Reconstructing Nuclear Fuel Cycle Operations with Nuclear Archaeology

Max Schalz*[1] and Malte Göttsche[1]

[1] Nuclear Verification and Disarmament Group, RWTH Aachen University, Aachen, Germany

Abstract

A key challenge for future nuclear disarmament treaties lies in verifying the completeness of fissile material baseline declarations. One approach for this is nuclear archaeology, which aims at reconstructing the past fissile material production of a country. It is a set of methods to infer operational production histories of nuclear facilities, typically combining forensic measurement data with simulation models of the examined facilities. Hence, nuclear archaeology methods usually apply to the facility level so far. They do not take into account fuel cycle-level information that may also be contained in declarations, such as material flows between facilities. To provide a platform for such fuel cycle analyses, we develop Bicyclus, an open-source Python3 module that couples nuclear fuel cycle simulations with an inference framework. The user models a fuel cycle in Cyclus, an open-source simulator, and inputs measurement data and rough estimates of the operational parameters to be reconstructed. Examples of parameters are the capacity factor of a reactor or the efficiency of a reprocessing facility. Following this, the software reconstructs those parameters using Bayesian inference and Markov Chain Monte Carlo algorithms, and yields estimates of the produced fissile material. Furthermore, Bicyclus can be used in a measurement-independent mode. Here, the user specifies ranges of values for uncertain operational parameters. Then, a Quasi-Monte Carlo method is used to efficiently sample this parameter space and to obtain aggregated fissile material estimates and uncertainties. We showcase our approach with a hypothetical nuclear fuel cycle for military purposes. First, we modelled the fuel cycle in Cyclus and generated synthetic measurements of the high-level reprocessing waste and the depleted uranium. Using these measurements in Bicyclus, we reconstructed both operational parameters key to the plutonium and HEU production as well as the overall fissile material production. Last, we performed a measurement-independent estimate of the produced fissile material.

1 Introduction

In a nuclear disarmament scenario, one important step is to determine a state's fissile material stockpile and to eventually safeguard it. During this process, uncertainties will inevitably remain and may well correspond to tens or hundreds of weapon-equivalents. Because of this, a robust uncertainty assessment is crucial to ensure that uncertainties are as low as possible and that the their origins are well understood [1].

Today, experts use available (open-source) data to get independent estimates, as done for example by the International Panel on Fissile Materials

*schalz@nvd.rwth-aachen.de

[2]. In a disarmament scenario, one could use other tools like nuclear archaeology to reconstruct the fissile material production of a state. By combining simulation models of the fissile material production processes with physical evidence, such as the isotopic composition of nuclear waste, nuclear archaeology can be used to verify a state's fissile material declaration [1].

To be successful, both methods have to take into account the complex processes of the nuclear fuel cycle (NFC) and disentangle the interlinked operations of the different facilities. Especially for larger NFCs, this is a difficult task. To overcome this problem, we propose to use NFC simulations. We introduce Bicyclus, a self-developed, open-source Python3 library, and use it to perform independent fissile material assessments as well as reconstruct the fissile material production of an example NFC.

2 Modelling the Nuclear Fuel Cycle

We use Cyclus, a modular, open-source simulator, to model the NFC [3]. In Cyclus, nuclear facilities are represented by so-called agents. These are independent from each other and their behaviour is solely governed by their respective internal states. Apart from 'doing things' with nuclear material, such as enriching it, agents can offer and request material in each timestep. Cyclus gathers these offers and requests, determines matching request-offer pairs and executes the trades by transferring the material between the facilities. Here, it takes into account any pre-set preferences, such as the user defining that a reactor prefers uranium oxide over mixed oxide fuel, or that the reactor's requests should get fulfilled first. Thus, the simulator effectively optimises material flows between facilities in each timestep. At the end of the simulation, a large amount of information is stored in an output file, including data on all material transfers and material compositions.

3 Investigating the NFC

In order to obtain sound fissile material assessments with Cyclus, we use two different statistical tools. First, we generate independent fissile material estimates using a quasi-Monte Carlo (QMC) sampling method. Then, in a second step, we can verify prior information or assumptions with actual measurement data in a Bayesian inference framework. While the actual work is similar in both tasks—repeatedly simulating the NFC with varying input parameters—the underlying concepts are fundamentally different, as will be explained in the following.

3.1 Assessing Fissile Material Production

Many unknowns arise when trying to determine fissile material production: how much natural uranium got mined, how long did the reactors run, how much enrichment capacity was available, etc. Such uncertainties must be propagated through the NFC to evaluate their impact on the final fissile material balance. Numerical methods are useful for this, especially for complex NFCs. By transforming uncertainties into probability distributions, e.g., uniform ranges where all parameter values are equally likely, and repeatedly simulating the NFC with different parameter combinations drawn from these distributions, we obtain aggregate fissile material estimates.

In this work, we use Sobol sequences, a QMC approach, to generate the parameter combinations [4]. Sobol sequences have the advantage of being so-called low-discrepancy sequences, meaning that they cover a high-dimensional parameter space much more efficiently than, e.g., standard grid sampling or pseudorandom sampling.

3.2 Verifying Inventories with Bayesian Inference

In the context of nuclear archaeology, we use Bayesian inference to reconstruct fissile material production with both measurement data and simulations.

Bayesian inference is based on Bayes' rule, $p(\theta|y) \propto p(\theta)p(y|\theta)$, where θ is a model input parameter and y a measurement. In the following, $p(\theta)$ is called the *prior distribution*, $p(y|\theta)$ the *likelihood* and $p(\theta|y)$ the *posterior distribution*.

In our application scenario, θ corresponds to a vector of nuclear facility parameters, such as an enrichment capacity or a reactor power, and y is a vector of measurements such as the isotopic composition or the total mass of nuclear waste. Thus, the prior reflects any knowledge on the nuclear facilities, e.g., as obtained through open-source data. The likelihood uses case-specific mathematical formulations to compare the measurements to corresponding simulated values, assuming certain parameter sets. Finally, the posterior describes the reconstructed facility parameters, taking into account both the prior knowledge and the measurements.

Solving an inverse problem with Bayesian inference is not trivial, as an analytical calculation of the posterior is only possible in few cases. Alternatively, the posterior can be numerically constructed using Markov chain Monte Carlo (MCMC) methods, where we repeatedly draw samples from a so-called transition distribution. This is done in such a way that with each step, the constructed posterior distribution approximates the true (unknown) posterior better and better [5].

4 Implementation

The final step is to combine Cyclus with the QMC and Bayesian inference methods into one coherent framework. To this end, we have developed Bicyclus, an open-source Python3 library [6, 7]. It offers two operation modes, the *forward mode* and the *reconstruction* or *inference mode*.

The forward mode corresponds to the QMC approach, where we generate independent fissile material estimates. Here, the user specifies the NFC in a template Cyclus input file and defines the parameter distributions for uncertain parameters. Then, Bicyclus uses SciPy's implementation of Sobol sequences to generate all input parameter sets and runs the large-scale simulations [8]. The fissile material estimates and any other information can subsequently be read out from the simulation output files using separate scripts and aggregate mean values and uncertainties can be determined.

Syntactically, the reconstruction mode is similar to the forward mode. The user first defines a template input file and the parameter distributions for uncertain parameters, i.e., the priors. Additionally, they need to define the likelihood: Which values should be extracted from a Cyclus output file, e.g., an isotopic composition vector $\vec{x}$ of a material, and how should they be compared to measurement data? For example, the user could use a Gaussian function as likelihood $\mathcal{L}$ such that

$$\mathcal{L} \propto \exp\left(-(\vec{x}_{\text{measurement}} - \vec{x}_{\text{simulation}})^2/2\sigma^2\right), \tag{1}$$

with a given uncertainty σ.

Having defined the prior as well as the likelihood, Bicyclus has all necessary information and can start preparing the sampling process. For this and all other MCMC-related operations, it uses PyMC, an open-source Python3 library [9]. It transforms priors and likelihood into PyMC functions, generates initial values and then starts the sampling process to obtain the posterior distribution. There, the following steps are performed repeatedly: 1. draw a parameter sample from the priors using PyMC, 2. generate a Cyclus input file, 3. run the simulation, 4. extract relevant simulation

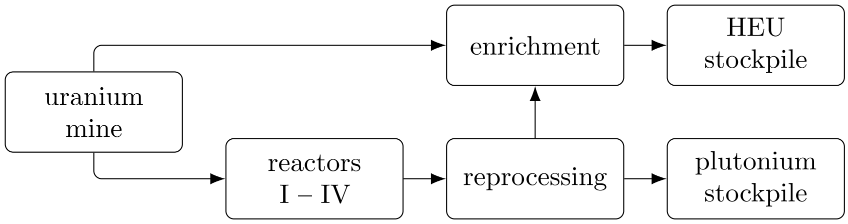

Figure 1: Nuclear fuel cycle and material flows used in the case study. Note that support facilities, such as material storages, were used in the simulation but are omitted in this scheme for clarity.

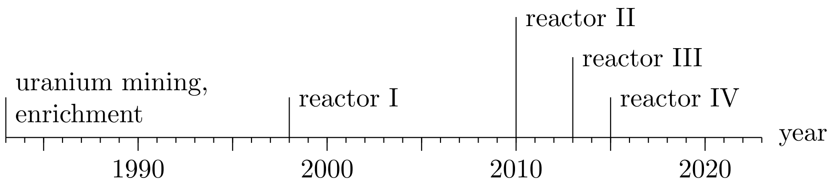

Figure 2: Timeline of the simulated NFC, showing the start of operations of each facility.

output and calculate the likelihood, 5. return the likelihood value to PyMC and go back to 1.

Apart from this core functionality, Bicyclus handles additional tasks such as logging, data visualisation, and it offers the possibility to store all of Cyclus' simulation output files. This feature is important, as it allows for an extensive analysis and postprocessing of data even after the actual inference process is finished. It will be used extensively in the following case study when reconstructing the fissile material production.

5 Case Study

In the following, we show how Bicyclus can be used to estimate fissile material production. First, we present the NFC used in this work. Then, we perform a measurement-independent assessment of the fissile material stockpiles and last, we use nuclear waste measurements to infer fissile material production.

5.1 Modelling a Military Fuel Cycle

The NFC used here is loosely based on Pakistan's military nuclear programme[1], following [10], and is shown schematically in Fig. 1. It can produce highly-enriched uranium (HEU) and plutonium, as well as recycle irradiated uranium which can then be enriched to weapon-grade levels. Each simulation covers a timespan of forty years and uses one day as simulation timestep. In order to increase realism and to have a dynamic simulation, some facility parameters are time-dependent, such as the enrichment facility's separative power, and the four reactors are deployed over the course of the simulation, see Fig. 2.

The uranium mine starts with a pre-existing natural uranium stockpile and produces additional uranium at varying rates during the simulation[2].

[1] This study does *not* aim at recreating Pakistan's programme. Some details, such as uranium imports, have been omitted for the purpose of simplicity.

[2] The initial stockpile of 339 t corresponds to the uranium production from 1971 to 1983 of, on average, 28.25 t per year. Later annual production varies between 23 to 50 t,

The enrichment facility is modelled as an ideal cascade [7], producing 90% enriched HEU with 0.3% enriched tails. The spent fuel composition is based on own reactor simulations and on data from [2] for a pressurised heavy water reactor (PHWR) running at a burnup of 1.2 MWd/kg.[3] The reprocessing facility separates the spent fuel into three outgoing material streams (plutonium, uranium and waste stream) with a 99% separation efficiency in plutonium and uranium streams.

Furthermore, we assume that plutonium is preferred over HEU, thus the simulation is configured to send natural uranium to the reactors first if requested, then send remaining natural uranium to the enrichment facility. It should be noted that this preference mechanism only takes into account the requests of the current timestep and it does not anticipate events, such as a reactor needing a large amount of uranium in the subsequent timestep.

Based on the structure of the NFC, we can—qualitatively—expect to observe the following three 'phases' over the course of a simulation:

1. HEU production: At the beginning, all natural uranium is enriched to HEU, within the given feed- and separative power-constraints.

2. HEU and plutonium production: Once reactor I starts operating, natural uranium resources are shared between enrichment and the reactor.

3. Natural uranium shortage: At some point during the deployment of reactors II to IV, the natural uranium demands of the reactors and the enrichment facility exceed the mine's production rate[4]. At this point, the enrichment facility starts using reprocessed uranium as enrichment feed. Additionally, if there is no uranium available, the reactors or the enrichment facility will temporarily suspend operations.

5.2 Quantities and Parameters of Interest

The overarching goal of the case study is to assess the total HEU and plutonium production given different input uncertainties and constraints. Specifically, we consider uncertainties on two parameters: the capacity factors of the reactors and the separative power of the enrichment facilities, both of which, in theory, influence the fissile material production linearly[5]. In practice however, this relationship may be more complex and questions arise, such as: How does the limited natural uranium supply change this result? How will the interplay between reactors and enrichment via the reprocessing facility affect the fissile material production? We investigate these questions and their influence on the final fissile material balance in the following using the forward and reconstruction modes of Bicyclus.

5.3 Independent Fissile Material Estimates

To obtain first estimates of the fissile material production, we use Bicyclus' forward mode and the parameter distributions from Table 1. We generate $2^{11} = 2048$ sets of parameter samples and run the corresponding simulations, producing the fissile material estimates shown in Fig. 3. The HEU production

settling at 45 t per year from 2010 onward [11].

[3]At this burnup, the plutonium in the spent fuel is still weapon-grade, i.e., it has a Pu-239 content of 93.8%.

[4]Refuelling all four reactor cores, each containing 9450 kg uranium, already consumes most of the yearly natural uranium production of 45 t. Given the low burnup and their powers ranging from 49 to 100 $\mathrm{MW_{th}}$ [12], they need to be refuelled more than once per year. This refuelling schedule can only be kept up for so long using excess natural uranium having accumulated over the years.

[5]The capacity factor equals the online time divided by the total cycle time, i.e., online plus offline time. The separative power is proportional to the number of centrifuges and is a measure of how much uranium can be enriched to a given level.

Table 1: Parameters varied throughout the case study. The parameters are uniformly distributed within the given ranges, the capacity factor is assumed to be identical for all four reactors and the separative power variation applies from 1999 onward. The ground truth only applies to the inference mode.

parameter	unit	range	ground truth
capacity factor	%	50 – 80	70
separative power	kg SWU/year	20 000 – 45 000	30 000

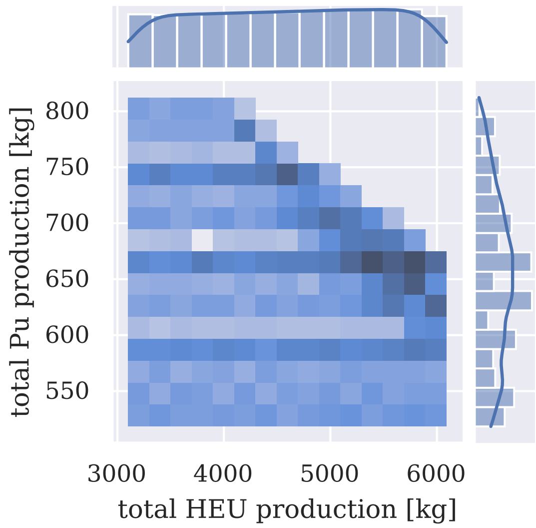

Figure 3: Two-dimensional histogram and marginal distributions of total HEU and plutonium produced in 2048 simulations, each covering 40 years. The solid lines in the marginals are kernel density estimates (KDEs).

spans nearly uniformly from 3 100 to 6 100 kg, while the plutonium production ranges from 520 to 810 kg.

Furthermore, these results show that it is not possible to maximise both HEU and plutonium production at the same time: A plutonium production of more than 670 kg implies there was an increased uranium consumption of the reactors, leading to a natural uranium shortage in the enrichment facility. To compensate this loss, the enrichment facility used reprocessed uranium as enrichment feed. Because of its lower enrichment level, this feed required more separative work per unit amount of produced HEU, resulting in a decrease in overall HEU production.

While the forward simulations yield important insights, such as proof of the natural uranium constraints, the fissile material estimates remain uncertain. These uncertainties can be significantly reduced if measurement data becomes available.

5.4 Inferring Fissile Material Production

Combining forward simulations and measurement data in the Bayesian framework allows to compare our prior assumptions and, ideally, to obtain aggregate, more precise fissile material estimates. In this case study, we use synthetic measurements which are generated by defining a ground truth (the 'true' parameter set), then running one Cyclus simulation with these values and using the resulting simulation output to obtain the measurements.

In the following, we will study the impact of different measurements on the reconstruction of the parameters, which in turn impacts the reconstruction

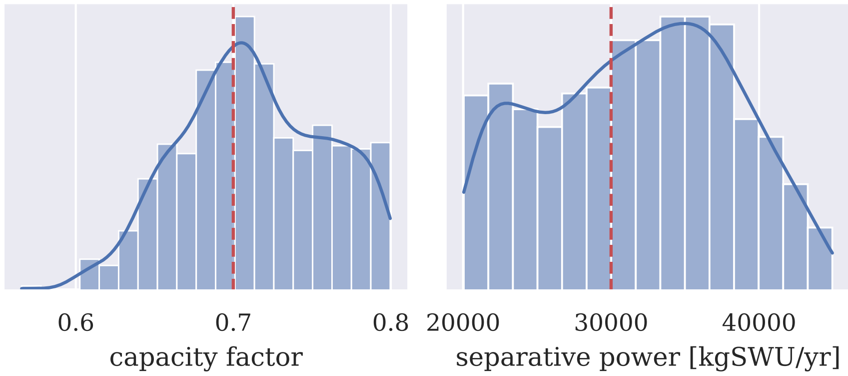

Figure 4: Posterior distributions obtained by measuring Cs-137 content in the reactor waste. The cutoffs are due to the fact that the posteriors cannot lie outside the range of the priors. The groundtruth values are shown as red, dashed line and the blue, solid lines are KDEs.

of the fissile material. Specifically, we first consider the amount of Cs-137 in the nuclear reactor waste, and second we additionally consider the mass of depleted uranium tails from the enrichment facility. These quantities have been chosen because of their relationship to the respective fissile material productions: Cs-137 is a fission product and its amount in the reactor waste thus scales with the reactor's total irradiation time; depleted uranium is a byproduct of uranium enrichment and, for a given enrichment process, scales linearly with the amount of HEU produced. While these measurements could be performed in an actual application, they cannot guarantee completeness (e.g., a malicious actor could hide tails containers or reenrich tails) and further studies are needed. This work focusses only on the simulations and the conceptual aspects.

Throughout the analysis, we use a Gaussian likelihood with 5% relative uncertainty, see Eq. (1), and the priors and ground truth from Table 1. All inferences use PyMC's `Slice` sampler with 15 independent chains, each with 100 tuning and 100 actual samples.

Using these settings and considering only the Cs-137 measurement, we obtain the posterior distributions shown in Fig. 4. The uncertainty on separative power remains essentially unchanged, while we have managed to reduce the uncertainty on the capacity factor with a posterior mean value of 0.713 ± 0.047 and a 95%-highest density interval (HDI) from 0.638 to 0.799. Nonetheless, uncertainties remain at a higher level, especially considering the right-hand flank of the posterior: Intuitively, we would have expected a better reconstruction of the capacity factor because of the physical correlation between itself and the caesium fission product.

This unexpected result can be explained by Fig. 5. Instead of the expected linear correlation, data points fan out for larger values of the capacity factor and we observe an additional dependency on separative power, where data points farther off the diagonal correspond to larger separative powers. This phenomenon is a consequence of the inference process: For a given capacity factor, different separative power values were drawn, which influences the natural uranium consumption of the enrichment facility and therefore also the amount of natural uranium available to the reactors. If the separative power is large, it can lead to a uranium shortage and as a consequence, reactor operations may temporarily be on hold. The reactors will not reach the desired (input) capacity factor and they will produce less spent fuel, i.e., less plutonium and less Cs-137.

To resolve this ambiguity, we add the depleted uranium mass measurement to the inference and we obtain the results shown in Fig. 6. We notice

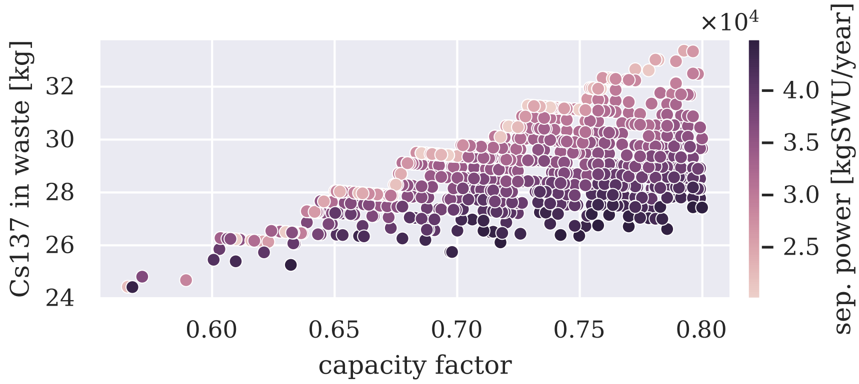

Figure 5: Using only the Cs-137 measurement in the inference results in ambiguities: one measurement can correspond to multiple capacity factors, and vice-versa.

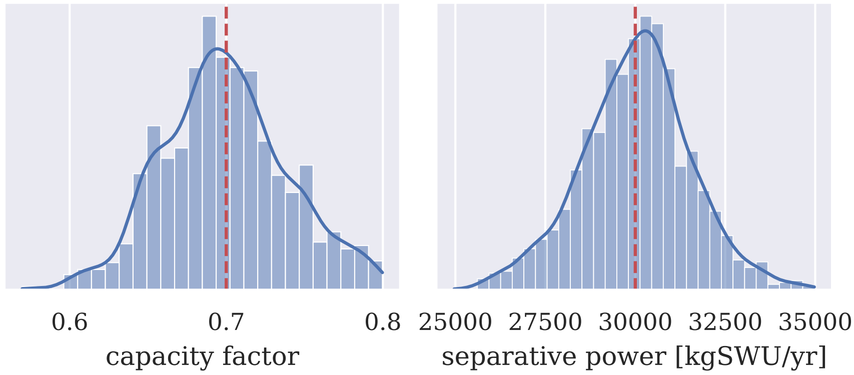

Figure 6: Posterior distributions obtained by measuring Cs-137 content in the reactor waste and the total uranium tails mass. The groundtruth values are shown as red, dashed lines and the blue, solid lines are KDEs.

that the separative power posterior now has a significantly smaller spread, peaking at around 30 300 kg SWU/year with a 95%-HDI from 26 710 to 32 890 kg SWU/year. Additionally, the reconstruction of the capacity factor improves: Its posterior becomes narrower and the right flank falls off much steeper compared to the caesium-only measurement. To explain this improvement, we refer back to Fig. 5: The depleted uranium measurement improves the separative power reconstruction drastically, which resolves the above-mentioned fanning out and the ambiguous mapping between Cs-137 and the capacity factor.

Having successfully reconstructed both parameter values, we can now evaluate the fissile material production corresponding to the posteriors using the stored Cyclus output files. We find that (4330 ± 180) kg HEU and (710 ± 34) kg plutonium were produced in this scenario, see Fig. 7. These results are in very good accordance with the true values of 4328 kg HEU and 709 kg plutonium.

Using the output files, we can reconstruct additional quantities, such as the specifics of the HEU enrichment. For instance, we can determine how much of each enrichment feed (natural and reprocessed uranium) was used as function of the sampled parameters, see Fig. 8. We clearly observe that higher separative powers consume more reprocessed uranium, which is due to the overall higher uranium consumption. Additionally, there is a more subtle dependency on the capacity factor: For a capacity factor larger than 0.7, reprocessed uranium gets used in the majority of cases. This arises from

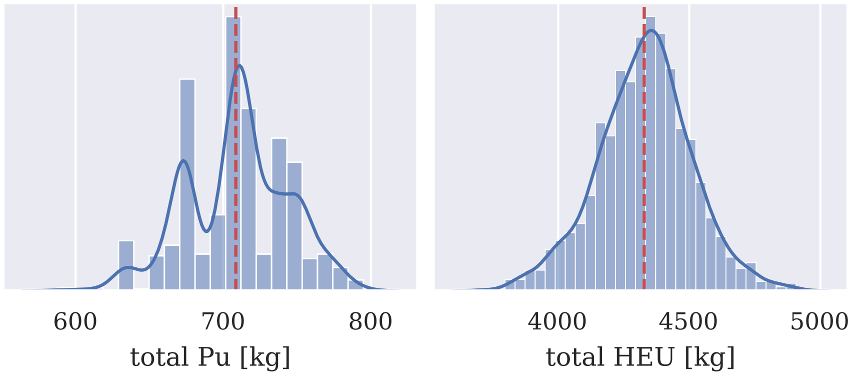

Figure 7: Fissile material estimates reconstructed using the Cs-137 and uranium tails measurements. The values corresponding to the groundtruth are shown as red, dashed lines and the blue, solid lines are KDEs.

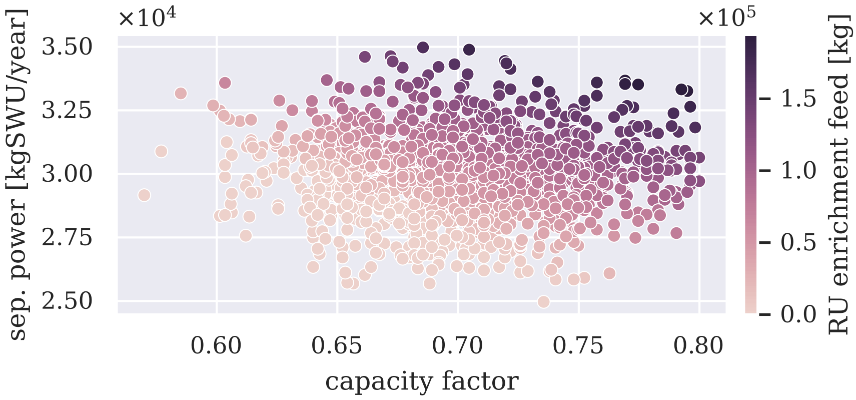

Figure 8: Total amount of reprocessed uranium used as enrichment feed

the reactors consuming most natural uranium, resulting in a shortage at the enrichment facility.

In a verification scenario, information such as material flows could be reconstructed and then be compared to archived data, such as shipping notes between facilities. Thus, this method could open up a multitude of possibilities to cross-check documentation with inference data and to verify any indications made by the entity under inspection.

6 Conclusion

This paper introduced Bicyclus, a software package aimed at investigating fissile material production in NFCs. We explained its working principle and showcased its use with an NFC with plutonium and HEU production capabilities. Here, we demonstrated how to obtain measurement-independent fissile material estimates, as well as how to reconstruct the fissile material production using measurement data. Last, we showed how the extensive Cyclus output data can be used to reconstruct additional quantities, which could be of interest in a verification scenario.

In a next step, we will investigate how accurately Cyclus depicts actual nuclear programmes. Especially in the early stages of a nuclear programme, the NFC may well be an inefficient cycle with losses and under development—opposing the picture drawn by Cyclus, which optimises the material flows in each timestep. Investigating such potential discrepancies will help understand model uncertainties, leading to a more robust overall uncertainty assessment.

Acknowledgements

We thank Lewin Bormann for his contribution to the initial version of what has now become Bicyclus. This work has been funded by the Volkswagen Foundation. Simulations were performed with computing resources granted by RWTH Aachen University under project `rwth0774`.

Reproducible Research

The source code presented and used here is open source and available at `github.com/Nuclear-Verification-and-Disarmament/bicyclus_inmm_2023`. Bicyclus is available at `github.com/Nuclear-Verification-and-Disarmament/bicyclus`.

References

[1] Alexander Glaser and Malte Göttsche. "Fissile Material Stockpile Declarations and Cooperative Nuclear Archaeology". In: *Verifiable Declarations of Fissile Material Stocks: Challenges and Solutions*. FM(C)T Meeting Series. UNIDIR, 2017, pp. 25–38.

[2] International Panel on Fissile Materials. *Global Fissile Material Report 2010. Balancing the Books: Production and Stocks*. Fifth annual report of the International Panel on Fissile Materials. Dec. 17, 2010.

[3] Kathryn D. Huff et al. "Fundamental concepts in the Cyclus nuclear fuel cycle simulation framework". In: *Advances in Engineering Software* 94 (Apr. 2016), pp. 46–59. ISSN: 0965-9978. DOI: `10.1016/j.advengsoft.2016.01.014`.

[4] I. M. Sobol'. "On the distribution of points in a cube and the approximate evaluation of integrals". In: *USSR Computational Mathematics and Mathematical Physics* 7.4 (1967), pp. 86–112. ISSN: 0041-5553. DOI: `10.1016/0041-5553(67)90144-9`.

[5] Andrew Gelman et al. *Bayesian data analysis*. Third Edition. Texts in statistical science. Boca Raton, FL: CRC Press, 2014. ISBN: 9781439840955.

[6] Max Schalz. *Bicyclus*. Source Code. 2022. URL: `https://github.com/Nuclear-Verification-and-Disarmament/bicyclus` (visited on 04/20/2023).

[7] Max Schalz, Lewin Bormann, and Malte Göttsche. "Using Fuel Cycle Simulators and Bayesian Inference in Nuclear Archaeology". In: *Transactions of the American Nuclear Society* 126.1 (June 2022), pp. 122–125. ISSN: 0003-18X. DOI: `10.13182/T126-37933`.

[8] Pauli Virtanen et al. "SciPy 1.0: fundamental algorithms for scientific computing in Python". In: *Nature Methods* 17.3 (Mar. 1, 2020), pp. 261–272. ISSN: 1548-7105. DOI: `10.1038/s41592-019-0686-2`.

[9] John Salvatier, Thomas V. Wiecki, and Christopher Fonnesbeck. "Probabilistic programming in Python using PyMC3". In: *PeerJ Computer Science* 2 (Apr. 6, 2016), e55. DOI: `10.7717/peerj-cs.55`.

[10] Zia Mian, A. H. Nayyar, and R. Rajaraman. "Exploring Uranium Resource Constraints on Fissile Material Production in Pakistan". In: *Science & Global Security* 17.2 (2009), pp. 77–108. DOI: `10.1080/08929880902975834`.

[11] NEA. *Uranium 2022: Resources, Production and Demand*. Paris: OECD Publishing, Feb. 2023. Previous editions were used, as well.

[12] Tamara Patton. "Combining Satellite Imagery and 3D Drawing Tools for Nonproliferation Analysis: A Case Study of Pakistan's Khushab Plutonium Production Reactors". In: *Science & Global Security* 20.2–3 (2012), pp. 117–140. DOI: `10.1080/08929882.2012.719383`.