

Covariance Applications with Kiwi

C.M. Mattoon^{1,a}, D. Brown², and J.B. Elliott¹

¹ Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore CA

² Brookhaven National Laboratory, P.O. Box 5000, Upton NY

Abstract. The Computational Nuclear Physics group at Lawrence Livermore National Laboratory (LLNL) is developing a new tool, named ‘Kiwi’, that is intended as an interface between the covariance data increasingly available in major nuclear reaction libraries (including ENDF and ENDL) and large-scale Uncertainty Quantification (UQ) studies. Kiwi is designed to integrate smoothly into large UQ studies, using the covariance matrix to generate multiple variations of nuclear data. The code has been tested using critical assemblies as a test case, and is being integrated into LLNL’s quality assurance and benchmarking for nuclear data.

1 Introduction

Covariance matrices for nuclear reaction data are receiving renewed interest from the international collaboration of cross section evaluators, resulting in a dramatic increase over the past few years in the amount and quality of covariances available in major nuclear reaction libraries (notably including the forthcoming ENDF/B-VII.1 library [1]). This increased interest in covariances comes in response to nuclear data users, who require good estimates of nuclear data uncertainties in order to provide realistic uncertainty estimates when simulating real-world designs.

At Lawrence Livermore National Laboratory, we are seeking to utilize these newly-available covariances in order to support the lab’s strong effort to improve uncertainty quantification. LLNL is in process of developing an extensive framework for Quantitative Measurement of Uncertainty (QMU) studies. This ‘UQ Pipeline’ has the ambitious goal of pulling together all measured data and uncertainties, and also accounting for model uncertainties and potential unknown systematic errors, in order to produce realistic uncertainty intervals on the final result of a simulation. The UQ pipeline promises to give more realistic answers to real-world problems including the performance and safety of weapons and nuclear reactors. It also increases the demand on the existing nuclear data infrastructure at LLNL, requiring more flexible handling of nuclear data as well as much more detailed uncertainties and covariances as input.

In this paper, we will present an overview of our efforts to meet the needs of large-scale sensitivity studies using the UQ pipeline. We will discuss three projects at LLNL that are focusing on improving the handling of nuclear data for UQ studies: Kiwi, a project for applying covariance matrices, Fudge2.0, which streamlines modifying and processing data for applications, and the new Generalized Nuclear Data format (GND), which is designed to modernize and simplify the storage and handling of nuclear data, replacing several older systems.

^a e-mail: mattoon1@llnl.gov

2 Introduction to Kiwi

LLNL’s Kiwi application is designed to be an interface between nuclear data covariances and UQ studies. Kiwi can be used to search LLNL’s ENDL library for available covariance matrices, and then used to make ‘realizations’ by varying the nuclear data based on a user’s request, within the limits of the covariance matrix. Kiwi can also be used to process the resulting realizations into the binary formats required for use by simulation codes, thus helping to automate the process of varying nuclear data inputs for UQ studies.

Kiwi was created by J. Pruet in 2007 to assist in various QMU studies in support of NNSA campaigns. In its original form, Kiwi used a very small (~ 6 material) library of cross section uncertainties either taken from the Neutron Standards project or from in-house estimates. These uncertainties were extended into full covariances inside Kiwi by assuming that each uncertainty is accompanied by a correlation matrix whose elements are given by:

$$R_{ij} = \exp(|E_i - E_j|/\Delta E) \quad (1)$$

where E_i is the mean energy of the i^{th} bin and ΔE is a user imposed correlation length. This version of Kiwi was used successfully in several other covariance studies by J. Pruet before the project was taken over by D. Brown in 2009. D. Brown extended Kiwi to read and sample from the full covariance matrices added to ENDL2009. These covariances were no longer ad-hoc constructions as in the original Kiwi, but rather were the full evaluated covariances from the ENDF/B-VII.0 and Low-Fidelity covariance libraries [2]. In 2010, D. Brown and T. Luu used this new version of Kiwi to perform a proof-of-principle study of the Jezebel critical assembly, and this study is given as an example project in the Kiwi source distribution.

The basic algorithm used by Kiwi is relatively simple. For a single variation, the code requires an (absolute) covariance matrix M containing the average uncertainties and covariances in each energy interval for a single reaction cross section, and also requires a vector V representing the user’s requested variation. For an $n \times n$ covariance matrix,

V is of length n , and each element of V contains the variation, in multiples of σ , requested for the corresponding energy interval. To generate the variation, Kiwi first obtains the eigenvalues λ and eigenvectors Λ of the matrix M , then creates the variation vector R :

$$R = \eta \cdot \Lambda^T, \quad (2)$$

where the elements of the vector η are computed by:

$$\eta_j = \sqrt{\lambda_j}(V \cdot \Lambda_j). \quad (3)$$

R contains the new variation, which is added to the original cross section to obtain the new cross section realization. The intermediate vector η can be thought of as the user's requested variation, rotated into the space spanned by the eigenvectors.

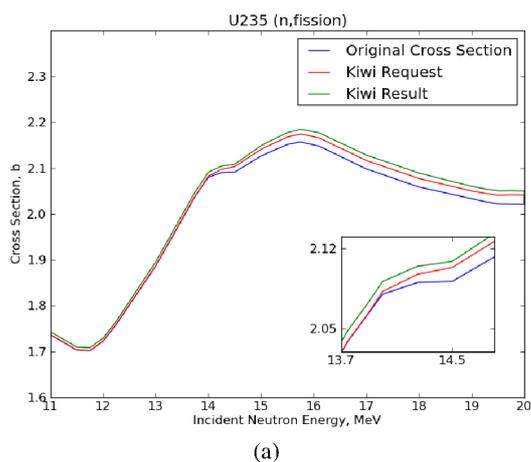
In addition to varying nuclear data, Kiwi must preserve the total cross section (for conservation of energy). By default, after computing the variation in a channel, Kiwi compensates by subtracting the same variation from the elastic cross section. While this is suitable for relatively small variations, large variations at high energy will require a different strategy to maintain energy balance, as the elas-

tic cross section may not be large enough to compensate a change in another channel.

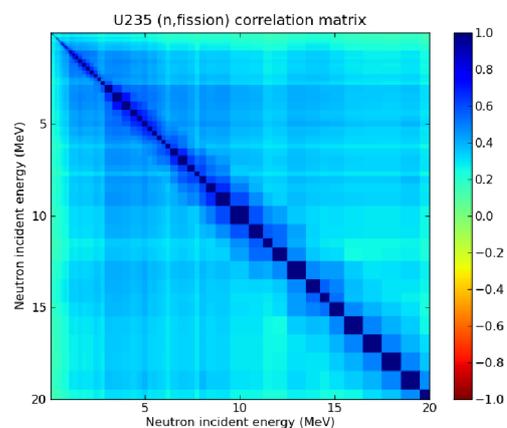
3 Kiwi in Action

Fig. 1 shows an example of typical Kiwi results. When the user requests a variation from Kiwi over an energy interval (shown in red), the Kiwi result (shown in green) is returned. As seen in the correlation matrix, the requested variation actually lies within a larger, strongly correlated region. In producing the requested variation, Kiwi therefore also increases the cross section in neighboring, correlated bins. The resulting variation thus spans a wider energy interval than initially requested, but is more realistic based on the covariance data.

Kiwi can give potentially surprising or confusing results, especially when strong anti-correlations are present in the covariance matrix. For example, in Fig. 2 the user has requested a cross section increased by one σ over a wide energy range that spans two anti-correlated regions. Even though the user requested the same positive variation throughout, the Kiwi results end up *decreasing* the cross section in one region while increasing in another. These

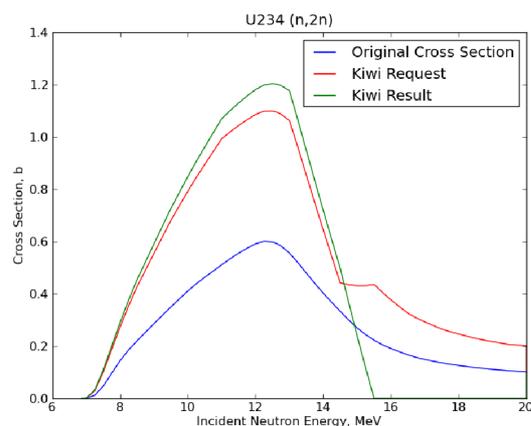


(a)

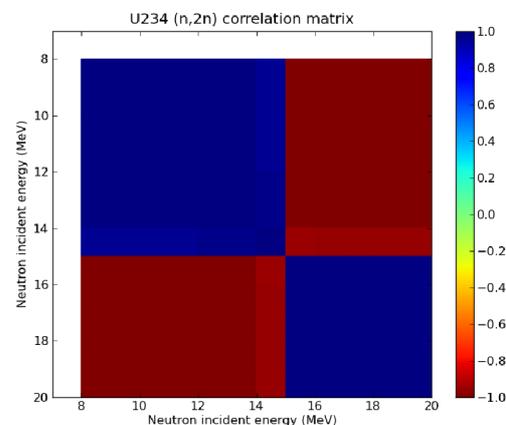


(b)

Fig. 1. Sample Kiwi result: the user requests a variation of $+1\sigma$ starting at 14 MeV (shown in red on Fig. 1(a)). Since the requested region is correlated, Kiwi returns a somewhat larger result, over a wider energy interval (shown in green, with greater detail in the inset). For comparison, the correlations are also shown in Fig. 1(b).



(a)



(b)

Fig. 2. Kiwi results should be treated with caution when strong anti-correlations are present in the covariance matrix. In this example, the user requests a variation of $+1\sigma$ throughout the entire energy region, yet due to anti-correlations seen in Fig. 2(b), the resulting variation is negative at some energies.

results arise directly from the anti-correlations in the covariance matrix, reflecting the fact that, according to the correlation matrix, an increase in the cross section in both regions simultaneously is unphysical. We note that this is the desired behavior when sampling a normalized probability distribution, where an increase in one region must be accompanied by an equal decrease at another energy to preserve the normalization. For sampling a single reaction cross section, however, the physical meaning of the result in Fig. 2 is less clear.

Fig. 2 demonstrates both that the Kiwi results depend strongly on the input covariance matrices, and also that strong anti-correlations in the matrix cause trouble when using Kiwi to obtain realizations of reaction cross sections. This example uses the U234(n,2n) covariance matrix from the ENDL-2009 library, which was originally taken from the ‘low-fidelity’ covariance estimates [2], produced in order to provide covariances for the many materials in major reaction libraries that do not yet contain evaluated covariance matrices. The strong anti-correlations in the low-fidelity covariance matrix appear to be the result of propagating model parameter uncertainties into a covariance matrix, without taking experimental data into account (no experimental data for U234 (n,2n) appear in the EXFOR database, although data for the competing (n,fission) channel are available). The Kiwi results suggest to us that the strong anti-correlations between (n,2n) at different energies should be treated with caution.

Since anti-correlations in the original covariance matrix lead to unexpected results, Kiwi provides alternate options for creating variations. The user may choose to receive a warning whenever a matrix containing anti-correlations is used to generate a realization. Kiwi also offers the option of disregarding the off-diagonal terms and computing the variation solely based on the diagonal terms (the variances). This option should be used with caution, however, since much of the information encoded in the covariance matrix will be lost. In the future, we hope that as new evaluations with covariances are produced, these work-arounds will become less important.

3.1 Integrating Kiwi into Sensitivity Studies

Before nuclear data can be used by deterministic or Monte Carlo production codes at LLNL, the data must first be converted into ENDL format, then processed and stored in the proper binary file formats. Thus, Kiwi must be able not only to generate nuclear data realizations, but also to handle the subsequent processing. Kiwi is therefore built upon LLNL’s code ‘Fudge’, which was written with the goal of simplifying the creation and processing of nuclear data libraries. Fudge was written in python over the past decade, and was designed without any covariance capabilities. Kiwi extends Fudge, inheriting the processing and data management capabilities and also adding the ability to handle covariances and generate realizations.

The full strength of Kiwi is seen when it is used in conjunction with the UQ pipeline to perform a sensitivity study. Tools such as Dakota [3] and PSUADE perform a self-directed search to explore the sensitivities of a problem to each input parameter of interest. Along with Kiwi, these tools can be used to explore the sensitivities of real-world applications to nuclear data, and produce realistic

estimates of how nuclear data impacts the overall uncertainty.

Applying Kiwi to a sensitivity study requires some initial effort from the user, who must decide first how to reduce the nuclear data into a set of input parameters, and second how to handle the model results. For example, the user may decide to vary a reaction channel by 1σ at all energies, or may split the channel into several energy regions varying independently.

4 Future Work

As previously mentioned, Kiwi is built upon the older code Fudge, which is used to make creating and processing nuclear data libraries simpler. Both Fudge and Kiwi were written to use data in LLNL’s ENDL format. In the meantime, however, a new project is underway at the lab to create a modern format intended to replace ENDL and ENDF, as well as multiple processed data forms. The ‘GND’ (Generalized Nuclear Data) format represents nuclear reaction information in a hierarchy intended to be both human-readable and also simple to integrate into production codes. GND is also designed with large-scale UQ studies in mind: variations in nuclear data inputs using the new format will be much simpler than the current format (at LLNL the full library must currently be rewritten to a new binary file whenever any cross section is modified, adding substantial overhead to UQ studies).

The improved GND format will also impact Kiwi. In particular, the new format will store many more types of covariance data than are currently available in the ENDL library. ENDL stores only reaction self-correlations, excluding all cross-correlation terms between reactions or between different materials. As these cross-terms become available, Kiwi must be extended to make use of them. This improvement may prove most important when deciding how to adjust cross sections while maintaining energy balance: if all reaction cross terms are available, the resulting full covariance matrix will encode all information needed to adjust one reaction channel while still maintaining the same total cross section.

5 Conclusions

The Kiwi code improves nuclear data handling at LLNL by permitting the use of covariance matrices when varying nuclear data for sensitivity studies. Kiwi has begun to be applied towards finding realistic uncertainty estimates for real-world applications. At the same time, several important improvements to the code are underway, primarily by expanding the types of covariance data that the code can use as input, and by making sensitivity studies simpler and more flexible.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore Laboratory under Contract DE-AC52-07NA27344.

References

1. M.Chadwick *et al.*, Nuclear Data Sheets **112** No. 12, (2011) 2887

2. R.C.Little *et al.*, Nuclear Data Sheets **109**, (2008) 2828
3. B.M.Adams *et al.*, “DAKOTA, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 5.0 Theory Manual”, Sandia Technical Report **SAND2010**, (2010) 2183