy.

3) Set up a run that will take a powder pattern with the accuracy calculated above. Start the run and make sure that you are actually collecting data. This is good practice because it is not unusual to spend hours setting up the spectrometer and the runs, and to forget to, for instance, open the beam shutter. This is the end of the part to do on the beam port floor.

4) You have been given the data file you measured, and access to the FULLPRO refinement program and manual. **Determine the structure of your material**, following the steps outlined in Chapter 6. We want to determine the crystal structure, the number of atoms in the unit cell, the positions of these atoms and their Debye-Waller factors. Give it your best shot, and make sure to look at some of the parameters as determined from your refinement such as the instrumental parameters (zero angle offset, incident wavelength, background, resolution), as well as how reasonable your Debye-Waller factors appear to be.

5) Try to identify what material you were looking at using any data base you can find on the internet.

Appendix O
Reflectometry Module

In this module we will align a sample on the reflectometer and measure its reflectivity profile. The reflectometer has been aligned for you, everything (including higher-order filters) is in place and the wavelength of the neutrons in the incoming beam is fixed at λ= 2.36 Å.

You have been provided with a sample of a Ni film of unknown thickness (to be determined in this module); the film is on a glass substrate, as shown in the photo. The goal is to determine the scattering length density profile of the sample as a function of distance to the air-film interface. The steps are to align the sample, to measure the on-specular and off-specular intensities, to convert the raw data to a reflectivity curve, and to fit a model reflectivity curve to the data. The last step will yield the desired scattering length density profile.
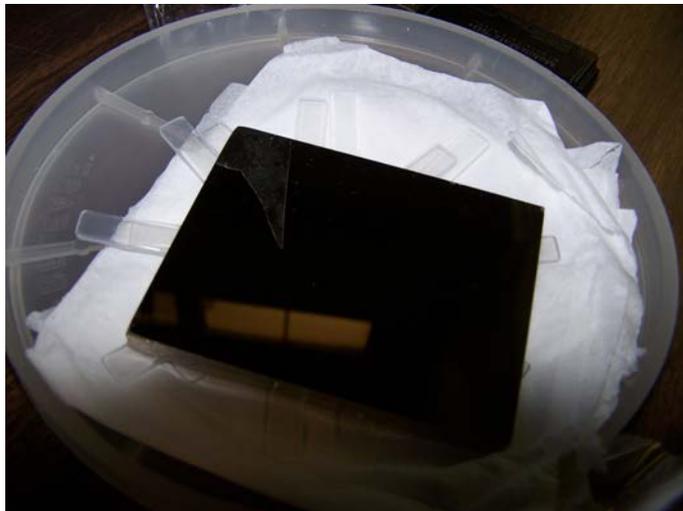
Fig. O.1 The sample consists of a thin nickel film on a fat glass substrate.

1) *Establish a reasonable starting position.* The alignment procedure is an iterative process. First, **mount the sample on the sample table**. To give yourself a good

starting point, check by eye that the sample is aligned with the beam (as defined by the slits). This is, of course, somewhat crude, but we are only trying to get a leg up for the next step where we use neutrons to align the sample in a more precise way. This should be just quick and dirty.

2) *Locate the sample.* That is, translate the sample laterally through the beam (refer to the instrument quick reference sheet for the correct drive commands as well as for the motor mnemonics). The sharpness of the edges that you observe depends on how close the θ-position is to the true zero position; this should improve with each iteration step [we will repeat steps 2) through 5)], unless you happened to have a very lucky starting position. Note that the sample holder backing consists of a sandwich of aluminum-cadmium-aluminum. **Determine where the sample surface is and move the the sample to that position. Sketch the translation profile and identify the features present and write down the position of the Ni surface**.

3) *Cut the beam.* Now your sample should be position so that it is just cutting the beam. By that we mean, if you back off the sample slightly the beam will pass by untouched, but if you push forward slightly it will begin to obscure the beam. This is the condition that we want for our measurements but the problem is that there are still some accumulated errors resulting from misalignments of the θ-position and the tilt under the sample. The next step is to **rock the θ-angle to determine the center position. Write down the apex of the scan, move to this position, and encode it to zero** (again, consult the quick reference sheet).

4) *Align the θ-position.* In order to gain accuracy in our θ-alignment, it is necessary to move to a non-zero 2θ-position and re-measure the θ-rocking curve. The reason behind this is to move away from the high background of the straight through beam and its associated air scattering and use actual specular reflection. So there is a trade-off, because going to higher 2θ moves you further from the high background but your reflected signal also drops off proportional to $1/q^4$. Choose a value of 2θ with this in mind and **measure the θ-rocking curve at this position. Mark down the center of mass, move to this position, and encode it to be where it was supposed to be** (how do you know where it should be?). This will be a longer scan because, of course, there are far fewer neutrons away from the straight through beam.

5) *Align the sample tilt position.* We have not yet touched the sample tilt. The way to align this is simply to **redo a rocking curve of the tilt axis** at the specular, non-zero $\theta - 2\theta$ position. This will also be a longer scan and may require repetition depending on your initial choice of scan range. **Move to the center of mass of the observed intensity**.

6) *Fine-tune the alignment.* Because the changes to tilt- and (especially) θ-offsets affect what the translation profile looks like, we must **repeat steps 2) through 5) until we are satisfied** that the alignment is sufficiently good. Do this and put some thought into how good is good enough for the alignment. And, as always, take careful

notes.

7) *Measure the reflectivity*. Well, actually, you will have to do some processing of your raw data in order to get to a reflectivity curve; in actuality, in this step you will **measure the specular and off-specular intensities. Write a macro** (with the aid of your instructor) to measure these. Write down the name of your macro data file in your notes. **Run the macro**. This concludes the in-class portion of the module.

8) *Convert the data into a reflectivity curve*. The definition of reflectivity was given in Chapter 8. How do you get from your raw data to a reflectivity curve? **Write down the procedure that you will use and implement it**.

9) *Refine the footprint correction*. The reflectivity below the critical angle (or equivalently, $q < q_{\text{critical}}$) should be flat as a function of angle down to where the straight through beam starts to come in and overwhelm our signal. Is this the case for your plot? If not, then consider how to adjust the footprint correction[1] in a more refined way. That is, consider a Gaussian beam (a beam that falls off in intensity according to a Gaussian profile with characteristic width W) impingent upon a sample surface whose length is L (measured along the straight through beam). From your notes during the alignment procedure, you should be able to come up with a beam width. **Write up a solution and apply it to your data**. Hint: your answer should take the form of an error function.

10) *Model the reflectivity curve*. Now that you have a nice flat reflectivity for angles below the critical angle, you are ready to model your reflectivity data. The NIST Center for Neutron Scattering supports a free-to-download reflectivity analysis software package called *reflpak*, which can be found at the following URL: http://www.ncnr.nist.gov/programs/reflect/data_reduction/software/index.html. This package will allow you to **model a reflectivity curve and ultimately determine the Ni film thickness and the scattering length density profile**. Note that it is necessary to convert the data into the appropriate format to be read into the software. Hand in a graph that shows the agreement between your model and the measured reflectivity profile.

---

[1]With footprint we mean what fraction of the incident beam is intercepted by the sample. The more we rotate this sample, the larger the fraction that is intercepted. Assuming a beam whose intensity is constant as a function of its width, geometry tells us that the fraction of the beam that is intercepted by the sample is proportional to $1/\sin\theta \sim 1/q$.

Appendix P
Triple Axis Spectrometer Module

In this module we will align a single crystal, and we will measure longitudinal and (for the hard-condensed matter scientists) transverse phonons. The TRIAX spectrometer has been aligned for you, everything (including higher-order filters) is in place, and the incoming neutron wave length is set to $\lambda = 2.37$ Å.

You have been given a single crystal of the mineral fluorite, CaF_2, as shown in the photo. The crystal structure of fluorite has been resolved using powder diffraction. The structure is face centered cubic (Fm3m, space group 225) (Fig. P.1). The room temperature lattice spacing is $a = 5.4626$ Å. The powder pattern relevant to our present TRIAX settings is shown in Fig. P.2.
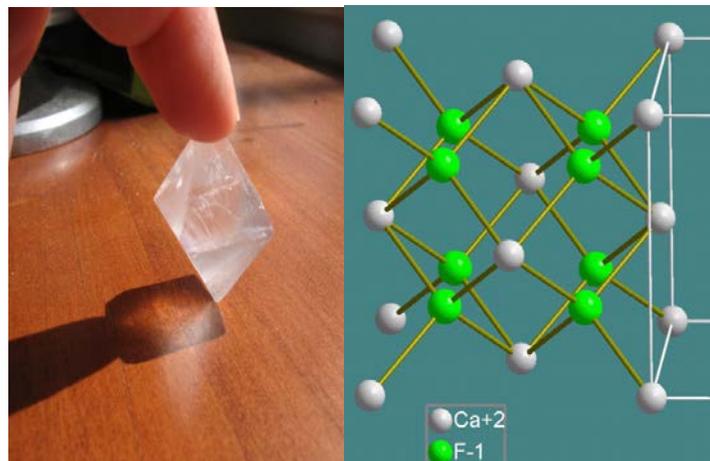
Fig. P.1 The left panel shows the single crystal of CaF_2, the right panel shows the positions of the atoms within the unit cell (image (right panel image source: Wikimedia Commons, image by NIMSoffice).

The first task is to align the crystal. This means that we will orient the crystal in such a way that we can transfer momentum along a specific crystallographic direction. Looking at the powder pattern, we see that the (family of) (220)-Bragg peak(s) is a strong one. This implies that phonons in this Brillouin zone will also be strong, and
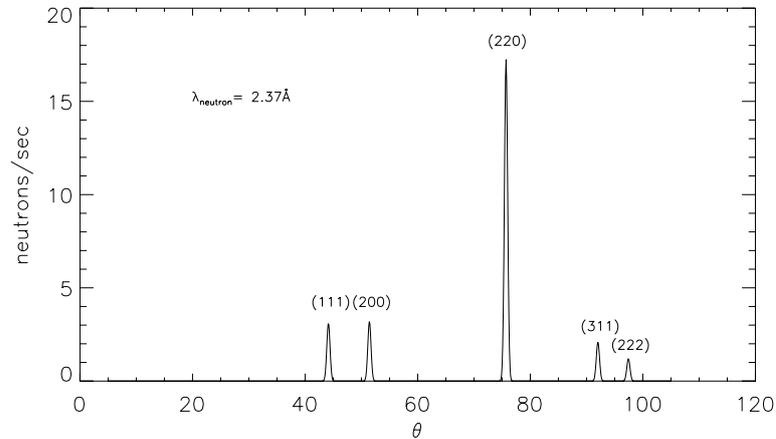
Fig. P.2 The powder pattern as a function of scattering angle θ for CaF$_2$ with incident neutron wave length λ= 2.37Å.

we will try to measure a longitudinal phonon at $(2+h, 2+h, 0)$. We can only do this if we have the crystallographic (110)-direction in our scattering plane.

1) To do after the lab: **Make a sketch of the positions of the atoms in the unit cell looking down along the c-axis (z-direction)**. Verify that it makes sense that the (220)-reflection is a very strong one (both F and Ca have positive scattering lengths).

Fluorite cleaves along (111)-planes, giving the crystal its distinct octahedral shape. This implies that if we place the crystal on one of its cleavage planes, that the (111) direction will be pointing straight up. This is good, because the (220)-direction is perpendicular to this [the (2,-2,0) direction really, but all these directions are equivalent in a cubic system], implying it will be in the scattering plane.

2) **Enter the crystal information (cubic, lattice spacing, etc.) into the SPICE software**. Tell the computer that one of the crystal directions in the scattering plane will be the (1,-1,0)-direction. Pick (1,1,-2) as the second vector that is in the scattering plane. Verify (for yourself) that the three vectors [(1,1,1), (1,-1,0) and (1,1,-2)] are perpendicular to each other).

3) **Place (and secure) the sample on the sample table at the right height for it to intercept the incoming neutron beam, and drive the detector to the correct angle for the (220)- Bragg reflection. Rotate the crystal (using a software command) until you notice a lot of counts in the detector (or on the rate meter)**. This is the (220) Bragg reflection.

4) Turn off the beam, and look at the crystal. Picture the (220) planes in your

mind's eye, and **reposition the crystal on the sample holder in such a way that the normal to the (220)-planes coincides with a tilt axis of the sample table**. In this way, we can adjust the tilt angles under the crystal in such a way that the two tilt adjustments are decoupled from each other. Once the crystal has been repositioned, **rotate the sample angle (using the computer) until you have found the (220) Bragg peak again**.

5) **Scan the sample angle, fit the data, and set the angle under the sample so that it is located at the peak of the intensity. Now redo the scan of the sample angle, while simultaneously varying the scattering angle by double the amount** (question for afterwards: why double?). For instance, this command could read 'scan s2 60 64 0.1 s1 40 42 0.05'. After the scan, **drive both motors to the peak position of the scan. Instruct the software that you are now sitting on top of the (2,-2,0) peak**.

6) While sitting on the peak, scan the sample tilt motor (sgu or sgl) to find the maximum intensity. If the sample tilt is off by more than a degree, then redo step 5) after driving the sample tilt motor to the peak position. You are now aligned on the (2,-2,0) Bragg peak, while the (equivalent) of the (111)-direction is pointing vertically up.

We now need to locate another Bragg peak in our scattering plane formed by the (1,-1,0) and (1,1,-2) directions. However, there are not many peaks to choose from as we learned from the powder pattern in Fig. P.2.

7) **Verify that the (2,0,-2) peak is in the scattering plane. Drive there, and repeat all the steps to get the optimal reflection for the (2,0,-2) reflection**. Update the software. Look at the powder pattern and tell the software to drive to another strong Bragg peak that is not one of the two peaks that we used for aligning the crystal. **Provide feedback on the software's feedback when you do this**. Instead, drive to the (1,1,-2) peak, and notice its absence. **Remove the PG filter, or ask the instrument scientist to remove it**. Notice the counts on the rate meter. What is happening here?

8) **Do an elastic scan of the (2,0,-2) Bragg peak in the longitudinal and transverse directions**. We will need this to see over what range (in q-space) the Bragg peak has been smeared out. Print the graphs.

9) Try to measure a longitudinal phonon. To do so, ensure that the momentum transfer is in the longitudinal direction (such as when the momentum transfer is (2.1,0,-2.1). **Pick a position in q-space that is outside of the resolution broadened Bragg peak, but that is still fairly close to the Bragg peak. Set up an inelastic scan** in such a way that a varying amount of energy is transferred to the sample at constant momentum transfer. Do this by varying the incident energy of the neutron, while keeping the final energy fixed. **Start with a scan that covers an**

energy transfer range of 0 $< E <$**8 meV**. If you observe a peak, put in a longer scan, if not, measure closer to the Bragg peak.

10) Optional, for those interested in hard-condensed matter. **Measure a transverse phonon** as well. Do this by typing in the correct software command and let it measure while you are no longer on the beam port floor.

11) To do after lab: **determine the longitudinal and transverse speed of sound for this mineral**. Compare with literature values.

Appendix Q
Final Projects

The following are examples of final projects that were given to the students during the first run of the course at MU, Spring 2012.

Final assignment for PHY xxx

The assignment consists of 4 parts:

- writing a proposal for beam time,
- obtaining the sample and (designing the) sample holders,
- doing the experiment,
- analyzing the data and writing a report,

The beam time proposal is due on xxx, the sample holder design is due April xxx, and the Final report is due on xxx.

To complete this assignment, you have at your disposal three instruments at MURR [the powder diffractometer, the reflectometer and the 3-axis spectrometer], a modest budget, and the physics machine shop. At any point you can ask Tom Heitmann, Helmut Kaiser, and Wouter Montfrooij for advice, or your advisor.

You will do the two projects in groups of two; all written parts will be submitted as a group, and they will be graded as a group.

The assigned topic is based on your research, and it should be viewed as a stepping stone to being able to do neutron scattering experiments that you will use in your research. In fact, after this course we will be happy to help you out with any future experiments (at least, that is what we say until the evaluations are in).

Topic: (all 5 projects that were given are listed A-E).

- (A) Loading activated carbon with N_2 might tell us something about what H_2 does in the pores. Load a sample with N_2 such that it will scatter 10% of the neutrons. Determine whether the N_2 molecules are deposited as molecules, or as atoms (dissociation). Determine whether N_2 forms a single layer, or multiple layers in the pores. Determine the average pore diameter from your results.
- (B) Loading activated carbon with N_2 might tell us something about what H_2 does

in the pores. Load a sample with N_2 such that it will scatter 10% of the neutrons. The dynamics of molecules in confined geometries will tell us something about their binding energies, and what brings about those binding energies. Set up a neutron scattering experiment that demonstrates that we can glean information about the dynamics of nitrogen molecules in pores of activated carbon. Try to make the experiment sensitive to changes (from bulk N_2 behavior) in short wavelength sound waves, in heat transport, and in the behavior of individual molecules. Report on what you can learn about the dynamics of N_2 in confinement, and how it differs from bulk behavior.

- (C) First, calculate the spin wave dispersion for a body centered cubic system of two magnetic atoms, where the one atom has a spin of magnitude 5/2, and the other has a spin magnitude of 3/2. Both species have their spin aligned along the c-axis, but they are pointing in opposite directions. Second, you have a single crystal of magnetite Fe_3O_4. This crystal is ferri-magnetic at room temperature. Measure the spin wave dispersion of the acoustic spin waves (in a suitable Brillouin zone) and determine the strength of the magnetic interaction between (some) Fe ions.

- (D) Determine the crystal structure, including any superstructure, of $TiMn_{0.25}S_2$s. Determine whether magnetic clusters exist at low temperature, and if so, measure their temperature dependence up to the point they disappear. Do this in such a way that shows whether the onset of magnetic clustering upon cooling is gradual (a second order transition), or sudden (a first order transition).

- (E) Determine the scattering length density of a single supported lipid bilayer on a Si substrate. Measure a second contrasting lipid sample and discuss what difference and similarities you are able to determine from the neutron data concerning the difference in lipids used as well as whatever you are able to say about the sample quality. Be as detailed as possible.

Proposal:

The proposal should include all the technical details: sample composition, instrument of choice for the experiment, estimate of how much beam time is required, dimensions of sample, what to use for a sample holder, any special instrument requirements (collimation for instance). Do you need to measure background runs? In short, everything that is relevant should be put in there, similar to the proposals that are included in the handouts. Based on this proposal, we will give you beam time on one of the instruments. The proposal is due on xxx.

Sample and sample holder:

To do the experiment, we need a sample, such as heavy water or some intercalated powder. We will also need to place the sample on the spectrometer, perhaps even in a cryostat or furnace. Likely we will have to design a suitable sample holder for this. Present a design for a sample holder to us (Tom, Helmut or Wouter), and once we think it will work, we will ask the physics machine shop to fabricate it. The due date on the sample holder design is xxx.

Experiment:
We will be there to make sure you do not make fatal mistakes, but essentially, it is all up to you. The experiment will take place in the period xxx-xxx.

Data Analysis:
It is a long process to go from the data as they roll off the instrument (raw data) to what we want to know about the sample. There (possibly) need to be corrections for background, attenuation, detector efficiency just to name a few. Your data analysis will not be perfect, but you should get as much information out of your data as possible. The final report, in journal article form, is due on xxx.

Index