https://doi.org/10.61092/iaea.c4vx-880h

IAEA
International Atomic Energy Agency

INDC(NDS)-0814
Distr. AC

INDC International Nuclear Data Committee

The Importance of Resonance Self-Shielding
Part 2

Dermott E. Cullen
1466 Hudson Way, Livermore, CA 94550, USA

Daniel Lopez Aldama
Center for Technology Applications and Nuclear Development
Havana, Cuba

Andrej Trkov
Jožef Stefan Institute
Ljubljana, Slovenia

August 2020

IAEA Nuclear Data Section
Vienna International Centre, P.O. Box 100, 1400 Vienna, Austria

Selected INDC documents may be downloaded in electronic form from
https://nds.iaea.org/publications
or sent as an e-mail attachment.
Requests for hardcopy or e-mail transmittal should be directed to
NDS.Contact-Point@iaea.org
or to:
Nuclear Data Section
International Atomic Energy Agency
Vienna International Centre
PO Box 100
1400 Vienna
Austria

Printed by the IAEA in Austria

August 2020

INDC(NDS)-0814
Distr. AC

The Importance of Resonance Self-Shielding
Part 2

Dermott E. Cullen
1466 Hudson Way, Livermore, CA 94550, USA

Daniel Lopez Aldama
Center for Technology Applications and Nuclear Development
Havana, Cuba

Andrej Trkov
Jožef Stefan Institute
Ljubljana, Slovenia

August 2020

Contents

Overview

Here using the TART Monte Carlo code we demonstrate the importance of neutron resonance self-shielding to calculate a neutron source problem, recently described as difficult to obtain consistent answers between Monte Carlo and deterministic codes. First, we present results for exactly the same system with and without self-shielding, in order to demonstrate the magnitude of the effect. Next, we present results obtained using three different methods to model self-shielding; Unshielded, Total Shielded and Multi-Band, to illustrate the magnitude of the differences between results. Lastly, based on the results in this report, we propose a future area of study to improve our self-shielding methods.

Details of our example problem

Part 1 of this paper [1] described the importance of resonance self-shielding in criticality calculations. Here in Part 2 we describe resonance self-shielding in source problem calculations. We have selected a problem described at a recent meeting on data for fusion neutrons applications as a difficult case to obtain agreement between Monte Carlo and deterministic calculations [2]. Below 20 MeV the FENDL-3.1d library is similar to ENDF/B-VIII.0, so we focus on the latter for convenience. We concentrate on one major concern that may lead to these differences in the results, namely the resonance self-shielding. In order to focus on self-shielding, we have modelled the problem to be as geometrically simple as possible: namely one, uniform sphere, of one material. Here is a complete description of our model.

Geometry: 100 cm radius uniform sphere (1 spatial zone)
Material: 57-La-139
Temperature: 293.6 K
Density: 6.2029 g/cm^3
Source: Monoenergetic isotropic at 20 MeV
Observable The neutron leakage spectrum from the surface of the sphere.

This completely describes the problem, hopefully in enough detail to allow any reader to uniquely reproduce this calculation. For 57-La-139 we used the ENDF/B-VIII.0 evaluation [3] as processed by the ENDF PREPRO codes [4], to,

1) LINEAR: Linearize all tabulated energy dependent cross sections.
2) RECENT: Reconstruct tabulated energy dependent cross sections from resonance parameters.
3) SIGMA1: Doppler broadened cross sections to 293.6 K.
4) GROUPIE: Calculated multi-group TART 616 group constant, and multi-band parameters in the unresolved resonance region (URR).

The 57-La-139 293.6 K SIGMA1 results are freely available on-line as part of POINT 2018 library [5] through the website **http://RedCullen1.net/Homepage.new**, for use by anyone who would like to reproduce our calculations; we encourage as many people as possible, using as many different transport codes (both Monte Carlo and deterministic) to try this simple test problem.

In this paper we first describe the 57-La-139 ENDF/B-VIII.0 evaluation in detail, which will help us better understand why we chose this evaluation, and how its features define our results. We next present results obtained for a number of models of self-shielded and unshielded cross sections, using the TART Monte Carlo code [6]. All of these results used exactly the geometrically simple problem defined above, and differ solely in how the neutron cross sections were defined:

1) Continuous energy cross sections, with multi-band unresolved resonance region self-shielding.
2) Multi-group cross sections, with and without self-shielding.
3) Multi-band (2-band) cross sections, with unresolved resonance region self-shielding.

Hopefully by the end of this paper the reader will appreciate the importance of neutron cross section self-shielding, and the magnitude (penalty) of the error introduced if self-shielding is not included. **We hope to show that self-shielding must be included in any realistic neutron transport calculation.** We close by using the results of this paper to suggest the subject of our next paper, closely related to this one, but specifically addressing modelling of the self-shielding in the unresolved resonance region.

57-La-139 ENDF/B-VIII.0 Evaluated Data

In order to understand the results of our test problem, and in order to allow any reader to reproduce the results presented here, we must uniquely define the evaluated data used. For this test problem we used only one evaluation, 57-La-139 ENDF/B-VIII, as shown here [5]. This evaluation has a resolved resonance region (RRR) up to 20 keV, and an unresolved resonance region from 20 keV to 100 keV. Note the large (n,2n) and (n,3n) cross sections near 20 MeV (our source energy), and the strong capture resonance near 72 eV; these play important roles in defining the leakage.

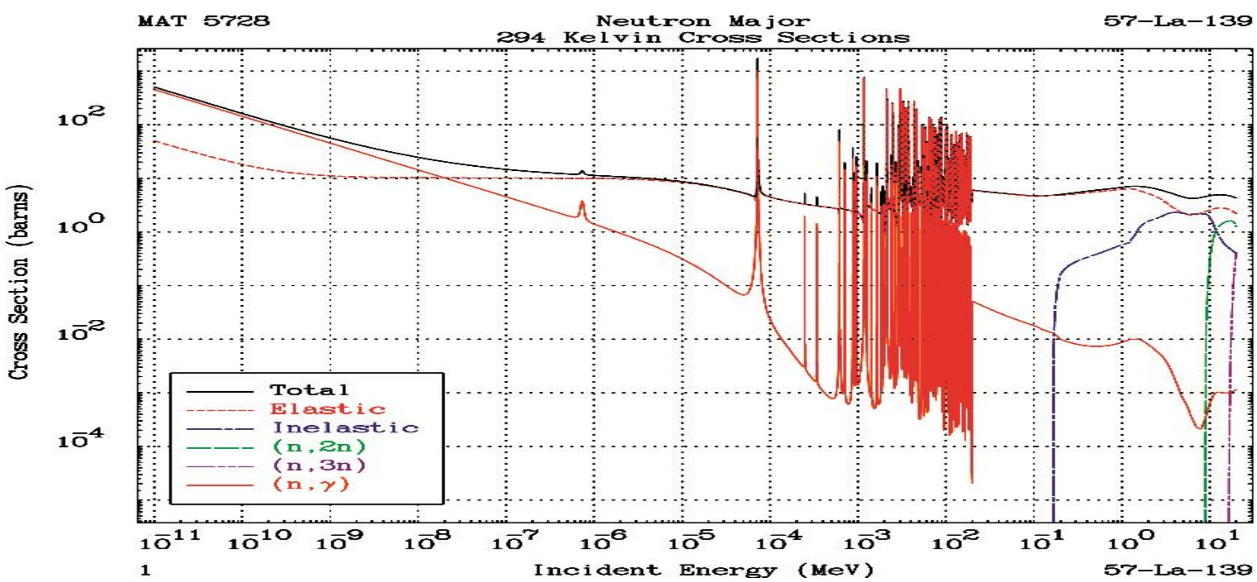

This evaluation was Doppler broadened [4] to room temperature, 293.6 K [5]. Temperature has a very important effect on the relatively narrow resonances in the resolved resonance region (RRR). As shown here the resonance peaks differ between old (0 K) and room temperature (293.6 K) data by up to over a factor of 14 (over 1400 %). In order to reproduce the leakage results presented here it is imperative to use exactly the evaluation Doppler broadened to 293.6 K [5]

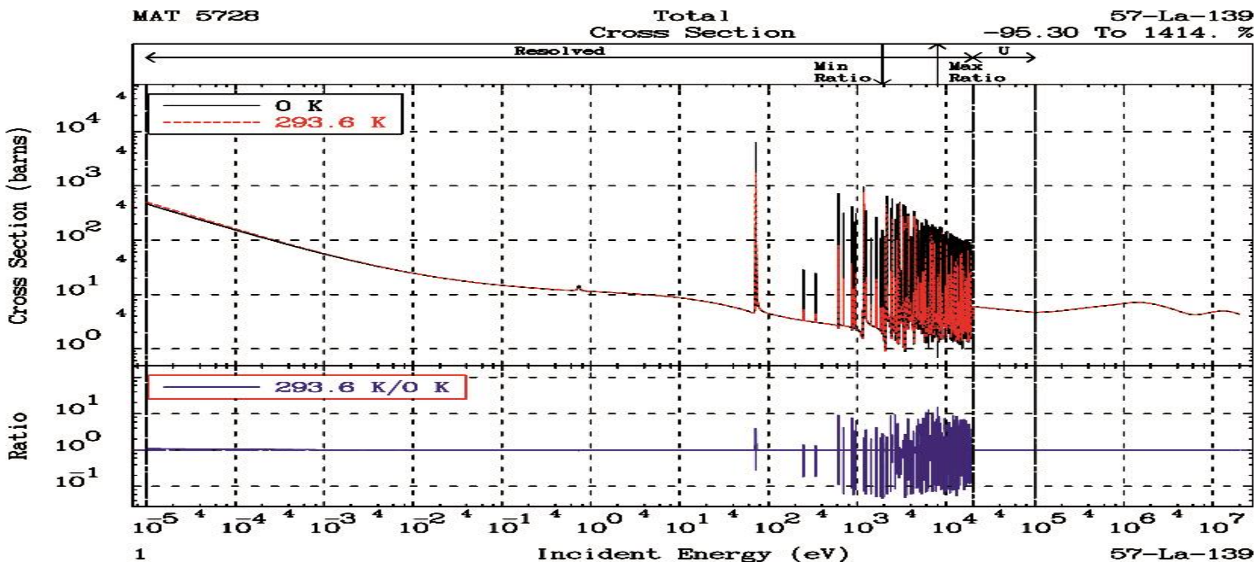

This evaluation also includes a classic design flaw that is worth noting, because it affects the results that we will be viewing. I say classic, because over the last 50 years I have seen the same defect in many ENDF formatted evaluations. With any ENDF formatted evaluation that includes a resolved resonance region it is important to include a sufficient number of resonances with a peak above the upper energy limit of the resolved resonance region, to allow the tail, or wings, of the resonances to extend down into the resolved region, and help to define a baseline, or the lower limit of the sum of resonances near the upper energy limit of the resolved resonance region. In this case the resolved resonance region extends up to 20 keV, **but there are no resonances defined above about 19 keV**; this is the classic ENDF error. The result is that in this evaluation the capture cross section approaching from below 20 keV, is physically far too small; based on the average capture just above 20 keV, the value just below 20 keV is a factor of about 1,000 too low = not 1,000 % - a factor of 1,000 – creating a unphysical "hole" in the capture cross section. The elastic minimum is "saved" by potential scattering, but even here the lack of resonances above 19 keV results in an elastic that is about a factor of two too low. I mention this here because when we later in this report look at the leakage we should expect to see a non-physical result just below the 20 keV resolved-unresolved resonance region boundary; I predict a "bump" in leakage due to this evaluation's cross sections being too low in this energy range; I will repeat: Unphysically too small. **We hope this serve as a WARNING to ALL evaluators: Your job includes extending the resonance sequence to establish baseline cross sections; failure to do so can lead to unrealistic results.**

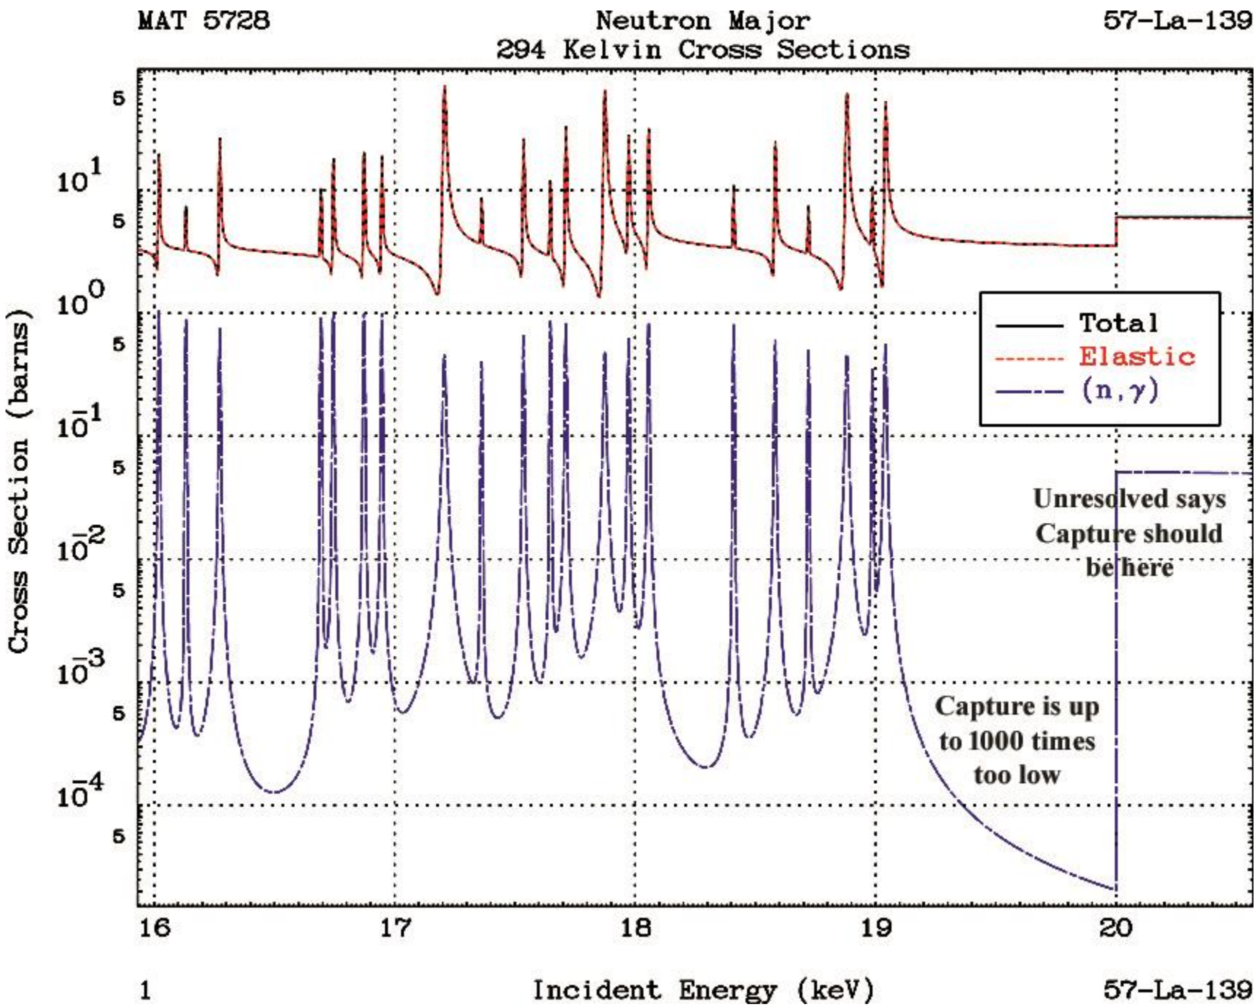
MAT 5728
Neutron Major
294 Kelvin Cross Sections
57-La-139
Total
Elastic
(n,γ)
Unresolved says Capture should be here
Capture is up to 1000 times too low
Cross Section (barns)
Incident Energy (keV)
1
57-La-139

Here is an overall view of the 57-La-139 total, elastic, and capture, TART 616 [7] multi-group averaged, unshielded and fully shielded cross sections, over the entire energy range [4]. I'll mention just a few points we should note because they will affect the leakage from our example sphere of 57-La-139. First is the very strong resonance near 72 eV; here we can see the unshielded total is close to 200 barns, and the shielded is close to 10 barns, 20 times smaller. Next, notice the unresolved region moderate self-shielding between 20 keV and 100 keV; moderate that is compared to the self-shielding in the resonance region below 20 keV.

Next, note the extremely small capture group average in the group just below 20 keV; again, let me stress that this is unphysical, and completely due to the flaw in the 57-La-139 ENDF/B-VIII.0 evaluation. I mention this because this "hole" in the cross section may show up as a "bump" in our calculated leakage; any such "bump" is not real, and strictly due to the flaw in this evaluation.

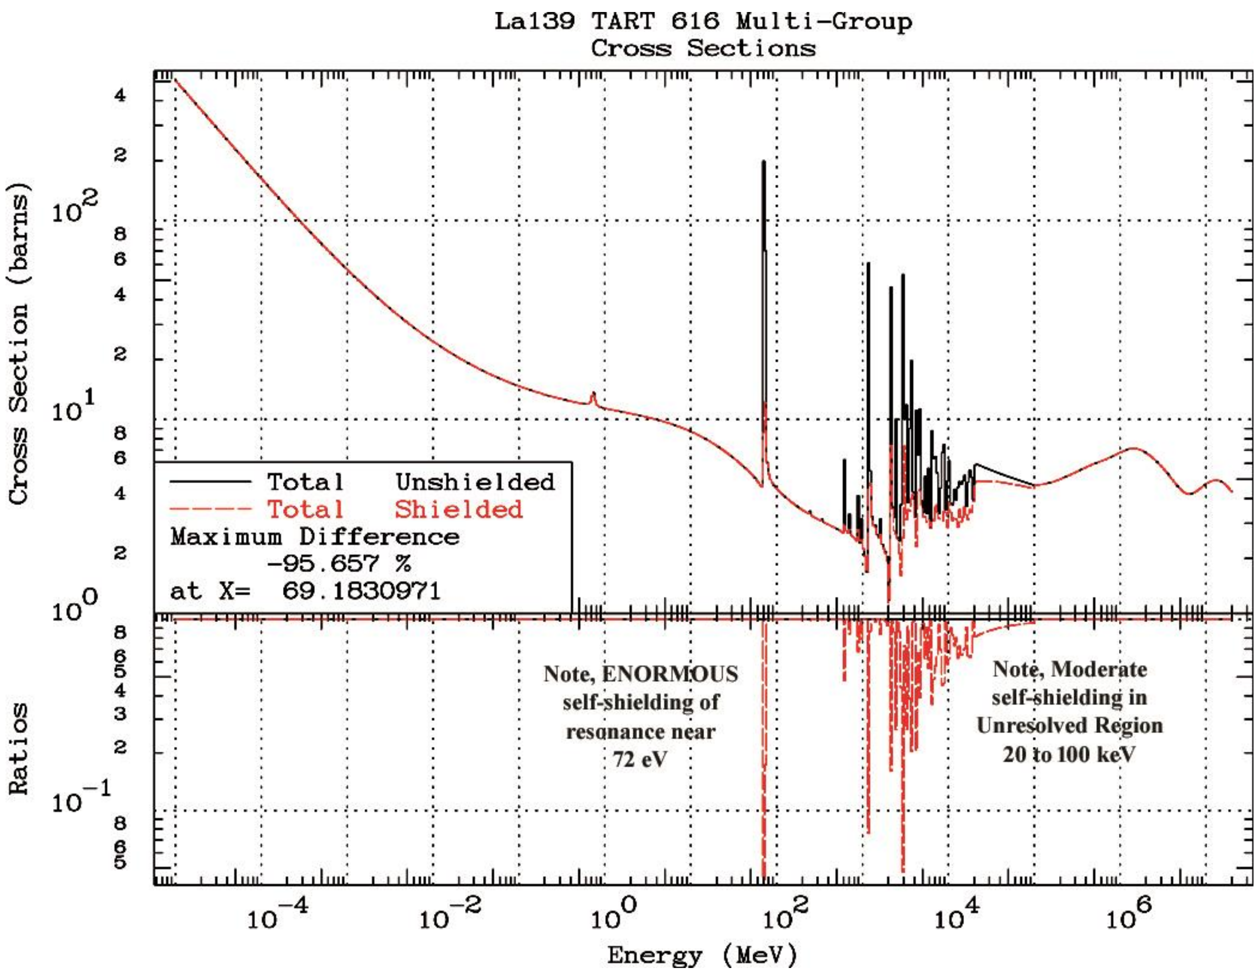

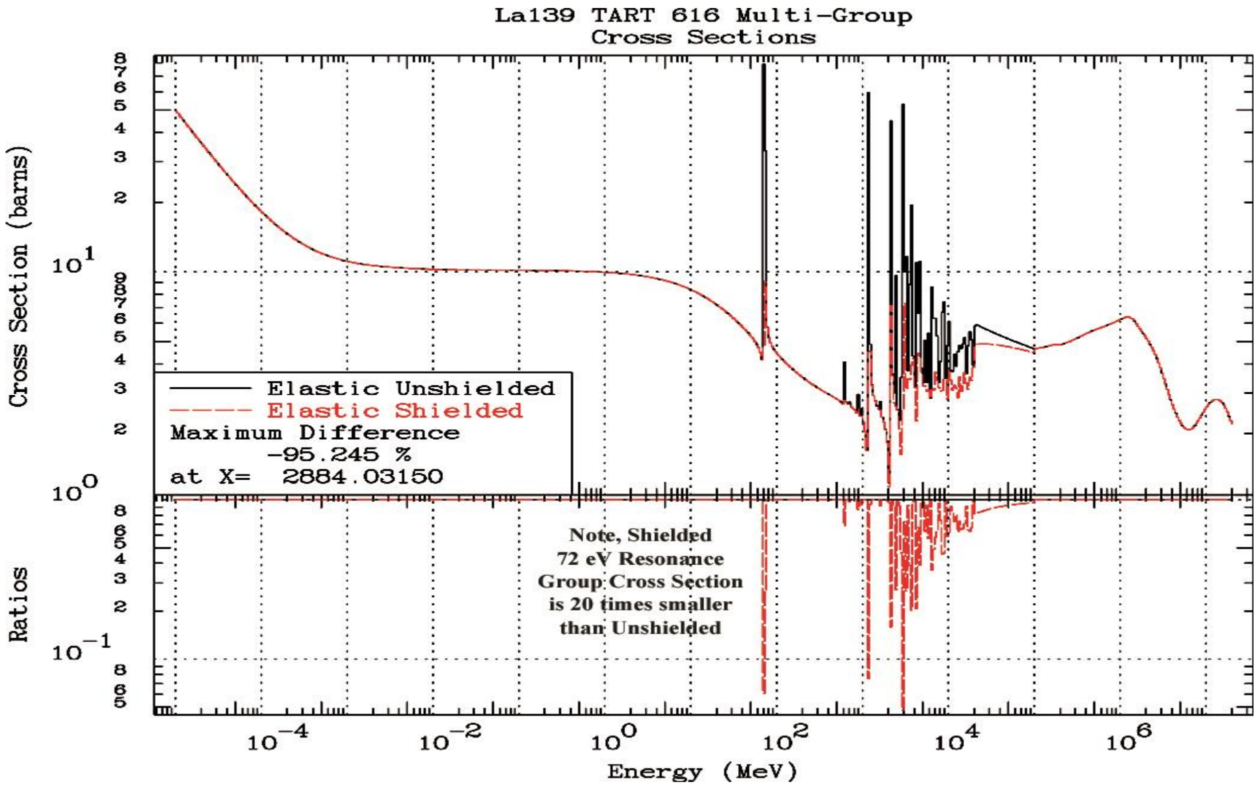
La139 TART 616 Multi-Group
Cross Sections
Cross Section (barns)
Ratios
Energy (MeV)
Elastic Unshielded
Elastic Shielded
Maximum Difference
-95.245 %
at X= 2884.03150
Note, Shielded
72 eV Resonance
Group Cross Section
is 20 times smaller
than Unshielded

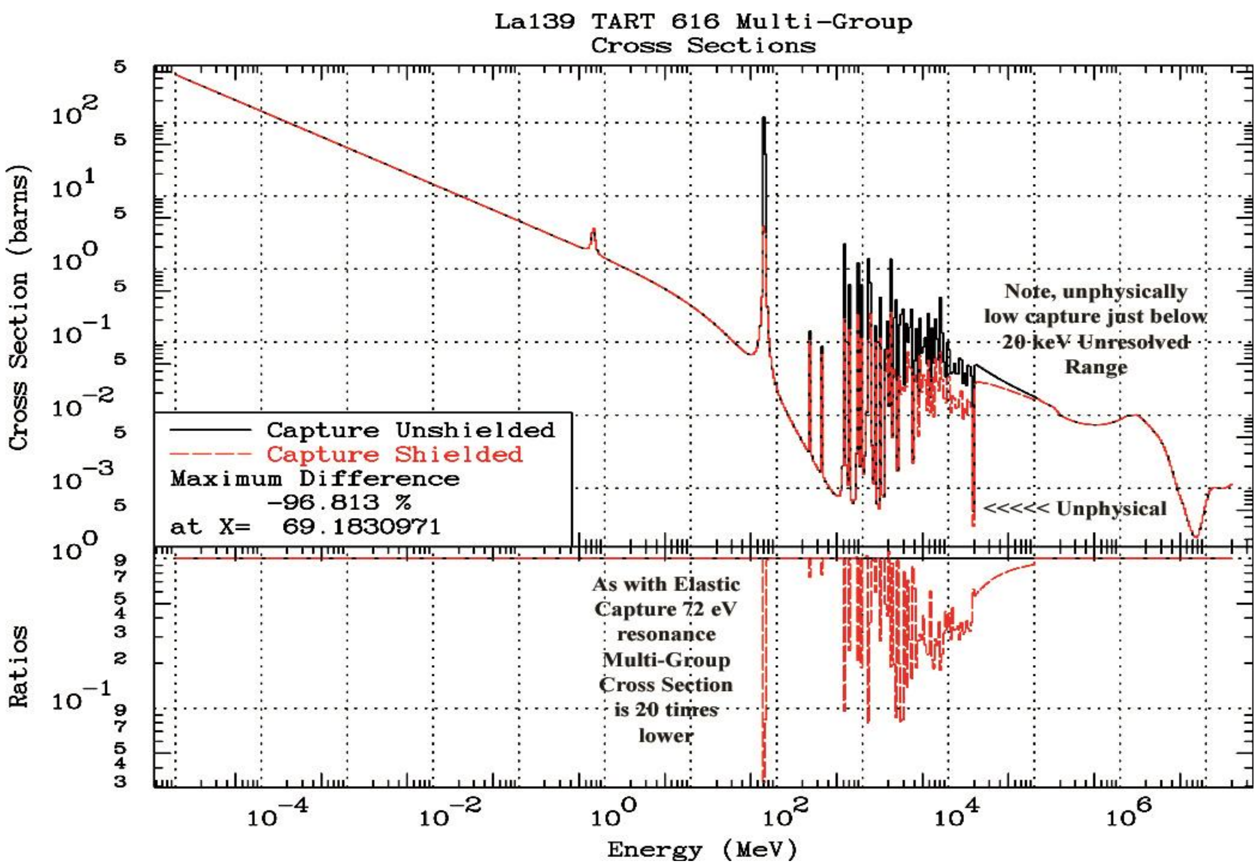
La139 TART 616 Multi-Group
Cross Sections
Cross Section (barns)
Ratios
Energy (MeV)
Capture Unshielded
Capture Shielded
Maximum Difference
-96.813 %
at X= 69.1830971
Note, unphysically
low capture just below
20 keV Unresolved
Range
<<<<< Unphysical
As with Elastic
Capture 72 eV
resonance
Multi-Group
Cross Section
is 20 times
lower

Another important feature of this evaluation are the spectra of (n,2n), and to a lesser degree (n,3n), secondary neutrons that will be produced by the monoenergetic 20 MeV we are using. In this evaluation these secondary neutron spectra are tabulated for a series of discrete incident neutron energies. Above roughly 400 keV of the secondary energy, the neutron shape of each spectrum is similar to a MeV-like Maxwellian, but below they are a rather crude linear function of secondary energy extending all the way down to zero energy.

These spectra will play a dominant role in defining the leakage spectrum. The total mean free path (MFP) for a neutron of 20 MeV is roughly 8 cm, so we expect few neutrons to transport 100 cm at 20 MeV and leak from the sphere at 20 MeV. We expect virtually all of the source neutrons to collide at 20 MeV, which will produce a spectrum of (n,2n) and (n,3n) secondary neutrons across the entire secondary neutron energy range; again I will mention that at higher energy this spectrum will be similar to a Maxwellian. The net effect of our monoenergetic 20 MeV source, will be similar to a Maxwellian neutron source created near the centre of the sphere (within a few MFPs), which will then transport and preliminarily scatter (elastic and inelastic) and be captured, as they find their way toward the surface of the sphere to define the leakage we measure in our test case.

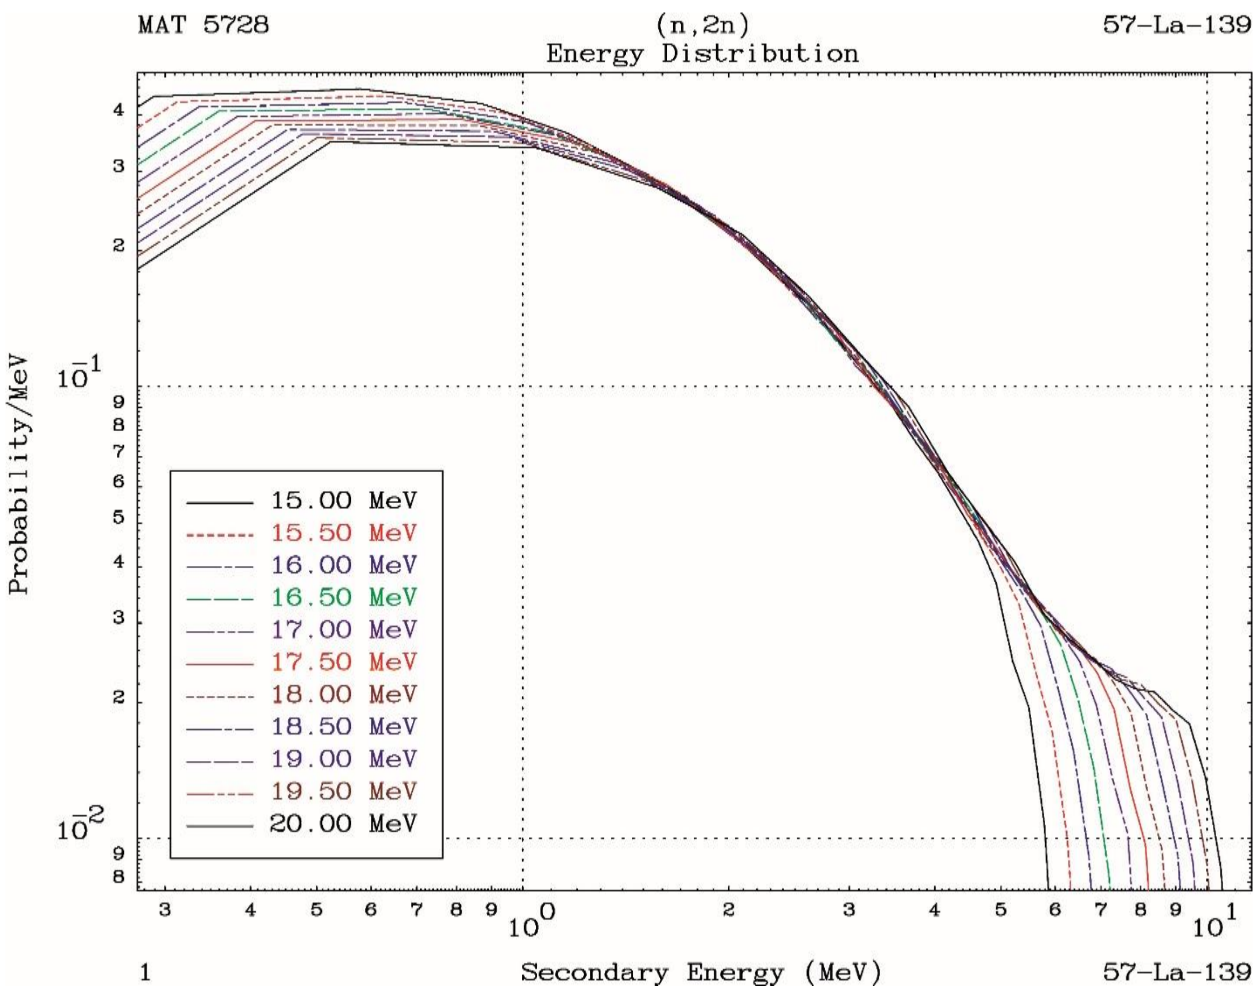

Here are preliminary results to illustrate the importance of La139 capture resonance near 72 eV; it is so strong and wide that few neutrons can scatter across it, greatly reducing the leakage spectrum below this energy, i.e., less than 1 in a million neutrons leak below roughly 100 eV, as we can see in the below figure. Because of this effect we will ignore the lower energy range for the remainder of this report.

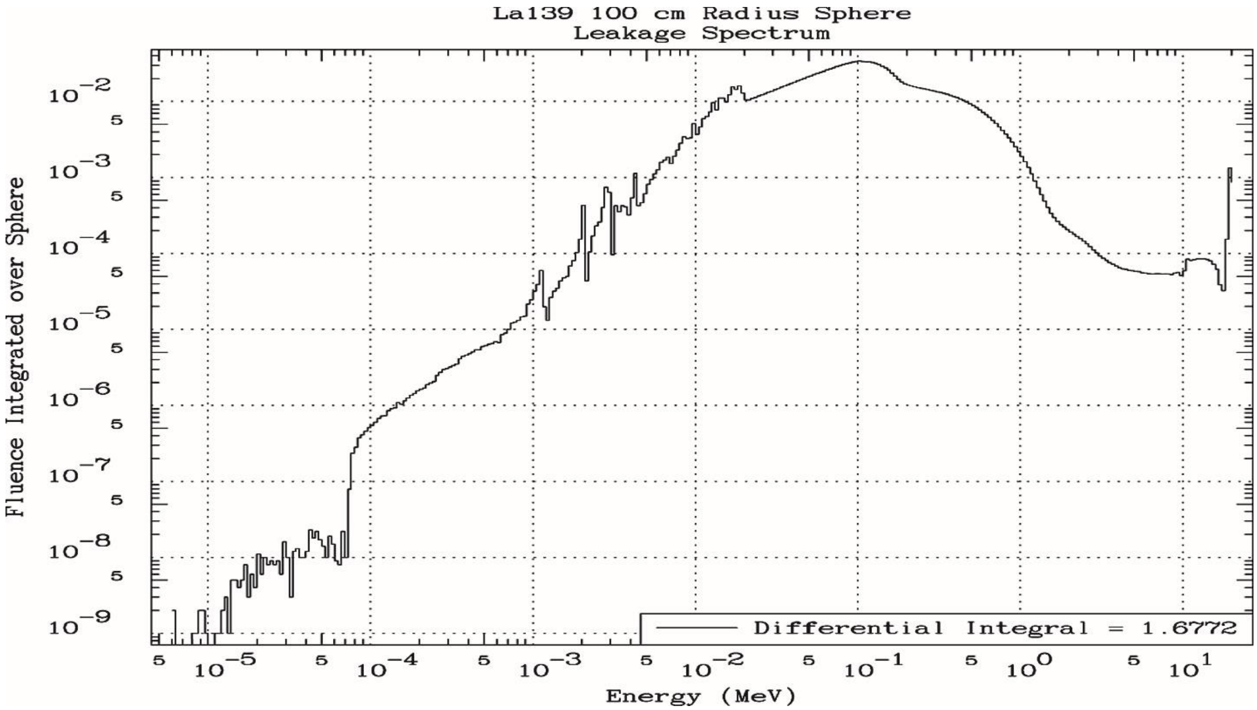

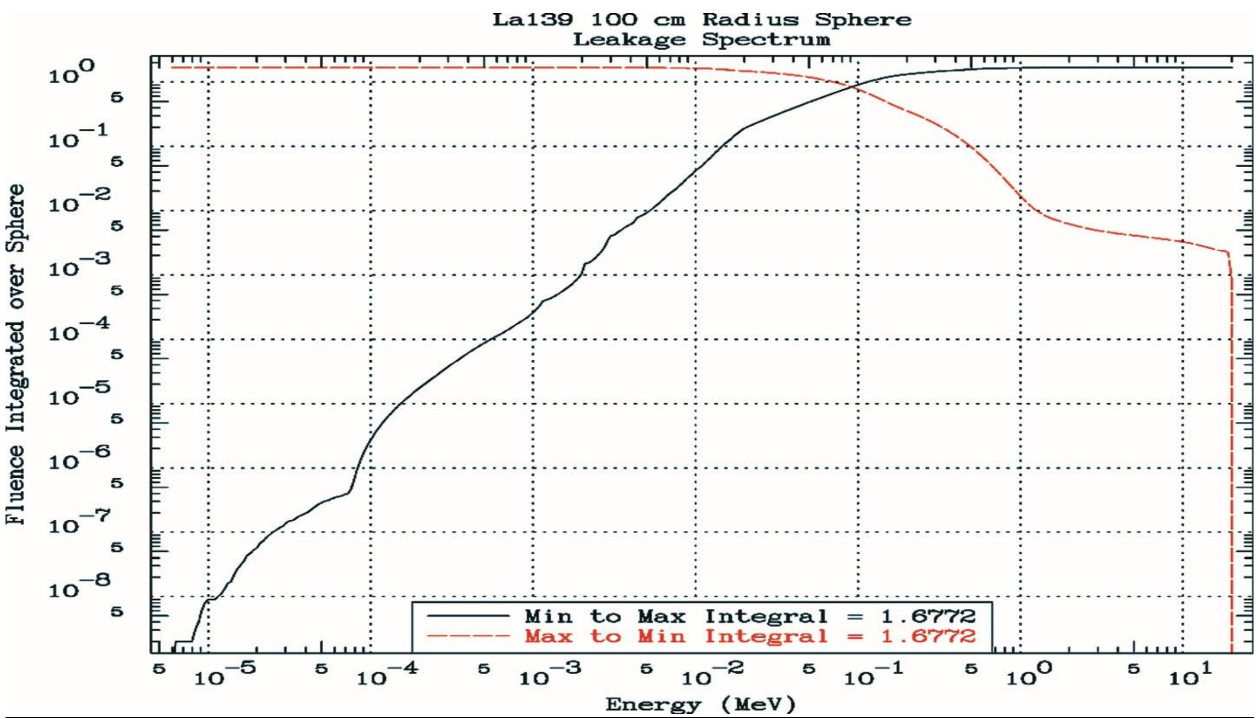

The lower part of the figure above represents cumulative integrals of the leakage spectrum from zero to E, and from E_{max} to E, respectively.

Here's a quick look comparing what's going on inside the sphere to the leakage on the surface. These two quantities are in different units; for comparison purpose, to show the energy dependent shape, each is normalized so that the integrals over energy and over space (the whole interior of the sphere or the surface of the sphere) are normalized to unity.

With the monoenergetic source at 20 MeV the high energy flux is quite high, but at very high energy the flux is close to the neutron source at the origin, so that very little leaks from the surface, i.e., low leakage. With decreasing energy the neutrons spread out from the source, and by 100 keV to 200 keV, reach a peak and then flux and leakage appear to be in equilibrium from that point on down to low energy. With spectra like these this would appear to be a perfect case to test the importance of resolved region (RRR), that extends up to 20 keV, and unresolved resonance region (URR) that extends from 20 keV up to 100 keV, i.e., near the peak of the leakage.

Note, the peaks and valleys of flux and leakage throughout the resolved resonance region, below 20 keV. If you compare these to the cross sections shown above you will see that the peaks in flux and leakage correspond to the valleys between resonances; this is classic neutron resonance self-shielding, where **the distance to collision (flux) varies inversely to the total cross section.**

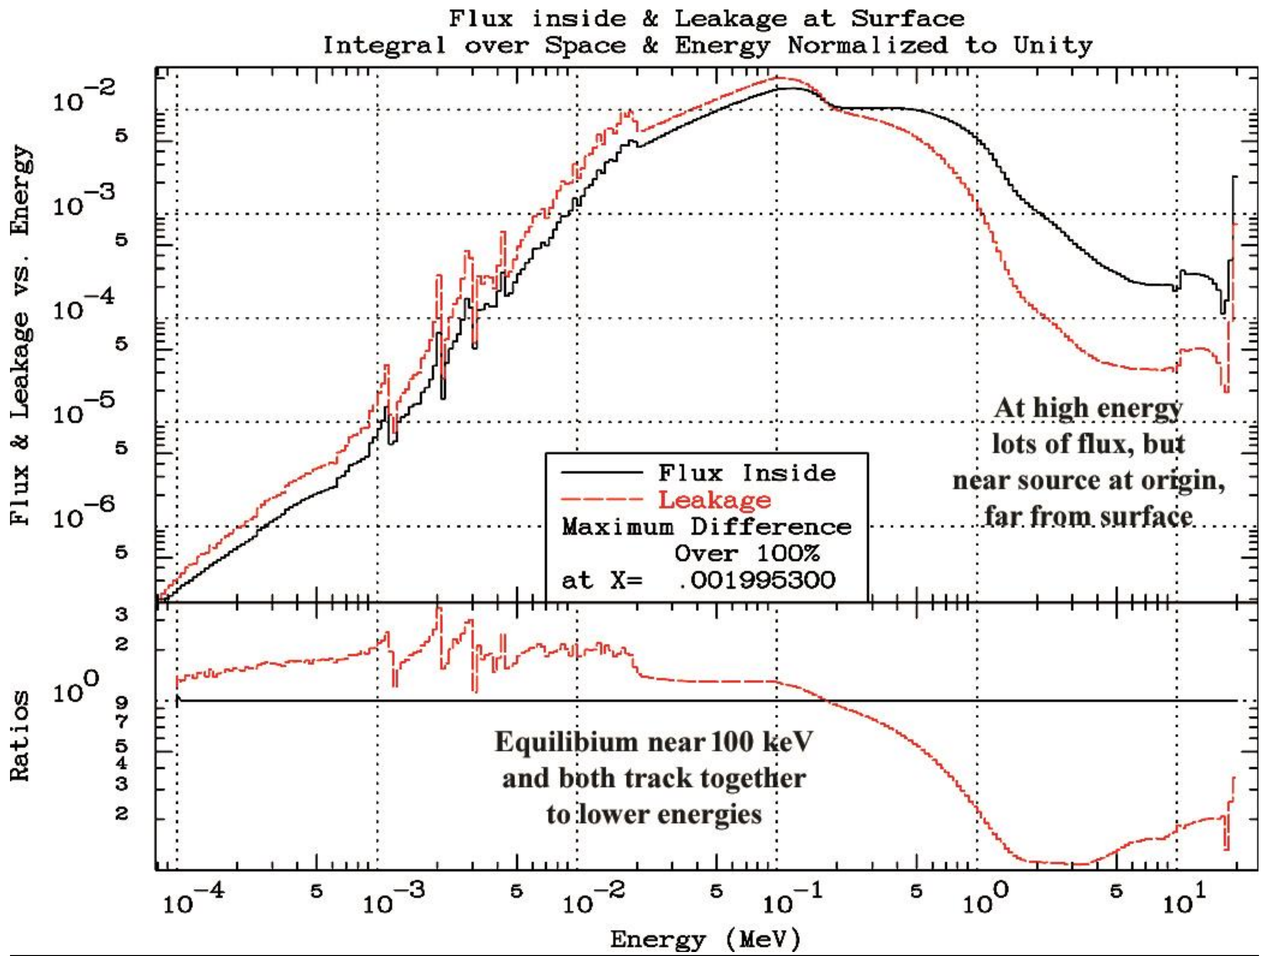
Flux inside & Leakage at Surface
Integral over Space & Energy Normalized to Unity
Flux & Leakage vs. Energy
Ratios
Energy (MeV)
Flux Inside
Leakage
Maximum Difference
Over 100%
at X= .001995300
At high energy
lots of flux, but
near source at origin,
far from surface
Equilibium near 100 keV
and both track together
to lower energies

Monte Carlo Statistical Uncertainty

Often users make the mistake of assuming that if they are using Monte Carlo the answers are as accurate as possible, because Monte Carlo allows us to use the most precise models, in terms of geometry and accuracy of the nuclear data, e.g., continuous energy cross sections rather than multi-group.

This assumption ignores the effect of the uncertainty in our Monte Carlo answers due to statistics; even the most precise model that will give us an accurate answer if we run our calculations long enough, may still include significant statistical uncertainty in any real case that we run a finite number of samples.

To illustrate this point below I present results for the leakage from the surface of our La139 sphere 100 cm radius, due to a point, monoenergetic 20 MeV neutron source at the origin. Here I present results using 100 million (10^{+8}) source neutrons compared to the results using ten time as many, 1,000 million (10^{+9}) source neutrons. I use the TART 616 tally bins, which are defined as 50 bins per energy decade, equally spaced in lethargy, which means that the leakage shown is the leakage integrated over each bin (not normalized by the width of the bin) and effectively represents the lethargy spectrum. In this form we can "see" how many counts (tallies) we predict in each bin per source neutron. Due to the large (n,2n) and (n,3n) cross sections at 20 MeV in La139, the integral of the leakage per source neutron exceeds unity, i.e., about 1.677. Combined with the capture within the sphere we actually end up tracking about twice the number of source neutrons.

Surely 100 million (10^{+8}) samples, to say nothing of 1,000 million (10^{+9}), are overkill for such a simple problem, and yet we still see difference in some bins of over 20%; even the integral slightly differs. Again, let me stress that both cases use exactly the same model and all differences shown here are solely due to statistics, i.e., 10^{+8} versus 10^{+9} samples.

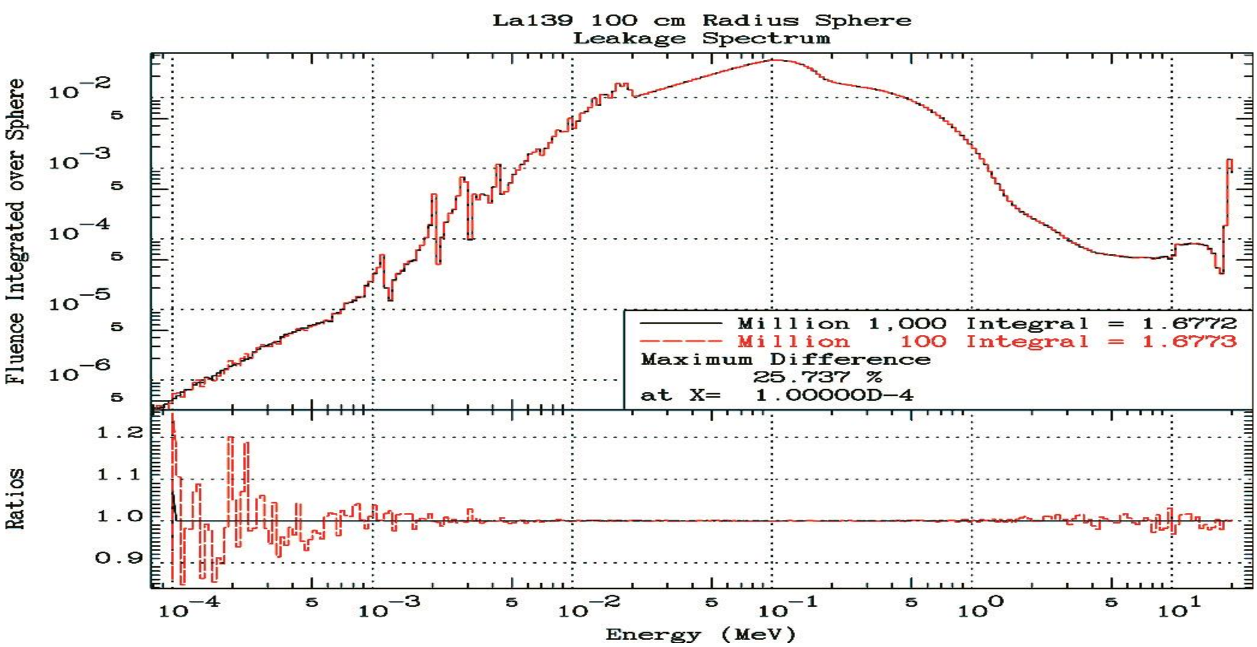

By looking at the integral(s) of the leakage we can get a better idea of why there is such a large difference due to statistics, even for such large sample sizes. Consider the integral from low energy to high. This shows that the probability of leakage below 1 keV is 10^{-4}, and below 100 eV is 10^{-6}. For a sample of 100 million (10^{+8}) this means only 100 neutrons will leak below 100 eV and 10,000 below 1 keV. Very crudely assuming that convergence varies as the reciprocal of the square root of the sample size, S/SQRT(N), we should expect: for 100 samples about 10% difference and 10,000 samples about 1%, which is very roughly what we see for the integral shown below between 100 eV and 1 keV. This result is for the integral; when we consider that the differential results shown above are further subdivided into 50 bins per energy decade, you should not be at all surprised to see differences of up to 20% in some bins.

From the combination of the two integrals shown below we can see the 99% (all but 1%) of the leakage occurs between about 7 keV and 1 MeV (where each integral crosses 0.01). Hopefully these differential and integral results shown here will help the reader focus on what is important AND how well we should expect different models to affect results shown later in this paper. **Monte Carlo is not perfect, and requires investing enough resources to insure accurate answers.**

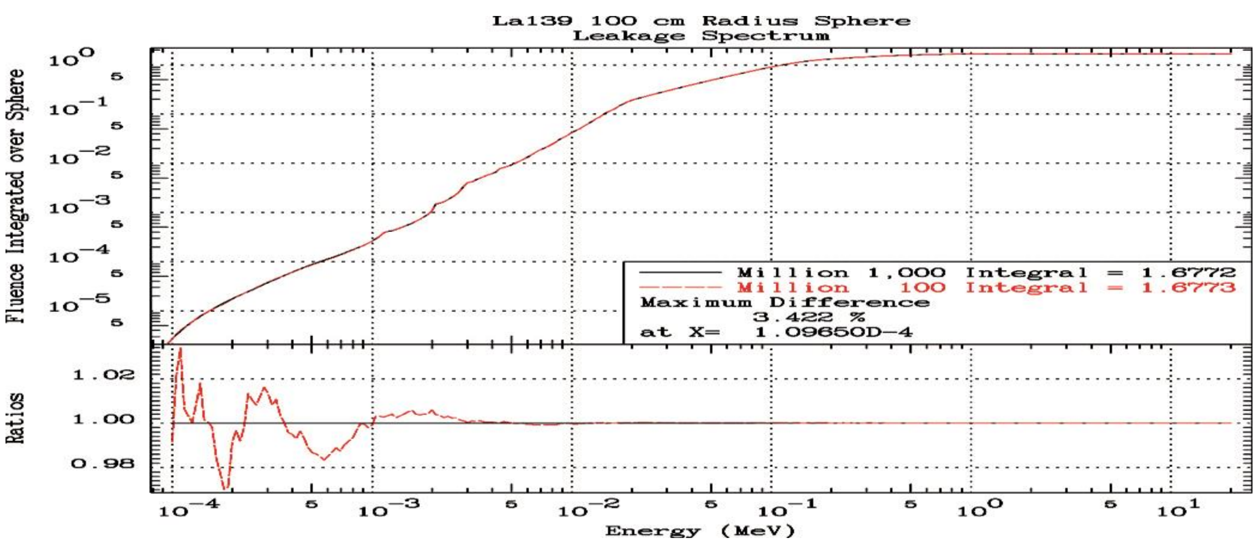

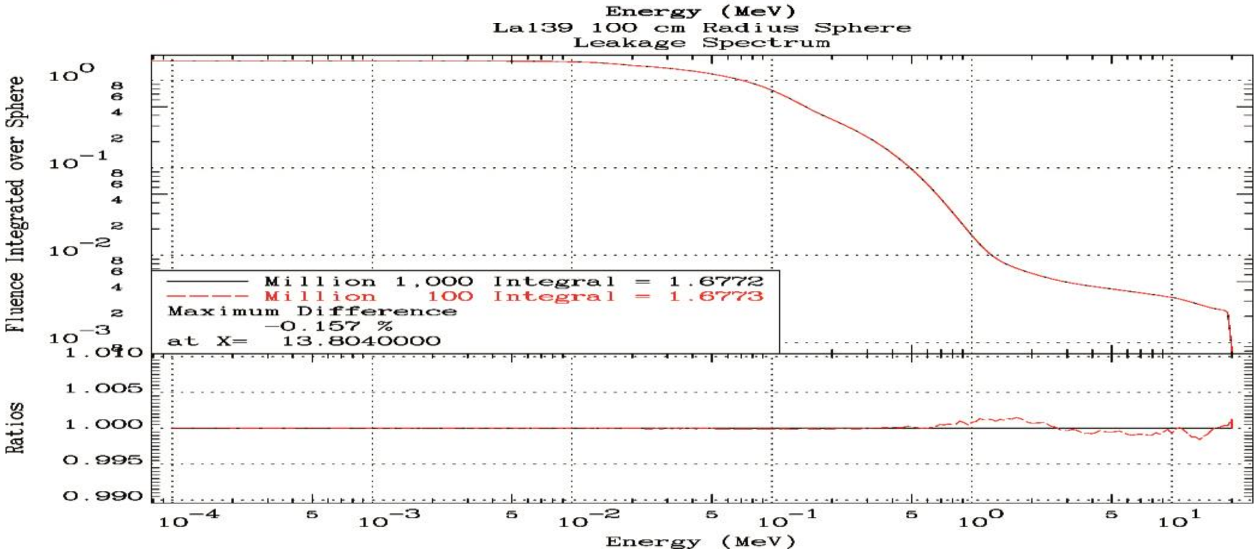

Model Dependent Convergence

The above results illustrate the speed of convergence using exactly the same model, versus the number of neutron samples. The above results are for about the worst case to achieve convergence, as far as the number of degrees of freedom we must sample, since it uses continuous energy cross sections. For example, in a 616 multi-group calculation our Monte Carlo code need only sample 616 different multi-group total cross sections. Whereas using continuous energy cross sections our code must sample many different cross sections, e.g., for fissile materials this can involve many thousands of energy dependent tabulated cross sections; in extreme cases even hundreds of thousands.

It is worth illustrating that the speed of convergence can be improved, using simpler models that still meets the needs. For example, below I compare results using 1,000 vs. 100 million neutron samples, first using continuous energy cross sections (the same as shown above), to results using multi-band (2 band) cross sections. Statistically both give the same integral answer to 4-digits, but the convergence of the multi-band is faster due to reducing the number of degrees of freedom it must sample. **This also illustrates that users of deterministic codes (multi-band) can potentially achieve the same accuracy as Monte Carlo codes.**

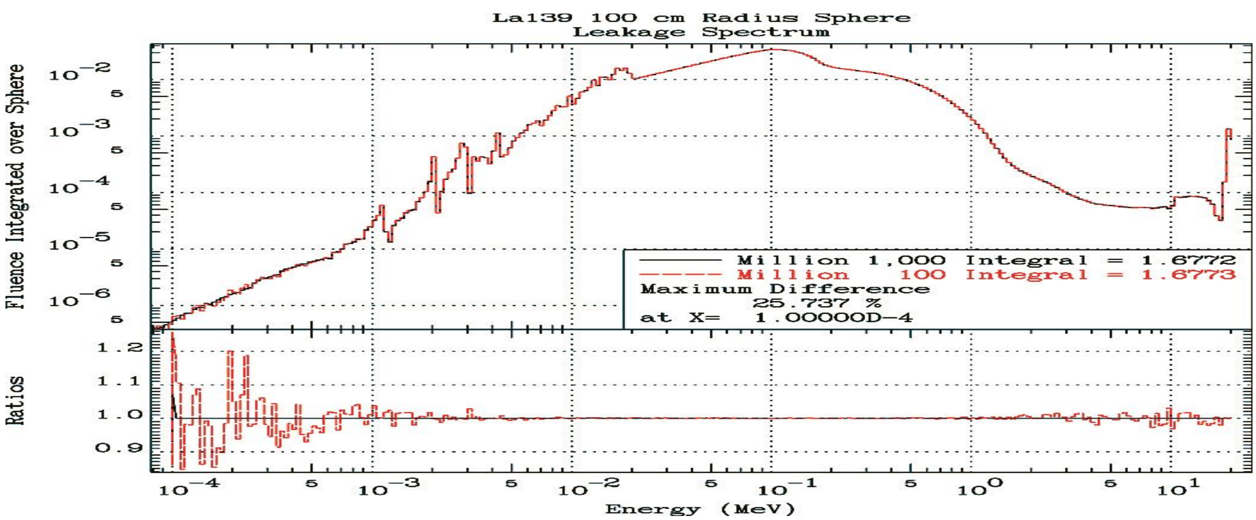

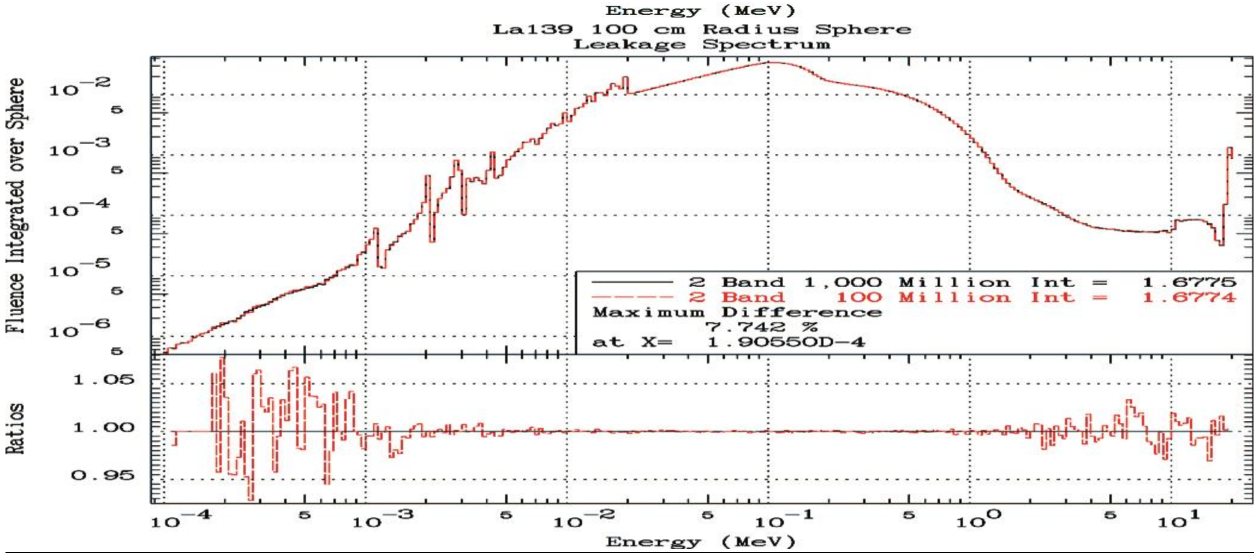

Continuous vs. Unshielded

ALL results below in this report used 1,000 million (10^{+9}) neutron samples to assure statistical accuracy. To clearly "see" how important self-shielding is, here I compare TART's BEST estimate of the correct answer, using continuous energy cross sections, and multi-band self-shielding in the unresolved region (URR), to TARTs answer using 616 group unshielded cross section. This is a very controlled test, using only one code, TART, one evaluation 57-La-139, one spatial model; everything is exactly the same except for how the cross sections are defined: Continuous versus unshielded group averages.

Starting from the monoenergetic neutron source at 20 MeV, we see close agree between the two results down to 100 keV; in this energy range there are no resonances and the generally "smooth" cross sections do not require self-shielding. Below 100 keV we see progressively more self-shielding effects; even by 20 keV, the lower energy limit of the unresolved resonance region (URR), the unshielded value is only 70 % of the shielded value (**clearly demonstrating the importance of self-shielding the unresolved resonance region)**. Below 20 keV in the resolved energy region we see ENORMOUS differences near the 57-La-139 resonances, where the unshielded cross sections can be over a factor of 10 too large = let me repeat that: not 10%, a factor of 10 too large, resulting in deep minimum in the leakage. The net effect is that the integrals of the leakage are: Continuous 1.6772, Unshielded 1.5646, **about 6.7% lower**, completely due to self-shielding. But with all of the fluctuations due to competition between scatter and capture it is difficult to judge which energy ranges are important, as far as defining the overall spectrum integral.

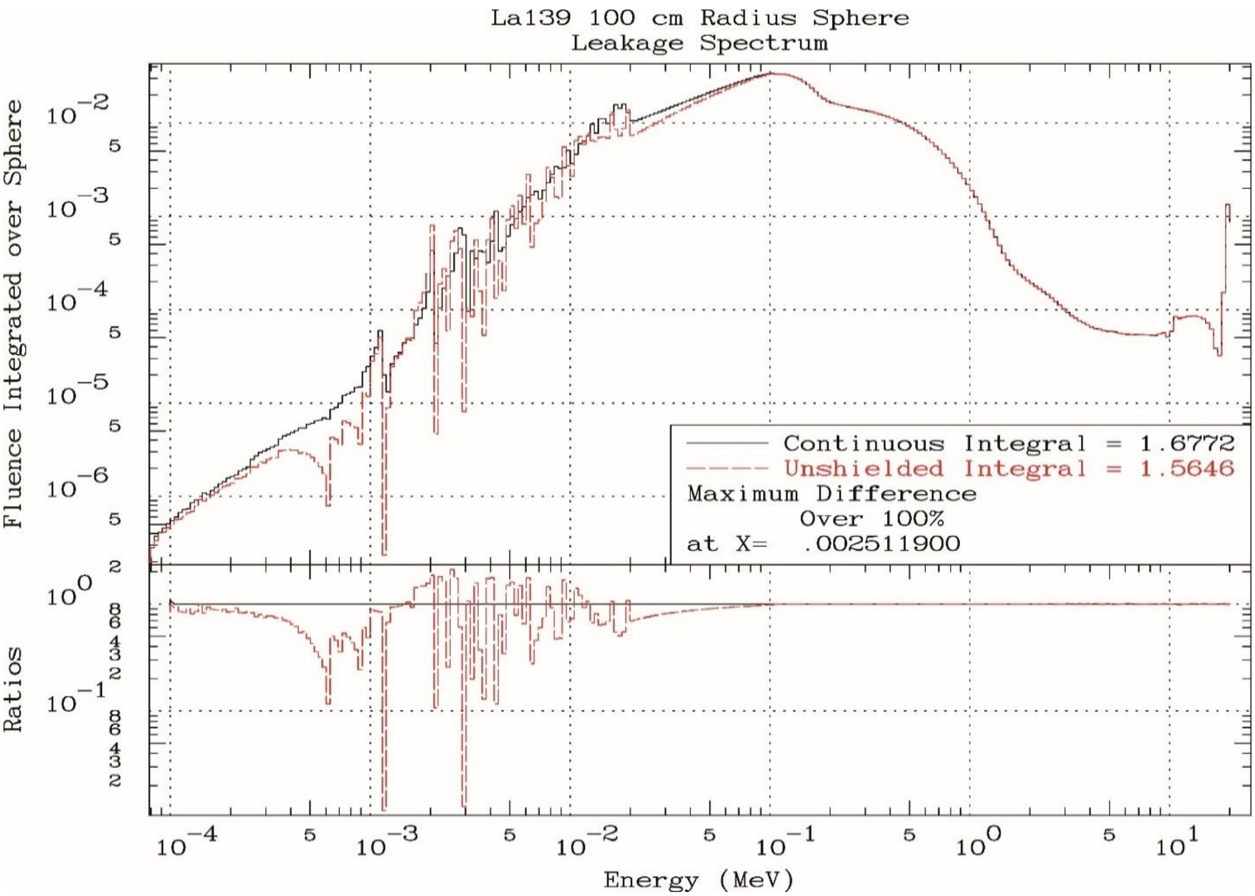

From the integrals we can see more clearly what energy ranges are important. From the integral from low to high energy we can again see many fluctuations at low energy, making it difficult to judge what the net effect of these fluctuations ultimately has on the overall integral. From the integral from high energy to low we see extremely close agreement from 20 MeV down to 100 keV, the upper energy limit of the unresolved resonance region (URR). Between 100 keV and 20 keV, the lower energy limit of the unresolved resonance region (URR) , the unshielded integral become lower than the continuous by 6.8 %, virtually the entire integral difference between the two spectra.

This difference clearly demonstrates the importance of self-shielding in the unresolved resonance region (URR), e.g., here it is responsible for virtually all of the difference in the two spectra.

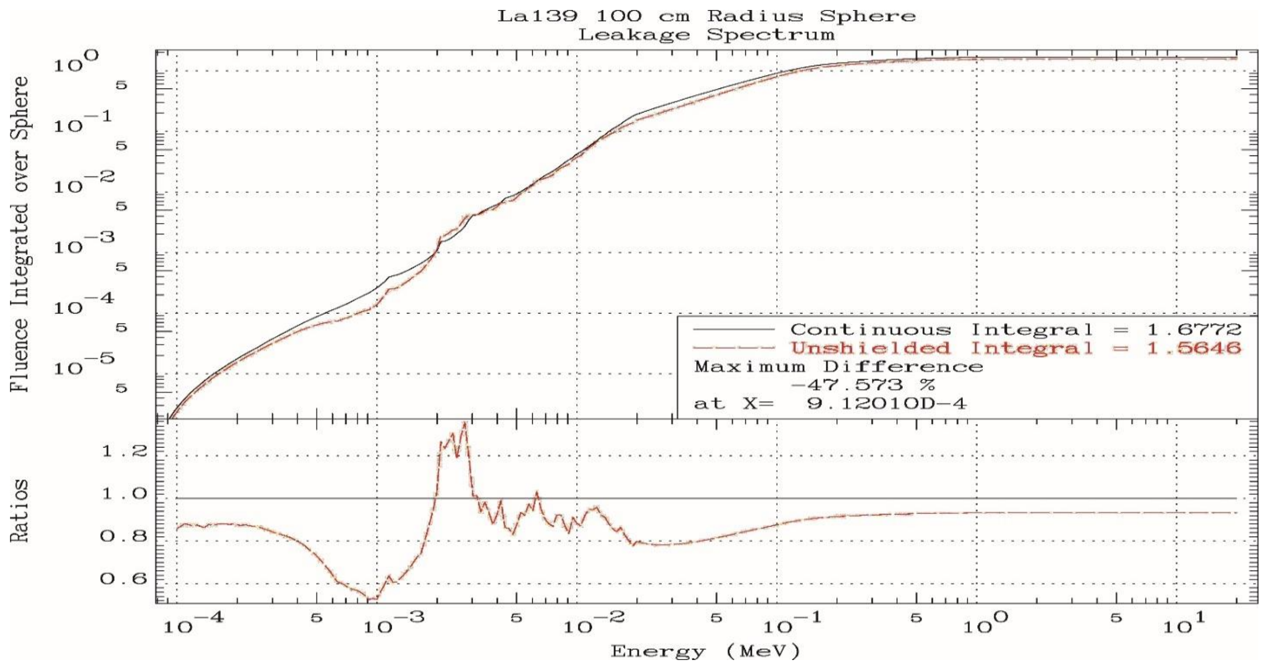

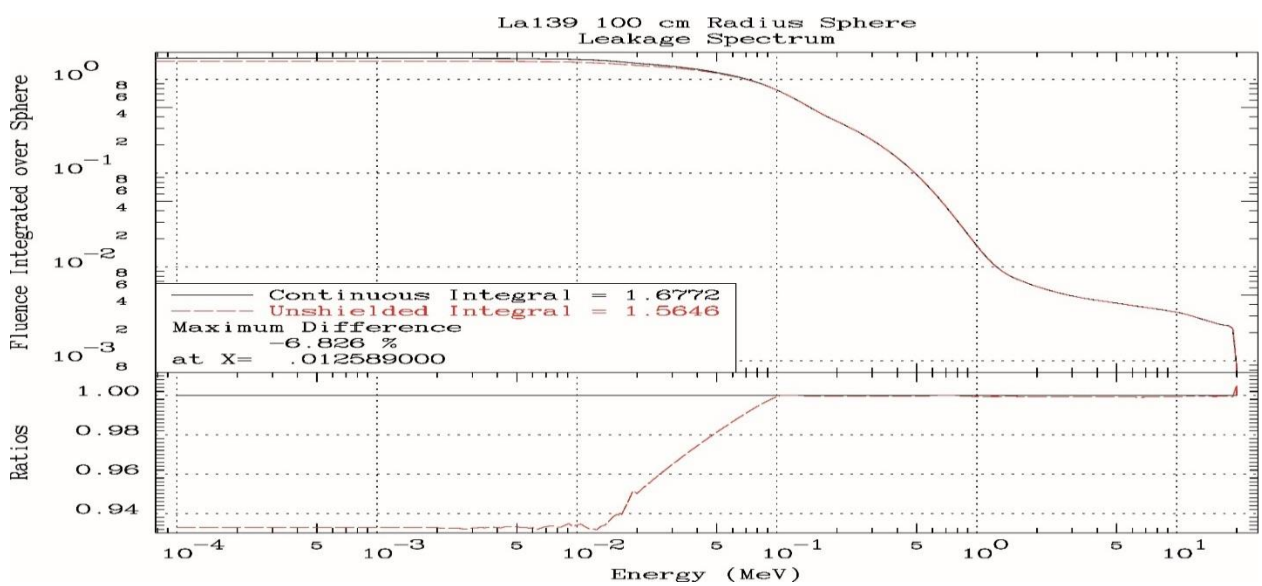

Continuous vs. Totally Shielded

The only difference between the above and below results is that for the above results **Unshielded** TART 616 group cross sections were used, and for the below results **Totally Shielded** (sigma0=0), TART 616 group cross sections were used, including in the unresolved resonance region.

Above 100 keV the unshielded and shielded results are similar, because there are no resonances and self-shielding is not required. Below 100 keV Self-shielding the cross sections definitely improves the results. In the unresolved region between 100 keV and 20 keV we start to see differences between the continuous and totally shielded results, but there is a definite improvement compared to using unshielded cross sections. Below 20 keV in the resolved region we still see differences in some tally bins up to about 37%, which is a major improvement over using unshielded where we saw differences in excess of a factor of 10 (1000%). There is a very definite improvement in the integral: Continue 1.6773, Totally Shielded 1.6738, **about 0.20% difference**; as a reminder, for Unshielded 1.5646, about 6.7% lower.

Depending on your application this 0.2% difference in the integral may to good enough to meet your needs.

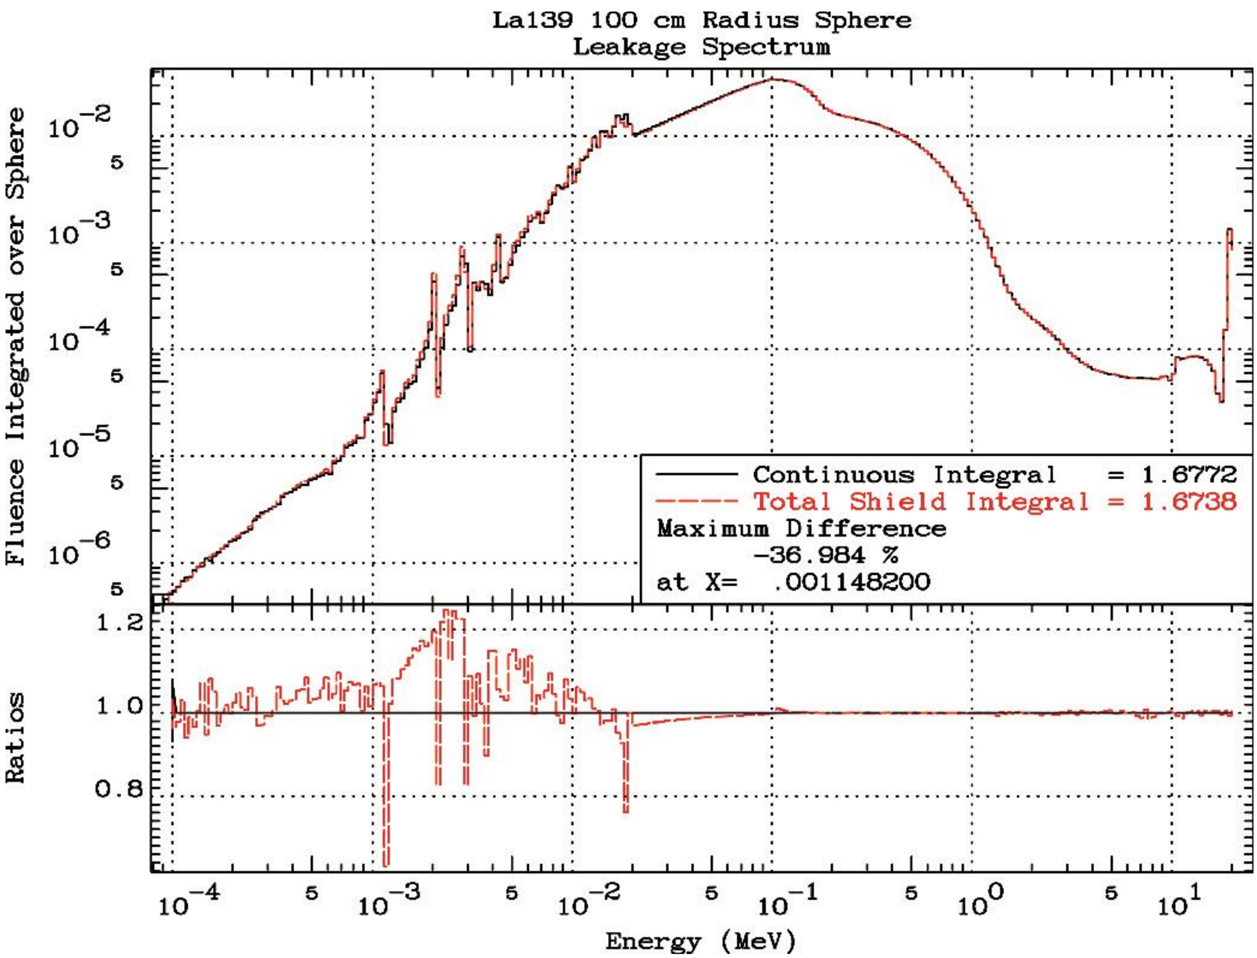

Continuous vs. Multi-Band (2-bands)

If the above totally shielded results do not meet your needs, you may need multi-band (2-band). Compared to the above totally shielded results, using multi-band results extends the agree down to 20 keV; this is because in the unresolved region the continuous model uses the same multi-band data. Below 20 keV the individual tally bin differences are reduced at about 12 % (shielded 37%) and multi-band does a good job tracking through the stronger resonance. The big improvement comes in the integral: Continuous 1.6772, Multi-Band 1.6775, only about **0.018% different; the integrals are statistically identical**. Here is a comparison of the three different results,

Method	Integral	% Difference
Continuous	1.6772	
Unshielded	1.5646	6.7%
Totally Shielded	1.6738	0.20%
Multi-Band	1.6775	0.018%

Hopefully one of these methods meets your needs. If not, sorry but deterministic methods probably will not work for use, i.e., you probably need Monte Carlo.

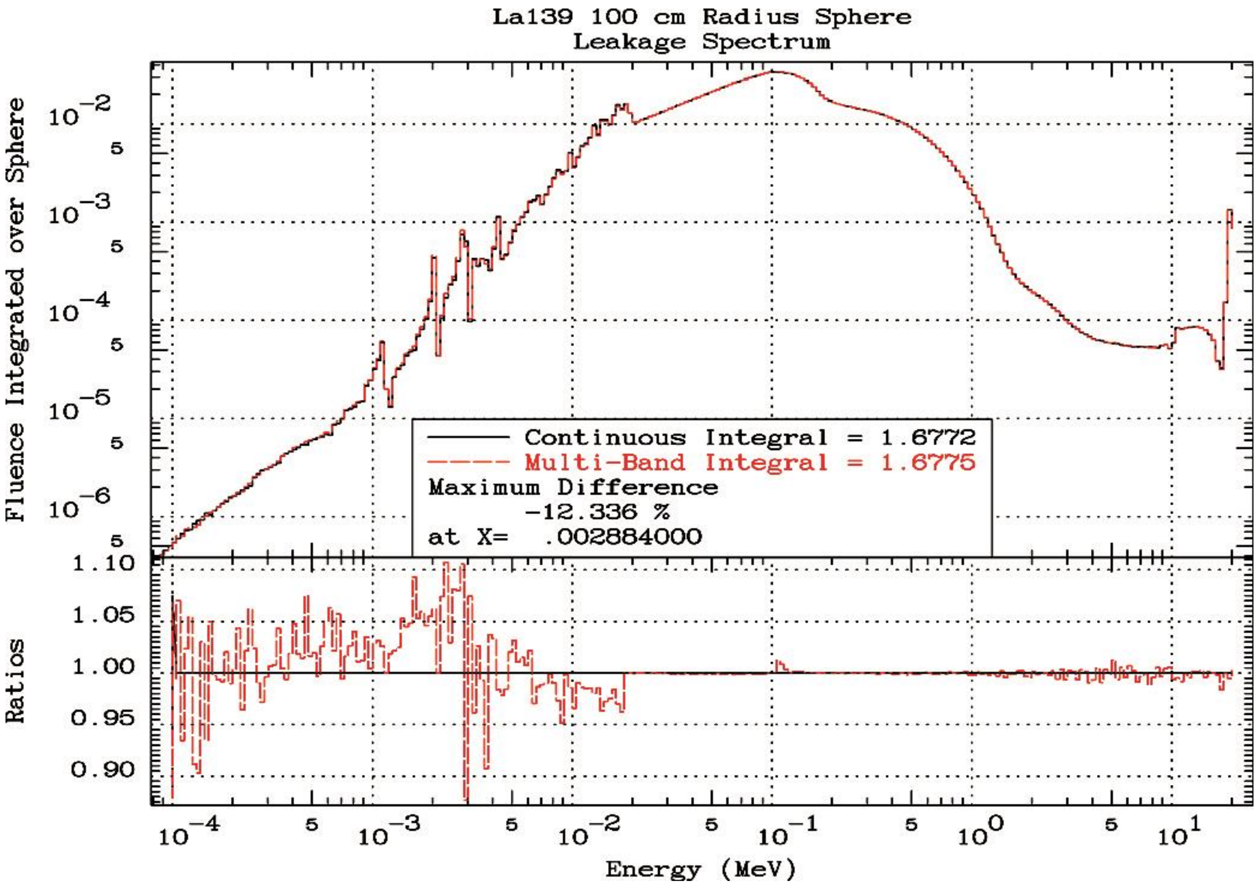

Continuous vs. Continuous without Unresolved Self-Shielding

As a hint at what we will address in our next paper, here we focus on the importance of the unresolved resonance region (URR). Far too many codes seem to ignore the URR or use a poor treatment to handle it. Here we use a TART option to turn off URR self-shielding. The below results are both based on our BEST continuous energy cross sections, and differ only that one includes URR self-shielding and the other does not. From the below figure we can see that, not shielding the cross sections in the URR between 20 keV and 100 keV, means the cross sections will be higher, resulting in a decrease in leakage, an increase in capture, (obvious 20 to 100 keV), and an increase in scatter causing a shift in the spectrum below 20 keV (the increase we see here).

What we see from the below results is how much of the difference is due to the URR self-shielding. For the integrals we find: Continuous 1.6772, No URR Shielding 1.6327, **2.65% lower**; compare this to Unshielded 1.5646, 6.7% lower. Roughly 1/3 of the difference is solely due to URR self-shielding. Hopefully this serves as a WARNING that regardless of how good you think your data and codes are, **if they do not include URR self-shielding you will not be able to produce accurate results.** Our next paper will focus on insuring we can accurately model unresolved resonance region (URR) self-shielding.

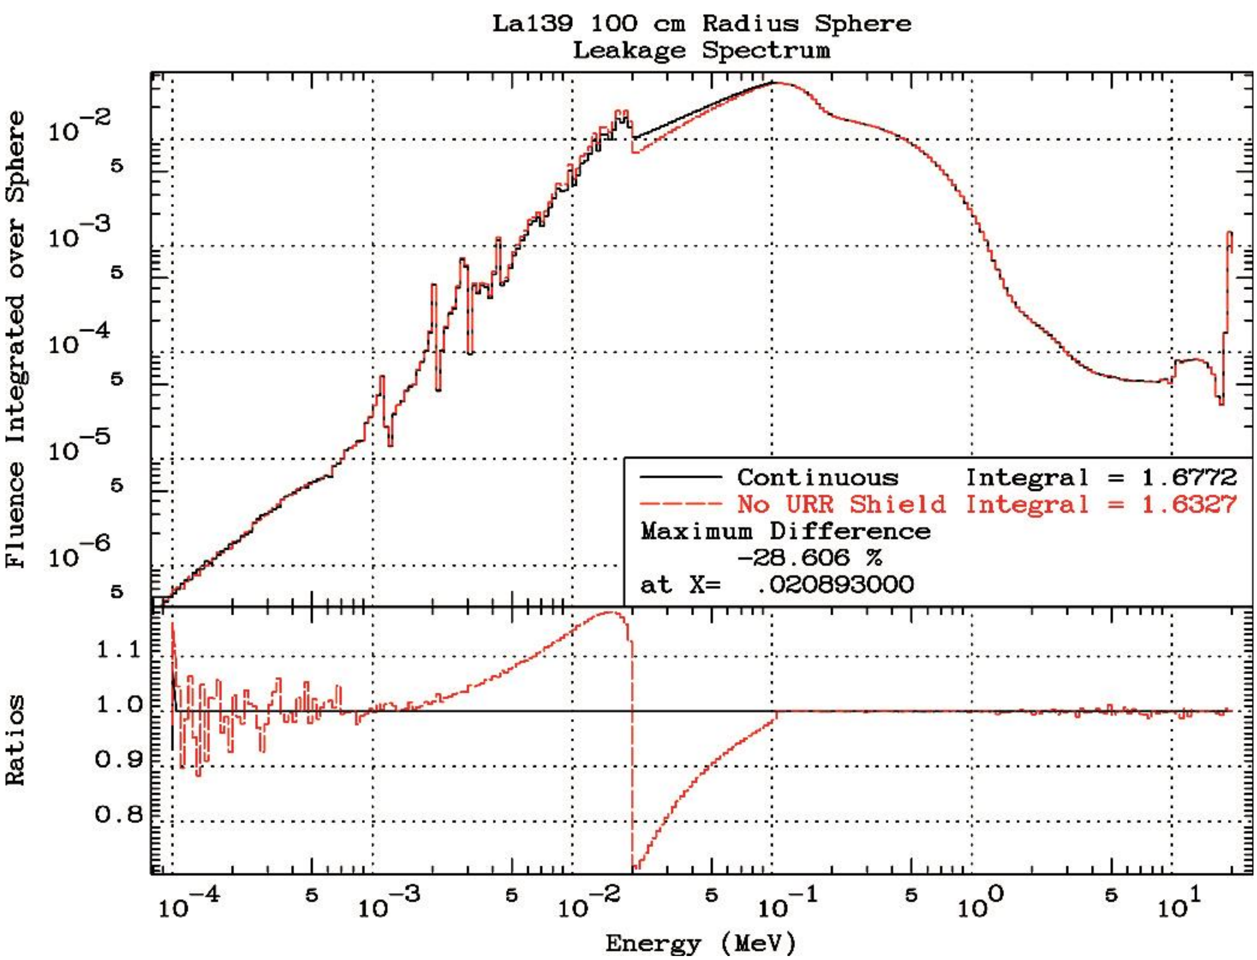

From the integrals of the spectra, particularly the one from max-to-min energy, we can see how the integral is the same down to top of the URR at 100 keV. It then drops between 100 keV and 20 keV, by about 5%. But then below 20 keV it rebounds, increasing back to the final 2.65% total integral difference.

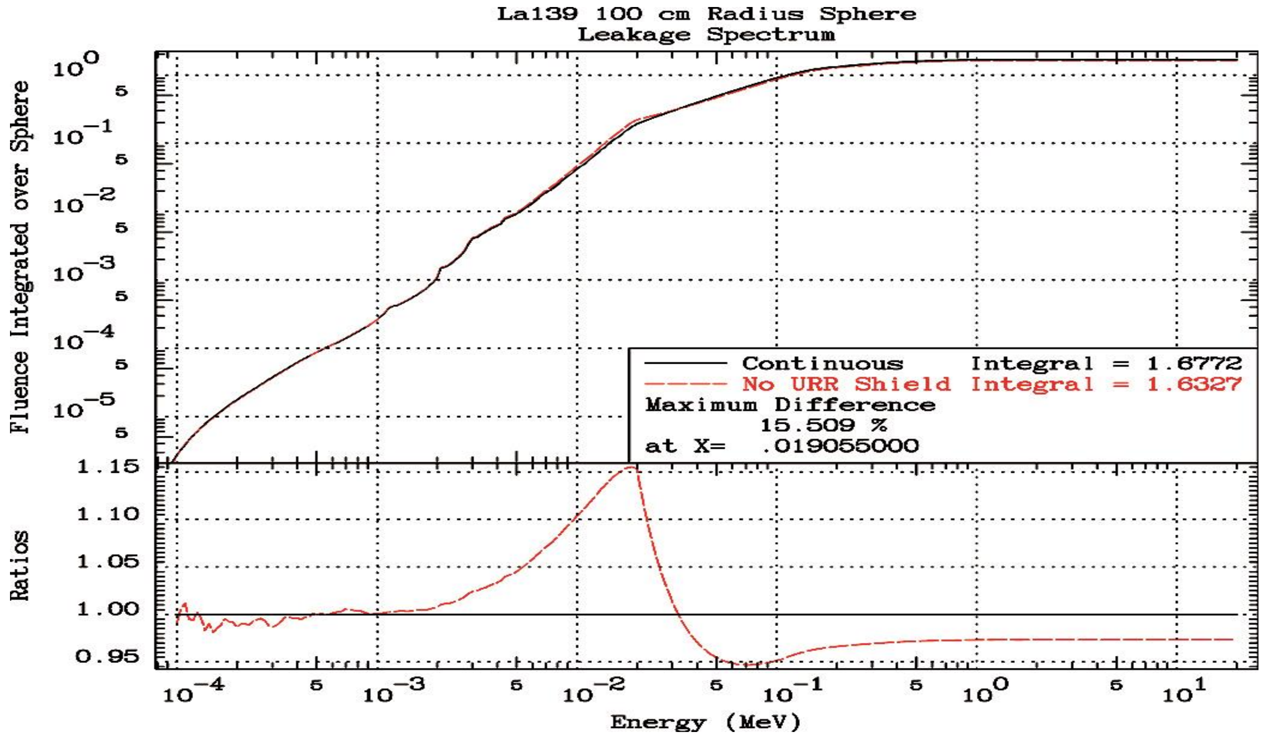

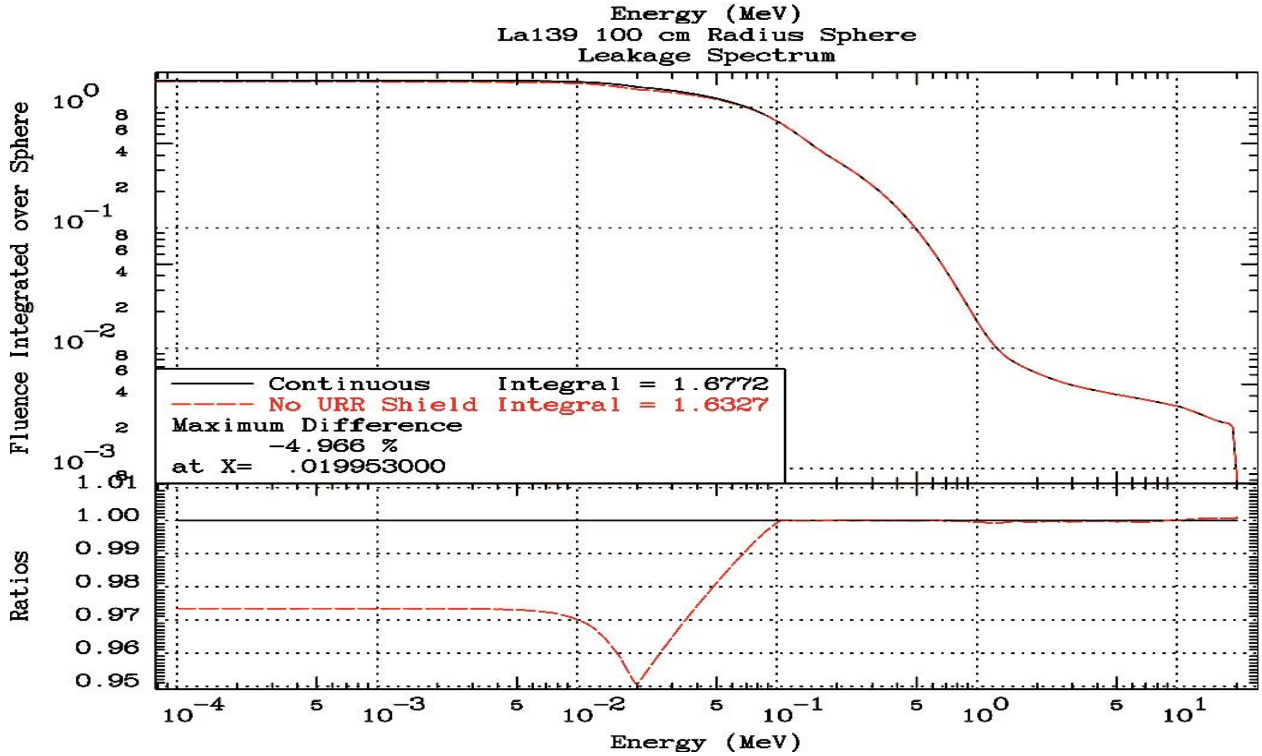

Conclusions

Here using the TART Monte Carlo code we demonstrated the importance of neutron resonance self-shielding to calculate a neutron source problem, recently described as difficult to obtain consistent answers between Monte Carlo and deterministic codes [2]. First, we presented results for exactly the same system with and without self-shielding, in order to demonstrate the magnitude of the effect. Next, we presented results obtained using three different methods to model self-shielding: Unshielded, Total Shielded and Multi-Band, to illustrate the magnitude of the differences between results. Lastly, based on the results in this report, we propose a future area of study to improve our self-shielding methods, namely improved unresolved resonance region (URR) self-shielding.

References

[1] D.E. Cullen: "**The Importance of Resonance Self-Shielding**", Report INDC(NDS)-0778, Nuclear Data Section, IAEA, Vienna, Austria, March 2019. FREELY available on-line at: https://www-nds.iaea.org/publications/indc/indc-nds-0778.pdf and http://RedCullen1.net/Homepage.new

[2] C. Konno: "New comments on data below 20 MeV in FENDL-3.1d", Presented at the Consultants' Meeting on the FENDL Library for fusion neutronics calculations, 15-18 October 2018, IAEA Headquarters, Vienna, Austria, "https://www-nds.iaea.org/index-meeting-crp/CM-FENDL-2018/docs/Konno_FENDL_r7_2017.pdf", Slide 38.

[3] D.A. Brown et al.: "**ENDF/B-VIII.0**: The 8th Major Release of the Nuclear Reaction Data Library with CIELO-project Cross Sections, New Standards and Thermal Scattering Data", Nuclear Data Sheets **148** (2018) 1-142.

[4] D.E. Cullen: "**PREPRO2019**: 2019 ENDF/B Pre-processing Codes (ENDF/B-VIII Verified)", Report IAEA-NDS-229, Nuclear Data Section, IAEA, Vienna, Austria, August 2019. FREELY available on-line at: https://www-nds.iaea.org/publications/iaea-nds/iaea-nds-0229.pdf and http://RedCullen1.net/Homepage.new

[5] D.E. Cullen: "**POINT2018**: ENDF/B-VIII Final Temperature Dependent Cross Section Library", Report IAEA-NDS-227, Nuclear Data Section (NDS), IAEA, Vienna, Austria, April 2018. FREELY available on-line at: https://www-nds.iaea.org/publications/iaea-nds/iaea-nds-0227.pdf and http://RedCullen1.net/Homepage.new

[6] D.E. Cullen: "**TART 2016**: An Overview of a Coupled Neutron-Photon 3-D, Combinatorial Geometry Time Dependent Monte Carlo Transport Code", Report: LLNL-SM-704560, Code Release: LLNL-CODE-708759, September 2016. Internationally available through code centers.

[7] **616 TART tally bins** are defined as 50 bins per energy decade, equally spaced in lethargy, i.e., the energy boundaries of successive bins are multiples of one another, from 10^{-5} eV to 20 MeV. With this tally bin scaling the spectra are the actual COUNT of the number of neutrons leaked per source neutron.

Nuclear Data Section
International Atomic Energy Agency
Vienna International Centre, P.O. Box 100
A-1400 Vienna, Austria

E-mail: nds.contact-point@iaea.org
Fax: (43-1) 26007
Telephone: (43-1) 2600 21725
Web: https://nds.iaea.org