Fundamental Concepts in the Cyclus Fuel Cycle Simulator Framework

Kathryn D. Huff^{a,*}, Matthew J. Gidden^b, Robert W. Carlsen^b, Robert R. Flanagan^c, Meghan B. McGarry^b,
Arrielle C. Opotowsky^b, Erich A. Schneider^c, Anthony M. Scopatz^d, Paul P.H. Wilson^b

^aUniversity of California - Berkeley, Department of Nuclear Engineering Berkeley, CA, 94720

^bUniversity of Wisconsin - Madison, Department of Nuclear Engineering and Engineering Physics, Madison, WI 53706

^cUniversity of Texas - Austin, Department of Mechanical Engineering, Nuclear and Radiation Engineering Program, Austin, TX 78758

^dUniversity of South Carolina, Nuclear Engineering Program, Columbia, SC 29201

Abstract

As nuclear power expands, technical, economic, political, and environmental analyses of nuclear fuel cycles by simulators increase in importance. To date, however, current tools are often fleet-based rather than discrete and restrictively licensed rather than open source. Each of these choices presents a challenge to modeling fidelity, generality, efficiency, robustness, and scientific transparency. The CYCLUS nuclear fuel cycle simulator framework and its modeling ecosystem incorporate modern insights from simulation science and software architecture to solve these problems so that challenges in nuclear fuel cycle analysis can be better addressed. A summary of the CYCLUS fuel cycle simulator framework and its modeling ecosystem are presented. Additionally, the implementation of each is discussed in the context of motivating challenges in nuclear fuel cycle simulation. Finally, the current capabilities of CYCLUS are demonstrated for both open and closed fuel cycles.

Keywords: nuclear fuel cycle, simulation, agent based modeling, nuclear engineering, object orientation, systems analysis

1. Introduction

As nuclear power expands, technical, economic, political, and environmental analyses of nuclear fuel cycles by simulators increase in importance. The merits of advanced nuclear technologies and fuel cycles are shaped by myriad physical, nuclear, chemical, industrial, and political factors. Nuclear fuel cycle simulators must therefore couple complex models of nuclear process physics, facility deployment, and material routing.

Indeed, the cardinal purpose of a dynamic nuclear fuel cycle simulator is to calculate the time- and facility-dependent mass flow through all or part the fuel cycle. Dynamic nuclear fuel cycle analysis more realistically supports a range of simulation goals than static analysis [1]. Historically, dynamic nuclear fuel cycle simulators have calculated fuel cycle mass balances and performance metrics derived

from them using software ranging from spreadsheet-driven flow calculators to highly specialized system dynamics modeling platforms.

To date, current tools are typically distributed under restrictive rather than open source licenses, having been developed in industrial contexts or using commercial software platforms. Additionally, having often been developed for customized applications, many possess inflexible architectures, never having been designed to enable new features or extensions. Finally, many model only fleet-level dynamics of facilities and materials rather than discrete resolution of those individual agents and objects. When the DOE-NE Fuel Cycles Technologies Systems Analysis Campaign developed requirements necessary in a next generation fuel cycle simulator, three main failure modes were associated with those software choices. First, they discourage targeted contribution and collaboration among experts. Next, they hobble efforts to directly compare modeling methodologies. Finally, they over-specialize, rendering most tools applicable to only a subset of desired simula-

*Corresponding Author

Email address: huff@berkeley.edu (Kathryn D. Huff)

tion fidelities, scales, and applications. Those three constraints were identified as presenting significant challenges to modeling fidelity, generality, efficiency, robustness, and scientific transparency in the field of fuel cycle analysis [2].

The CYCLUS nuclear fuel cycle simulator framework and its *modeling ecosystem*, the suite of agents and other physics plug-in libraries compatible with it, incorporate modern insights from simulation science and software architecture to solve these problems. These modern methods simultaneously enable more efficient, accurate, robust, and validated analysis. This next-generation fuel cycle simulator is the result of design choices made to:

- support access to the tool by fuel cycle analysts and other users,
- encourage developer extensions,
- enable plug-and-play comparison of modeling methodologies,
- and address a range of analysis types, levels of detail, and analyst sophistication.

CYCLUS is a dynamic, agent-based model, which employs a modular architecture, an open development process, discrete agents, discrete time, and arbitrarily detailed isotopic resolution of materials. Experience in the broader field of systems analysis indicates that agent-based modeling enables more flexible simulation control than system dynamics, without loss of generality [3]. Furthermore, openness allows cross-institutional collaboration, increases software robustness [4, 5], improves the strength and quality of results through peer review [6, 7, 8, 9, 10], and cultivates an ecosystem of modeling options. This ecosystem is *modular*, being comprised of dynamically loadable, interchangeable, plug-in libraries of fuel cycle component process physics that vary in their scope, depth, and fidelity. This modularity allows users and developers to customise CYCLUS to analyze the cases that are of interest to them rather than any custom application the simulator was originally developed to address. Additionally, that customizability allows users and developers to address those cases at the level of fidelity necessary for their application. The fundamental concepts of the CYCLUS nuclear fuel cycle simulator capture these modern insights so that novel challenges in nuclear fuel cycle analysis can be better addressed.

1.1. Background

Nuclear fuel cycle simulators drive research development and design (RD&D) by calculating ‘metrics’, quantitative measures of performance that can be compared among fuel cycle options. The feasibility of the technology development and deployment strategies which comprise a fuel cycle option, the operational features of nuclear energy systems, the dynamics of transitions between fuel cycles, and many other measures of performance can be expressed in terms of these metrics. For example, economic feasibility is often measured in levelized cost of electricity (LCOE), a combination of fuel and operating costs normalized by electricity generation, while environmental performance might be measured by spent fuel volume, radiotoxicity, or mined uranium requirements. A meta-analysis of fuel cycle systems studies identified over two dozen unique quantitative metrics spanning economics and cost, environmental sustainability and waste management impacts, safety, security and nonproliferation, resource adequacy and utilization, among others [11]. With few exceptions, these metrics are derived from mass balances and facility operation histories calculated by a fuel cycle simulator. For example, where nuclear waste repository burden is derived from ejected fuel masses, water pollution or land use can be derived from facility operational histories (as in [12]).

However, methods for calculating those metrics vary among simulators. Some model the system of facilities, economics, and materials in static equilibrium, while other simulators capture the dynamics of the system. Similarly, while some simulators discretely model batches of material and individual facilities, others aggregate facilities into fleets and materials into streams. Some simulators were designed to model a single aspect of the fuel cycle in great detail while neglecting others. For example, a simulator created for policy modeling might have excellent capability in economics while capabilities for tracking transformations in material isotopes and the effects of isotopes on technology performance are neglected. The Code for Advanced Fuel Cycles Assessment (CAFCA)[13] simulator is problem-oriented in this way, having elected to neglect isotopic resolution in favor of integral effects.

Historically, domestic national laboratories have driven development and regulated the use of their own tools: the Verifiable Fuel Cycle Simulation Model (VISION)[14], Dynamic Model of Nuclear Development (DYMOND)[15], and Nuclear Fuel Cycle Simulator (NFCSim)[16, 17]. Internationally,

other laboratories have created their own as well, such as Commelini-Sicard (COSI)[18, 19, 20, 21] and ORION[22]. Finally, some simulators initiated in a national lab setting have been continued as propriety, industry-based simulators, such as Dynamic Analysis of Nuclear Energy System Strategies (DANESS)[23]. Outside the national laboratories, researchers have created new nuclear fuel cycle simulation tools when existing tools were not available or not sufficiently general to calculate their metrics of interest. With limited access to the national laboratory tools and a need to customize them for research purposes, universities and private industry researchers have “reinvented the wheel” by developing tools of their own from scratch and tailored to their own needs. Examples include CAFCA[24] and Dynamic Analysis of Nuclear Energy Systems Strategies (DESAE)[25, 26, 17].

CYCLUS emerged from a line of tools seeking to break this practice. Its precursor, Global Evaluation of Nuclear Infrastructure Utilization Scenarios (GENIUS) Version 1 [27, 28], originated within Idaho National Laboratory (INL) and sought to provide generic regional capability. Based on lessons learned from GENIUS Version 1, the GENIUS Version 2 [29, 30] simulator sought to provide more generality and an extensible interface to facilitate collaboration. The CYCLUS project then improved upon the GENIUS effort by implementing increased modularity and encapsulation. The result is a dynamic simulator that treats both materials and facilities discretely, with an architecture that permits multiple and variable levels of fidelity. Using an agent-based framework, the simulator tracks the transformation and trade of resources between autonomous regional and institutional entities with customizable behavior and objectives. Each of these concepts (agent-based, resource tracking, and regional as well as institutional entities) will be described in their own sections (sections 2.2, 2.3.1, and 2.2.2 respectively). Together, they provide a capability for extension and reuse beyond that pursued by any existing fuel cycle simulator.

1.2. Motivation

The CYCLUS paradigm enables targeted contribution and collaboration within the nuclear fuel cycle analysis community to achieve two important goals: lower the barrier for users to include custom nuclear technologies and facility types in their fuel cycle analyses while improving the ability to compare simulations with and without those custom concepts.

This essential capability is absent in previous simulators where user customization and extensibility were not design objectives. While the *modular and open architecture* of CYCLUS is necessary to meet these goals, it is not sufficient. *Agent interchangeability* is also required to facilitate direct comparison of alternative modeling methodologies and facility concepts. With this concept at its core, CYCLUS provides a platform for users to quickly develop the capabilities at a level of detail and validation necessary for their unique applications. Finally, CYCLUS is applicable to a broader range of fidelities, scales, and applications than other simulators, due to the flexibility and generality of its *agent-based modeling (ABM)* paradigm and *discrete, object-oriented approach*.

This structure recognizes that specialists should utilize their time and resources in modeling the specific process associated with their area of expertise (e.g., reprocessing and advanced fuel fabrication), without having to create a model of the entire fuel cycle to serve as its host. CYCLUS supports them by separating the problem of modeling a physics-dependent supply chain into two distinct components: a simulation kernel and archetypes that interact with it. The kernel is responsible for supporting the deployment and interaction logic of entities in the simulation. Physics calculations and customized behaviors of those entities are implemented within *archetype* classes.

Ultimately, modeling the evolution of a physics-dependent, international nuclear fuel supply chain is a multi-scale problem which existing tools cannot support. They have either focused on macro effects, e.g., the fleet-level stocks and flows of commodities, or micro effects, e.g., the used-fuel composition of fast reactor fuel. Each focus has driven the development of specialized tools, rendering the task of answering questions between the macro and micro levels challenging within a single tool. In contrast, the open, extensible architecture and discrete object tracking of CYCLUS allow the creation and interchangeability of custom archetypes at any level of fidelity and by any fuel cycle analyst.

1.2.1. Open Access and Development Practices

The proprietary concerns of research institutions and security constraints of data within fuel cycle simulators often restrict access. Use of a simulator is therefore often limited to its institution of origin, necessitating effort duplication at other institutions and thereby squandering broader human resources. License agreements and institutional ap-

proval are required for most current simulators (e.g. COSI6, DANESS, DESAE, EVOLCODE, FAMILY21, NFCSim)[31], including ORION, and VISION. Even when the source code is unrestricted, the platform on which it relies (e.g. VENSIM) is often restricted or costly. The MIT CAFCA software, for example, relies on the commercially licensed VENSIM product as a platform. CYCLUS, on the other hand, is written in C++ for which freely available development tools and an open standard are available. Further, CYCLUS relies only on open source, freely available libraries. As such, it provides fully free and open access to all users and developers, foreign and domestic.

Moreover, both technical and institutional aspects of the software development practices employed by the CYCLUS community facilitate collaboration. Technically, CYCLUS employs a set of tools commonly used collaborative software development that reduce the effort required to comment on, test and ultimately merge individual contributions into the main development path. For many of the simulation platforms adopted by previous simulators, there were technical obstacles that impeded this kind of collaboration. Institutionally, CYCLUS invites all participants to propose, discuss and provide input to the final decision making for all important changes.

1.2.2. Modularity and Extensibility

Modularity is a key enabler of extending the scope of fuel cycle analysis within the CYCLUS framework. Changes that are required to improve the fidelity of modeling a particular agent, or to introduce entirely new agents, are narrowly confined and place no new requirements on the CYCLUS kernel. Furthermore, there are few assumptions or heuristics that would otherwise restrict the algorithmic complexity that can be used to model the behavior of such agents.

For example, most current simulators describe a finite set of acceptable cycle constructions (once through, single-pass, multi-pass). That limits the capability to create novel material flows and economic scenarios. The CYCLUS simulation logic relies on a market paradigm, parameterized by the user, which flexibly simulates dynamic responses to pricing, availability, and other institutional preferences.

This minimal set of mutual dependencies between the kernel and the agents is expressed through the dynamic resource exchange (DRE) that provides a level of flexibility that does not exist in other fuel cycle simulators. It creates the potential for novel agent archetypes to interact with existing archetypes

as they enter and leave the simulation over time and seek to trade materials whose specific composition may not be known *a priori*.

1.2.3. Discrete Facilities and Materials

Many fuel cycle phenomena have aggregate system-level effects which can only be captured by discrete material tracking [2]. CYCLUS tracks materials as discrete objects. Some current fuel cycle simulation tools such as COSI [26, 32, 24], FAMILY21[26], GENIUS version 1, GENIUS version 2, NFCSim, and ORION also possess the ability to model discrete materials. However, even among these, the ability to model reactor facilities individually is not equivalent to the ability to model distinct activities. COSI, for example, has some support for modeling reactors individually, but according to a recent benchmark [33], it models many reactors operating in sync. That is, refuelling and discharging occur simultaneously for all reactors. While CYCLUS allows this type of fleet-based aggregation of reactor behavior, CYCLUS also enables operations in each facility to vary independently of any others in the simulation.

Similarly, the ability to model disruptions (i.e. facility shutdowns due to insufficient feed material or insufficient processing and storage capacity) is most readily captured by software capable of tracking the operations status of discrete facilities [2]. Fleet-based models (e.g. VISION) are unable to capture this gracefully, since supply disruptions are modeled as a reduction in the capacity of the whole fleet. All of the software capable of discrete materials have a notion of discrete facilities, however not all handle disruption in the same manner. DESAE, for example, does not allow shutdown due to insufficient feedstock. In the event of insufficient fissile material during reprocessing, DESAE borrows material from storage, leaving a negative value [26]. The CYCLUS framework does not dictate such heuristics. Rather, it provides a flexible framework on which either method is possible.

A final benefit of the discreteness of facilities and materials is their power when combined. The ability to track a material's history as it moves from one facility to another is unique to CYCLUS. While some current simulators track materials in discrete quanta, they do not necessarily preserve the identity of each quantum as the materials move around the fuel cycle. When coupled with CYCLUS' individual facility modeling, this capacity becomes distinct from what other fuel cycle simulators are able to do.

So, while FAMILY21 and COSI can identify whether a batch being discharged from a reactor originated in mixed oxide (MOX) fabrication rather than fresh uranium oxide (UOX) fabrication, CYCLUS can go further, tracking which of the fresh batches contained material from a particular discharged batch. By extension, CYCLUS can also report which individual facilities the batch passed through and in which it originated. The ability to track a single material through the simulation, though it might be split, transmuted, or merged with other materials along the way, allows CYCLUS to answer more data-rich questions that previous simulators have been unable to ask. For example, it allows precise tracking of specific material diversions, so queries about nonproliferation robustness in a facility can be levied either in the context of a single event or a series of nefarious acts.

2. Methodology and Implementation

A modular, agent-based modeling (ABM) approach is ideal for modeling the coupled, physics-dependent supply chain problems involving material routing, facility deployment, and regional and institutional hierarchies which arise in fuel cycle analysis. Additionally, the choice to build CYCLUS on open source libraries in modern programming languages enables both remote and multiprocess execution on a number of platforms. This section begins by describing the general design features that make CYCLUS both flexible and powerful: *cluster-ready software* and *dynamically loadable libraries*. The ABM framework is then described, focusing on its implementation and benefits in a fuel cycle context. A discussion of the time-dependent treatment of discrete resources follows, focusing on the DRE. Support for users and developers via the CYCAMORE library of archetypes and the experimental toolkit are also presented. Lastly, the methods for quality assurance are outlined.

2.1. Modular Software Architecture

The architecture of CYCLUS allows developers to define nuclear fuel cycle processes independent of the simulation logic. To achieve this, *agents* are developed which represent facilities, institutions, and regions comprising the nuclear fuel cycle. These agents are created using the CYCLUS framework application programming interface (API), a set of functions and protocols which assist in agent development and specify how agents should be defined.

This encapsulated ‘plug-in’ design choice provides two major benefits. First, analysts can take advantage of the simulation logic API and archetype ecosystem when they apply CYCLUS to their specific problem. A modeler can focus on creating or customizing nuclear facility, institution, resource, and toolkit models within their specific area of technical expertise. Second, because CYCLUS uses a modular archetype approach, comparing two archetypes is straightforward. For example, if an analyst would like to compare the effect of using different models to determine the input fuel composition for fast reactors, fuel fabrication archetypes can be developed and interchanged while keeping the rest of the models used in the simulation fixed.

2.1.1. Cluster-Ready Software

Many fuel cycle simulators rely on commercial, off-the-shelf (COTS) and Windows-only software that limits performance on and compatibility with resource computing infrastructures (e.g. cluster or cloud computing). This constrains the possible scope of simulations and increases the wall-clock time necessary to conduct parameterized sensitivity analyses and other multi-simulation studies. For example, large scale sensitivity analyses to quantify the dependence of fuel cycle outcomes are only feasible in a massively parallelized environment. To enable such research, CYCLUS, is primarily written in C++ and relies on libraries supported by Linux and UNIX (including Ubuntu and OSX) platforms, which are flexible and support parallelization.

Cyclopts [34], a proof-of-principle CYCLUS-enabled application on such a large compute system, uses UW-Madison’s HTCondor high-throughput computing (HTC) infrastructure to study DRE performance and outcomes. For example, an investigation of the effect of solution degeneracy, a commonly observed phenomenon in scenarios with individual facilities and basic (e.g., commodity-only) preference definitions, was performed for three different fuel cycles: once-through, MOX fast reactor-thermal reactor cycles, and MOX-thorium oxide (ThOX) recycle in fast reactors-thermal reactor cycles. Objective coefficients were generated based on two factors: commodity-facility pairings and facility location. The population of the possible values of commodity-facility pairings is always small ($\mathcal{O}(10)$). The size of the population of possible values of the location effect was investigated for values of zero, ten, and infinity (i.e., any real number). The study additionally included an investigation of problem-

scaling behavior in order to quantify the magnitude and rate-of-increase of DRE solution times as a function of the simulation-entity population for each fuel cycle[35]. In total, Cyclopts has run over 10^5 jobs, comprising more than 60,000 compute hours. The HTC infrastructure has separately been utilized to run and collect information from full CYCLUS simulations running in parallel on 10^3 machines reliably for order 10^5 simulations.

2.1.2. Dynamically Loadable Libraries

The CYCLUS architecture encourages efficient, targeted contribution to the ecosystem of archetype libraries. With CYCLUS, a researcher can focus on generating an archetype model within their sphere of expertise while relying on the contributions of others to fill in the other technologies in the simulation. Similarly, individual developers may explore different levels of complexity within their archetypes, including wrapping other simulation tools as loadable libraries within CYCLUS.

CYCLUS achieves this behavior by implementing generic APIs and a modular architecture via a suite of dynamically loadable plug-in libraries (pictured in Figure 2.1.2). By anticipating the possible classes of information required by the simulation kernel, the CYCLUS APIs facilitate information passing between the plug-in agents and the core framework. Though common in modern software architecture, such a plug-in paradigm has not previously been implemented in a nuclear fuel cycle simulator. It allows the core CYCLUS framework to operate independently from the plug-in libraries, and the dynamically loadable plug-ins to be the primary mechanism for extending CYCLUS' capabilities independent of the core.

An additional benefit is the ability for contributors to choose different distribution and licensing strategies for their contributions. Users and modelers control the accessibility of their archetypes and data sets (See Figure 2.1.2). In particular, since the clean plug-in architecture loads libraries without any modifications to the CYCLUS kernel, closed-source archetypes can be used with the simulator alongside open source archetypes without transfer of sensitive information. This architecture allows closed-source libraries (e.g., those representing sensitive nuclear processes and subject to export control) to be developed and licensed privately.

Finally, dynamically loadable libraries enable CYCLUS to easily handle varying levels of simulation complexity. Hence a single simulation engine can be

Figure 1: The CYCLUS core provides APIs that abstract away the details in the kernel and allow the archetypes to be loaded into the simulation in a modular fashion.

Figure 2: The CYCLUS framework enables fully open, partially open, and fully closed collaborations[36].

used by both users interested in big-picture policy questions as well as users focused on detailed technical analyses. They merely choose their preferred level of modeling depth from among the available libraries in the ecosystem.

2.2. Agent-based Paradigm

CYCLUS implements an agent-based modeling paradigm. ABM enables model development to take place at an agent level rather than a system level. In the nuclear fuel cycle context, for example, an analyst can design a reactor agent that is entirely independent from a fuel fabrication agent. Defining the behavior of both agents according to the API contract is sufficient for them to interact with one another as bona fide agents in the simulation. The two archetype libraries can be used in the same simulation without any shared knowledge, allowing modelers to construct a simulation from building blocks of many types and origins.

Furthermore, the ABM paradigm is more flexible and intuitive (from both a developer and user perspective) than the system dynamic approach used in current simulators. System dynamics is a popular approach for modeling nuclear fuel cycles [37, 23, 13, 24]. Formally however, system dynamics models are simply a strict subset of agent-based models [3]. That is, any system dynamics model can be translated into an agent-based model. ABM techniques therefore enable a broader range of simulations in a more generic fashion.

2.2.1. Agent Interchangeability

ABM is inherently object-oriented because agents represent discrete, independently acting objects. Figure 2.1.2 illustrates the modular nature of CYCLUS archetypes. The core of the CYCLUS simulator creates a set of classes on which agent plug-ins are based. Agent plug-ins utilize the generic core APIs that define agent-to-agent interaction as well as agent-to-environment interaction. For example, they use the resource exchange paradigm API for trading resources with one another. For the archetype developer, these interfaces provide enormous power simply. The API abstracts away details unnecessary to specifying the archetype behavior, while providing all necessary functionality for interacting with the CYCLUS simulation kernel. For the user, multiple archetypes that inherit the same APIs are interchangeable in a simulation.

Critically, this novel functionality enables the comparison between agent implementations. For example, an archetype that inherits the **Region** interface, as in Figure 3, is interchangeable with any other Region agent.

Figure 3: Inheriting CYCLUS class interfaces, such as the **Agent**, **Facility**, **Institution**, and **Region** classes, abstracts away unnecessary details while exposing powerful functionality. In the above example, the **Dummy** archetype simply inherits from **Region** in order to become a bona fide Region-type **Agent**.

In this way, a researcher can directly compare two different reactor modeling implementations (perhaps the imaginary classes **DetailedReactor** and **SimpleReactor**) simply by exchanging the two corresponding archetypes. That is, two reactor archetypes both inheriting from the **Facility** class are indistinguishable from a simulation perspective. This can be done with any agent type, where agents can be “Regions,” “Institutions,” or “Facilities.”

2.2.2. Regions, Institutions, and Facilities

CYCLUS provides a novel representation of entities in the nuclear fuel cycle that reflects the reality in international nuclear power: facilities implementing individual fuel cycle technologies, institutions managing those facilities, and regions providing geographical and political context for institutions and facilities. While few simulators have provided any notion of static regional effects [2, 31], CYCLUS allows for both regions and institutions to be first-class agents in simulated fuel cycles. The fundamental interactions for each entity are implemented in a corresponding archetype class in CYCLUS, i.e., the **Region** class, **Institution** class, and **Facility** class. Archetype developers can then build on the provided functionality by inheriting from the appropriate class. CYCLUS implements a Region-Institution-Facility (RIF) relationship through a parent-child hierarchy where regions are the parents of institutions which are, in turn, the parents of

facilities. In other words, RIF hierarchies form a directed acyclic graph (DAG)¹, with regions as root nodes and facilities as leaf nodes.

Two primary consequences arise from this structure. First, institutions are nominally responsible for deploying and decommissioning facilities. Accordingly, advanced logic regarding building and decommissioning can be implemented on top of those behaviors inherited from the **Institution** interface. Second, the **Facility** class implements the **Trader** interface to participate in resource exchange, and institutions and regions, respectively, can adjust the resource flow preferences of their managed facilities. Importantly, this capability allows for the modeling of preferential regional trading of resources (e.g., tariffs) as well as preferential institutional trading (e.g., long-term contracts).

2.3. Discrete Objects

CYCLUS models facilities, institutions, regions, and resources as discrete objects. A discrete resource model allows for a range of modeling granularity. In the macroscopic extreme, it is equivalent to time-stepped continuous flow. In the microscopic extreme, the model is capable of representing arbitrarily small material objects at isotopic resolution. In this way, CYCLUS is applicable across the full range of modeling fidelity.

Fleet-based, lumped-material models do not distinguish between discrete facility entities or materials. However, some questions require resolution at the level of individual facilities and materials. As a result, many detailed performance metrics cannot be captured with previously existing fleet-based models. For example, meaningful models of spent nuclear fuel storage transport, and disposal strategies, require representation of discrete casks and their varying isotopic compositions.

For all of the reasons that the ABM paradigm in 2.2 enables novel simulations, multiple use cases require that these agents, such as the regions, institutions, and facilities in CYCLUS, must be represented as discrete objects. For instance, tracking of individual shipments is only viable if materials and resources are tracked as discrete objects.

¹ DAGs are common graph theoretic structures whose most important feature is a lack of *cycles*, i.e., there is a single path from the root node to any other node in the graph.

2.3.1. Resources and Materials

Another such use case seeks to capture system vulnerability to material diversion. Provenance and trade-history of distinct materials is the fundamental information unit in such studies, and so this type of analysis requires discrete simulation of a target facility and the individual materials modified within it. Material risk analysis, therefore, demands that both facilities and materials should be discretely modeled objects like those in CYCLUS.

In CYCLUS, agents can transfer discrete resource objects among one another. Cyclus supports two types of resources:

- materials: these represent typical nuclear materials with nuclide compositions;
- products: these can represent any user-defined measure: carbon credits, build permits, employees, etc..

All operations performed on every resource object (splitting, combining, decay, etc.) are tracked in detail as they are performed. This information includes the agent that created each resource when it was introduced into the simulation. The parentage of each resource is also tracked. This makes it possible to follow the history of resources as they are transferred between agents.

The CYCLUS kernel has built-in experimental support for decay calculations. Materials store the time since their last decay and agents are free to invoke the decay function on them as desired to decay them to the current simulation time. CYCLUS can operate in 3 decay modes, with 1 additional mode likely to be added in a future release:

- “manual” (currently implemented) is the default mode where agents decay materials when requested by an archetype,
- “never” (currently implemented) globally turns off all decay. The Material decay function does nothing,
- “lazy” (currently implemented) decays material automatically whenever its composition is observed (e.g. when an agent queries information about a material’s ²³⁹Pu content),
- “periodic” (future) automatically decays all materials in a simulation with some fixed frequency.

When decay is invoked, a material checks to see if it contains any nuclides with decay constants that are significant with respect to the time change since the last decay operation. If none of the decay constants are significant, no decay calculation is performed and the material remains unchanged. This error does not accumulate because the next time the material's decay function is invoked, the time change will be larger.

CYCLUS has no notion of “tracked” versus “untracked” nuclides. In CYCLUS, the composition of a material is represented by an arbitrarily large list (potentially thousands) of nuclides. Agents are free to treat nuclides present in materials any way they please - including ignoring them. It is the responsibility of archetype developers to choose how to handle potentially full-fidelity compositions.

In large simulations, many material objects may change frequently. Material decay can also contribute significantly to such changes. In order to help avoid unnecessary runtime performance and database size impacts, compositions in CYCLUS have some special features. In particular, compositions are immutable once created. This allows multiple material objects to hold references to the same composition safely. Additionally, new compositions resulting from decay are cached and used to avoid redundant decay calculations. Figure 4 illustrates how this decay history cache works. Composition immutability in concert with decay history caching help eliminate many redundant calculations in addition to reducing the total number of composition entries recorded in the database.

2.3.2. Dynamic Resource Exchange (DRE)

The CYCLUS simulation paradigm allows discrete agents, based on archetypes about which the kernel has no knowledge, to enter the simulation at arbitrary times and trade in discrete resources. These resources are not defined *a priori*. Therefore, the logic engine defining resource interaction mechanisms among agents is crucial. The DRE, described in detail in [35], is that critical logic engine on which CYCLUS simulations are built. Supporting the general CYCLUS philosophy, facilities are treated as black boxes and a supply-demand communication framework is defined.

The DRE consists of three steps: supply-demand information gathering, resource exchange solution, and trade execution. Importantly, each step is agnostic with respect to the exchange of resources in

Figure 4: A simple decay line cache holding compositions A, B, C, and a yet-uncomputed composition D. B comes from decaying A 1 time step. C comes from decaying B 2 time steps, etc. Black represents the cache for this particular composition chain. Blue indicates decay operations that can be satisfied by cache lookups. If A needs to be decayed 1 time step (A +1), a quick lookup returns the previously computed B. Decaying B 3 time steps requires a decay calculation to compute a new composition D, but subsequent requests such as decaying A 4 time steps will not require any recalculation.

question, i.e., the same procedure is used for both Materials and Products.

The information-gathering step begins by polling potential consumers. Agents define both the quantity of a commodity they need to consume as well as the target isotopics, or quality, by posting their demand to the market exchange as a series of *requests*. Users may optionally parameterize the agent to associate a collection of demand constraints with each collection of requests. Collections of requests may be grouped together, denoting *mutual* requests that represent demand for some common purpose. For example, a reactor may request UOX and MOX fuel mutually, indicating that either will satisfy its demand for fuel.

Suppliers then respond to the series of requests with a *bid*. A bid supplies a notion of the quantity and quality of a resource to match a request. Suppliers may add an arbitrary number of constraints to accompany bids. For example, an enriched UOX supplier may be constrained by its current inventory of natural uranium or its total capacity to provide enrichment in Separative Work Units (SWUs). It attaches such constraints to its bids.

Any potential resource transfer, i.e., a bid or a request, may be denoted as *exclusive*. An exclusive

transfer excludes partial fulfilment; it must either be met fully or not at all. This mode supports concepts such as the trading of individual reactor assemblies. In combination with the notion of mutual requests, complex instances of supply and demand are enabled.

Finally, requesting facilities, institutions and regions may apply *preferences* to each potential request-bid pairing based on the proposed resource transfer. Facilities can apply arbitrary complex logic to rank the bids that they have received, whether based on the quantity available in each bid or on the quality of each bid, and the consequent implications of the physics behavior of that facility. In addition, an institution can apply a higher preference to a partner to which it is congenial; similarly, a region may negate any transfers of material which have a higher uranium enrichment than is allowable.

Figure 5: The flow graph representing three suppliers (left), two requesters (right), and the potential flows of various commodities among them. The second consumer makes two different requests. Meanwhile, the second supplier can supply the commodities requested by both consumers and has provided two bids accordingly [35].

Given a full definition of supply and demand, represented in Figure 5 as a flow graph, the DRE may be solved either optimally using a mathematical program or approximately by a simulation-based heuristic [35]. If any trade is denoted as exclusive, e.g., if an analyst desires an assembly-fidelity model, then either a heuristic must be used or exchanges must be represented as a mixed-integer linear program (MILP). If no exclusive trades exist, a linear program (LP) model may be used. The LP formulation in CYCLUS is of the form given by Gidden

[35]:

$$\min_x z = c^\top x \quad (1)$$

$$\text{s.t. } Ax \leq b \quad (2)$$

$$x_{i,j} \in [0, \tilde{x}_j] \quad \forall i \in I, \forall j \in J \quad (3)$$

where

I = set of supply nodes

J = set of request nodes

$c_{i,j}$ = unit cost of flow from i to j

$x_{i,j}$ = flow from node i to node j

$a_{i,j}^k$ = coefficient for constraint k between i and j

$b_{s|r}^k$ = RHS for constraint k for a given supplier or requester

$\tilde{x}_j$ = requested quantity of request node j .

In practice, LPs solve much faster than MILPs. However, both capabilities exist in CYCLUS in order to provide users with the desired level of fidelity.

Trades between agents are initiated by the CYCLUS kernel after a solution to the DRE is found. For each trade, the supplying agent is notified of its matched request and provides an associated resource to the exchange. All supplied resources are then sent to the corresponding requesting agents.

In CYCLUS, the DRE is executed at each time step. Therefore, if a facility's request for a resource is not met at a given time step, it may offer a request in the following time step. Because agent behavior may change arbitrarily, the exchange executed at any given time step may be unique in a simulation.

The DRE is a novel simulation concept in the nuclear fuel cycle domain. It provides a flexible supply-demand infrastructure, supporting dynamic flows of resources between agents, even as those agents enter and leave the simulation, and even when those agents are defined by archetypes of arbitrary complexity. Trading between agents can be affected by both the proposed quality of a resource and agent relationships through the use of preferences. Accordingly, a wide range of possible effects can be modeled, from capacity-limited fuel supply to international trade agreements.

2.4. Simulation Support

So that users and developers can build working simulations in the shortest time possible, the

CYCLUS ecosystem provides fundamental building blocks: basic archetypes and a toolkit of commonly needed functions. The CYCAMORE library provides a suite of fundamental Region, Institution, and Facility archetypes, while the CYCLUS toolkit provides assistance to developers.

2.4.1. Cycamore

CYCAMORE [38], the CYCLUS additional module repository, provides a fundamental set of dynamically loadable libraries providing agent archetypes for basic simulation functionality within CYCLUS. Since CYCLUS relies on external archetypes to represent the agents within a simulation, CYCAMORE provides the basic archetypes a new user needs to get started running simple simulations. These archetypes support a minimal set of fuel cycle simulation goals and provide, by example, a guide to new developers who would seek to contribute their own archetypes outside of CYCAMORE.

As of version 1.0, CYCAMORE contains one region archetype, two institution archetypes, and four facility archetypes. Three additional facilities were added in version 1.3. Short descriptions of these functions can be found in Table 1.

Section 3.2 will demonstrate how the current CYCAMORE release provides basic functionality enabling simple fuel cycle analyses. As future contributions are vetted, the capabilities in CYCAMORE may grow.

2.4.2. Toolkit

In addition to the core functionality of the CYCLUS kernel, which is focused on the set of capabilities needed to implement an agent-based simulation with DRE, a toolkit is provided to assist developers and users with related simulation and nuclear engineering tasks. The toolkit is an actively developed part of CYCLUS, with a primarily forward-looking focus on supporting interesting *in situ* metric analysis tools.

Simulation Tools. A series of utility classes are provided to support demand-constrained agent facility deployment. For example, symbolic function representations of linear, exponential, and piecewise functions are supported via the `SymbFunctionFactory` class. Such functions are used with other toolkit classes to determine commodity demand (e.g., power demand) from user input. Four mix-in

classes provide the basis for in-simulation deployment determination: `CommodityProducer`, `CommodityProducerManager`, `Builder`, `BuildingManager`. The `CommodityProducer` class provides an interface for querying the *prototypes* which have the capacity to produce a given commodity. The `CommodityProducerManager` provides an interface for registering `CommodityProducers` and querying the current capacity (supply) of a commodity. The `Builder` class provides an interface for querying which prototypes can be built and interacts with the `BuildingManager`, which orders prototypes to be built. The `BuildingManager` uses a simple minimum cost algorithm to determine how many of each prototype, y_i , to build given a demand (Φ), capacities (ϕ_i), and costs (c_i). Here i indexes I available prototypes which perform a similar function, and the demand, capacity and cost carry prototype-specific units which are defined by the developer.

$$\begin{aligned} \min \sum_{i=1}^N c_i y_i \\ \text{s.t. } \sum_{i=1}^N \phi_i y_i \geq \Phi \\ y_i \in [0, \infty) \quad \forall i \in I, \quad y_i \text{ integer} \end{aligned} \quad (4)$$

Nuclear Engineering Tools. The CYCLUS toolkit provides two useful interfaces for querying physical parameters of `Material` objects. First, the `MatQuery` tool provides a basic querying API, including the atom and mass fractions of nuclides, the number of moles of a nuclide in a material, etc. (i.e., a `Composition`), in a material. The `Enrichment` tool provides an API for determining enrichment-related parameters of a material, including the SWU and natural uranium required to enrich a material provided knowledge of feed, product, and tails assays.

Toolkit Extensions. In addition to those that already exist, new tools will emerge from the archetypes developed by the community. As these tools gain adoption between projects and demonstrate their utility to the developer community, they will be considered for screening and adoption into the kernel as toolkit extensions. Likely extensions include

- fuel cycle metrics calculators,
- supportive data tables,
- policy models,

Entity	Archetype	Functionality
Facility	EnrichmentFacility	This facility enriches uranium at a specified capacity.
Facility	Fab*	This facility fabricates fuel material based on separated streams.
Facility	Reactor*	A reactor model that handles batch refueling, based on pre-determined recipes of compositions. It requests any number of input fuel types and transmutes them to static compositions upon discharge.
Facility	Separations*	This facility separates an incoming material into specified streams.
Facility	Sink	This facility is capable of accepting a finite or infinite quantity of some commodity produced in the simulation.
Facility	Source	This facility generates material of the composition and commodity type specified as input.
Institution	ManagerInst	The manager institution manages production of commodities among its facilities by building new ones as needed.
Institution	DeployInst	This institution deploys specific facilities as defined explicitly in the input file.
Region	GrowthRegion	This region determines whether there is a need to meet a certain capacity (as defined via input) at each time step.

Table 1: The Archetypes in CYCAMORE seek to cover a large range of simple simulation use cases [38]. Facilities added in version 1.3 are marked with *.

- and economic models.

2.5. Quality Assurance

To garner the trust of a broad user and archetype developer community, the CYCLUS project must implement a strategy to assure the ongoing quality of the software it provides. Multiple strategies, collectively known as *quality assurance (QA)*, have been developed by the scientific software development community to mitigate structural and algorithmic errors in software. These include *verification and validation (V&V)* [39], testing, and others.

Nuclear engineering software quality is often governed by Nuclear Quality Assurance - 1 (NQA-1), an American Society of Mechanical Engineers (ASME) specification whose latest revision appeared in 2009 [40]. CYCLUS has adopted an *agile* development process [41], interpreting NQA-1 in a manner similar to the process adopted by the Department of Energy (DOE) within Nuclear Engineering Advanced Modeling and Simulation (NEAMS) [42] or by the PyNE toolkit [43].

Validation of simulators like CYCLUS, that are intended to give forecasts of system behavior in uncertain futures, is not well defined as there are no experimental benchmarks upon which to rely. Instead, code-to-code comparisons with fuel cycle simulators that use other modeling paradigms are underway as the best approach to establish confidence that CYCLUS produces correct answers. [44]

Verification of CYCLUS implementation relies on software development best practices such as version control, testing, continuous integration, documentation, and style guidelines to ensure reliability and reproducibility. Sections 2.5.1-2.5.3 discuss in greater detail the software development components that comprise the CYCLUS verification strategy, which allows the simulation kernel and physics models to be tested explicitly and separately. Each of these approaches on its own is a valuable addition to QA but it cannot be the entire answer to the requirements imposed by NQA-1. Taken together and strictly adhered to, they present a fortress to protect against poorly designed or otherwise undesirable code.

2.5.1. Version Control

Automated version control, provided by git, a well-established *distributed version control* tool [45], is at the heart of the QA strategy because it records the full history of any change, along with metadata such as the author, a timestamp, and an accompanying message. This makes it possible to identify when changes were introduced, how they are related to other changes, and who made those changes. If necessary, it also facilitates reversing individual changes from a long and intricate set of changes. Version control also enables a code review process described in section 2.5.3 below.

2.5.2. Testing

Automated software *testing* is the first line of defense against errors in implementation. Testing directly compares the actual results of running software versus the expected behavior of the software. In CYCLUS, three categories of tests are defined: unit tests, integration tests, and regression tests. Before a proposed code change is allowed into CYCLUS, the change must be covered by a test, either new or existing, and all tests must pass.

Unit Tests. Unit tests verify behavior of the smallest code *unit*, typically a single function or a class. CYCLUS uses the Google Test framework [46] as a harness for running unit tests. Sufficient unit tests are required for any proposed change to the CYCLUS code base. Currently, CYCLUS implements over 450 unit tests and CYCAMORE implements 85. These cover approximately 65% of their respective code bases, and these numbers are expected to grow over time.

Integration Tests. Integration tests combine multiple elements of the CYCLUS interface and test that they work correctly with each other. In CYCLUS and CYCAMORE, integration tests are performed by running sample simulations for scenarios verifying that results match predictions. A set of standard input files are run, then the output is inspected and compared via Nose [47], a Python test framework.

Regression Tests. Regression tests ensure significant unintended changes do not occur over the course of CYCLUS development. Regression tests are implemented similarly to integration tests. In this category, however, the comparison is done against the output of the same input file when run with a previous version of CYCLUS, typically the last released version. In some sense, regression tests are ‘dumb’ in that they do not care about the contents of a simulation being correct, only whether or not it changed.

Continuous Integration. Continuous integration (CI) is the idea that software should be automatically tested and validated as it is being developed, rather than as a final stage in a longer development cycle. The results of this testing are shown to code reviewers prior to reviewing the software changes themselves. The CYCLUS project uses the CircleCI [48] service for continuous integration. When a code change is submitted for review, CircleCI builds a

version of the CYCLUS source code that includes the requested changes and runs the complete test suite, reporting back whether or not those steps were successful. If CI was not successful, the code author(s) must first identify and resolve the problems. These steps are performed for all incoming code prior to inclusion, so broken code never enters the main software development branch.

2.5.3. Code Review

Automated testing including CI is a necessary but not sufficient component of the CYCLUS QA system: it keeps bad code out of CYCLUS. However, CYCLUS will always require human eyes and hands to let good code in.

The main CYCLUS repository is hosted remotely and publicly on the GitHub website [49], in part because it provides tools that facilitate a culture of frequent and thorough code review. A number of policies exist to ensure that a proposed set of changes, known as a *pull request*, adhere to the projects QA standards. Every pull request must be reviewed and accepted by a member of the CYCLUS core team that was not involved in authoring the changes. In addition to reviewing the algorithmic design in the source code, the reviewer relies on tests to ensure correct behavior, and requires the authors to adhere to the style guide and provide sufficient documentation. Only after the QA standards have been met are the proposed changes merged into the software. This step has been repeatedly shown to improve code quality [4].

Once CI is successful, a code reviewer inspects not only the changes that are proposed to the software itself, but also the changes that have been proposed to tests.

Because of their long term benefit to the maintainability of the project, documentation and coding style are also reviewed as part of this process. The API must be documented as required by the CYCLUS QA policy. In CYCLUS, this information is aggregated together into static websites with the Doxygen [50] and Sphinx [51] tools, and can be accessed at <http://fuelcycle.org>. CYCLUS also strictly enforces the use of the Google C++ Style Guide [52] for all software contributions. This means that all developers of CYCLUS write CYCLUS code in the same way. This homogenization may be a hurdle to new developers but ultimately improves code legibility and, therefore, robustness [4].

3. Demonstration

The success of the CYCLUS simulator can be measured in many ways. The most compelling are early successes of the community-driven development model and demonstrations of its fundamental simulation capabilities. Promising growth of the CYCLUS ecosystem at multiple institutions indicates that a fuel cycle simulator can advance in this community-driven development paradigm. Additionally, simulation results for both once-through and more complex recycling scenarios demonstrate that CYCLUS possesses the fundamental fuel cycle simulation capabilities to contribute to the field.

3.1. Ecosystem

The CYCLUS community-driven software development model seeks to break the proliferation of specialized simulators. It instead leverages the collective expertise of fuel cycle researchers toward a single, more extensible, tool. Through the targeted contributions of those researchers, an ecosystem of capabilities should emerge. The CYCLUS ‘Ecosystem’ is the collection of tools, calculation libraries, archetypes, data, and input files intended for use with the CYCLUS simulator. Members of the ecosystem include:

- the archetypes provided in the CYCAMORE [38] repository
- the archetypes created by researchers
- isotopic composition data
- historical facility deployment data
- the CYCLUS graphical user interface (GUI) tool Cyclist
- fundamental analysis tools in the CYCLUS toolkit
- tools for CYCLUS optimization, parallelization, and development

Taken together, these form an ecosystem of capabilities. Over time, this ecosystem will grow as archetype developers, kernel developers, and even users contribute capabilities developed for their own needs. Indeed, the long-term vision for the CYCLUS framework predicts an ever-expanding ecosystem of both general and specialized capability extensions.

Already, the ecosystem is growing. Early cross-institutional contributions to the ecosystem demonstrate a significant achievement by the CYCLUS framework and provide the basis for a community-driven development model.

3.1.1. Supplementary Projects

A number of projects and tools outside of the core simulation kernel have been developed to improve the scope and the diversity of the capabilities in the CYCLUS ecosystem. Table 2 lists the tools and projects developed under close integration with the CYCLUS kernel. These tools are used to ease development and simulation design (Cycstub, Cycic, Ciclus), data visualization and analysis (Cyclist, Cyan), and remote execution (Cloudlus).

3.1.2. Archetype Contributions

It is expected that the most common type of contribution to CYCLUS will be contributions of new archetypes. Researchers will be driven to create a new archetype when a need arises, such as to improve the fidelity of simulation, or to represent a novel reactor type, an innovative reprocessing strategy, or a particular governmental or institutional behavior. The real-world utility of CYCLUS can in part be measured by the breadth and diversity of archetypes being developed in this way.

Early progress has been promising. Many archetypes external to the CYCAMORE library (Table 3) have been [59, 60] or are being [61, 62, 63] developed for contribution to the CYCLUS ecosystem. These archetypes provide the first examples of developer-contributed capabilities. They add to the fundamental CYCAMORE archetypes by providing physics-based reactors, separations logic, fuel fabrication processing, storage facilities, and expanded institutional paradigms. The existence and diversity of these contributed archetypes illustrate the power and potential of the community-based development approach that CYCLUS has taken.

3.2. Simulations

To illustrate the flexibility of CYCLUS, this section will discuss the creation and results of a range of representative fuel cycle simulations. Previous benchmarks between CYCLUS and system dynamics simulators for more complex problems, including transition analyses, have been reported elsewhere and CYCLUS has demonstrated satisfactory performance [64, 65, 35, 66]. The examples presented here

Name	Description	Ref.
Cycic	Input control embedded in Cyclist	[53]
Cyclist	Interactive data exploration environment	[54]
Ciclus	Continuous integration scripts for CYCLUS	[55]
Cycstub	Skeleton for clean slate module development	[56]
Cyan	CYCLUS analysis tool	[57]
Cloudlus	Tools for running CYCLUS in a cloud environment	[58]

Table 2: Many tools have been developed outside of the scope of the CYCLUS kernel for improved user, developer, and analyst experiences with CYCLUS.

Name	Description	Ref.
Bright-lite	A physics-based reactor archetype and fuel fabrication archetype	[61]
Nuclear Fuel Inventory Model	A flexible, ORIGEN-based, reactor analysis module	[62]
CommodConverter	A simple commodity converting storage facility archetype	[60]
MktDrivenInst	An institution that controls deployment based on commodity availability	[63]
SeparationsMatrix	A facility for elemental separations of used fuel	[59]
StreamBlender	A facility for fuel fabrication from reprocessed streams	[59]

Table 3: A diverse set of archetypes under development reflect the diverse needs of researchers at various institutions. These archetypes, contributed outside of the CYCLUS core and CYCAMORE libraries are the first demonstration of community-driven development in a fuel cycle simulator.

are not the limit of CYCLUS’ capabilities, which are extended with each new addition to the ecosystem, as discussed above. Rather, these simulations are designed to illustrate that CYCLUS matches the capabilities of any recipe-based simulator and that variations on a single CYCLUS simulations can be run with small changes to the input specification. For simplicity of the current demonstration, many simplifying assumptions have been made with respect to material compositions, fuel transmutation, among others. The three fuel cycles examined are:

1. No Recycle (once through)
2. 1-pass MOX Recycle
3. Infinite-pass MOX Recycle

For each of these fuel cycles, a 1,100 month single-reactor CYCLUS simulation was run in addition to a 10-reactor simulation with staggered refueling times. As the number of staggered-cycle reactors increases the system converges toward continuous material flow results. CYCLUS flexibility allows this transition to be examined.

Facility	Parameter	Value
LWR	type	LWR
cycamore::Reactor	batches	3 [batches/cycle]
	batch size	20 [MTHM]
	cycle length	18 [months]
	refuelling time	2 [months]
	requests	recycled MOX (1st preference)
	requests	enriched UOX (2nd preference)
Spent Fuel Fabrication	requests	depleted uranium
cycamore::Fab	requests	separated fissile material
	offers	recycled MOX
Separations	requests	all spent fuel types
cycamore::Separations	offers	separated fissile material
	efficiency	0.99
	Pu separation capacity	$6.0E4$ [kg/month]
Repository	requests	all waste
cycamore::Sink	capacity	∞
UOX Fabrication	offers	UOX
cycamore::Source	capacity	∞
DU Source	offers	depleted U
cycamore::Source	capacity	∞

Table 4: Facility Configurations for Example Simulations

As described in Table 4, the facilities used in these simulations are:

- **Reactor** (`cycamore::Reactor`): This is a reactor model that requests any number of input fuel types and transmutes them to static compositions when they are burnt and discharged from the core. The reactor was configured to model a light water reactor with a 3 batch core operating on an 18 month cycle with a 2 month refuel time. The batches in the core each contain 20 metric tons of heavy metal (MTHM), where heavy metals are actinide elements like thorium, uranium, and plutonium. The reactor was also configured to take in either enriched UOX fuel or recycled MOX fuel.
- **Spent Fuel Fabrication** (`cycamore::Fab`): This facility requests depleted uranium and separated fissile material and mixes them together into recycled MOX fuel that it offers to requesters. The two streams are mixed in a ratio in order to match simple neutronics properties of the target fuel as closely as possible. The method used is based on a variation “equivalence method” originally developed by Baker and Ross [67]. This technique has also been used in the COSI fuel cycle simulator developed by the French Commissariat à l’Energie Atomique et aux Energies Alternatives (CEA).
- **Separations** (`cycamore::Separations`): This facility takes in all kinds of spent fuel and separates it into plutonium and waste streams with some efficiency (0.99 was used for these simulations). Up to 60,000 kg of Pu can be separated per month.
- **Repository** (`cycamore::Sink`): This is an infinite-capacity facility that can take in all types of material including separations waste streams and spent reactor fuel.
- **UOX Fabrication** (`cycamore::Source`): This facility provides enriched UOX fuel as requested. This facility has infinite production capacity.
- **DU Source** (`cycamore::Source`): This facility provides depleted uranium as requested. This facility has infinite production capacity.

To model each of the three fuel cycles, only simple adjustments to the input file specification were necessary. Specifically, only the commodity types and

trade preferences for each of the facilities needed to be altered from one simulation to another.

For all cases, the reactor was configured to request recycled MOX fuel with a higher preference than new UOX fuel. For the No Recycle case, the **Reactor** was set to offer all spent fuel as waste. For the 1-pass recycle case, the **Reactor** offered spent UOX into a spent fuel market, but spent MOX is still offered as waste. In the Infinite-Pass Recycle case, the **Reactor** offers all burned fuel into a spent fuel market. The separations facility requests spent fuel with a higher preference than the repository resulting in preferential recycle. All these preferences are easy to adjust and CYCLUS dynamically handles supply constraints and non-uniform preference resolution. It is notable that separations and recycle fuel fabrication capacity are deployed identically in all simulations. In the once through case, the recycling loop never acquires material, and so reactors always receive fresh UOX fuel. The DRE ensures everything operates smoothly in all cases.

CYCLUS’ discrete materials make single-pass recycle straightforward to implement. The reactors keep track of fuel as discrete batches. A reactor remembers where it received each batch from. If a batch was received from the recycled fuel fabrication facility, it does not offer it to separations and instead offers it as a waste commodity which is only requested by the repository. The material flows for the Single-Pass Recycle case are shown in Figure 6.

Changing the scenario from a 1-pass fuel cycle to an infinite-pass fuel cycle requires only a one-word change in the input file regarding the output commodity for the spent MOX fuel of the **Reactor**:

```
<fuel>
  <incommodity>mox</incommodity>
-  <outcommodity>waste</outcommodity>
+  <outcommodity>spent_fuel</outcommodity>
  <inrecipe>mox_fresh_fuel</inrecipe>
  <outrecipe>mox_spent_fuel</outrecipe>
</fuel>
```

This results in the material flows in Figure 7. A similarly trivial change was used to switch from the No Recycle to a 1-pass fuel cycle. Note that because the reactors always transmute fuel into fixed compositions, the error in isotopic compositions is larger for the Infinite-Pass Recycle case.

Figure 8 shows the full-system plutonium buildup for No Recycle (once through), Single-Pass Recycle, and Infinite-Pass Recycle variations of the one-reactor scenario described above. The figure was generated directly from CYCLUS output data. After

Figure 6: 1-pass MOX recycle scenario material flows.

Figure 7: Full MOX recycle (multi-pass) fuel cycle material flows.

several batch cycles (near month 300) in the 1-pass and infinite-pass cases, enough separated fissile material accumulates in the fuel fabrication facility to generate a full recycled batch. When this batch is transmuted, more plutonium is burned than created. This results in a drop in the total fuel cycle system plutonium inventory. This pattern repeats roughly every 10 cycles (200 months) for the Single-Pass Recycle case and every 9 cycles (180 months) for the Infinite-Pass Recycle case. Because the 1-pass recycle scenario does not re-recycle material, it takes the fabrication facility 2 cycles longer to accumulate a full batch of fissile material.

Because facilities are represented individually and transact discrete materials as discrete events, realistic non-uniform patterns in facility behavior that affect total system behavior are observed using CYCLUS.

Figure 8: System plutonium buildup with one reactor.

Figure 9 shows plutonium buildup for the 10-reactor simulations of the No Recycle, Single-Pass Recycle, and Infinite-Pass Recycle scenarios. As the number of reactors (with staggered refueling) increases, the behavior of the system approaches a more steady average reminiscent of continuous material flow models.

The fundamental capabilities of demonstrated in these simulations qualify CYCLUS and its ecosystem to model the breadth of scenario types expected of nuclear fuel cycle simulator tools. These examples further show the flexibility provided by the DRE logic within the CYCLUS framework and provide an example of the resolution made possible by discrete-facility and discrete-material modeling in fuel cycle simulation.

Figure 9: System plutonium buildup with staggered refueling for many reactors.

4. Conclusions

The CYCLUS nuclear fuel cycle framework presents a more generic and flexible alternative to existing fuel cycle simulators. Where previous nuclear fuel cycle simulators have had limited distribution, constrained simulation capabilities, and restricted customizability, CYCLUS emphasizes an open strategy for access and development. This open strategy not only improves accessibility, but also enables transparency and community oversight. Furthermore, the object-oriented ABM simulation paradigm ensures more generic simulation capability. It allows CYCLUS to address common analyses in a more flexible fashion and enables analyses that are impossible with system dynamics simulators.

Similarly, the fidelity-agnostic, modular CYCLUS architecture facilitates simulations at every level of detail. Simulations relying on arbitrarily complex isotopic compositions are possible in CYCLUS, as are simulations not employing any physics at all. Physics is introduced through optional system-wide radioactive decay of materials and through the use of physics-enabled facility libraries. To support calculation of physical processes in nuclear facilities, the CYCAMORE library provides models employing basic physics for core fuel cycle facilities and extension libraries from the community support more detailed simulations. Indeed, agents of such varying fidelity can even exist in the same simulation. Researchers no longer need to reinvent the underlying simulator framework in order to model a simulation focused on the aspects of the fuel cycle relevant to their research.

Furthermore, when the capabilities within CYCLUS, CYCAMORE, and the rest of the ecosystem are insufficient, adding custom functionality is simplified by a modular, plug-in architecture. A clean, modern API simplifies customization and independent archetype development so that researchers can create models within their domain of expertise without modifying the core simulation kernel. Throughout the CYCLUS infrastructure, architecture choices have sought to enable cross-institutional collaboration and sustainable, community-driven development. The ecosystem of capabilities, already growing, may someday reflect the full diversity of use cases in the nuclear fuel cycle simulation domain.

5. Acknowledgements

This work has been supported by a number of people. The authors would like to thank software contributors Zach Welch and Olzhas Rakhimov. Additionally, the authors are grateful to the many enthusiastic members of the CYCLUS community and our Nuclear Energy University Programs (NEUP) collaboration partners at the University of Texas, University of Utah, and University of Idaho.

Direct support for the CYCLUS project has been received from the Nuclear Regulatory Commission (NRC) Faculty Development program, the NEUP Research & Development Program, the National Nuclear Security Administration (NNSA) Consortium for Verification Technology, and the University of Wisconsin (UW). In addition students contributing to CYCLUS have received fellowship support from the Argonne National Laboratory (ANL) Lab Grad program, the NEUP Fellowship program, the National Science Foundation (NSF) Graduate Fellowship program, and the Department of Homeland Security (DHS) Nuclear Forensics Graduate Fellowship program.

- [1] S. J. Piet, B. W. Dixon, J. J. Jacobson, G. E. Matthern, D. E. Shropshire, Dynamic Simulations of Advanced Fuel Cycles, Nuclear Technology 173 (3) (2011) 227–238.
URL <http://epubs.ans.org/?a=11658>
- [2] K. Huff, B. Dixon, Next Generation Fuel Cycle Simulator Functions and Requirements Document, Tech. Rep. fcrd-sysa-2010-000110, Idaho National Laboratory (Jul. 2010).
- [3] C. M. Macal, To agent-based simulation from system dynamics, in: Simulation Conference (WSC), Proceedings of the 2010 Winter, IEEE, 2010, pp. 371–382.
URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5679148

- [4] J. Cohen, Modern Code Review, Making Software: What Really Works, and Why We Believe It (2010) 329–336.
URL <http://books.google.com/books?hl=en&lr=&id=DxuGi5h2-HEC&oi=fnd&pg=PA329&dq=cohen+modern+code+review&ots=0VrtsgOdqP&sig=-XrFj4mAA2MB-20qZX0Fx8Felzo>
- [5] M. Fagan, Design and code inspections to reduce errors in program development, in: Software pioneers, Springer, 2002, pp. 575–607.
URL http://link.springer.com/chapter/10.1007/978-3-642-59412-0_35
- [6] D. L. Donoho, A. Maleki, I. U. Rahman, M. Shahram, V. Stodden, Reproducible Research in Computational Harmonic Analysis, Computing in Science & Engineering 11 (1) (2009) 8–18, 00110. doi:10.1109/MCSE.2009.15.
URL <http://scitation.aip.org/content/aip/journal/cise/11/1/10.1109/MCSE.2009.15>
- [7] D. C. Ince, L. Hatton, J. Graham-Cumming, The case for open computer programs, Nature 482 (7386) (2012) 485–488, 00142. doi:10.1038/nature10836.
URL <http://www.nature.com/doifinder/10.1038/nature10836>
- [8] V. Stodden, The Scientific Method in Practice: Reproducibility in the Computational Sciences, SSRN Electronic Journal00037. doi:10.2139/ssrn.1550193.
URL <http://www.ssrn.com/abstract=1550193>
- [9] J. M. Wicherts, M. Bakker, D. Molenaar, Willingness to Share Research Data Is Related to the Strength of the Evidence and the Quality of Reporting of Statistical Results, PLoS ONE 6 (11) (2011) e26828. doi:10.1371/journal.pone.0026828.
URL <http://dx.doi.org/10.1371/journal.pone.0026828>
- [10] M. Petre, G. Wilson, Code Review For and By Scientists, arXiv:1407.5648 [cs]00000 arXiv: 1407.5648.
URL <http://arxiv.org/abs/1407.5648>
- [11] M. Flicker, E. A. Schneider, P. Campbell, Evaluation Criteria for Analyses of Nuclear Fuel Cycles, in: Transactions of the American Nuclear Society, Vol. 111 of Fuel Cycle Options Analysis – III, American Nuclear Society, La Grange Park, IL 60526, United States, Anaheim, CA, United States, 2014, pp. 233–236.
URL <http://epubs.ans.org/?a=36342>
- [12] C. Poinssot, S. Bourg, N. Ouvrier, N. Combernoux, C. Rostaing, M. Vargas-Gonzalez, J. Bruno, Assessment of the environmental footprint of nuclear energy systems. Comparison between closed and open fuel cycles, Energy 69 (2014) 199–211. doi:10.1016/j.energy.2014.02.069.
URL <http://www.sciencedirect.com/science/article/pii/S0360544214002035>
- [13] L. Guerin, M. Kazimi, Impact of Alternative Nuclear Fuel Cycle Options on Infrastructure and Fuel Requirements, Actinide and Waste Inventories, and Economics, Technical Report MIT-NFC-TR-111, Massachusetts Institute of Technology, Cambridge, MA, United States (Sep. 2009).
- [14] J. Jacobson, A. Yacout, G. Matthern, S. Piet, D. Shropshire, R. Jeffers, T. Schweitzer, VERIFIABLE FUEL CYCLE SIMULATION MODEL (VISION): A TOOL FOR ANALYZING NUCLEAR FUEL CYCLE FUTURES, Nuclear Technology 172 (2010) 157–178.
- [15] A. M. Yacout, J. J. Jacobson, G. E. Matthern, S. J. Piet, A. Moissyetsev, Modeling the Nuclear Fuel Cycle, in: The 23rd International Conference of the System Dynamics Society, Boston, 2005.
URL <http://www.inl.gov/technicalpublications/Documents/3169906.pdf>
- [16] E. A. Schneider, C. G. Bathke, M. R. James, NFCSim: A Dynamic Fuel Burnup and Fuel Cycle Simulation Tool, Nuclear Technology 151 (1) (2005) 35–50.
- [17] C. Allan, International Atomic Energy Agency, International Project on Innovative Nuclear Reactors and Fuel Cycles, Guidance for the application of an assessment methodology for Innovative Nuclear Energy Systems INPRO Manual: overview of the Methodology., rev.1 Edition, no. IAEA-TECDOC-1575 in IAEA-TECDOC, International Atomic Energy Agency, Vienna, 2008.
URL http://www-pub.iaea.org/MTCD/Publications/PDF/TE_1575_web.pdf
- [18] L. Boucher, J. P. Grouiller, "COSI" : A Simulation Software for a Pool of Reactors and Fuel Cycle Plants, in: Fuel Cycle and High Level Waste Management, Beijing, China, 2005.
- [19] L. Boucher, J. P. Grouiller, "COSI": The Complete Renewal of the Simulation Software for The Fuel Cycle Analysis, in: Fuel Cycle and High Level Waste Management, Vol. 1, ASME, New York, NY, USA, Miami, FL, United states, 2006.
- [20] M. Meyer, L. Boucher, New Developments on COSI 6, the Simulation Software for Fuel Cycle Analysis, in: Proceedings of GLobal 2009, Paris, France, 2009.
- [21] C. Coquelet-Pascal, M. Meyer, R. Girieud, R. Eschbach, C. Chabert, C. Garzenne, P. Barbraut, L. Van Den Durpel, T. Duquesnoy, M. Caron-Charles, B. Carlier, J. C. Lefevre, Comparison of Different Scenarios for the Deployment of Fast Reactors in France - Results Obtained with COSI, in: Proceedings of GLobal 2011, Makuhari, Japan, 2011.
- [22] A. Worrall, R. Gregg, Scenario Analyses of Future UK Fuel Cycle Options, Journal of Nuclear Science and Technology 44 (3) (2007) 249–256. doi:10.1080/18811248.2007.9711279.
URL <http://www.tandfonline.com/doi/abs/10.1080/18811248.2007.9711279>
- [23] L. Van Den Durpel, A. Yacout, D. Wade, T. Taiwo, U. Lauferts, DANESS V4.2: Overview of Capabilities and Developments, in: Proceedings of Global 2009, Paris, France, 2009.
- [24] L. Guerin, B. Feng, P. Hejzlar, B. Forget, M. S. Kazimi, L. Van Den Durpel, A. Yacout, T. Taiwo, B. W. Dixon, G. Matthern, L. Boucher, M. Delpech, R. Girieud, M. Meyer, A Benchmark Study of Computer Codes for System Analysis of the Nuclear Fuel Cycle, Technical Report, Massachusetts Institute of Technology. Center for Advanced Nuclear Energy Systems. Nuclear Fuel Cycle Program, electric Power Research Institute (Apr. 2009).
URL <http://dspace.mit.edu/handle/1721.1/75245>
- [25] E. A. Andrianova, V. D. Davidenko, V. F. Tsubul'skii, Desae program for systems studies of long-term growth of nuclear power, Atomic energy 105 (6) (2008) 385–390.
URL <http://www.springerlink.com/index/1581v3n5111475u8.pdf>
- [26] K. A. McCarthy, B. Dixon, Y.-J. Choi, L. Boucher, K. Ono, F. Alvarez-Velarde, E. M. Gonzalez, B. Hyland, Benchmark Study on Nuclear Fuel Cycle Transition Scenarios-Analysis Codes.
URL <http://www.iaea.org/inis/collection/>

- NCLCollectionStore/_Public/44/089/44089401.pdf
- [27] M. L. Dunzik-Gougar, C. A. Juchau, K. Pasamehmetoglu, P. P. H. Wilson, K. M. Oliver, P. J. Turinsky, H. S. Abdel-Khalik, R. Hays, T. E. Stover, Global Evaluation of Nuclear Infrastructure Utilization Scenarios (GENIUS), in: GLOBAL 2007: Advanced Nuclear Fuel Cycles and Systems, September 9, 2007 - September 13, 2007, GLOBAL 2007: Advanced Nuclear Fuel Cycles and Systems, American Nuclear Society, Boise, ID, United States, 2007, pp. 1604–1611.
- [28] R. Jain, P. P. H. Wilson, Transitioning to global optimization in fuel cycle system study tools, in: 2006 Winter Meeting of the American Nuclear Society, Nov 12 - 16 2006, Vol. 95 of Transactions of the American Nuclear Society, American Nuclear Society, Albuquerque, NM, United States, 2006, pp. 162–163.
- [29] K. M. Oliver, P. P. Wilson, A. Reveillere, T. W. Ahn, K. Dunn, K. D. Huff, R. A. Elmore, Studying international fuel cycle robustness with the GENIUSv2 discrete facilities/materials fuel cycle systems analysis tool, in: Proceedings of GLOBAL 2009, GLOBAL 2009: Advanced Nuclear Fuel Cycles and Systems, Paris, France, 2009.
- [30] K. D. Huff, K. M. Oliver, P. P. Wilson, T. W. Ahn, K. Dunn, R. Elmore, GENIUSv2 Discrete Facilities/Materials Modeling of International Fuel Cycle Robustness, in: Transactions of the American Nuclear Society, Vol. 101 of Nuclear Fuel Cycle Codes and Applications, American Nuclear Society, Washington D.C., United States, 2009, pp. 239–240.
- [31] C. A. Juchau, M. L. Dunzik-Gougar, J. J. Jacobson, Modeling the Nuclear Fuel Cycle, Nuclear Technology 171 (2) (2010) 136–141.
URL <http://epubs.ans.org/?a=10778>
- [32] G. Grasso, S. Monti, M. Sumini, NEA-WPFC/FCTS benchmark for fuel cycle scenarios study with COSI6, Tech. rep., Report RSE/2009/136 (2009).
URL http://www.enea.it/it/Ricerca_sviluppo/documenti/ricerca-di-sistema-elettrico/nucleare-fissione/studi-accordi/rse136.pdf
- [33] L. Boucher, Benchmark Study on Nuclear Fuel Cycle Transition Scenarios Analysis Codes, Tech. Rep. NEA/NSC/WPFC/DOC(2012)16, OECD, Nuclear Energy Agency (Jun. 2012).
- [34] M. Gidden, Cyclopts 0.0.10, Figshare-<http://dx.doi.org/10.6084/m9.figshare.1288959>. doi: <http://dx.doi.org/10.6084/m9.figshare.1288959>.
URL <http://figshare.com/articles/cyclopts/1288959>
- [35] M. J. Gidden, An Agent-Based Modeling Framework and Application for the Generic Nuclear Fuel Cycle, Ph.D. thesis, University of Wisconsin, Madison, WI, United States (Mar. 2015).
- [36] R. W. Carlsen, M. Gidden, K. Huff, A. C. Opatowsky, O. Rakhimov, A. M. Scopatz, Z. Welch, P. Wilson, Cyclus v1.0.0, Figshare-<http://dx.doi.org/10.6084/m9.figshare.1041745>. doi: <http://dx.doi.org/10.6084/m9.figshare.1041745>.
URL http://figshare.com/articles/Cyclus_v1_0_0/1041745
- [37] J. J. Jacobson, R. F. Jeffers, G. E. Matthern, S. J. Piet, B. A. Baker, J. Grimm, VISION User Guide-VISION (Verifiable Fuel Cycle Simulation) Model, Tech. rep., Idaho National Laboratory (INL) (2009).
URL [http://www.inl.gov/technicalpublications/](http://www.inl.gov/technicalpublications/Documents/4363869.pdf)
- Documents/4363869.pdf
- [38] R. W. Carlsen, M. Gidden, K. Huff, A. C. Opatowsky, O. Rakhimov, A. M. Scopatz, P. Wilson, Cycamore v1.0.0, Figshare-http://figshare.com/articles/Cycamore_v1_0_0/1041829. doi: http://figshare.com/articles/Cycamore_v1_0_0/1041829.
URL http://figshare.com/articles/Cycamore_v1_0_0/1041829
- [39] B. Boehm, Software risk management, Springer, 1989.
- [40] ASME, NQA-1a-2009, Addenda to ASME NQA-1-2008, Quality Assurance Requirements for Nuclear Facility Applications, no. NQA-1a-2009 in Nuclear Quality Assurance, American Society of Mechanical Engineers, New York, NY, 2009.
- [41] C. Larman, Agile and Iterative Development: A Manager's Guide, Addison-Wesley Professional, 2004.
- [42] NEAMS, NUCLEAR ENERGY ADVANCED MODELING AND SIMULATION (NEAMS) SOFTWARE VERIFICATION AND VALIDATION (V&V) PLAN REQUIREMENTS, Tech. rep., U.S. Department of Energy, Office of Nuclear Energy (Jul. 2013).
URL <http://www.energy.gov/sites/prod/files/2013/09/f2/NEAMS%20Software%20Verification%20and%20Validation%20Plan%20Requirements%20Version%200.pdf>
- [43] E. Biondo, A. Scopatz, M. Gidden, R. Slaybaugh, P. P. Bates, Cameron Wilson, Quality Assurance within the PyNE Open Source Toolkit, in: American Nuclear Society 2014 Winter Meeting, ANS, 2014.
- [44] K. D. Huff, M. Fratoni, H. Greenberg, Extensions to the Cyclus Ecosystem In Support of Market-Driven Transition Capability, in: Transactions of the American Nuclear Society, American Nuclear Society, Anaheim, CA, United States, 2014.
- [45] Software Freedom Conservancy, Git, Software Freedom Conservancy, 2014.
URL <http://git-scm.com>
- [46] G. Inc, googletest - Google C++ Testing Framework, Technical Report, Google Inc. (2008).
URL <https://code.google.com/p/googletest/>
- [47] J. Pellerin, Nose is nicer testing for python, Technical Report, nose (2007).
URL <https://nose.readthedocs.org/en/latest/>
- [48] P. Biggar, CircleCI, <https://circleci.com/> (2015).
URL <https://github.com/circleci/frontend>
- [49] L. Dabbish, C. Stuart, J. Tsay, J. Herbsleb, Social coding in GitHub: transparency and collaboration in an open software repository, in: Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work, ACM, 2012, pp. 1277–1286.
- [50] D. van Heesch, Doxygen: Source code documentation generator tool, URL: <http://www.doxygen.org>.
- [51] G. Brandl, Sphinx Documentation.
URL <http://sphinx-doc.org/sphinx.pdf>
- [52] B. Weinberger, C. Silverstein, G. Eitzmann, M. Mentoai, T. Landray, Google C++ style guide, <http://google-styleguide.googlecode.com/svn/trunk/cppguide.xml>.
- [53] R. Flanagan, E. A. Schneider, Input Visualization for the Cyclus Nuclear Fuel Cycle Simulator: Cyclus Input Control, in: Proceedings of GLOBAL 2013, Salt Lake City, UT, United States, 2013.
- [54] Y. Livnat, H. Gur, K. Potter, R. Flanagan, Cyclist, github.com/SCI-Utah/cnome (Oct. 2014).
URL github.com/SCI-Utah/cnome

- [55] A. Scopatz, Z. Welch, M. Gidden, Ciclus, <https://github.com/cyclus/ciclus> (Oct. 2014).
URL <https://github.com/cyclus/ciclus>
- [56] R. W. Carlsen, M. Gidden, K. Huff, A. C. Opotowsky, O. Rakhimov, A. M. Scopatz, Z. Welch, P. Wilson, Cyc-Stub, <https://github.com/cyclus/cycstub> (Jun. 2014).
URL <https://github.com/cyclus/cycstub>
- [57] R. W. Carlsen, CyAn, Figshare-
[Http://dx.doi.org/10.6084/m9.figshare.1041836](http://dx.doi.org/10.6084/m9.figshare.1041836). doi:
<http://dx.doi.org/10.6084/m9.figshare.1041836>.
URL http://dx.doi.figshare.com/articles/CyAn_Cyclus_Analysis_Tools/1041836
- [58] R. W. Carlsen, Cloudlus, <https://github.com/rwcarlsen/cloudlus> (Oct. 2014).
URL <https://github.com/rwcarlsen/cloudlus>
- [59] K. D. Huff, StreamBlender, Figshare-
[Http://dx.doi.org/10.6084/m9.figshare.1134648](http://dx.doi.org/10.6084/m9.figshare.1134648). doi:
<http://dx.doi.org/10.6084/m9.figshare.1134648>.
URL <http://figshare.com/articles/StreamBlender/1134648>
- [60] K. D. Huff, CommodConverter, Figshare-
[Http://dx.doi.org/10.6084/m9.figshare.1134647](http://dx.doi.org/10.6084/m9.figshare.1134647). doi:
<http://dx.doi.org/10.6084/m9.figshare.1134647>.
URL <http://figshare.com/articles/CommodConverter/1134647>
- [61] R. Flanagan, Bright-lite, <https://github.com/bright-dev/bright-lite> (Oct. 2014).
URL <https://github.com/bright-dev/bright-lite>
- [62] S. Skutnik, Development of an ORIGEN-based Reactor Analysis Module for Cyclus, Figshare-
[Http://dx.doi.org/10.6084/m9.figshare.1291144](http://dx.doi.org/10.6084/m9.figshare.1291144). doi:
<http://dx.doi.org/10.6084/m9.figshare.1291144>.
URL http://figshare.com/articles/Development_of_an_ORIGEN_based_Reactor_Analysis_Module_for_Cyclus/1291144
- [63] K. D. Huff, MktDrivenInst, Figshare-
[Http://dx.doi.org/10.6084/m9.figshare.1134650](http://dx.doi.org/10.6084/m9.figshare.1134650). doi:
<http://dx.doi.org/10.6084/m9.figshare.1134650>.
URL <http://figshare.com/articles/MktDrivenInst/1134650>
- [64] D. Djokic, A. M. Scopatz, H. R. Greenberg, K. D. Huff, R. P. Nibbelink, M. Fratoni, The Application of CYCLUS to Fuel Cycle Transition Analysis, in: Proceedings of Global 2015, LLNL-CONF-669315, Paris, France, 2015, p. 5061.
- [65] A. M. Scopatz, Non-judgemental Dynamic Fuel Cycle Benchmarking, arXiv:1511.09095 [physics]ArXiv:1511.09095.
URL <http://arxiv.org/abs/1511.09095>
- [66] M. Gidden, P. Wilson, K. Huff, R. Carlsen, Once-Through Benchmarks with C YCLUS, a Modular, Open-Source Fuel Cycle Simulator, in: Proceedings of the 2012 ANS Winter Conference, San Diego, CA, 2012.
- [67] A. Baker, R. Ross, Comparison of the Value of Plutonium and Uranium Isotopes in Fast Reactors, in: Proc. Conf. on Breeding, Economics and Safety in Large Fast Breeder Reactors, ANL-6972, 1963, pp. 7–10.