

The Performance of Neutron Scattering Spectrometers at
a Long-Pulse Spallation Source

by

Roger Pynn

Los Alamos National Laboratory, Los Alamos, New Mexico 87545, U.S.A.

Introduction — The Simple-Minded Approach

In this document I consider the performance of a long pulse spallation source for those neutron scattering experiments that are usually performed with a monochromatic beam at a continuous wave (CW) source such as a nuclear reactor [1]. This class of experiments can be represented generically by the upper part of Figure 1. A monochromatic neutron beam with wavevector \mathbf{k}_i falls on a sample and the scattered beam, with wavevector \mathbf{k}_f , is analysed in some way to obtain information about the sample.

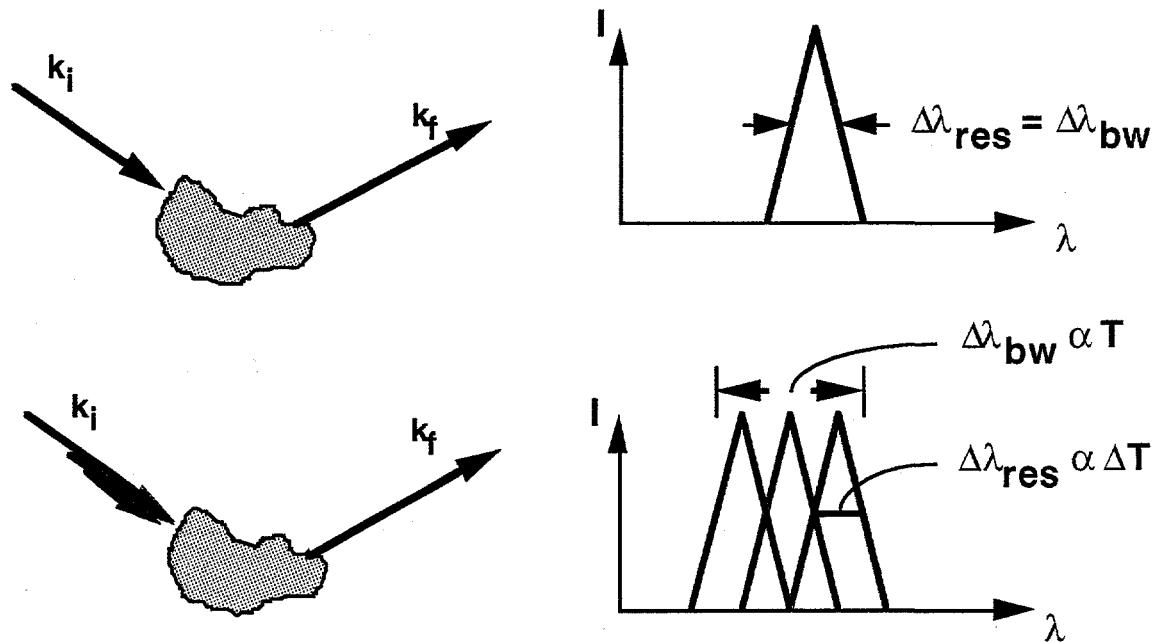


Figure 1: The upper part of the figure represents a generic scattering experiment carried out with a monochromatic beam. The graph to the right represents schematically the intensity of neutrons incident on the sample as a function of neutron wavelength. The lower part of the figure indicates that the same experiment can be performed with several different incident wavelengths that can be labeled by the neutron time of flight.

If one were to perform the same experiment at a pulsed source, one could use neutrons with several incident wavelengths, as depicted in the lower panel of Figure 1, each labeled by its time of flight to a detector.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

How many different wavelengths can be used, in principle, at a pulsed source? It is fairly clear from Figure 1 that the answer is simply the accessible wavelength bandwidth ($\Delta\lambda_{bw}$) divided by the full-width at half maximum of the incident wavelength resolution ($\Delta\lambda_{res}$). If we do nothing special to the beam from a pulsed source, the incident wavelength resolution $\Delta\lambda_{res}$ is proportional to the pulse duration, ΔT , while the accessible wavelength bandwidth is proportional to the time between pulses, T . In fact:

$$\Delta\lambda_{res} \approx 4 \cdot \Delta T / L \quad (1)$$

$$\Delta\lambda_{bw} \approx 4 \cdot T / L \quad (2)$$

where L is the distance in meters from the source to the neutron detector, times are measured in milliseconds and wavelengths in Ångstrøms. Eqns (1) and (2) tell us that we can make $1/c$ times more measurements with the same incident wavelength resolution at a pulsed source of duty factor $c = \Delta T/T = \Delta\lambda_{res}/\Delta\lambda_{bw}$ than we can at a CW source. Since this same factor relates average to peak flux — i.e. $I(av) = c I(peak)$ — we are led to the conclusion that the peak flux of a spallation source should be compared to the average flux of a CW source.

Caveats

This, of course, was the state of play before the Shelter Island meeting in 1984 [2] and has long been known to be too simplistic a view of the comparison between CW and pulsed sources. The Shelter Island workshop taught us that one has to consider each type of experiment — powder diffraction, SANS etc — separately and make a detailed comparison of the time taken at pulsed and CW sources to take the same data. Let us examine the reasons for this conclusion.

There are several obvious things that were swept under the rug in the simple argument given above. For example, it is clear that the argument only holds if the same secondary spectrometer is used at both sources and if the collimation of the incident beam is also the same. Even if these conditions are fulfilled, several other factors affect the comparison, as Figures 2 and 3 show.

Suppose, for example, that the data obtained from all of the different wavelengths at a pulsed source are not useful, either because the dynamic range of the wavevector transfer Q is too large at the pulsed source or, as depicted in Figure 2, because the vector Q lies in parts of the reciprocal space of a single-crystal sample where there is no useful information. In this case, it matters not that we have simultaneously measured many data points at a pulsed source because some of the information we have obtained is not required. A similar argument, not depicted in Figure 2, pertains to inelastic scattering experiments: the many wavelengths used at a pulsed source may probe energy transfers that are not of interest.

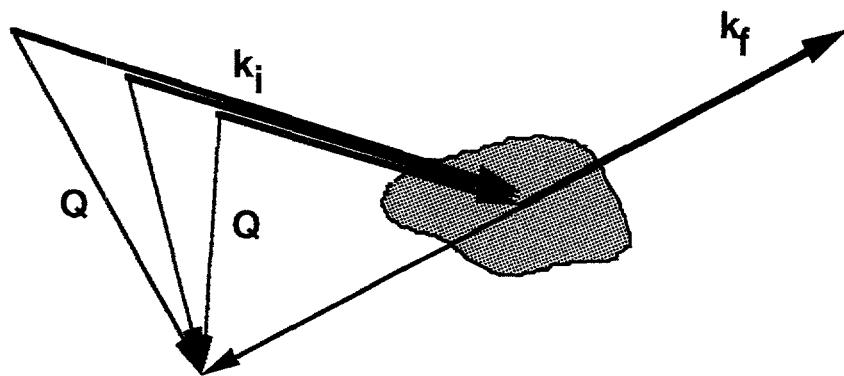


Figure 2: When different incident wavelengths are used in an experiment, different wavevector transfers, Q , are accessed in the sample if the scattered wavevector is not changed.

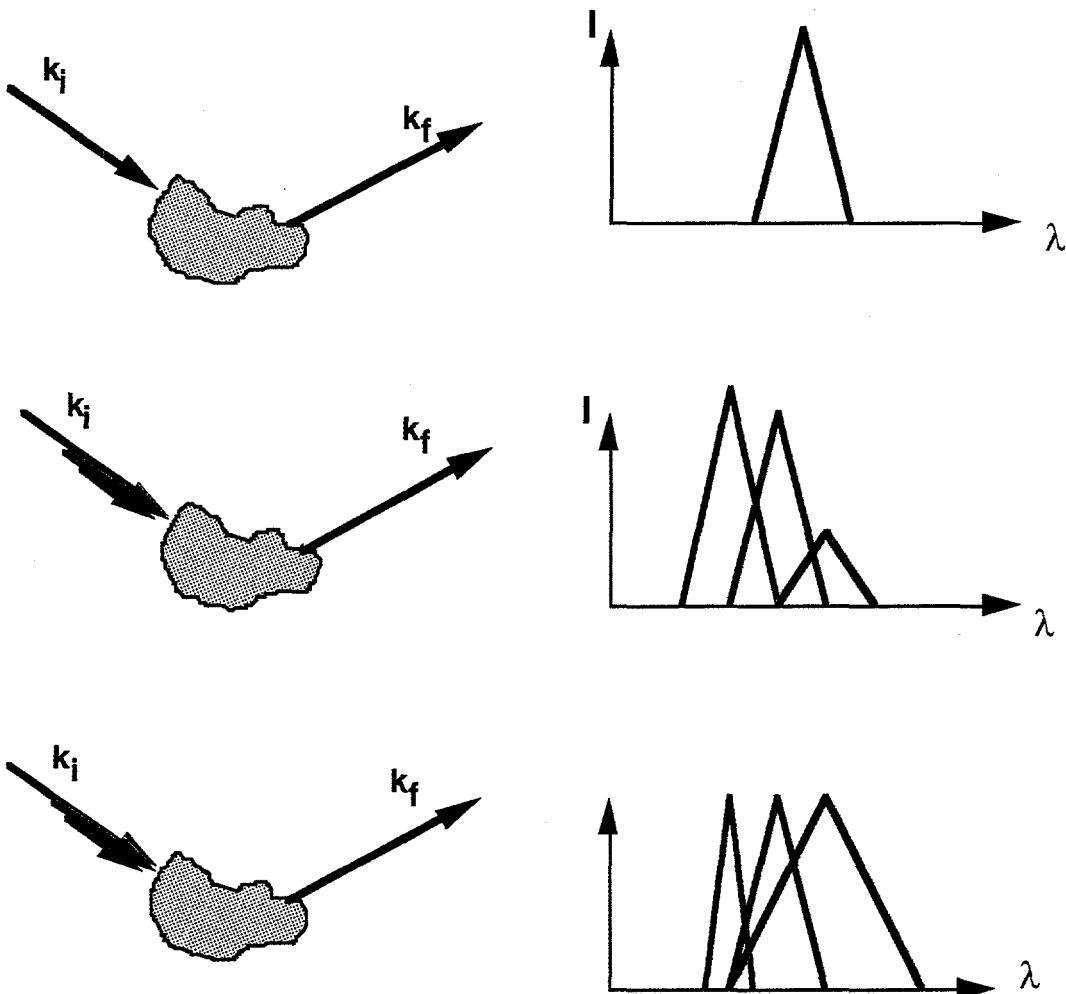


Figure 3: The upper panel represents a generic experiment using monochromatic neutrons with a particular wavelength spread. The second panel indicates that when different incident neutron wavelengths are used, the incident neutron flux may not be the same for each incident wavelength. In the lower panel an experiment is performed with different incident wavelengths but the resolution varies with that wavelength across the wavelength band used.

One way of restating the relative performance of pulsed and CW sources that accounts for these problems is to say that, for monochromatic beam experiments, a pulsed source gains a factor of $1/c$ with respect to its mean flux by using time of flight methods provided the **required** incident wavelength resolution, $\Delta\lambda_{\text{res}}$ (i.e. the wavelength resolution we would use at a CW source) is related to the **required** incident band width, $\Delta\lambda_{\text{bw}}$ (i.e. the bandwidth that contains useful information) by the expression

$$\Delta\lambda_{\text{res}} = c \cdot \Delta\lambda_{\text{bw}} \quad (3)$$

Figure 3 illustrates two more potential problems. If the wavelength resolution changes across the wavelength bandwidth, the number of resolution elements that fit within the accessible bandwidth is more difficult to calculate because the effective duty factor also varies across the wavelength band used. In addition, it may turn out that measurements at each wavelength are not made with the required resolution. In fact, for a large region of the energy spectrum of neutrons used at a short-pulsed source such as LANSCE, IPNS or ISIS, the pulse duration scales as the reciprocal of neutron wavelength, rather than being a constant. While this effect is often useful — for example, in assuring that the fractional d-space resolution of a powder diffractometer at a short-pulsed source is independent of the actual d-spacing — it nevertheless destroys the simple argument given above for the comparison between pulsed and CW sources.

Less severe, but still a problem for detailed comparison, is the fact, illustrated in the middle panel of Figure 3, that the neutron spectrum at a pulsed source is not independent of neutron energy (nor does it even vary linearly with energy). Measurements that employ different incident wavelengths are thus made with different neutron intensities incident on the sample and have different statistical precision after a given measurement time. Each measurement does not have equal weight and so one must do better than simply counting the number of measurements made to compare relative source performances.

Several potential problems are not depicted in Figures 2 and 3. For example, if the neutron pulse is not a simple shape with well-defined sides like the triangles depicted in Figures 1 and 3 — if, for example, it has a long tail on one side — it is not possible in general to use a simple Rayleigh criterion to say how many resolution elements can be fitted within the band width. Even more insidious is the fact that monochromatisation at a CW source is often performed by Bragg reflection from a crystal, thereby introducing a correlation between the wavelength and direction of travel of a particular neutron. This correlation, which is absent when time of flight is used to determine neutron wavelength, has important ramifications, for example, in determining the resolution of a triple-axis spectrometer, especially for higher incident neutron energies. This so-called focusing effect may be advantageous or not, depending on the experiment to be performed.

Comparison of Short-Pulse and CW Sources

The shorter the pulse length of a source, the more important are the caveats described above in the comparison of CW and pulsed sources. To understand this, consider the pulse durations obtained with a high-intensity water moderator at a traditional short pulse spallation source (SPSS) such as LANSCE that are shown in Figure 4. From this Figure we see that, for a 20 Hz source, the effective duty factor, c , varies between a value of about 10^{-3} at low neutron energies to less than $2 \cdot 10^{-4}$ at high neutron energies. Because c is so

small, a large bandwidth is required if eqn (3) is to be satisfied for any useful wavelength resolution. Even if it is usable (i.e. even if eqn 3 is satisfied), this large bandwidth makes it far more likely that the problems depicted in Figures 2 and 3 will occur and make the comparison of pulsed and CW sources more complex than the simple-minded approach outlined at the beginning of this report.

As Figure 4 also shows, the pulse duration at a short pulse spallation source does indeed depend on neutron energy, so the problem depicted in the lower panel of Figure 3 renders the simple $1/c$ comparison false.

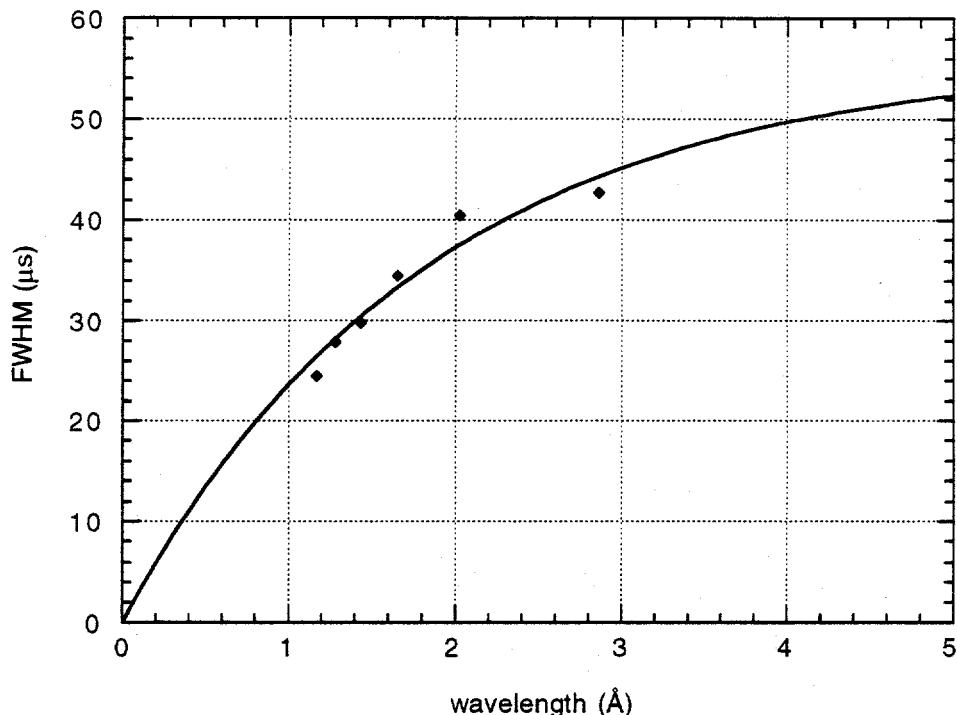


Figure 4: Pulse durations for the high-intensity water moderator at LANSCE

These arguments show that, even for the experiments that favour short-pulse sources, such as powder diffraction, it is not always true to say that the peak flux of a spallation source should be compared to the average flux of a CW source. Although one can use the entire bandwidth usefully, so that eqn (3) is rather well satisfied, the Devil hides in the details, such as those expressed by Figure 3.

Comparing Long-Pulse and CW Sources

Although simple comparisons between CW and short-pulse sources are very difficult to make using the simple-minded approach outlined at the beginning of this report, the situation is not as bad for long-pulse spallation sources (LPSS). In this case, c is between 0.05 and 0.1, so the incident bandwidth only needs to be 10 to 20 times the required resolution for eqn (3) to be satisfied and for the peak flux to be the relevant parameter for performance. Because a narrower band width is used at an LPSS, none of the problems

sketched in Figures 2 and 3 are as serious at an LPSS as they are at a SPSS. Furthermore, at a long-pulse source, the neutron pulse length tends to be almost independent of neutron energy, so that the effective duty factor and wavelength resolution are also almost constant, especially across the small wavelength band width that would be used at an LPSS.

In order to minimise the neutron pulse duration at a short pulse source, target designers go to great trouble to include neutron poisons and absorbers that prevent low-energy neutrons from getting into the moderators long after the proton pulse has struck the target [3]. While these efforts successfully keep the pulse short, they do so at the expense of its intensity. At a long-pulse source, where the proton pulse lasts a millisecond instead of a fraction of a microsecond, these absorbers can be eliminated to produce the full neutron flux that a coupled moderator has to offer. In the thermal neutron region, the average flux of a spallation source with coupled moderators is four to five times higher than the average flux of the same source equipped with decoupled moderators [4].

Figure 5 shows what can be achieved in this respect. While the average unperturbed thermal flux of a 1 MW spallation source with coupled moderators is about the same as that of an optimised 10 MW reactor [5], the performance of a cold neutron source can be greater because the lower heating of a spallation source allows cold-neutron moderators to be placed in a higher neutron flux. Using neutron transport codes that have been

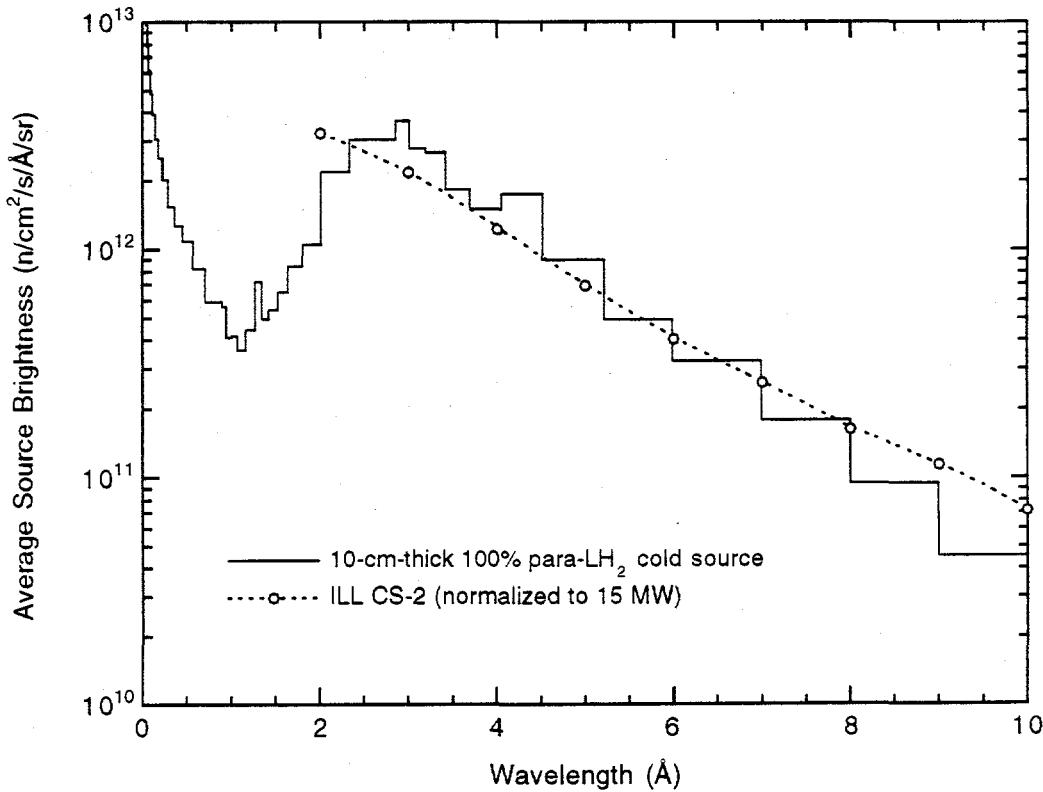


Figure 5: The calculated average thermal flux of a 1 MW, 60 Hz long-pulsed spallation source with 1 millisecond proton pulses is compared to the calculated flux of the second ILL cold source normalised to 15 MW reactor power.

extensively bench-marked [6], Gary Russell at Los Alamos has designed a neutron production target with cold moderators whose brightness for neutrons in the $2 - 10 \text{ \AA}$ range is essentially the same as would be obtained with the new ILL cold source with the ILL reactor power scaled to 15 MW.

Potentially, Figure 5 implies that if one can use time-of-flight methods to gain the full factor of $1/c$ ($\approx 10 - 20$) offered by the naive argument at the beginning of this report, a 1 MW LPSS would perform as well as a 150 - 300 MW reactor, provided the flux from such a reactor scaled linearly with its power over the range between 60 MW (the ILL power) and 300 MW. The question is, how much of the factor $1/c$ can one realistically hope to gain at a 1 MW LPSS?

The Performance of Specific Spectrometers at an LPSS

Roughly speaking, the argument we shall use to establish the performance of LPSS instrumentation goes as follows. We consider a spectrometer at an LPSS that has the same incident-beam wavelength resolution as one would use at a CW source, with time-of-flight replacing the monochromating element used at the CW source. Equal wavelength resolution is achieved by adjusting the length of the total flight path, which then fixes the incident neutron bandwidth (cf eqns 1 and 2). We then check to see whether eqn (3) is satisfied. The extent to which the usable bandwidth differs from that given by eqn (3) determines one factor by which the optimal gain ($1/c$) needs to be decreased. A second such factor derives from a comparison of the incident-beam collimation implied by the spectrometer length and the collimation that would be used at a CW source. Of course, we may also choose to match the resolution of a CW spectrometer by using a chopper to reduce the pulse length of a pulsed source. In this case, the duty factor, c , is reduced but so is the effective average flux, by the same factor. Provided eqn (3) is satisfied when a chopper is used to reduce ΔT , it is thus still the peak flux of the spallation source which determines performance.

The Effective Duty Factor of Long-Pulse Sources

In eqn (2), I assumed that the full time, T , between the beginning of successive pulses is available for counting. While this is a good approximation at an SPSS where $\Delta T \ll T$, it is not always adequate for long-pulse sources, as the time-distance diagram in Figure 6 demonstrates. To define the accessible wavelength band and avoid frame overlap, a band-defining chopper is needed on many spectrometers, such as the SANS and reflectometer considered below. At the beginning and end of each measurement frame, this chopper introduces a penumbra because the detector cannot "see" neutrons from the whole pulse duration. If the band-defining chopper is located half way between the neutron source and the detector, the effective counting time (the time during which the detector is fully illuminated on average) is reduced from T to $T - \Delta T$. In addition to this penumbra effect, the band-defining chopper also takes a finite time to open, shadowing part of the neutron beam as it opens and closes across the beam. For a source with $\Delta T \approx 1 \text{ msec}$, this closure time effectively reduces the counting frame by another $\Delta T \text{ msec}$ to about $T - 2\Delta T$, instead of the value T assumed in eqn (2). For a 60 Hz source with 1 msec pulses, the combination of these two effects imply that the potential gain from pulsing the source is reduced from 17 $\{ = T/\Delta T \}$ to 15 $\{ = (T - 2\Delta T)/\Delta T \}$.

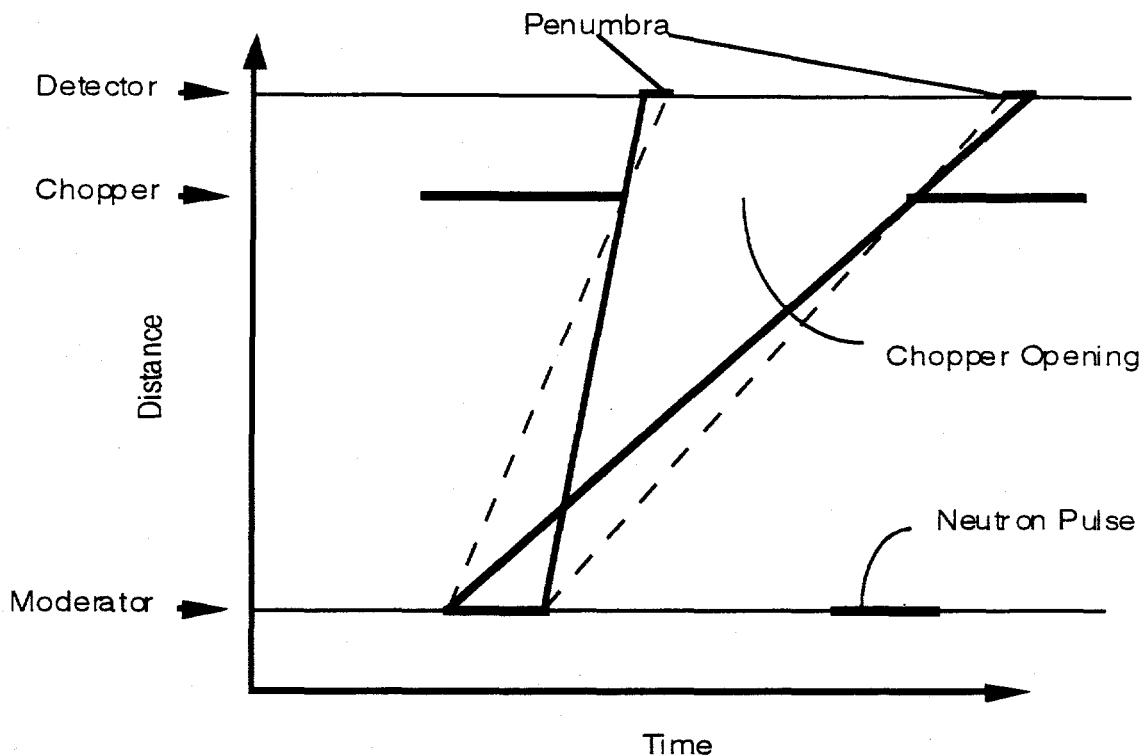


Figure 6: Distance-time diagram for a typical spectrometer. The vertical axis represents distance from the neutron source while the horizontal axis is time. The band-defining chopper results in a significant penumbra at the detector for a LPSS because the pulse duration is 5 - 10% of the time between pulses.

Low Energy Inelastic Scattering

The simplest case one can imagine is that of a chopper spectrometer for inelastic scattering such as the IN5 at ILL. In this case, one simply takes the entire spectrometer and places it in front of the LPSS without modification. The spectrometer sees the peak flux of the spallation source and so gains the full factor of $1/c$, provided the repetition rate of the pulsed source is the same as one would choose at a CW source. In fact, for a 60 Hz LPSS the repetition rate is somewhat lower than one might choose to use on IN5, where 120 Hz is likely to be more suitable. Thus the gain factor that results from pulsing the source is reduced to about $17/2 \approx 8$ compared with the full factor of $1/c$ (≈ 17 for a source with 1 msec pulses). Since the average flux of the 1 MW LPSS is about 4 times less than the ILL (cf Figure 5), this means that a low energy chopper spectrometer (such as IN5 at the ILL) would perform about twice as well at a 1 MW LPSS as it would at ILL.

Small Angle Scattering

A second relatively simple case is SANS. Here we assume for the moment that we are comparing the LPSS spectrometer to a spectrometer like D11 at the ILL with its velocity selector set up to give $\Delta\lambda/\lambda = 0.1$ and a total spectrometer length of L (m). For a 60 Hz LPSS, the time between pulses is given by $T = 1000/60$ msec. Thus, a spectrometer of

length L (m) will, according to eqn (2), have an available band width of $200/3L$ Ångstrøms. Including the effect of the band-defining chopper described above reduces this band width to about $60/L$ Ångstrøms. The required resolution of 0.1λ fits into this bandwidth $600/L\lambda$ times, so this is the gain obtained by pulsing the source. For a 20 m (total length) SANS instrument using 10 Å neutrons, this implies a gain factor from source pulsing of a factor of about 3. In light of the calculations displayed in Figure 5, a 20 m SANS instrument at a 1 MW LPSS would thus perform about 75% as well as an optimised D11 at the ILL in the sense that a measurement made with a spectrometers of total length 20 m and neutrons of 10 Å wavelength would take about 25% longer at the 1 MW LPSS.

From eqn (3), one sees that the wavelength resolution at the LPSS would be about 0.5 Å for 20 m machine at a source with 1 msec pulses, implying that a resolution better than that of D11 can be obtained if required, especially for neutrons scattered towards the edges of the detector. None of the problems (cf Figures 2 and 3) that make it difficult to compare SANS at CW and low duty factor pulsed sources are relevant to the comparison of CW and LPSS because the total wavelength band width used is scarcely greater at the LPSS than it is at a CW source.

Of course, this comparison assumes that the D11 velocity selector has perfect peak transmission and that D11 could be moved to the second ILL cold source, which provides a higher flux than the source actually viewed by D11. In the past these two factors have resulted in a D11 performance which is almost a factor of three less than that used to derive the comparison described above. I have been told that the recent replacement of the Brunhilde velocity selector has significantly reduced this penalty [7].

Reflectometry

One might think the performance comparison outlined above for SANS would also apply to reflectometry, but this turns out not to be the case. A reflectometer at a CW source normally uses a monochromator crystal to deflect a beam downwards on to the horizontal surface of a reflecting sample. The degree of monochromatisation achieved depends on the collimation before and after the monochromator and, in general, on the monochromator mosaic spread. For an instrument placed on a supermirror guide using a wavelength of 6 Å and a high quality pyrolytic graphite monochromator, the FWHM wavelength spread in the incident beam is about 1.5% (i.e. about 0.1 Å). From eqn (1), we see that this level of monochromatisation can be achieved at an LPSS with 1 msec pulses by using a 40 m total flight path. The incident flight path, just like that at the CW source, can be made of supermirror guide since the path length uncertainty introduced by this means has an insignificant effect on wavelength resolution. A 40 m instrument will use a 1.7 Å band width according to eqn (2), all of which is useful for reflection measurements. Since eqn (3) is satisfied by the required resolution and the useful bandwidth, the reflectometer gets the full gain from time of flight techniques, a factor of 15 for our hypothetical 6% duty factor source once the effect of the band-defining chopper is included. Thus, for the conditions considered here, such a spectrometer at an LPSS would capture data about 4 times faster than an equivalent spectrometer at the ILL and with the same resolution. Once again, in computing this factor we have ignored the less than perfect reflectivity of the monochromator used at a CW source and the fact that a reflectometer at such a source usually requires some type of finite-transmission filter to eliminate higher order scattering

from the monochromator. Both of these factors would, of course, tilt the comparison even further in favour of the LPSS.

The key difference between the reflectometer and the small angle scattering instrument considered above is in the required incident wavelength resolution — 1.5% at 6 Å for the reflectometer and 10% at 10 Å for SANS. The length of the SANS spectrometer is set by the required collimations and the desire to use a reasonably large sample. From eqn (1), such a long spectrometer automatically has a good wavelength resolution at a long pulse source — much better than the 10% required. Because of this, one cannot achieve the full gain from time of flight methods (i.e. eqn 3 cannot be satisfied). The sharper wavelength resolution needed by the reflectometer is more naturally matched to a long pulses source, allowing eqn (3) to be satisfied and the full time of flight gain to be realised.

The incident beam collimation used for reflectometers is typically ± 1 mrad, or about $\pm 5\%$ of the angle of incidence on the sample whose reflectivity is to be measured. On the face of it, this would seem to imply that the wavelength resolution used for reflectometry at a CW source (1.5% FWHM) is too good because it is not matched to the angular resolution of the incident beam. If this were the case, the performance of a reflectometer at a CW source could be improved substantially by degrading the wavelength resolution and increasing the neutron intensity. This would make the reflectometer look more like the SANS case and would indicate that the gain factor of 15 calculated above for reflectometry was over-estimated. However, the mismatch between incident angular collimation and wavelength resolution for reflectometry is deliberate. Because the angles of incidence and reflection are equal for specular reflection, the actual angular resolution is determined by the pixel size at the neutron detector rather than by the incident beam collimation. This resolution is typically 1% or less, implying that wavelength resolution in this range really is needed.

The Devil is in the Details

Arguments like those given above can be repeated for any spectrometer. One can carry out the *gedanken* experiment of transferring an instrument from a CW source to an LPSS, and replacing the monochromating element at the CW source by time of flight. Simple arguments can then be used to quantify the gains or losses from such a strategy. These arguments work because the useful wavelength band width is small. A similar analysis, for a pulsed source with an intermediate pulse length (≈ 100 μ sec), is documented in the proceedings of the Maria Laach workshop [8] held in support of the German SNQ project.

A summary of the results which these simple arguments give for the comparison of a 1 MW, 6% duty factor LPSS with the ILL second cold source is given in the following table for several instrument types.

gain with respect to
ILL cold source #2

| | |
|---|------|
| SANS (total length = 20 m, $\lambda = 10 \text{ \AA}$) | 0.75 |
| (total length = 10 m, $\lambda = 10 \text{ \AA}$) | 1.5 |
| (total length = 20 m, $\lambda = 5 \text{ \AA}$) | 3 |
| NSE (total length = 20 m, $\lambda = 10 \text{ \AA}$) | 0.75 |
| Diffuse scattering (e.g. D7 at ILL) | 4 |
| Inelastic TOF (e.g. IN5 at ILL) | 2 |
| Reflectometry | 4 |
| Backscattering — single monochromator | 0.5 |
| — several monochromators (e.g. MUSICAL) | 2 |
| — chopper monochromator | 4 |
| Laue diffraction for large unit cells | 2 |
| Triple Axis Spectrometer (e.g. IN14 at ILL) | 0.25 |

Even at this level, however, one has to be careful and not lose sight of details. For example, at an LPSS it will almost certainly be necessary to install T-zero choppers on all spectrometers to attenuate fast neutrons generated when protons strike the neutron production target. These choppers will need to be outside the bulk shield, 6 m or so from the target. At a 60 Hz source, the T-zero choppers will effectively block neutrons in the 10 \AA to 15 \AA range in addition to the fast neutrons they are designed to attenuate. If neutrons in this part of the spectrum are needed for a particular spectrometer, simple minded comparisons of the type detailed above will clearly not work.

There are also details that are not captured by the simple calculations. For example, coupled moderators give rise to neutron pulses with notoriously long tails, as shown in Figure 7.

The effects of these tails on spectrometer performance needs to be assessed and mitigating strategies need to be explored. Mitigation by the inclusion of neutron poisons in the moderators or by the use of composite reflectors will inevitably reduce the total neutron flux, as Figure 7 shows. Appropriate choppers can be used to eliminate pulse tails but careful consideration needs to be given to their design if they are to be placed in high radiation fields close to the neutron moderators where they are most useful.

The case of the triple axis spectrometer given in the above table also needs to be considered more carefully. While the naive "gain" factor of 0.25 may be correct for experiments with a high signal to noise ratio, the ability to gate the detector at a pulsed source can significantly reduce many sources of background. As I showed at the workshop on

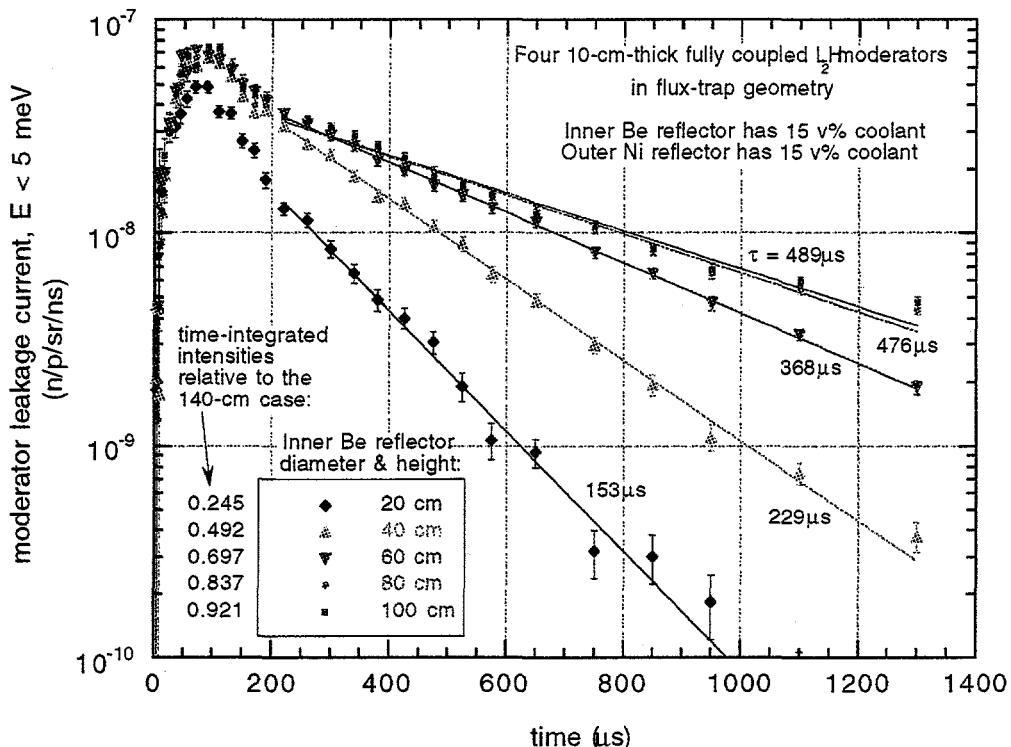


Figure 7: Neutron pulse shapes obtained with a coupled liquid hydrogen moderator for various dimensions of the surrounding beryllium reflector. For this calculation, the proton pulse was assumed to be a delta function in time.

instrumentation for the proposed European Spallation Source (ESS), this effect can be quite dramatic for hard experiments — a factor of 7 in favour of the LPSS for signal-to-noise ratios of 0.1, for example [9]. Further issues that needs to be considered for the triple axis spectrometer are the elimination of higher order neutrons by time of flight and the performance of augmented spectrometers such as RITA that are being pursued at Risø and elsewhere. In the latter case, if time of flight is used as a monochromating element, the crystal analyser can, in principal, be configured to perform a constant-Q scan if this is desired.

In making the comparisons listed in the above table, I have assumed that a 1 MW LPSS really can produce the same cold neutron flux as a 15 MW reactor. However, the results presented in Figure 5 are obtained from a calculation which does not include all of the engineering details of the target and moderators. While the largest effects (such as dilution of the target and the reflector by coolant) are included, smaller perturbations (such as absorption by the material of the target and moderator canisters) are not. The effect of these engineering realities, as well as possible gains from better target design need to be considered in detail.

Finally, it is not at all obvious that the optimum design of a spectrometer at an LPSS imitates the design at a CW source, even though such imitation makes comparisons between the two types of sources much easier. One might find, for example, that the performance of a SANS spectrometer could be improved at an LPSS by shortening the flight path, much as one does for low-Q diffraction at a short pulse spallation source. To

determine the advantages of using such strategies, it is necessary to consider the requirements of different experiments in detail.

Conclusion

The first conclusion I want to draw is that comparison of the performance of neutron scattering spectrometers at CW and pulsed sources is simpler for long-pulsed sources than it is for the short-pulse variety. Even though detailed instrument design and assessment will require Monte Carlo simulations (which have already been performed at Los Alamos for SANS and reflectometry), simple arguments are sufficient to assess the approximate performance of spectrometers at an LPSS and to support the contention that a 1 MW long-pulse source can provide attractive performance, especially for instrumentation designed for soft-condensed-matter science. Because coupled moderators can be exploited at such a source, its time average cold flux is equivalent to that of a research reactor with a power of about 15 MW, so only a factor of 4 gain from source pulsing is necessary to obtain performance that is comparable with the ILL. In favourable cases, the gain from pulsing can be even more than this, approaching the limit set by the peak flux, giving about 4 times the performance of the ILL. Because of its low duty factor, an LPSS provides the greatest performance gains for relatively low resolution experiments with cold neutrons. It should thus be considered complementary to short pulse sources which are most effective for high resolution experiments using thermal or epithermal neutrons.

References

1. F. Mezei, Editorial in *Neutron News* Vol. 5, No. 3; F. Mezei, Proc. ICANS XII. (RAL 94 - 025: Rutherford Appleton Lab. report); G. Bauer, Editorial in *Neutron News* Vol. 6, no. 1.
2. *Scientific Opportunities with Advanced Facilities for Neutron Scattering*, CONF-8410256 (available from the U.S. Dept. of Commerce)
3. G. J. Russell, H. Robinson, G. L. Legate, R. Woods, E. R. Whitaker, A. Bridge, and K. J. Hughes, Proc. ICANS IX (ISBN 3-907998-01-4: Schweizerisches Institut fur Nuklearforschung report; eds: F. Atchison and W. Fischer)
4. see ref 3
5. G. S. Bauer, Proc. ICANS VII (AECL-8488, Atomic Energy of Canada Ltd report; ed: S. O Schriber)
6. T. O. Brun, C. A. Beard, L. L. Daemen, E. J. Pitcher, G. J. Russell, and W. B. Wilson, Proc. ICANS XII, (RAL 94 - 025: Rutherford Appleton Lab. report)
7. Peter Lidner, private communication
8. Proceedings of the Workshop on Neutron Scattering Instrumentation for SNQ, Maria Laach, 3-5 Sept. 1984 (ISSN 0366-0885)
9. R. Pynn, in *Instrumentation and Techniques for the European Spallation Source* (RAL-92-040: Rutherford Appleton Lab. report)

Acknowledgments

Several friends and colleagues have contributed to my education in the matters discussed in this report. I would especially like to acknowledge conversations with Gunther Bauer and

Feri Mezei, whose help was indispensable. Several of my colleagues at Los Alamos either helped to clarify my ideas or corrected errors in the early versions of this report. In this respect I thank Rex Hjelm, Andre Michaudon, Rob Robinson and Phil Seeger.