INL/RPT-23-74725
Revision 0

A Multidisciplinary Approach to Integrated Energy Systems: Advanced Nuclear Plants with Thermal Storage for Dynamic and Flexible Operation in Diverse Markets

January 2024

Rami M. Saeed, Václav Novotný, So-Bin Cho, Amey Shigrekar, Courtney M. Otani, Jakub Toman, Aidan Rigby, Daniel Mikkelson

Idaho National Laboratory

IES
Integrated Energy Systems

INL is a U.S. Department of Energy National Laboratory operated by Battelle Energy Alliance, LLC

INL/RPT-23-74725
Revision 0

A Multidisciplinary Approach to Integrated Energy Systems: Advanced Nuclear Plants with Thermal Storage for Dynamic and Flexible Operation in Diverse Markets

Rami M. Saeed, Václav Novotný, So-Bin Cho, Amey Shigrekar, Courtney M. Otani, Jakub Toman, Aidan Rigby, Daniel Mikkelson
Idaho National Laboratory

January 2024

Idaho National Laboratory
Integrated Energy Systems
Idaho Falls, Idaho 83415

http://www.ies.inl.gov

Prepared for the
U.S. Department of Energy
Office of Nuclear Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Page intentionally left blank

ABSTRACT

Energy supply, distribution, and demand are continuously evolving in accordance with the integration of new power generation sources. In the United States, there is a progressive shift toward incorporating fluctuating energy sources such as wind and solar into the energy mix, alongside the distributed generation. Simultaneously, we are seeing a revolution in consumer technologies such as electric vehicles, energy devices for both residential and commercial use, and industrial processes—all contributing to a fundamental change in energy consumption dynamics. These scenarios have led to huge demand for enhanced flexibility from nuclear power plants (NPPs), requiring them to operate with unprecedented adaptability. Furthermore, advanced NPPs (A-NPPs) could potentially apply to scenarios in which power generation flexibility is prioritized over its current and conventional baseload generation capacity for traditional demand patterns.

The present work covers several configurations for coupling nuclear energy production with thermal energy storage (TES). Previous work evaluated systems entailing a simple nuclear-TES configuration that used steam as a heat source for the TES of three different advanced reactors (ARs), but its heat diversion ratio (HDR) for charging was limited. The present work proposes a new configuration that features full decoupling of the nuclear reactor and the power cycle, using reactor coolant (i.e., gas) as a heat source for the TES and offering a theoretically unlimited HDR. This configuration is, however, more complex, as both high- and low-temperature TES (HT-TES and LT-TES) systems are required in order to cover the extended temperature range. Although this configuration imposes certain limitations on the discharge system configuration, it also presents the opportunity to design a more efficient system. For this case in particular, a steam reheat cycle is advantageous. Conventional TES systems take advantage of nominal efficiency when discharging, whereas the reheat cycle in the fully decoupled TES system affords a significant efficiency increase over conventional cases. Material selection, component/equipment sizing, and costing analyses were performed, followed by economic dispatch and size optimization.

Dynamic models were developed for the decoupled system in order to generate insights into aspects of off-design operation (e.g., drops in cycle efficiency), and these models enable exploration of potential mitigation strategies for addressing such drops. The models also highlight the need to design fail-safes that ensure continuous operational stability and minimize the impact of power ramping on the reactor parameters.

Page intentionally left blank

ACKNOWLEDGEMENTS

This manuscript was authored at Idaho National Laboratory by Battelle Energy Alliance LLC under contract no. DE-AC07-05ID14517 with the U.S. Department of Energy (DOE). This work was prepared for the U.S. DOE Office of Nuclear Energy (DOE-NE) via funding from the Integrated Energy Systems (IES) program.

The authors would like to acknowledge Paul Talbot, Mohammad Abdo, Dylan McDowell, Aaron Epiney, and Shannon Bragg-Sitton for their invaluable support and insights in evaluating and improving this work.

Page intentionally left blank

CONTENTS

FIGURES

TABLES

Page intentionally left blank

ACRONYMS

A-NPP	Advanced nuclear power plant
APEA	Aspen Process Economic Analyzer
BOP	Balance of plant
CHX	Charge heat exchanger
CSP	Concentrated solar power
EDR	Exchanger Design and Rating
ERCOT	Electric Reliability Council of Texas
HDR	Heat diversion ratio
HERON	Holistic Energy Resource Optimization Network
HPT	High-pressure turbine
HTGR	High-temperature gas-cooled reactor
HT-TES	Hight-temperature thermal energy storage
IES	Integrated energy system
LMP	Locational marginal prices
LPT	Low-pressure turbine
LT-TES	Low-temperature thermal energy storage
LWR	Light-water reactor
LZ	Load zone
MCT	Model costing technique
MISO	Midcontinent Independent System Operator
NPP	Nuclear power plant
NPV	Net present value
PJM	PJM Interconnection, LLC (formerly Pennsylvania–New Jersey–Maryland Interconnection)
RAVEN	Risk Analysis Virtual Environment
ROM	Reduced-order model
RTE	Round-trip efficiency
TES	Thermal energy storage

Page intentionally left blank

1. INTRODUCTION

In recent years, advanced nuclear reactors have gained attention as a potential solution for meeting the growing demand for clean, reliable energy. However, challenges such as the intermittent nature of renewable energy sources and the inflexibility of traditional nuclear power plants (NPPs) have sparked interest in integrating advanced nuclear reactors with thermal energy storage (TES) systems. Therefore, the aim of this research is to facilitate flexible power generation across diverse energy markets and to explore alternative applications of nuclear energy—applications that move beyond conventional electricity generation.

Flexible power generation is particularly valuable for advanced NPPs (A-NPPs) operating in environments where the benefits of baseload generation are outweighed by the ability to implement demand following production. TES systems are crucial for enabling NPPs to swiftly adapt to market fluctuations. Additionally, they allow A-NPPs to participate in multi-commodity markets by complementing the decarbonization pathways for industrial processes, thereby enhancing the economic competitiveness of NPPs. To illustrate this, Figure 1 showcases one possible thermally coupled integrated energy system (IES) architecture.

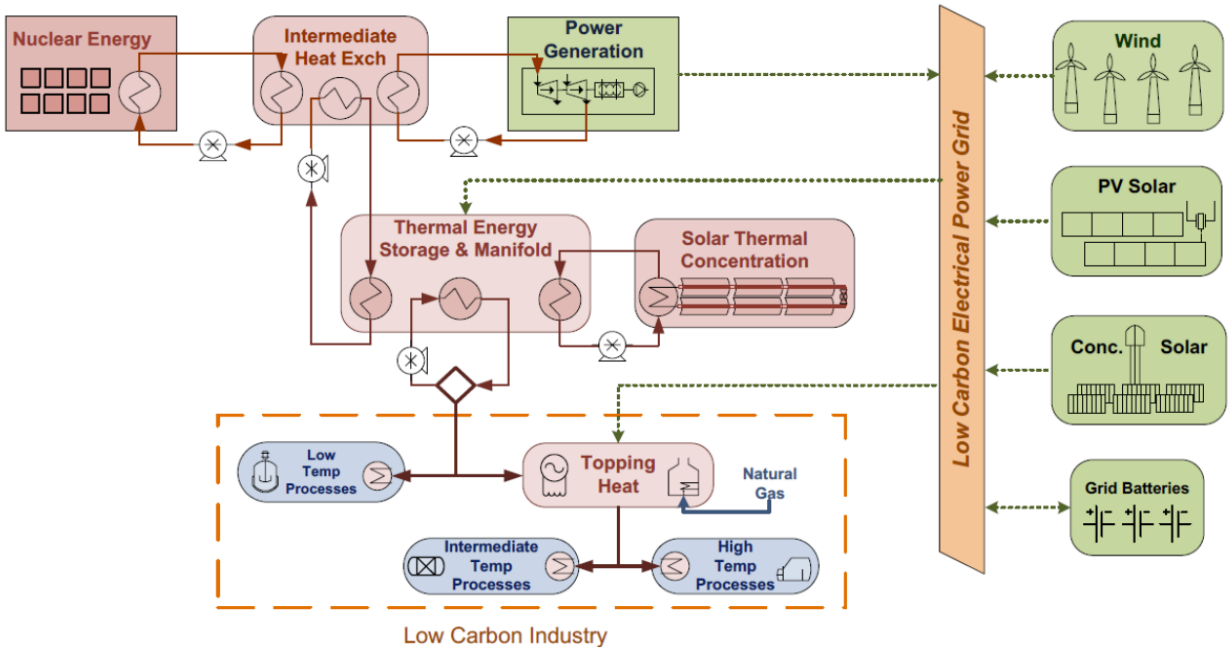

Figure 1. General architecture for a thermally coupled IES.

In a thermally coupled IES configuration, the nuclear reactor serves as a source of baseload power or heat, operating at a high capacity factor to cover both the operational and capital costs. A TES is employed to dampen subsystem dynamics/transients or defer energy delivery to a later time. This enables the reactor to maintain steady-state design conditions, thus improving performance and reducing maintenance expenses. By dispatching varying amounts of heat to thermal storage, the power generated from the balance of plant (BOP) can be adjusted, allowing for flexible power generation. The heat recovered from storage can either be directly utilized within the power generation cycle or directed to an industrial process via an IES. Note that the presented configuration represents one potential scenario, and that the specific need for and potential benefits of energy storage may differ in other IES architectures.

In previous work, Idaho National Laboratory assessed the compatibility of various TES technologies being considered for integration with A-NPPs [1]. This evaluation considered qualitative and merit indicators in order to identify suitable options. Of the technologies evaluated, two-tank molten-salt, solid-media, and latent-heat TES systems were deemed best suited for coupling with a broad range of A-NPPs, as these systems can accommodate different reactor sizes and operating temperatures. These preliminary findings served as a foundation for additional work in evaluating three different advanced reactor (AR) systems in various markets [3]. That work considered a TES charged by steam from the steam generators of the baseline A-NPPs. This approach minimizes the impact to the nuclear island but simultaneously limits efficiency, due to exergy (the potential of heat to be converted into work) loss in heat transfer processes between single-phase fluids and condensing or evaporating steam.

The present work further builds on this area of study and evaluates additional novel nuclear-TES coupling techniques. It considers tapping the primary loop for heat transfer to storage, and focuses on improving IES system performance by maximizing these systems' flexibility and efficiency in dispatching heat and generating power. A journal article and additional models are published in [4].

2. BACKGROUND AND HOLISTIC OVERVIEW

This study utilizes modeling capabilities and development know-how afforded by the Department of Energy's IES program, including process design and thermodynamic modeling, sizing and costing of system components, dynamic off-design modeling, and economic dispatch and size optimization. Idaho National Laboratory's set of tools for conducting this comprehensive range of tasks is known as the Framework for Optimization of Resources and Economics.

2.1 Approach to System Design and Evaluation in HYBRID

By evaluating multiple configurations and answering the following questions, the present study aims to identify optimal ways coupling NPPs with IES:

- What are economically and technically viable options for coupling IES to NPPs in specific grid energy systems?
  - *Technical Performance*
  - *Economic Analysis*
- What are the optimal strategies for coupling IES technologies with NPPs?
  - *Component Design*
  - *Thermodynamic Evaluations*
- What is the statistically ideal nuclear-IES system size for various markets?
  - *Portfolio Optimization*
- What is the governing control scheme for IES?
  - *Real-Time Optimization*
  - *Dynamic Modeling*

These questions are addressed using a comprehensive approach. First, the potential nuclear-TES coupled systems are evaluated, focusing on their ability to facilitate flexible and hybrid plant operations. Next, each coupling design's grid-wide economics and stochastic optimization is thoroughly analyzed. This involves assessing the systems' economic viability and optimizing their performance based on price signals from PJM Interconnection LLC (PJM), the Electric Reliability Council of Texas (ERCOT), and the Midcontinent Independent System Operator (MISO). Lastly, the optimized system designs and sizes are modeled for evaluation under transient operation, considering dynamic factors, off-design

performance parameters, and feedback within the IES. These three stages (see below) are presented in Figure 2.

1. **Thermal design and process modeling:** Aspen HYSYS process models were created to assess the technical value propositions of different approaches for coupling nuclear systems with TES. A high-temperature gas-cooled reactor (HTGR) design was evaluated in regard to two different coupling approaches involving TES and various forms of energy output boosting. This process also included the generation of cost functions derived from the fully balanced nuclear-TES models, in consideration of various system sizes.
2. **Stochastic optimization using HERON:** This process involved evaluating price signals from the three different energy markets by using an autoregressive moving average (ARMA) model, which was then utilized for stochastic optimization of the nuclear-coupled energy storage system by employing the Risk Analysis Virtual Environment (RAVEN) tool and its dispatch optimization plugin, the Holistic Energy Resource Optimization Network (HERON). The three aforementioned markets were explored, along with various parameters of the analysis.
3. **Dynamic modeling and control feasibility:** Taking the output from the previous step, this phase focused on development of the transient state analysis, assessment of off-design performance, control characterization of dynamic systems, and the feedback generated in integrated systems. Success in this area was achieved by implementing Modelica models within the HYBRID repository. The developed models were used to explore certain operating limits of the physical systems, to implement fail-safes, and to demonstrate system operation with dispatch profiles from optimized models in HERON.

The examination of different thermal energy bypass ratios, or heat diversion ratios (HDRs), was intended to foster an understanding of thermal power extraction's impact on overall system efficiency. Parametric studies were conducted by varying the HDRs away from the steam cycle.

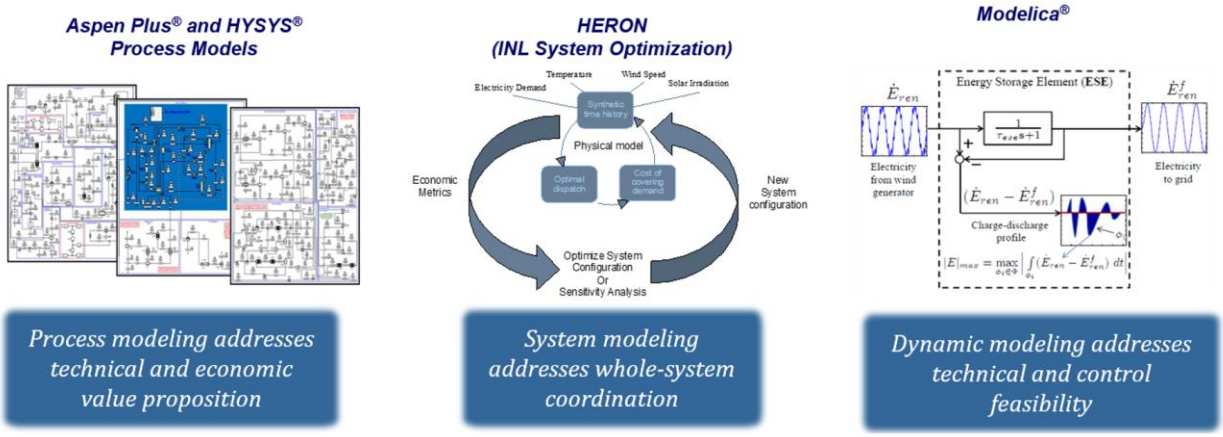

Figure 2. Adopted approach to nuclear-TES system architecture design and evaluation.

2.2 Mapping TES Systems for Nuclear Coupling

In previous work, a comprehensive analysis was conducted on 10 TES technologies, with a particular focus on their potential integration with advanced nuclear system technologies. Each was carefully explored in regard to 10 specific engineering questions. In the analysis, these engineering questions were defined as the following figures of merit: technology readiness level, temperature compatibility, energy density, size, cycle frequency, ramp time, realignment frequency, geographic needs, environmental impact, and interventions. The 10 figures of merit were then evaluated within a 3 x 3 matrix that included low-, medium-, and high-temperature A-NPPs, with each of these categories being further divided into

three subgroups: small-, medium-, and large-sized NPPs. The study emphasized tailoring TES technology choices to specific advanced nuclear system categories, rather than applying a generalized approach.

Recommendations for high-temperature advanced nuclear systems leaned toward molten-salt sensible heat storage, as this type of storage is applicable to most reactor sizes and temperatures. Table 1 lists the top-level technology recommendations and ranks each category of A-NPPs. It summarizes the entire technical space that was navigated; however, the actual study also provided an in-depth discussion on the maturity levels, pertinent performance data, and existing research gaps pertaining to these various TES systems. This information serves as a valuable resource for further enhancing these technologies.

Table 1. Top-level technologies, and the rankings for each category of A-NPP.

	Low Temperature (<350˚C)			Medium Temperature (350-650˚C)			High Temperature (>650˚C)		
Technology	<25 MWth Microreactors	25-750 MWth SMRs/medium	>750 MWth Larger NPPs	<25 MWth Microreactors	25-750 MWth SMRs/medium	>750 MWth Larger NPPs	<25 MWth Microreactors	25-750 MWth SMRs/medium	>750 MWth Larger NPPs
Two-tank Molten salt	15	16	16	15	16	16	13	14	14
Two-tank Thermal oils	13	13	13	9	9	9	9	9	9
Thermocline Molten salt	11	12	12	11	12	12	9	10	10
Thermocline Thermaloil	10	11	11	7	8	8	7	8	8
Hot/cold Water Tank	8	9	9	8	9	9	8	9	9
Underground Storage	6	6	6	6	6	6	6	6	6
Thermochemicals	6	6	5	5	5	4	5	5	4
Latent Heat Storage	14	14	14	14	14	14	12	12	12
Solid-media Storage	13	14	14	13	14	14	12	13	13
Steam accumulators	10	10	9	9	9	8	9	9	8

Blue: ranked highest; White: medium ranking; Red: ranked lowest

The analysis revealed that selecting an optimal technology depends on the unique operational conditions of the NPP and the inherent characteristics of the reactor type. Two-tank molten-salt systems were found to be uniquely well suited across the range of different reactor categories (e.g., various operating temperatures and sizes). For the present work, this versatility made it a preferred choice—or at least an attractive initial demonstration case—for studying TES integration with A-NPPs. Some technologies were found extremely suitable in relation to a specific A-NPP category (e.g., latent heat storage systems, which exhibited a particularly exceptional ranking within the low- and medium-temperature small-sized reactor category [microreactors]). Thanks to its demonstrated adaptability under diverse conditions, the two-tank molten-salt sensible heat storage system was utilized in the present work as the foundational basis for exploring a wide range of coupling methods, approaches, and electricity markets.

2.3 Nuclear-IES Thermal Coupling Design

This section explores and contrasts the different ways that advanced nuclear reactors can be coupled with TES systems (i.e., direct and indirect thermal coupling), highlighting their advantages and drawbacks. Such a comprehensive overview of these coupling methods helps foster the development of the most efficient and cost-effective TES systems for A-NPPs. The primary focus is on HTGR systems, mainly because they have the highest primary circuit temperatures, thereby showcasing the limits and the high potential of IES for enabling efficient conversion cycles while eventually also meeting the high-temperature thermal demand in industries.

Previous work primarily focused on a coupling approach in which steam from the steam generator is utilized as the charging fluid for the TES system [3]. Figure 3 illustrates two examples of this coupling method, whose primary advantage is that it eliminates the need for extensive reengineering of the nuclear island thermal extraction process, thus simplifying the licensing and permitting procedures. However, one drawback of this approach is the high loss of exergy (i.e., the potential for heat to be converted into work) associated with heat transfer between a single-phase fluid and a fluid undergoing a phase change—in this

case, water. This adverse effect is visible in the Q-T diagram of the heat transfer both to and from storage (see Figure 4), highlighting the decrease in the temperature (and thus pressure) of the steam generated from the TES system, as compared to using live steam to charge the TES. As a result, electricity production efficiency when discharging the TES can drop dramatically.

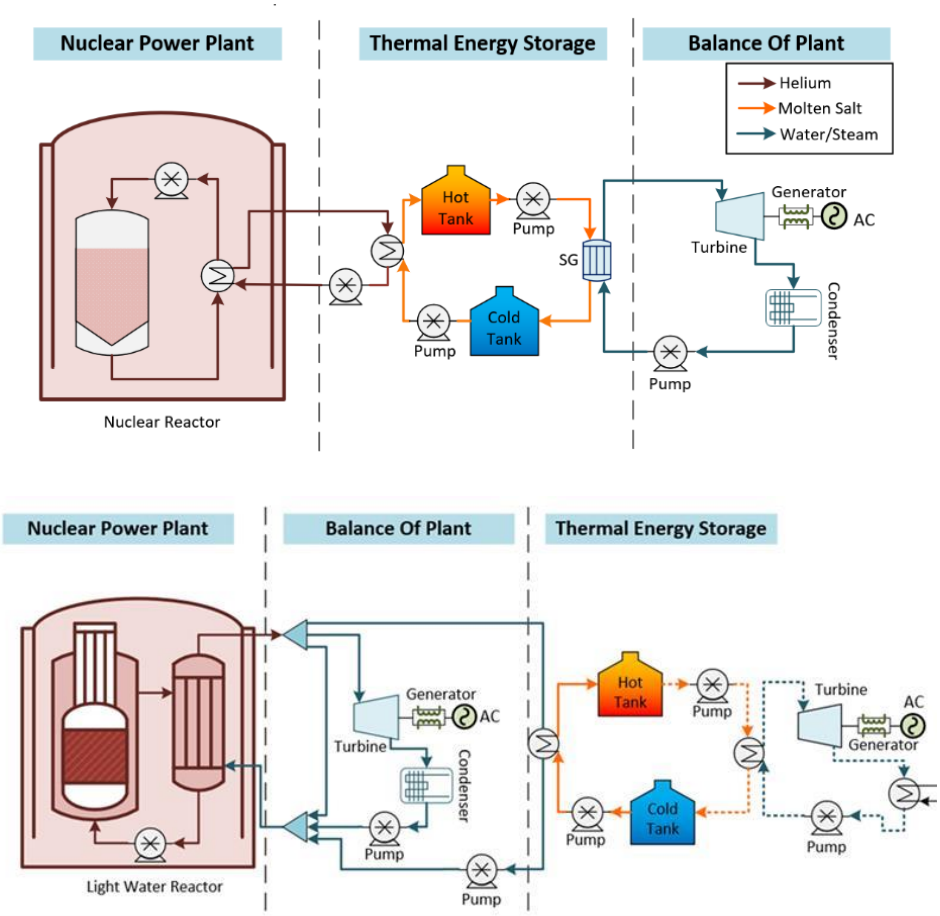

Figure 3. Nuclear-TES coupling approach in which steam us used for charging, with all energy passing through an intermediate loop (top), and featuring the option of routing a portion of the energy through the TES system and dedicated BOP (bottom).

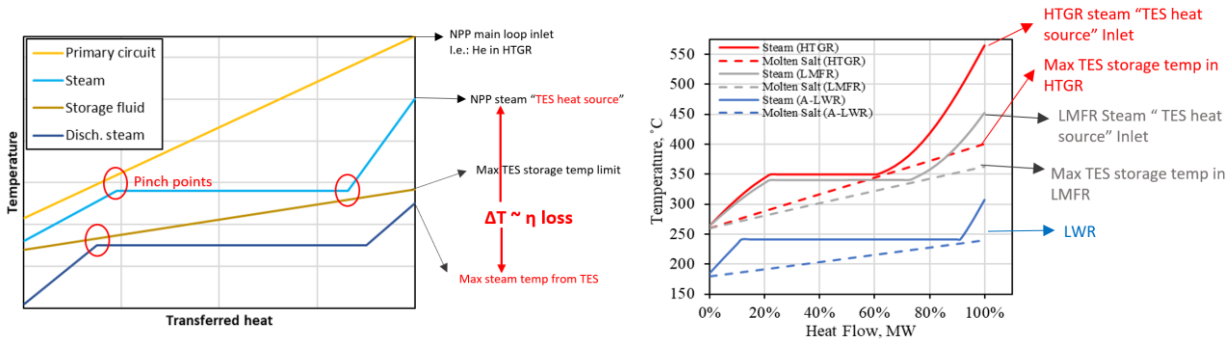

Figure 4. Heat transfer Q-T diagram of the coupled system approach in which steam (with phase change) is used for charging, showing both a general illustration (left) as well as specific results from previously explored systems (right).

To overcome the exergy loss limitation, one option is to integrate—into the TES—phase change materials whose melting temperatures are close to the steam saturation temperature, while also including sensible heat storage. A similar approach was tested for small-scale systems with heat pumps and an organic Rankine cycle [5]. However, implementing it in nuclear systems is considered unfeasible due to the issues and associated costs involved, as well as the difficulty of attaining balanced uniform performance in large-scale systems.

Losses can also be diminished by utilizing supercritical steam. Though this approach is thermodynamically promising, it is not being actively pursued in AR designs with high commercialization potential. The technically feasible alternative is to modify the heat removal system so that it directly transfers heat from the primary loop to a heat transfer fluid and eventually to storage. (Note that molten salt can act as both a heat transfer fluid and a storage fluid.) Figure 5 shows an ideal configuration for this system. The temperature of the storage medium and the steam produced during discharging is seen to remain near that of the high-temperature fluid in the primary loop. The steam cycle maintains the same parameters, regardless of whether directly operating from the primary loop heat source or discharging from the TES.

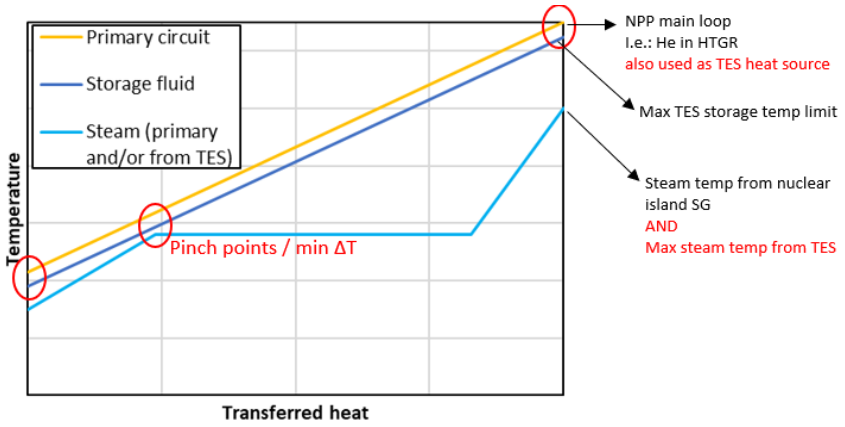

Figure 5. General heat transfer Q-T diagram for the system that uses primary loop fluid for TES charging.

When considering HTGR parameters, the TES should allow for temperatures above 700°C in order to eventually be able to supply high-temperature process heat yet simultaneously ensure that the temperature of the He as it exits the primary loop and is returned to the reactor core is at the nominal design point of 250°C. This high-temperature range poses a challenge for the storage system design. One option worth considering is a packed bed with a gaseous heat transfer fluid [6]. Alternatively, there are proprietary systems now under experimental development. For example, the Hyme system [7] is said to utilize a hydroxide salt and a specific corrosion control technology over a temperature span of 200°C–700°C, and the company is currently constructing their first pilot installations. However, detailed cost and design information on these systems is not publicly available.

One way to navigate the aforementioned limitations is via a system design in which only a portion of the heat at highest temperature is routed to the TES, whereas the remaining heat is used in a baseline BOP. Figure 6 shows this configuration for two discharging options: a dedicated BOP and a BOP with an oversized turbine. Note that in both cases, the steam is supplied at maximum temperature, regardless of whether the energy is passing through the TES system. The limit of the HDR into the TES system corresponds to the condition when the heat source temperature at the steam generator inlet (i.e., after mixing with bypass for TES charging) is insufficiently high to provide steam at nominal temperatures/pressures. Note that the baseline steam generator thus actually operates in off-design mode when charging the TES, transferring less heat but at the same time reflecting a lower temperature difference between the primary and secondary fluids.

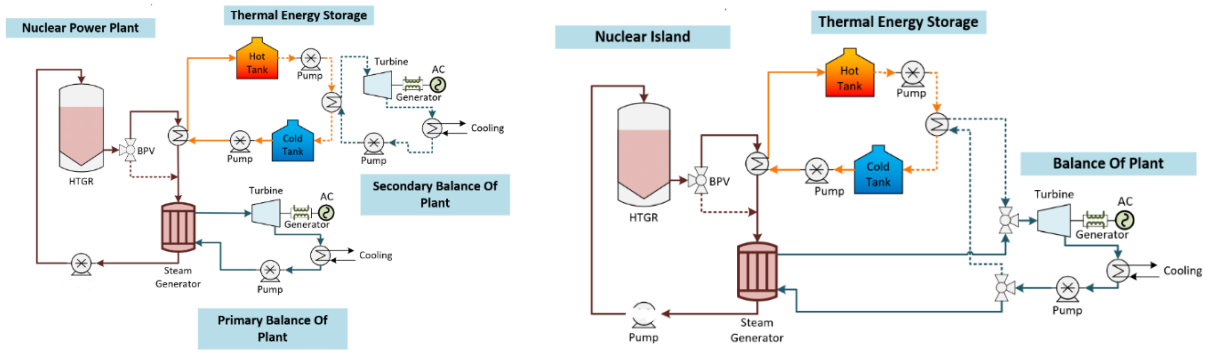

Figure 6. Proposed configurations of the TES-integrated nuclear plant that utilizes single-phase heat-transfer charging with a dedicated BOP (left), and an oversized BOP (right) to accommodate the discharging power increase.

To achieve higher HDRs, a fully decoupled system has been proposed (see Figure 7). As mentioned, using a single readily available storage fluid to achieve the full temperature range of the cold and hot fluid is unfeasible. Therefore, the system is split into two parts. For the high-temperature TES (HT-TES), a molten-salt mixture of $MgCl_2$-KCl was chosen. This mixture is limited by its 422°C melting point. Hence, to store all the incoming heat from the helium, which must return to the reactor at 250°C, a second, low-temperature TES (LT-TES) is introduced downstream of the helium loop. In this case, a HITEC molten salt would suffice. Having two separate storage systems is expected to slightly elevate the cost (the same overall amount of heat but in two smaller systems) as well as introduce several design challenges (e.g., the balancing and distribution of heat into the steam cycle). It is technically unfeasible to implement the boundary (or breaking point) between the HT-TES and LT-TES in the middle of single heat exchanger—especially one that involves boiling the working fluid (evaporator). To address this, a reheat cycle was suggested, and would benefit from a suitable ratio of heat input among the distinct components (the LT-TES supplies the preheater and evaporator, and the HT-TES supplies the superheater and reheater). Such a configuration could potentially improve the cycle efficiency.

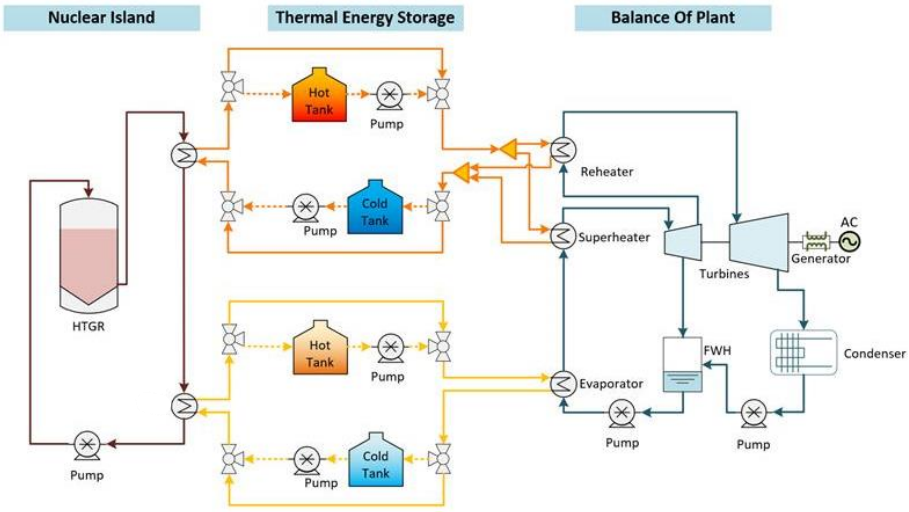

Figure 7. Nuclear-TES plant in a configuration that fully decouples the nuclear island from the power conversion cycle, using two TES systems.

2.4 Stochastic Dispatch Optimization

The growing volatility of electricity prices has been directing interest toward the potential grid services and business models that TES can offer. These include energy arbitrage, capacity reserve, and frequency regulation, as well as non-electric applications such as hydrogen production. The goal is to strategically align storage operations with market signals such as demand fluctuations and electricity price variations. The analysis conducted within the context of the present work evaluated the potential revenue that could be generated using a fully decoupled nuclear-TES system, thus substantiating its effectiveness through empirical evidence. The scope of this analysis was specifically confined to investigating energy arbitrage opportunities within the electricity sector. To this end, we gathered hourly locational marginal prices (LMPs) spanning a duration of 4 years (i.e., 2018 to 2021). These LMPs were employed to generate synthetic profiles of 30 years' worth of price signals for the ERCOT, PJM, and MISO markets, using the ARMA algorithm contained in RAVEN [8]. Under the assumption of a price-taking TES that operates with perfect foresight in the real-time market, we used HERON to derive profit maximizing capacities and TES charge/discharge profiles. Figure 8 presents the workflow applied to this analysis.

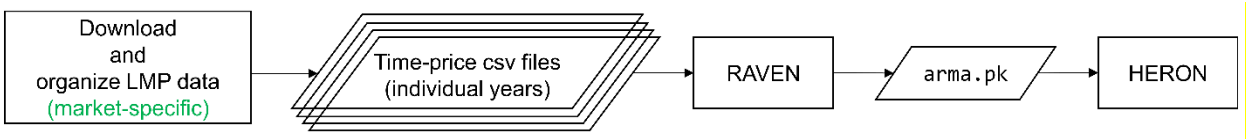

Figure 8. Illustrative flow chart of the stochastic technoeconomic analysis [8].

2.5 Off-design and Transient System Characterization

Application of the nuclear-IES system intrinsically requires flexibility—as the systems operate in terms of a wide, variable range of outputs (power, heat, manufacturing, and chemical batch processes), in conjunction with intermittent input from renewables—while other primary energy sources such as nuclear need to operate under a steady, nominal, undisturbed regime. HERON's economic dispatch optimization cannot directly account for the behavioral complexities observed in the physical system. Thus, to enable a complex assessment, transient models were developed to showcase the system's ability to operate according to the dynamic requirements, eventually find the limits of such operation, and assess off-design performance. The Modelica language was used to model the systems that were originally developed in Aspen HYSYS, assessed for physical limits and partial admission performance, while also using the HERON optimization and dispatch profiles as input. The results can serve as feedback for modifying the system design as needed, but are mainly intended for modifying the HERON input parameters so as to more closely match the system's overall performance in real life.

Our previous work included development of Modelica models for light-water reactor (LWR) designs [3]. It also demonstrated off-design and transient operation, as well as capabilities for reliable dispatch from the physical system. The previous work on HERON dispatch performance for the HTGR model reported in [8] indicated significant economic benefits in comparison to both the lead fast reactor and LWR designs. HTGRs present a novel transient problem as a result of their slow thermal load following characteristics and high-temperature/high-pressure capabilities on the nuclear plant secondary loop. Thus, extending the transient modeling work performed for the LWR design to the HTGR design was expected to be particularly informative. An independent BOP model was selected as the candidate for a detailed control, efficiency, and physical limitations analysis performed in Modelica. This aspect in particular directly continues the work reported in our previous study [3], and is thus included to complement the previous work covered in APPENDIX A of the present report.

2.6 Holistic Overview of the Project Workspace and Status

In summary, the research work revolves around investigating different nuclear-TES coupling configurations. Initially, a more straightforward nuclear-TES setup was examined, in which steam served as a heat source for the TES of three different ARs (two-phase heat transfer, partially decoupled). The coupling in each of these scenarios was evaluated, and a cost analysis was conducted, followed by technoeconomic HERON analyses in different markets. Dynamic models were developed to elucidate off-design operation and system behavior during transients. This approach was limited by its restricted HDR during charging, as well as a significant drop in efficiency. To overcome this, a more complex configuration was explored, featuring complete decoupling of the nuclear reactor and power cycle (fully decoupled). This configuration employed reactor coolant (gas) for TES, offering full heat diversion potential. This complex arrangement required both HT-TES and LT-TES systems, with a steam reheat cycle for improved efficiency. Table 2 lists the technical spaces that were comprehensively studied for the nuclear-TES coupling configurations, as applied to different reactor types.

Table 2. Technical spaces evaluated for nuclear-TES coupling configurations.

Nuclear-TES coupling	Reactor	Technical Space	Reference
Partially decoupled, two-phase heat transfer	Advanced LWR	Coupling methodology and steady-state physical models	Chapter 3 [3]
		Technoeconomic, costing, and HERON analyses	Chapter 4 [3], and [8]
		Transient/dynamic modeling	Chapter 6 [3]
	HTGR	Coupling methodology and steady-state physical models	Chapter 3 [3]
		Technoeconomic, costing, and HERON analyses	Chapter 4 [3], and [8]
		Transient/dynamic modeling	Present, APPENDIX A
	Liquid-metal fast reactor	Coupling methodology and steady-state physical models	Chapter 3 [3]
		Technoeconomic, costing, and HERON analysis	[8]
		Transient/dynamic modeling	—
Fully decoupled, single-phase heat transfer, with reheat cycle	Advanced HTGR	Coupling methodology and steady-state physical models	Present, Chapter 3
		Cost analysis	Present, Chapter 4
		Technoeconomic and HERON analyses	Present, Chapter 5
		Transient/dynamic modeling and off-design	Present, Chapter 6

3. DEVELOPMENT AND ANALYSES OF PHYSICAL MODELS

Design and operation of nuclear-TES systems require careful consideration of thermodynamic conditions in order to ensure optimal performance and efficiency. This subsection provides developers of these systems with valuable insights into the critical factors and criteria involved in selecting the

thermodynamic conditions for each component. It also explores the fundamental principles that govern the balancing of these systems, enabling developers to achieve optimal configurations. By understanding the intricacies of thermodynamic design and system balancing, developers can unlock the full potential of nuclear-TES systems, thus maximizing their efficiency and effectiveness in meeting energy demands or delivering heat.

3.1 Partially Decoupled Single-Phase Heat Transfer System

The partially decoupled system offers a relatively simplified configuration that, for the HTGR case, takes advantage of the wide temperature margin between the helium and steam. The primary loop is branched to the charge heat exchanger (CHX), upstream of the original steam generator. The presented model demonstrates that heat diversion to this charging branch produces the following limitation. However, the temperature at the steam generator entrance—after a portion of the heat has been transferred to the CHX—is too low to meet the required steam load parameters.

3.1.1 Input Parameters and Assumptions

Table 3 lists the assumptions and design parameters upon which the analyses were based and the system was modeled. The helium loop parameters were derived from the nominal HTGR design values, and the steam parameters correspond to the Xe-100 system. Note that the overstated isentropic efficiency of the turbine was calculated such that it thermodynamically compensates for the simplified feedwater heating in the process models. This approach was adopted for simplicity in modeling, and was carried over from our previous works.

Table 3. Design parameters of the partially decoupled system.

Parameter	Design value	Parameter	Design value
$T_{He,hot}$	750°C	$\dot{m}_{He}$	78.5 kg/s
$T_{He,cold}$	250°C	$p_{steam,HP}$	16.4 MPa
$T_{steam,SH}$	565°C	p_{cond}	7 kPa
T_{FW}	195°C	η_{turb}	96.94%
$\Delta T_{min,HX}$	5 K	η_{pump}	90%
$T_{cold,HT-TES}$	500°C	$\Delta p_{max,HX}$	3% (abs. p)

3.1.2 BOP and Nuclear-TES Analysis

The BOP configuration uses the original steam generator, along with an additional heat exchanger upstream of the steam generator on the helium side, to divert some of the high-temperature heat into the TES. The stored heat can be easily used to generate steam at parameters identical to those of a standard steam cycle without TES. The steam cycle is modeled in a simplified manner, with feedwater heating being achieved by mixing high-pressure steam with the condensate behind the feedwater pump. The model, created in Aspen HYSYS, is illustrated in Figure 9. The values represent a discharging mode in which the power output is increased to 133%. The cycle retains its 42% efficiency regardless of operating mode. (The constant cycle efficiency is valid at this modeling level, as further discussed in the off-design and transient modeling activities covered in Section 6.)

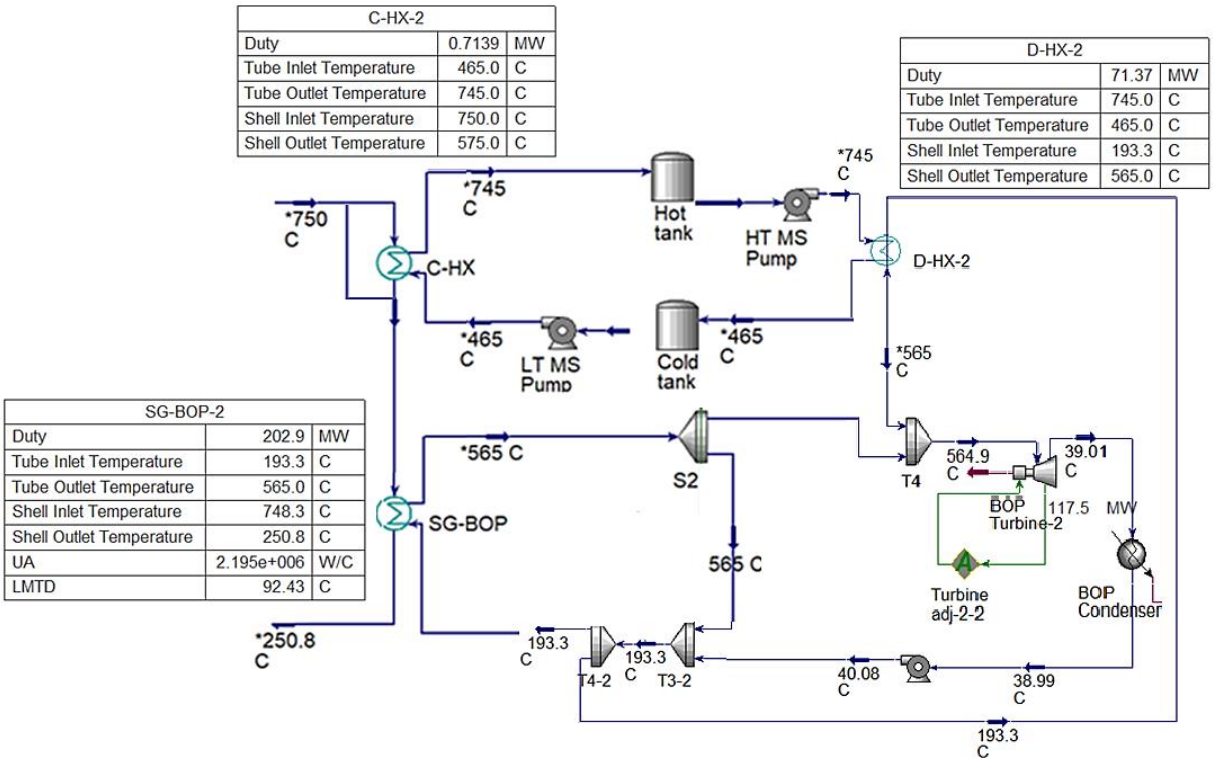

Figure 9. Aspen HYSYS model of the partially decoupled configuration.

3.2 Fully Decoupled Single-Phase Heat Transfer System with Reheat Cycle

Full decoupling enables designs featuring an unlimited (unrestricted) HDR and affords flexibility in the steam cycle design. This unique feature allows for exploring a completely new cycle configuration, distinct from the baseline Xe-100 design, thereby opening up new possibilities for optimizing BOP performance and achieving enhanced efficiency.

3.2.1 Input Parameters and Assumptions

The assumptions and design parameters for the fully decoupled system are summarized in Table 4. The helium loop parameters were again derived from the nominal HTGR Xe-100 design values. The steam parameters are not bound by the Xe-100 steam cycle design, as a completely standalone and decoupled system is considered. Instead, they correspond to modern sub-critical coal-fired plants. Numerous design choices were made via iterative processes in order to optimize the system, including the pinch point and heat distribution ratio between the LT-TES and HT-TES. Proper selection of these parameters balanced the need for the highest reheat cycle power (the optimum intermediate pressure would correspond to about 4 MPa) and a simplified configuration with a single feedwater heater, while also satisfying the thermal balance between the HT-TES and LT-TES. The trade-off is that there is now a need to consider the requirements for including the feedwater temperature's effect on the steam generator pinch point and to introduce a reasonable pressure drop on the steam extraction line in order to afford sufficient margin during flexible operation (around 10 K on the saturation temperature).

Table 4. Design parameters of the system.

Parameter	Design value	Parameter	Design value
$T_{He,hot}$	750°C	$T_{cold,HT-TES}$	500°C
$T_{He,cold}$	250°C	$T_{hot,LT-TES}$	470°C
$T_{steam,SH}$	565°C	$\dot{m}_{He}$	78.5 kg/s
$T_{steam,RH}$	565°C	$p_{steam,SH}$	18 MPa
$T_{sat,RH}$	175°C	p_{cond}	7 kPa
($p_{sat,RH}$)	(0.89 MPa)	η_{turb}	90%
T_{FW}	175°C	η_{pump}	80%
$\Delta_{Tmin,HX}$	5 K	$\Delta p_{max,HX}$	3% (abs. p)

The intermediate temperature of the helium at a position between the HT-TES and LT-TES is not a design choice; rather, it is dictated by the required heat distribution between the preheating and boiling stages and the superheating and reheating stages. The salt temperatures must remain within the freezing point limits; hence, the LT-TES cold tank temperature stems from the pinch-point analysis conducted on the corresponding CHXs. Furthermore, the heat exchanger type considered was a shell and tube, with the higher-pressure fluid always on the tube side. Thermodynamically, the reactor heat can support even higher cycle parameters, increasing the pressure (to supercritical) and temperature (steam cycles up to 650°C–700°C) design values. Such parameters are possible since multiple coal plants have been built and operated at these parameters, though not without issues in terms of cost and reliability, especially in flexible operation modes [9].

3.2.2 BOP and Nuclear-TES Analysis

As already mentioned, to fulfill the requirement of achieving the desired upper temperature for TES while still maintaining a fully decoupled system within the limitations specified in the previous section, the most viable solution is to employ two TES plants at two distinct storage temperatures. As the LT-TES will operate at below 250°C, HITEC salt was selected for use with it. High temperatures are achieved in the HT-TES by using $MgCl_2$-KCl, which has a melting point of 422°C.

Using two distinct TES loops (i.e., the HT-TES and LT-TES)—with the heat being transferred to/from each of them at a similar rate—coupled to a simple steam cycle would result in the boundary between the HT-TES and LT-TES to be located in the middle of the evaporator. However, this arrangement presents technical challenges and is impractical. Thus, an alternative approach involving a reheat cycle is proposed. The reheat cycle enables a better heat distribution among discrete components, such that the heat input for the preheater and evaporator is comparable to that for the superheater and reheater. Moreover, the reheat cycle generally offers greater thermal efficiency than does a simple steam cycle.

Obviously, a reheat cycle affords several advantages when used in configurations that involve a TES. One possible reason why this cycle is not typically considered in conventional nuclear plant designs (except for steam-steam reheat for limiting steam wetness in LWRs) is the increased complexity that arises when coupling the reheat cycle directly to the primary loop, which operates inside containment and often at high pressure (nuclear safety considerations). Unlike conventional plant designs, a TES-coupled plant offers a solution, in that the low-pressure heat transfer fluid in the TES loop can fully decouple the BOP cycle from the high-pressure primary side. Thus, the reheat cycle can be employed between the low-pressure TES plant and the BOP without posing technical difficulties. Note that the reheat cycle is also commonly employed in concentrated solar power (CSP) plants, especially ones utilizing a solar tower configuration and a molten-salt as the heat transfer fluid [10],[11].

A simplified thermodynamic model of the cycle was developed, considering a single open feedwater heater (deaerator) heated via controlled or uncontrolled extraction after the high-pressure turbine (HPT), upstream of the reheat process. To fulfill the thermal balance and prescribed temperature limits, the Q-T curves and design parameters were carefully tuned and calibrated. For example, the temperature of the feedwater entering the preheater must be about 165°C. At higher feedwater temperatures, efficiency can be improved by enabling greater feedwater heating through turbine extraction; however, this presents pinch-point challenges (e.g., cooling the salt in the LT-TES to sufficiently low temperatures and achieving the prescribed helium temperature limit of 250°C). Additionally, at lower feedwater temperatures, an excessive amount of heat would be transferred from the LT-TES as compared to the HT-TES, leading to unnecessary losses in efficiency. The details of the heat transfer are well visualized in the Q-T diagram in Figure 10 (see Figure 11 for the Aspen Models and CycleTempo diagrams). To ensure an optimal heat distribution during system operation, a specific control method was implemented in the subsequent dynamic models described in Section 6.3.

Note that certain parameters for the fully balanced system may slightly differ based on the fluid package considered or the fluid properties of helium, the storage medium, and water. The resulting thermodynamic parameters of the system during design steady-state operation (no charging or discharging) are listed in Table 5.

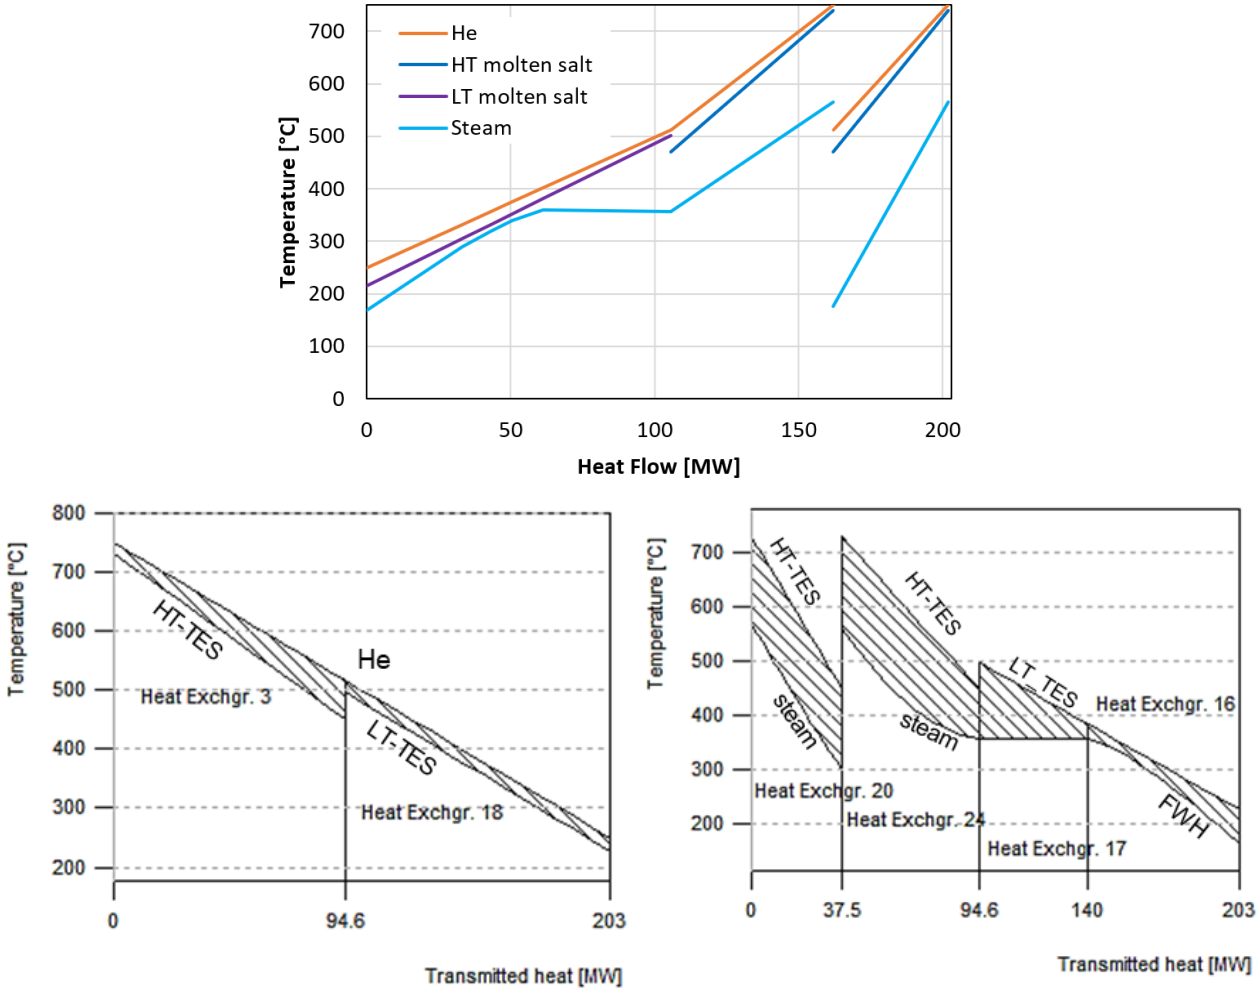

Figure 10. Q-T diagram of the heat transfer within the reheat cycle system (top) corresponding to the Aspen Model in Figure 11, as well the He-salt heat exchangers (bottom left) and salt-to-steam set of heat

exchangers (right) corresponding to the CycleTempo in Figure 11. Heat exchanger numbers and their locations are shown in the model schematic in Figure 11.

Table 5. Nominal steady-state parameters of the decoupled system with the reheat cycle, during nominal operation. The naming of the streams corresponds to the process flow diagram in Figure 11.

Stream	p [kPa]	T [°C]	h [kJ/kg K]	x [-]	m [kg/s]
He Loop					
He in	5840	750	3784	—	78.5
HT he out	5840	512	2550	—	78.5
He out	5840	250	1189	—	78.5
HT-TES					
HT MS out	150	745	363.5	—	304.7
HT H-MS in	150	745	363.5	—	304.7
S1	150	745	363.5	—	125.2
HT C-MS out-A	150	470	45.6	—	125.2
S2	150	745	363.5	—	179.5
HT C-MS out-B	150	470	45.6	—	179.5
HT C-MS in	150	470	45.6	—	304.7
LT-TES					
LT H-MS out	150	500	-1281	—	311.7
LT H-MS in	150	500	-1281	—	311.7
LT C-MS out	150	215	-1623	—	311.7
LT C-MS in	150	215	-1623	—	311.7
Steam Cycle					
G	701	39	162.57	—	59.3
12	701	162	697.24	0	59.3
FW (= G-2)	19000	165	719.91	—	59.3
Sat liquid	18525	359	1755.27	0	59.3
A (sat vapor)	18050	357	2507.80	1	59.3
B	18050	565	3460.03	—	59.3
C	893	178	2774.32	-1	47.5
D	848	565	3622.43	—	47.5
E	7	61	2604.69	-1	47.5
F	6.86	39	161.78	0	47.5
11 (Extraction)	893	178	2774.32	-1	11.8

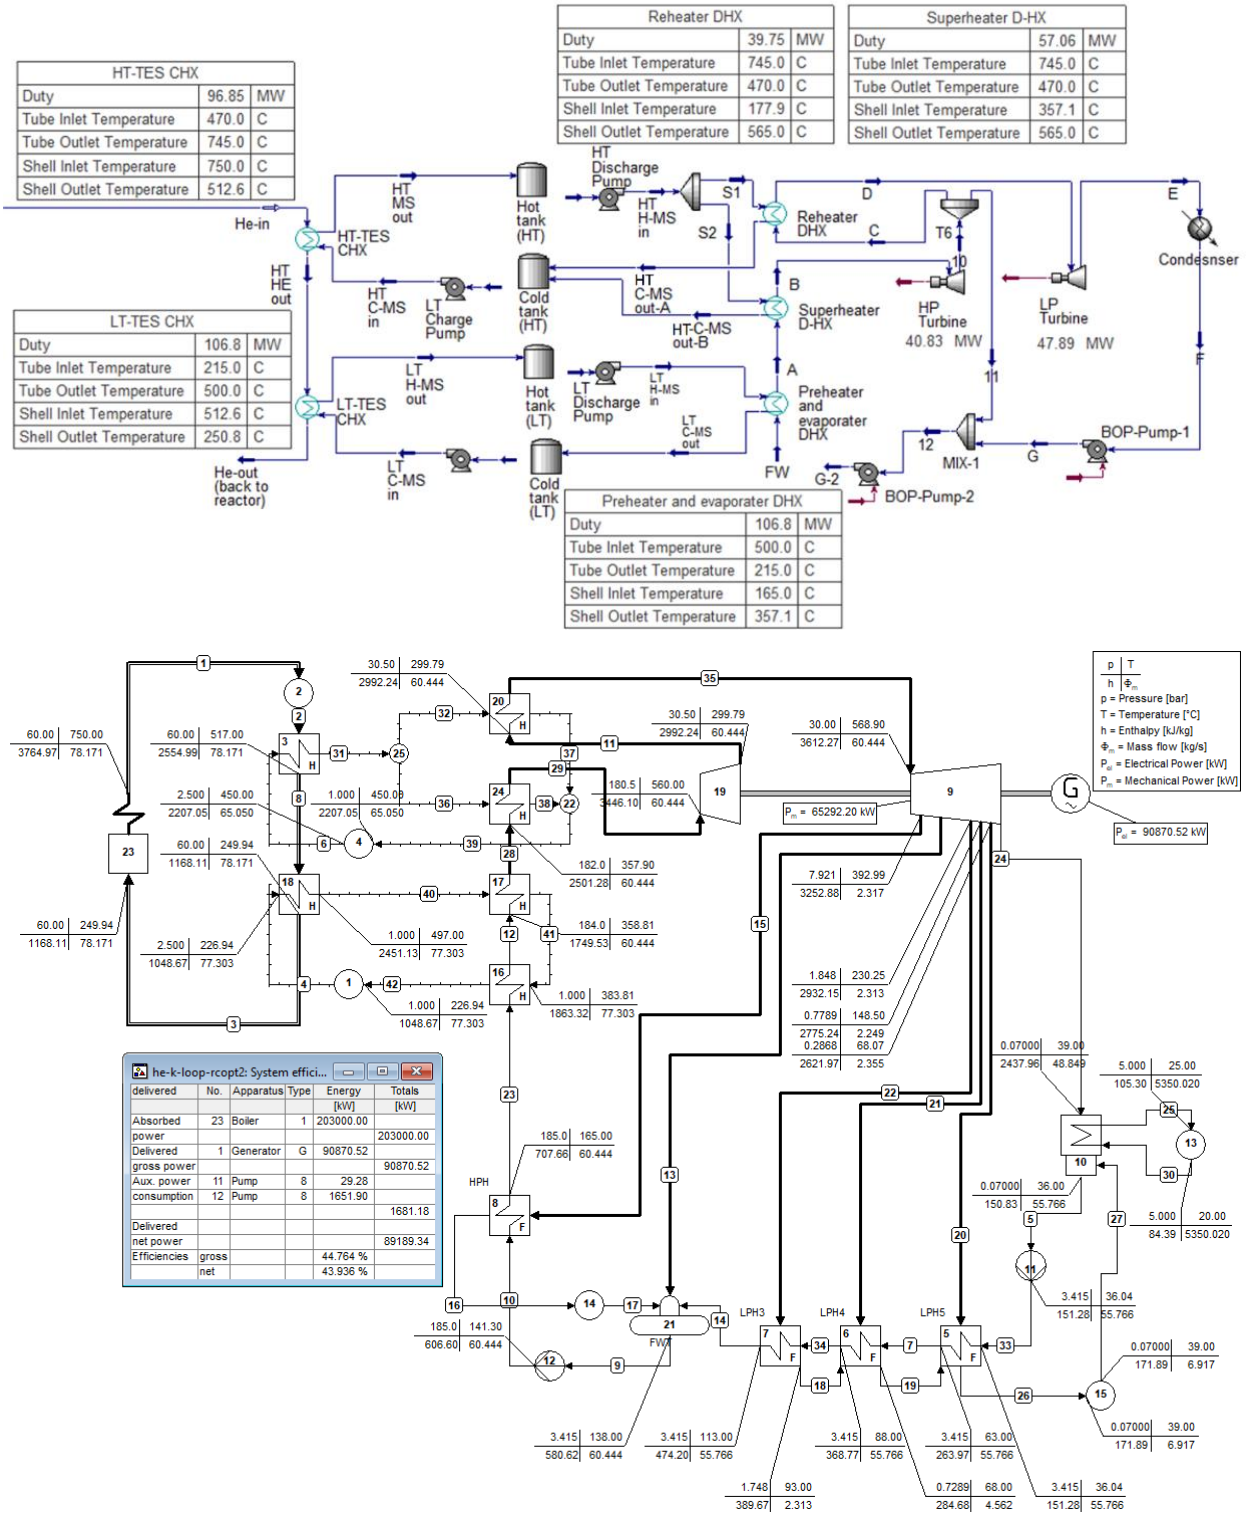

Figure 11. Aspen HYSYS model (top) with simple reheat and CycleTempo (bottom) with full feedwater heating train for the fully decoupled configuration.

The cycle achieves 43.1% efficiency, independent of whether the plant is undergoing charging, nominal, or discharging operations (assuming constant pressures, as covered in the dynamic modeling

discussion in Section 6.3). This efficiency is higher than that of the conventional HTGR BOP (i.e., 42%), regardless of considering the standard turbine efficiency of 90% (as compared to 97% for the baseline HTGR model). Note that the proper multistage feedwater heating design typically seen in CSP systems (at least five feedwater heaters) affords the potential for even further efficiency enhancements.

Per the Q-T diagram, a substantial size reduction of the HT-TES CHX can be achieved, allowing for higher temperature differences between the primary and secondary fluids. The selection of the HT-TES upper temperature partly stemmed from demonstrating the technical limits of such a system by enabling the HT-TES to operate and store heat at the highest possible temperature. Besides the ability to discharge higher quality heat, the higher storage temperature also helps minimize the storage system size by utilizing a higher temperature difference between the cold and hot tanks. For power-only applications and the considered steam parameters, this means a high potential for further size and cost optimizations.

3.3 Comparison of System Parameters

For an overall view of the thermal and electrical installed capacities and efficiencies of the explored system, the parameters of the reference and TES-integrated decoupled plants are summarized in Table 6. The TES system coupled with the HTGR (from previous work) is included for comparison. Note that the current systems offer constant plant efficiency, which cannot said of the prior configuration. The HDR of the partially decoupled system is limited to around 33%. The advantages of the fully decoupled system clearly suggest that reactor system developers should consider an intermediate heat transfer loop when dealing with steam cycle conversion systems.

Table 6. Overall parameters of the reference and TES-integrated decoupled plants.

Design	Reactor output (MWth)	Heat dispatch to TES (MW_{th})	Heat dispatch to BOP (MW_{th})	Net electricity (MWe)	Plant thermal efficiency (%)	Turbine efficiency
Reference plant (no TES)	203.6	N.A.	203.6	86	42.2%	97%
Steam-charged TES nuclear plant (prev. work)	203.6	203.6	203.6 + boosting	86 + boosting	39.5-42.2%	97% [3]
Partially decoupled TES nuclear plant	203.6	203.6 - HDR	203.6 + boosting	86 + boosting	42.2%	97%
Fully decoupled TES nuclear plant	203.6	203.6	203.6 + boosting	88 + boosting	43.1%	90%

4. TECHNOECONOMIC ASSESSMENT

This section examines the development of cost functions for the fully decoupled system. Using the current information—obtained from steady-state models—for the coupled nuclear-TES cases affords sufficiently detailed data for creating cost functions that reflect the system cost as a function of IES size (i.e., TES boosting). The developed cost functions in this section were then incorporated into HERON's stochastic optimization process for each of the three different electricity markets. This integration provides region-specific direction for optimizing the superset capacities to meet the distinct needs of each regional market.

4.1 Methodologies

Section 4.1.1 covers the cost function development and equipment sizing methodologies. Section 4.1.2 discusses the cost functions created by utilizing the most recent publicly accessible data from Aspen Process Economic Analyzer (APEA)-V11, or by employing the module costing technique (MCT).

4.1.1 Cost Functions

Fixed costs and variable costs are economic drivers that play a crucial role in IES cases, particularly in terms of how they scale with changes to system size during the subsequent HERON optimization. Hence, the cost functions were formulated using the following scaling equation:

$$Y = A \cdot (D / D')^{x}$$

where

Y is the installed cost of the equipment of interest

A is the reference installed cost for the equipment, corresponding to capacity or size *D*'

D is the scaled equipment size determined in the optimization or that must be costed

D' is the reference equipment size that will be fixed and assigned to each piece of equipment

x is the exponential scaling factor (<1 implies economy of scale).

In developing cost functions, *A*, *D'*, and *x* were assigned as constants for each piece of equipment, enabling derivation of the installed cost (*Y*) of the equipment at any scaled size (*D*).

Cost functions based on these models were developed using the latest publicly available published data from vendors, data taken from APEA-V11, or data generated via MCT, as discussed in the following section. Once a costing database was created, the cost functions were developed via the following steps:

[1] The equipment and installation costs for all components that appeared in the nuclear-TES models (e.g., turbine, heat exchangers, condensers, pumps, tanks, and energy storage materials) were tabulated to form a database of component costs at discrete component sizes.

***Note**: The installed cost for each piece of equipment reflects the direct and indirect costs, which include estimates for the following cost elements: equipment and setting, piping, civil, structural steel, instrumentation, electrical installation, insulation, and paint.*

[2] A separate cost function (equation) for each component type was developed (installed cost as a function of equipment size). These cost functions are presented in Section 4.2.

The individual equipment subsets were grouped to form three major TES (superset) models (see Figure 12):

a. Charge model: A group of charge-loop-specific equipment, specific to the single-phase heat transfer fully decoupled nuclear-TES case. The charge superset herein replaces the conventional/reference NPP steam generator and always charges the TES at a constant rate that aligns with the reactor heat. The cost of the high-pressure steam generator in the reference case (i.e., no TES) is assumed to be comparable to that of the He-to-low-pressure molten-salt heat exchangers (HT-TES CHX and LT-TES CHX). Therefore, the CHX assumes no additional cost in comparison to the conventional plant design. The cost difference of the charging system is therefore due solely to the charging pumps, which operate at a fixed charge rate that aligns with the reactor power output, having a fixed cost irrespective of the scaling of the TES system.
b. Storage model: A group of storage-loop-specific equipment (molten-salt material and holding tanks), with a reference system size (*D'*) that relates to the storage capacity in MWh_{th}.
c. Discharge model: A group of discharge-loop-specific equipment (molten-salt pumps from the HT-TES and LT-TES, preheater with evaporator, superheater, reheater, turbine, condenser, condensate, and main feedwater pump), with a reference system size (*D'*) that relates to the thermal power (in MW_{th}) consumed by the turbine. Specific to the single-phase fully decoupled nuclear-TES case, the cost of this superset model is the difference between the cost of a larger

steam Rankine cycle designed to achieve a given amount of heat boosting and the total cost of a conventional plant (0% boosting, no TES, steam produced directly from the steam generator).

[3] A cost function was created for each of the three superset models by using the individual cost functions for each piece of equipment falling under that superset model. The cost functions of the three superset models can indicate the additional cost realized by adding TES as a function of system size, as detailed in Section 4.2.

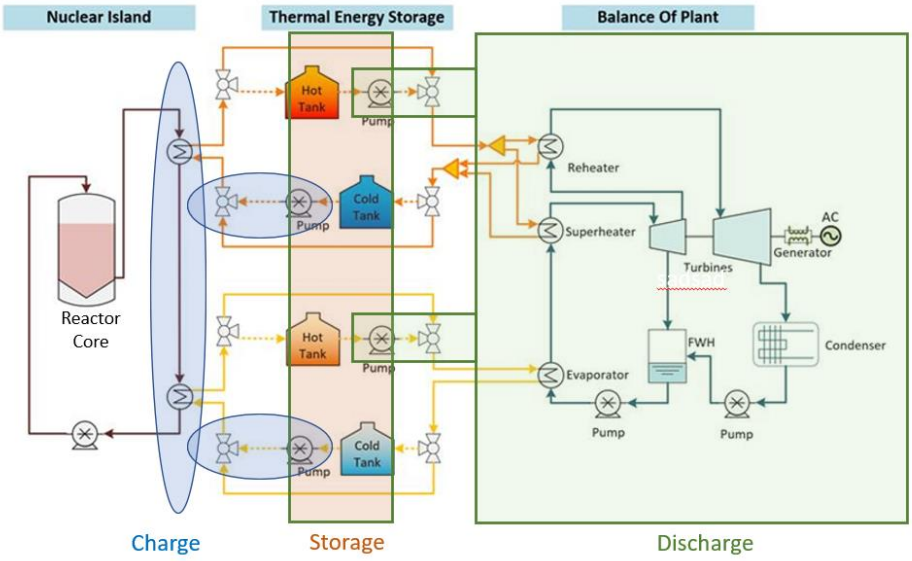

Figure 12. Component groupings for deriving the three superset cost functions.

4.1.2 APEA and MCT Costing Methods

Two methods were used to acquire the cost data for the presented systems. The first was to access the latest publicly available data from APEA-V11, an Aspen suite module specifically designed for analyzing and assessing the economic aspects and costs of various process plants and components. Second, when the equipment being investigated—or the operating conditions thereof—fell outside the APEA library or database, MCT was used to acquire the capital cost of such equipment. Per Turton et al. [12], MCT works by assigning the bare module cost for each piece of equipment. The cost of a specific device includes both the direct and indirect capital costs, as well as all additional items such as piping, installation, and instrumentation and control elements. These costs are factored in by applying a multiplication factor for each component. First, the capital cost is estimated based on the base condition, under the assumption that the equipment is made from regular carbon steel and operates under ambient pressure. Next, the calculated costs are adjusted to account for the nominal operational pressure and the material used in the actual system being considered. Finally, the calculated costs are further adjusted by using the Chemical Engineering Plant Cost Index [13] to obtain the net present market value of the equipment.

4.2 Cost Analysis

4.2.1 Material Cost Data

The methodology for acquiring the raw material cost data entailed a comprehensive literature review of published data, cost analysis reports, and market surveys. The raw material costing data were then cross referenced and verified through global suppliers, taking into account both relative (different multipliers for different material prices) as well as absolute ($/kg) prices.

Table 7 shows the relative prices of the considered Ni alloys in comparison to the reference material, stainless steel 316L. Several values for Alloy 800, Alloy 625, and Inconel 625 were taken directly from the suppliers. The rest of the values were calculated based on additional inputs such as absolute prices or relative prices in reference to stainless steel 309. The relative price for Incoloy 800 (or its derivatives) falls within the range of 2.4 to 3.9, with a median value of 3.51. The relative price for Inconel 625 (or its derivatives) falls within the range of 6.8 to 8.6, with a median value of 7.75. To maximize accuracy, the refence material factors in the MCT for carbon and stainless steel were then calibrated using heat exchangers previously costed by APEA.

Table 7. Overview of high-temperature material costs.

Material	Relative price to 316	Relative price to 309	Price $/kg	Year	Ref.
SS 316	1.00	0.75	5.04	2019	[14]
Alloy 800	3.65	—	—	2012	[15]
RA800AT	2.40	1.8	—	n.a.	[16]
Incoloy 800 / H / HT	3.37	—	17	n.a.	[17]
Incoloy 800 Price per kg in India	3.85	—	19.4	2019	[18]
Incoloy 800 HT Price per kg in India	3.37	—	17	2019	[18]
Alloy 800 H Price in the U.S.	4.76	—	24	2019	[19]
Alloy 625	7.75	—	—	2012	[15]
Inconel 625	6.80	5.1	—	n.a.	[16]
Inconel 625 Price per kg in India	8.57	—	43.18	2019	[18]

The manufacturability of nickel alloys in terms of machining and welding is often reported as being comparable or slightly higher than that of stainless steel 316. Thus, the final cost of manufacturing heat exchangers, tanks, and other components can be reasonably approximated by multiplying the cost of manufacturing them from stainless steel 316 by a material cost multiplier.

4.2.2 Heat Exchanger Costing

Before obtaining the equipment cost and installed cost for the heat exchangers, their relevant size and geometry must be determined for the system being investigated. This applies to both the plant's charge and discharge heat exchangers. The costing process for heat exchangers entails the following steps:

[4] Sizing: The sizing of the heat exchangers was carried out, using the Exchanger Design and Rating (EDR) in Aspen HYSYS, as a rigorous shell-and-tube model. The operating conditions for the heat exchanger's two inlet streams and single outlet stream were defined based on the system being studied, and the sizing was accomplished by adjusting the allowable pressure drop on the hot and cold sides until four specific targets were met: (1) limiting the pressure drop ratio for each stream to less than 10% of the total inlet pressure, (2) achieving the most cost-effective solution, (3) limiting the Excess Surface Area in the EDR geometry window to a range of 0–5%, and (4) limiting the Δp-ratio (shell-side/tube-side pressure drop ratio) in the EDR geometry window to a range of 0.85–1. All heat exchangers in this study were sized based on Tubular Exchanger Manufacturers Association-type BEM heat exchangers.

[5] Geometry: Based on the analysis conducted in the previous step, an optimal geometry for the heat exchangers was selected, stream conditions were revised as needed, and the surface area of the heat exchangers was reported and used as input for the MCT model or APEA.

[6] Costing: MCT (described above), as applied to shell-and-tube heat exchangers, uses the heat transfer surface area, tube-side pressure, and material combination of the shell and tubes as inputs,

and ultimately adjusts the cost from the reference year to the desired value by referring to the Chemical Engineering Plant Cost Index. For Incoloy 800 heat exchangers, the material factor followed the relative median price for stainless steel, as listed in Table 7. The heat exchangers in the HT-TES plant are considered to be made of Incoloy 800, whereas both heat exchangers in the LT-TES plant are made of stainless steel.

4.2.3 Tank and Storage Media Cost Data

The HITEC heat-transfer salt demonstrates excellent stability, providing several years of reliable performance at temperatures of up to approximately 850°F (454°C) when free from contaminants. The operating conditions studied in this research remained below the maximum recommended temperature limit (538°C) for HITEC, but it is important to note that, when operating near the upper limit, the salt may undergo a gradual thermal breakdown in closed systems. To mitigate this, it is advisable to use an inert gas blanket or actively control the atmosphere above the molten salt via the oxygen content [20]. Over extended periods, nitrogen gas can slowly evolve from HITEC, and the freezing point of the salt mixture gradually increases. Therefore, the design cold-tank temperature in this work was kept at well above the freeze point. To comprehensively understand the long-term performance of HITEC at near its maximum operating temperature range, further analyses and experimental work are required. Note that other companies have suggested the use of sodium hydroxide (NaOH) as a stable molten salt at temperatures of up to 700°C [7]. However, limited technical information on this alternative is available in the public domain. While this study primarily focuses on the feasibility of system design rather than material science, conservative design criteria were employed, even though detailed evaluations of material characteristics were not conducted.

For increased confidence in the cost estimates for the TES systems, the results of our costing analysis were compared to those of other similar systems typically adopted in the CSP industry. These system costs are typically described in units of $/kWh, even though the value taken from this approach can be imprecise, as it will vary with the actual operating temperature difference, density, and cost of the selected storage materials. In contrast, [21] is useful in that it provides the specific cost of a hot tank as 1293 $/m^3 (3 x D = 43 m, H = 12 m, stainless steel), a cold tank as 740 $/m^3 (3 x D = 41 m, H = 12 m, carbon steel), and the specific cost contribution from installed power capacity as 25.17 $/kW$_{th}$ (or from installed storage capacity as 22.73 $/kWh$_{th}$). The systems that utilize solar salt are, interestingly enough, becoming cheaper over time, mainly as a result of reported cost decreases for storage materials. A summary of the TES system and molten-salt storage material costs is provided in Table 8 and Table 9.

Table 8. Overview of TES system capital costs.

Source, year	SAM 2013	SAM 2017	SAM 2019	SAM 2022	Gen 3 CSP 2021
Unit cost ($/kWh$_{th}$)	27.33	24	22	22	39.5
Note	566/288°C sol. salt	—	290/550°C, sol. salt	290/574°C sol. salt	720°C/500°C $MgCl_2$-NaCl
Ref.	[21]	[22]	[23]	[24]	[26]

Table 9. Overview of molten-salt storage material costs.

Salt cost	Solar salt	Hitec	Hitec XL	$MgCl_2$-KCl
Price $/kg	0.75, 1.3, 1.77, 1.81	1.93	1.66	0.75
Ref.	[23],[25],[3],[21]	[25]	[25]	[26]

The MCT costing method, described above, was used to obtain the cost of the tanks. The tanks are of the API fixed-roof type, with a volume of up to 30,000 m^3, and operate at near-ambient pressure. For this type of tank, the parameters input to the MCT model are the volume and material factor, with the adjusted

material factor again being used. The HT-TES hot tank must be made of Incoloy 800. The HT-TES cold tank and LT-TES hot tank are stainless steel, and the LT-TES cold tank is carbon steel.

4.2.4 BOP Equipment Costing

The turbine equipment costs were obtained from APEA, taking into account the 22 MWe maximum modeling capacity for turbines. To accommodate any equipment with a capacity exceeding this limit, the turbines are divided into multiple sections, each with a capacity of 20 MWe or less. Downsizing a turbine involves adjusting the operating conditions and reducing the mass flow rate accordingly. By analyzing the data points from turbines whose output is 20 MWe or less, a scaling constant (x) can be derived. This scaling constant is then utilized to estimate the cost of turbines whose output exceeds 20 MWe, utilizing the cost function derived from the data points. In addition, condensers are modeled as heat exchangers in Aspen HYSYS, and their sizing is performed in EDR via an approach similar to that used for the charge/discharge heat exchangers. However, in this case, the cost of the condensers was acquired from APEA instead of via the MCT model. This was done to prioritize APEA as a costing method whenever the operating conditions align with the APEA's database and capabilities. In the case of condensers, APEA was best suited to provide accurate cost estimations. The pumps for the molten-salt loop and BOP were also mapped and costed through APEA.

4.3 Cost Functions

This subsection discusses the cost functions for the nuclear-TES use case. The data points retrieved from the steady-state models that were developed using different system sizes and coupling options provide sufficient resolution for creating cost functions for the various component and system sizes. Note that the present work primarily focuses on the fully decoupled case, due to its greater application flexibility. The data collected from the steady-state models, which were created to assess the models under different boosting values, offer sufficient resolution to establish the cost functions for different component and system sizes.

4.3.1 Cost Function Results for the Individual Equipment

First, the costing data for the various sized equipment were acquired from the MCT model or APEA, then used to create a cost function for each piece of equipment. The costing data for the molten-salt storage media and storage tanks were the only components analyzed outside of APEA or MCT, following the methodology outlined in Section 4.2. Table 10 summarizes the cost function constants for the various components in the fully decoupled nuclear-TES use case, along with the superset model under which each piece of equipment falls. Figure 13 combines the cost function curves for the specific components. The green line in the turbine data indicates when the data began to be extrapolated after reaching 20 MWe (APEA's cost database limit).

Table 10. Summary of cost function constants for the fully decoupled HTGR-TES use case components.

Equipment	A	D'	X	Error	Notes
Turbine	2,237,600.0	11.26	0.7192345	0.4%	—
Condenser	4,577,000.0	92.18	0.6415093	9.3%	—
Cond. Feedwater Pump	1,063,000.0	0.195	0.4689908	4.3%	101 to 230 kPa
HT-TES CHX	—	—	—	—	Replacement of SG
LT-TES CHX	—	—	—	—	Replacement of SG
Preheater and Evaporator	2,862,475.5	213.60	0.5124595	5.8%	—
Superheater	4,254,510.9	57.06	0.5411958	6.0%	—
Reheater	2,331,326.1	39.7	0.5011288	5.3%	—

Equipment	A	D'	X	Error	Notes
BOP Pump (>400 kWe)	1,034,400.0	1.369	0.5078757	8.2%	—
HT Cold Salt Pump	85,900.0	0.14977	0.5384690	9.8%	—
HT Hot Salt Pump	1,239,400.0	0.16215	0.8125600	2.8%	—
LT Cold Salt Pump	234,200.0	0.07284	0.7066847	2.4%	—
LT Hot Salt Pump	98,300.0	0.08169	0.8024849	9.9%	—
HT-TES Tanks	8,705,590.1	581.1	0.7000000	2.4%	SS 316 and Incoloy 800
LT-TES Tanks	2,708,324.8	640.8	0.7000000	2.9%	SS 316
HT Molten Salt	4,607,064.0	6581520	0.8500000	6.9%	$MgCl_2$-KCl
LT Molten Salt	11,903,449.0	6732720	0.8500000	6.9%	HITEC

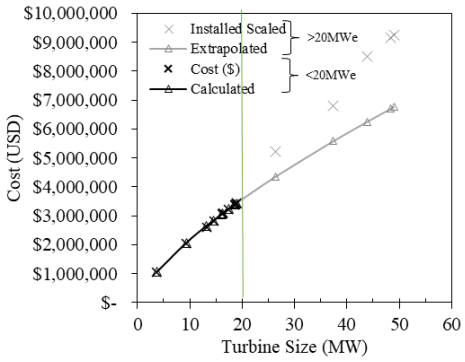

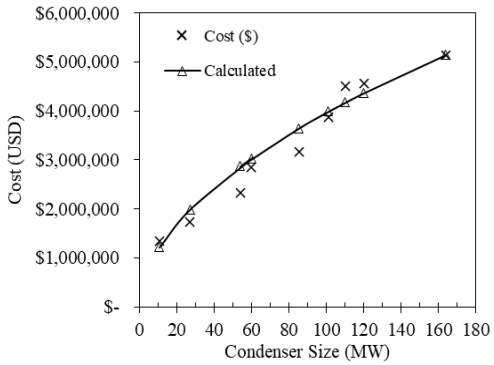

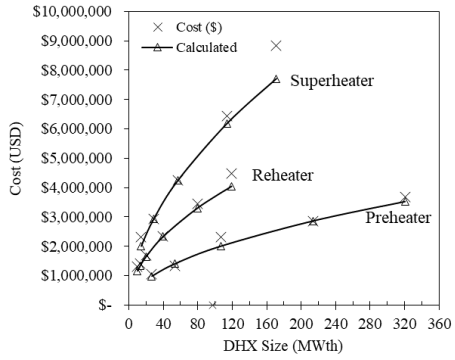

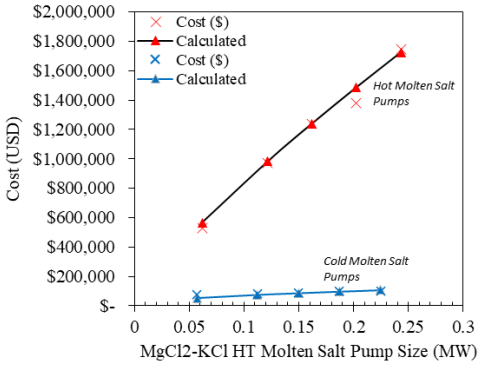

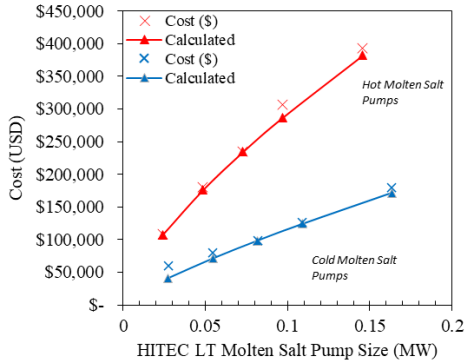

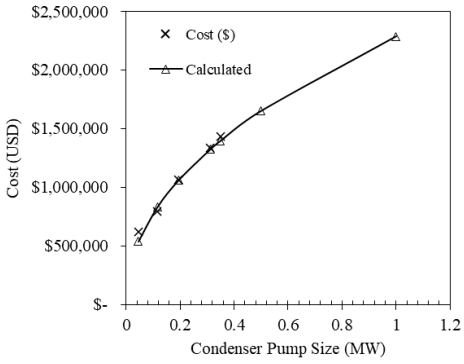

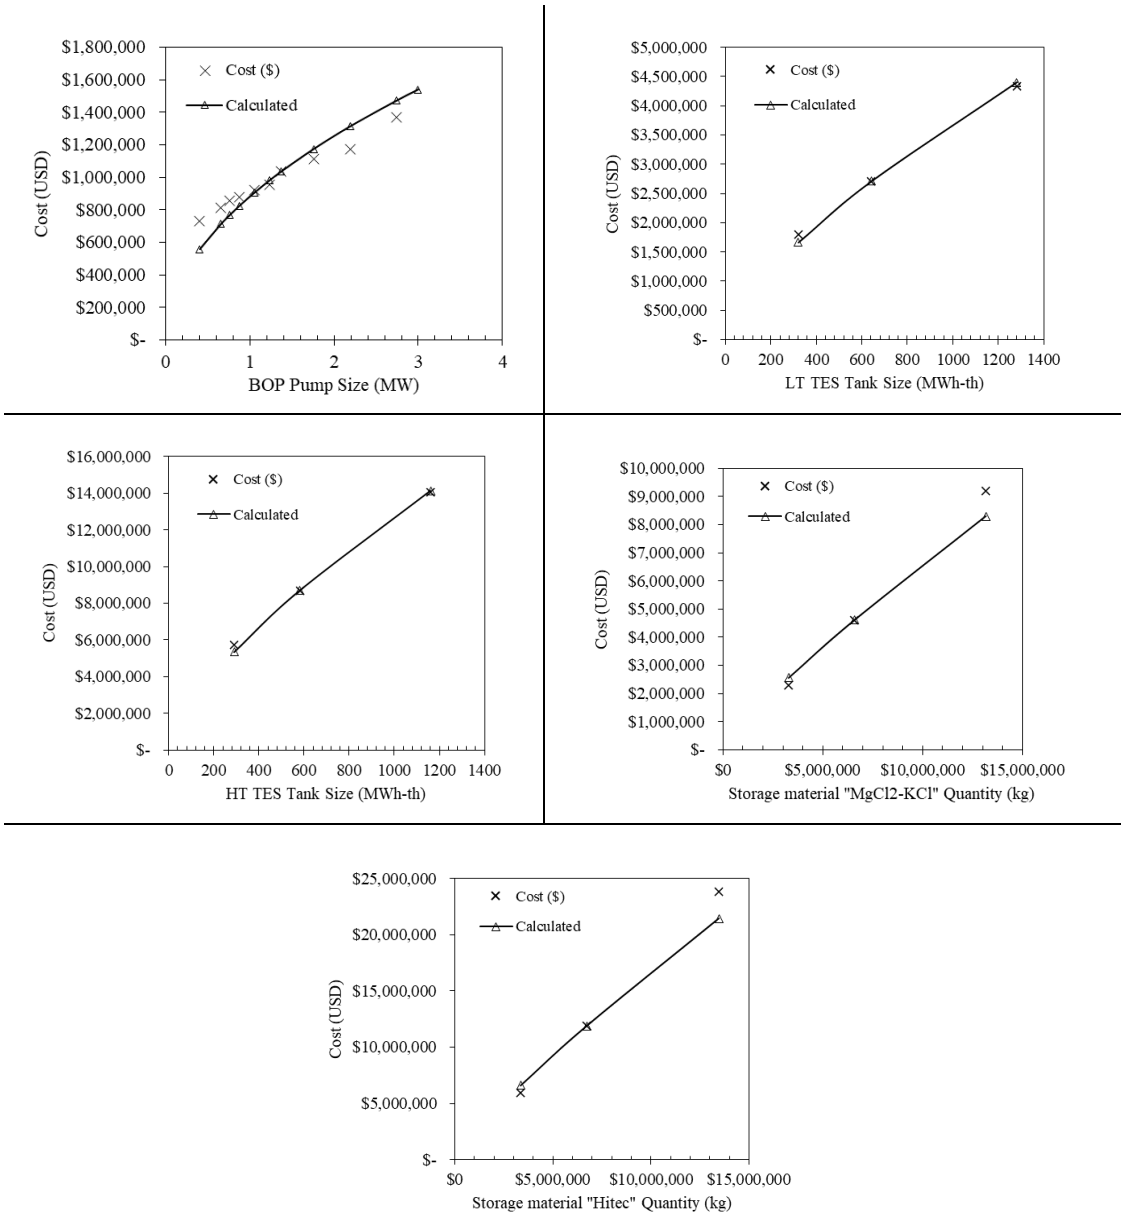

Figure 13. Cost function curves for the various fully decoupled nuclear-TES use case components.

4.3.2 Superset Cost Functions

As outlined in the previous section, the cost functions derived from the individual equipment components were utilized to construct three overarching models, defined herein as the superset models: the (1) charge model, (2) storage model, and (3) discharge model. Each of these superset models was created by combining the cost functions of the individual equipment subsets that fell within their respective categories. Figure 14 represents the combined curves of the storage and discharge superset cost functions, presenting them as a function of system size. The secondary x-axis in the discharge superset model establishes a relationship among system size, proportional energy boosting ratio, and cost.

summarizes the cost function constants for the three superset models in the fully decoupled nuclear-TES use cases, along with the mean absolute percentage error for each cost function. Note that

the discharge system is exclusively described as $Y = A \cdot (D / D')x + C$, wherein C = -16,661,065. The constant comes as a result of replacing the original BOP with the corresponding cost.

Table 11. Cost function constants for the fully decoupled nuclear-TES coupling superset models.

Superset Model	A	D'	X	Mean absolute percentage error
Charge	234,181.4	203.6 MW_{th}	0	0.0%
Equation: $Y = 234{,}181.4 \cdot (D / 203.6)^{0}$	—	—	—	—
Storage	22,258,933.6	916.2 MWh_{th}	0.78654	0.1%
Equation: $Y = 22{,}258{,}933.6 \cdot (D / 916.2)^{0.78654}$	—	—	—	—
Discharge	33,486,397.9	305.4 MW_{th}	0.63050	0.5%
Equation: $Y = 33{,}486{,}397.9 \cdot (D / 916.2)^{0.630504} - 16{,}661{,}065$	—	—	—	—

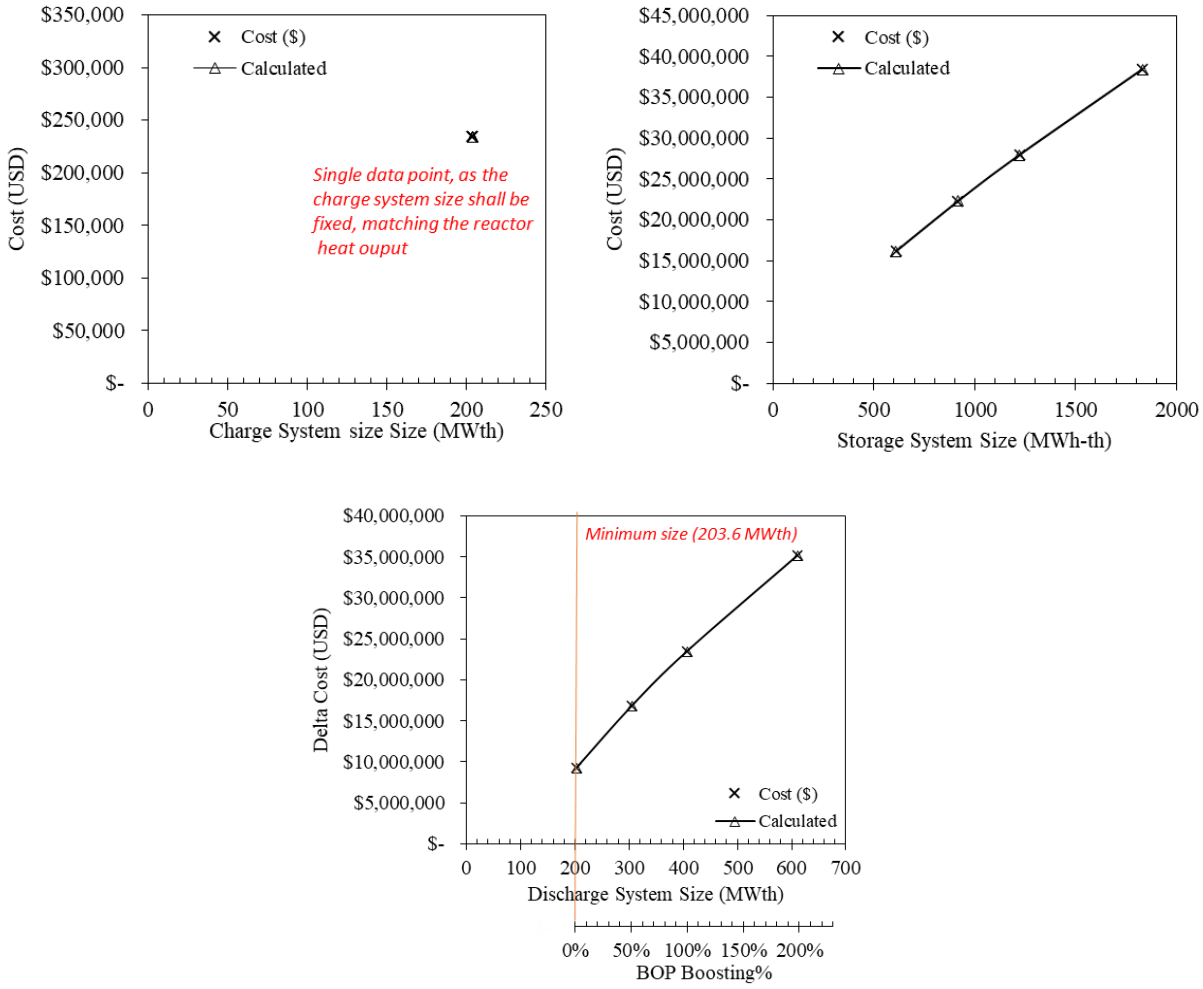

Figure 14. Cost functions curves for the fully decoupled nuclear-TES coupling superset models.

5. SYNTHETIC ELECTRICITY MARKET DATA GENERATION AND HERON ECONOMICS OPTIMIZATION

In this study, market prices are captured in synthetic histories, which allow for stochastic optimization. Synthetic histories are time-dependent models that can be stochastically sampled and are created to include a range of stochastically possible scenarios not necessarily present in the time series data upon which they were trained. In this study, synthetic histories retain the same statistical properties as the historical data they were trained on, but provide different time-dependent realizations in each sample. As a result, the synthetic histories, when sampled to generate a representative synthetic price signal, are distinct from the historical price signal while maintaining virtually the same distribution of prices.

5.1 Selection of Historical Price Signals

We sought to identify locations with heightened potential for power arbitrage and to then conduct economic optimization at those locations. This approach assumes that a rational market actor would seek an economically advantageous location at which to site a grid-connected energy storage system.

Due to their publicly accessible data, the three regions in which we searched for suitable locations (i.e., nodes) were operated by the following balancing authorities (i.e., market operators, referred to in the present report simply as "markets"): PJM, ERCOT, and MISO. A node is generally the lowest-level element of an electric grid monitored by the operator, and typically corresponds to a physical location.

Under the premise that high variability coupled with high prices would be economically advantageous, we sought a location with a high standard deviation, high mean, and high 95th percentile LMP within the real-time market of each studied region. These regions had readily available LMP data, and the data set (i.e., historical price signal) for the selected location in each was then used to generate synthetic histories.

Ultimately, the highest 95th percentile price was used to select each node from the three market regions, as outlined below. Of the three nodes, two also featured the highest standard deviation within their respective regions.

Real-time hourly[a] market LMP data for 2018–2021 were downloaded from the publicly accessible website of each market operator. For PJM, 17 nodes geographically separated throughout the PJM footprint were somewhat arbitrarily preselected for data acquisition. Load zones (LZs) were used in lieu of nodal data for ERCOT, due to immediate access to their LMP data at the time of download, and included the North, West, and South LZs. Once the LMP data for a single PJM node or a single LZ were compiled, the mean, standard deviation, and 95th percentile prices were calculated over the period of 2018–2021[b]. The data structure made available by MISO allowed for downloading the LMPs for all nodes throughout the period of 2018–2021, thus enabling ranking of all the nodes. A purpose-written Python script listed all the nodes that remained active over the entire 4-year period, then ranked them by 95th percentile price. For each market, we selected the node or LZ with the highest 95th percentile price.

---

[a] While the real-time market in each regional balancing authority operates on a smaller time increment, hourly data were used due to their availability and the practical considerations involved (e.g., downloading many files manually, and the inability to filter for a single node in PJM's DataMiner2-archived data covering the earlier parts of 2018–2021). PJM and MISO directly provide hourly averaged data, while quarter-hourly data, among other types, are located on ERCOT's public website. Quarter-hourly data from ERCOT were averaged to create an hourly dataset.

[b] In the case of ERCOT, the LMPs of February 14–19, 2021, were excluded from the data so as to remove the effect of the exceptionally high prices during the Texas blackout event of 2021. In the data set used for synthetic history generation, hourly prices for these dates were substituted with linearly interpolated values spanning from the last hour of February 13 to the first hour of February 20 ($1,190.35 and $17.43, respectively).

Table 12 gives the LMP means and standard deviations for the selected nodes (PJM and MISO) and the selected LZ (ERCOT). Their historical data are plotted in Figure 15 through Figure 17, respectively.

Table 12. LMP statistics (in $/MWh) for the selected nodes and LZ, for the period of 2018–2021.

	PJM (Brandon Shores)	ERCOT (West LZ)	MISO (node in Northern States Power service area)
Node/zone ID	50660	WEST	NSP.HATFIHAT1
Mean	34.03	38.58	32.32
Standard deviation	33.30	146.72	33.99
95th percentile	75.91	106.90	77.49

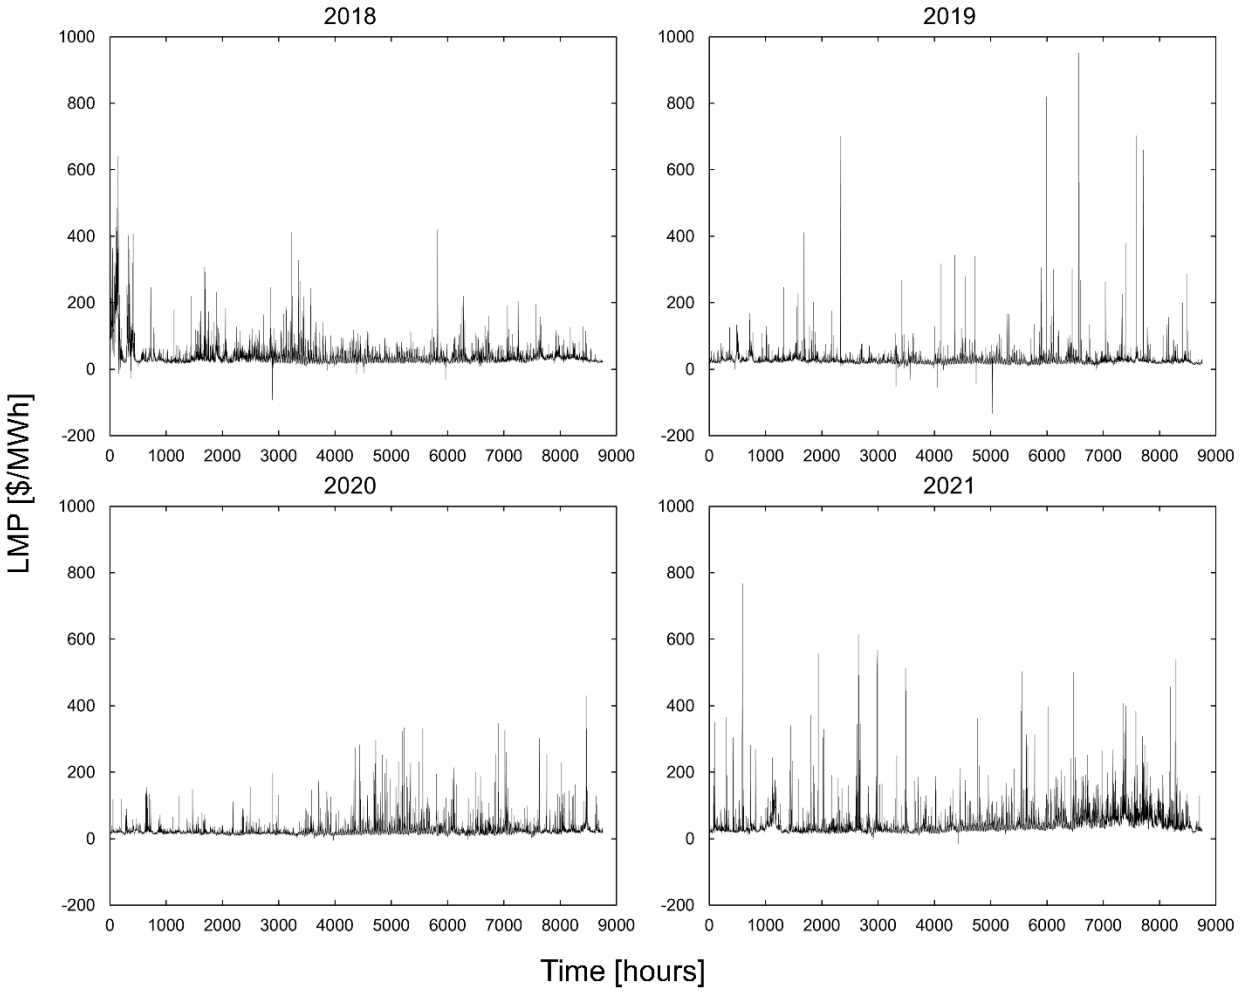

Figure 15. Historical real-time market LMPs at PJM's Brandon Shores node (hourly-averaged).

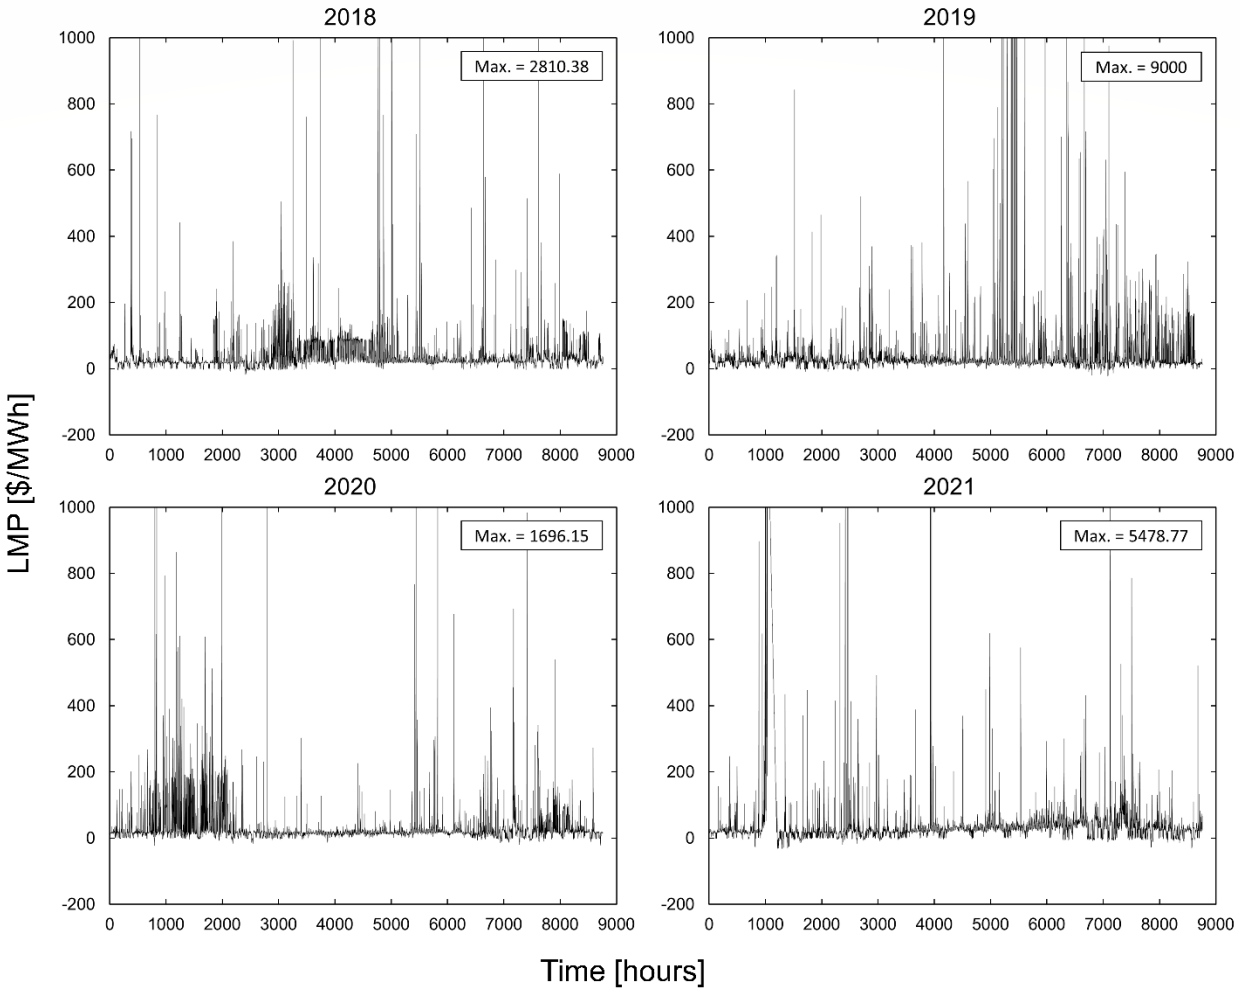

Figure 16. Historical real-time market LMPs for ERCOT's West LZ (hourly-averaged from quarter-hourly data).

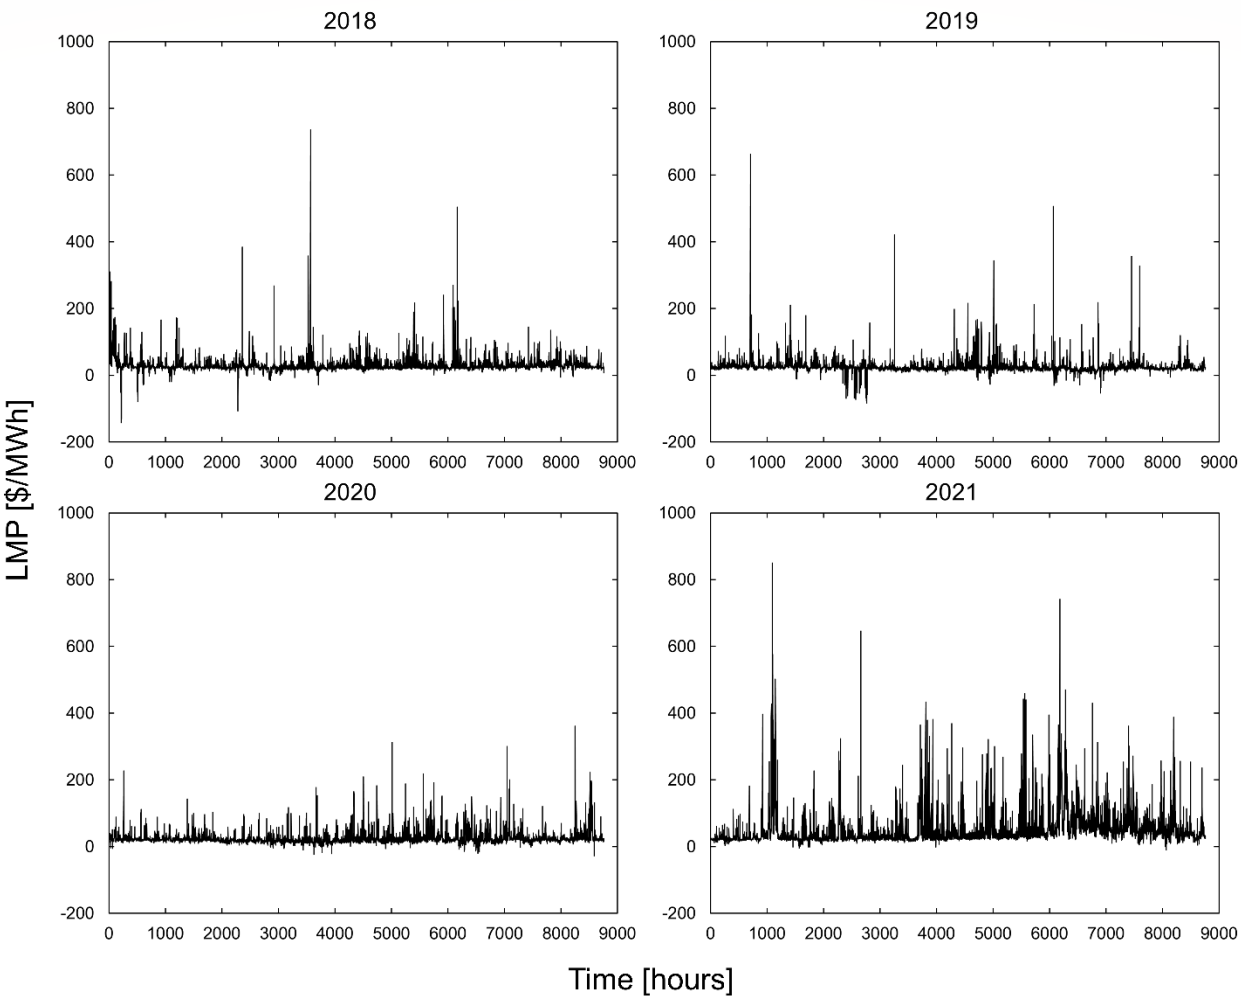

Figure 17. Historical real-time market LMPs at MISO's NSP.HATFIHAT1 node (hourly-averaged).

5.1.1 Synthetic History Generation

The overall workflow for synthetic history generation, shown schematically in Figure 18, consisted of data collection and formatting for input to Idaho National Laboratory's RAVEN software, which was configured to generate a synthetic history object (arma.pk) for subsequent use in HERON.

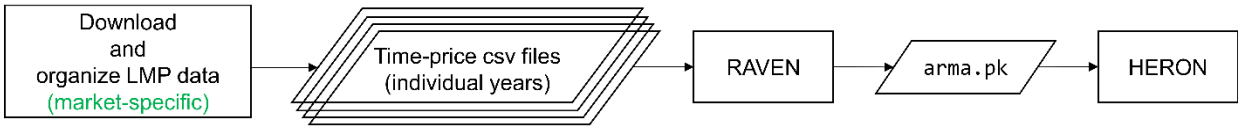

Figure 18. Simplified workflow for synthetic history generation.

RAVEN generated synthetic LMP histories via its ARMA reduced-order model (ROM), with parameters chosen largely through user judgement, as aided by visual examination of plots and statistical measures (described below). These parameters, listed in

Table 13, were varied within the Models block of the RAVEN input file to arrive at trained models. The models produced synthetic histories, each consisting of 20 clusters, representing each year of the history. The clusters had a uniform length known as the pivot length (hereafter referred to as the "segment length").

Parameters developed based on the LMP data for PJM's Brandon Shores node were later directly applied to the ERCOT and MISO data. Of the parameters listed in Table 13, only the Fourier modes were partially based on a computational algorithm rather than purely selected by the user. In a process prior to ARMA, RAVEN's <FastFourierTransform> post-processor identified the Fourier modes of the PJM Brandon Shores data set and recorded them in a comma-separated values file (see ARMA_fft.xml in APPENDIX B). The resulting modes were sorted by their corresponding amplitude, and a subset of those were chosen subjectively for equaling common measures of time essential to the economy (year, half-year, week, or day)[c]. Modes of 8760, 4380, 24, and 12 hours were selected for inclusion in all parameter trials, while modes of 1095 and 168 hours were included in trials in an additive manner (none; 1095; or 1095 and 168). Because the solution did not converge for simultaneous non-zero powers of the autoregressive (AR) and moving average (MA) portions of the model, the only value pairs considered were (0, 1) and (1, 0), respectively.

Histograms, binned by price, were used to visually compare one year of synthetic data generated by the ARMA analysis to historical data for that same year. The tail ends of the distribution were primarily examined so that they would match the historical values, due to their economic significance. This visual process allowed for judging the appropriateness of the trialed values, along with the means, minimums, and maximums.

The final parameter values chosen, shown in the second column of Table 13, were input to the Models block of the RAVEN input file (an input file is shown in full in APPENDIX C). Because the power of the autoregressive portion of the model was ultimately set to zero, every model adopted is a moving average model rather than a full ARMA. Two synthetic history models were adopted for each market: one with a segment length of 24 hours, and one with a segment length of 120 hours.

In mapping historical years to synthetic years when generating the 30-year synthetic history objects, the LMP data set for the node or LZ was first split into four files, each containing a year (8760 hours)[d] of data. To build a synthetic history file, the ROM was configured to fit each synthetic year to a corresponding historical year in a repeating sequence: the ROM modeled the first 4 years (2018–2021) of synthetic data on their corresponding historical years, then modeled future years on those same historical years in repeating order until reaching 2047, as summarized in

Table 14. The result of running the ROM once was a 30-year synthetic history object. For each market, two segment lengths were used in running the ROM, producing two synthetic history objects (for a total of six objects, each saved as an arma.pk file) destined for use in the next stage of the technoeconomic analysis.

The first six years of realizations of the PJM Brandon Shores and ERCOT West synthetic histories are plotted in Figure 19 and Figure 20, respectively, each with a 24-hour segment length. Figure 21 plots the first four years of the MISO 24-hour-segment-length synthetic history, and Figure 22 plots the same for a 120-hour-segment-length synthetic history. Lastly, Figure 23 provides a detailed workflow chart.

[c] While mode 168 did not appear directly, nearby values did; thus, mode 168 was included to represent a week.

[d] Note that February 29 was deleted from each 2020 historical data set, shortening the leap year to the length of a common year.

Table 13. Parameters trialed and adopted for synthetic history generation by RAVEN via the ARMA ROM. Two parameter sets result under "Adopted values": one with a segment length of 24 hours, and one with a segment length of 120 hours.

Parameter <Input file property>	Values trialed	Adopted (final) values
Clusters (segments) <n_clusters>	—	20
Autoregressive power <P>	0, 1, 2 *while <Q> = 0*	0
Moving average power <Q>	0, 1, 2 *while <P> = 0*	1
Fourier modes <Fourier>	1095, 168	{8760,4380,2190,1095,168,24,12}
Segment length <subspace pivotLength="…>	24, 120, 146, 219	24, 120
Preservation of input continuous distribution function <preserveInputCDF>	False, True	True

Presence in this table does not imply that all combinations of parameters were trialed. Fourier modes were always trialed with the full set of all the following modes: 8760, 4380, 2190, 24, and 12.

Table 14. Historical years used for synthetic history generation, ordered by synthetic history year. This table reflects the contents of the ARMA_pointer.csv file.

Synthetic history year	Historical year
2018	2018
2019	2019
2020	2020
2021	2021
2022	2018
2023	2019
2024	2020
2025	2021
…	…
2047	2019

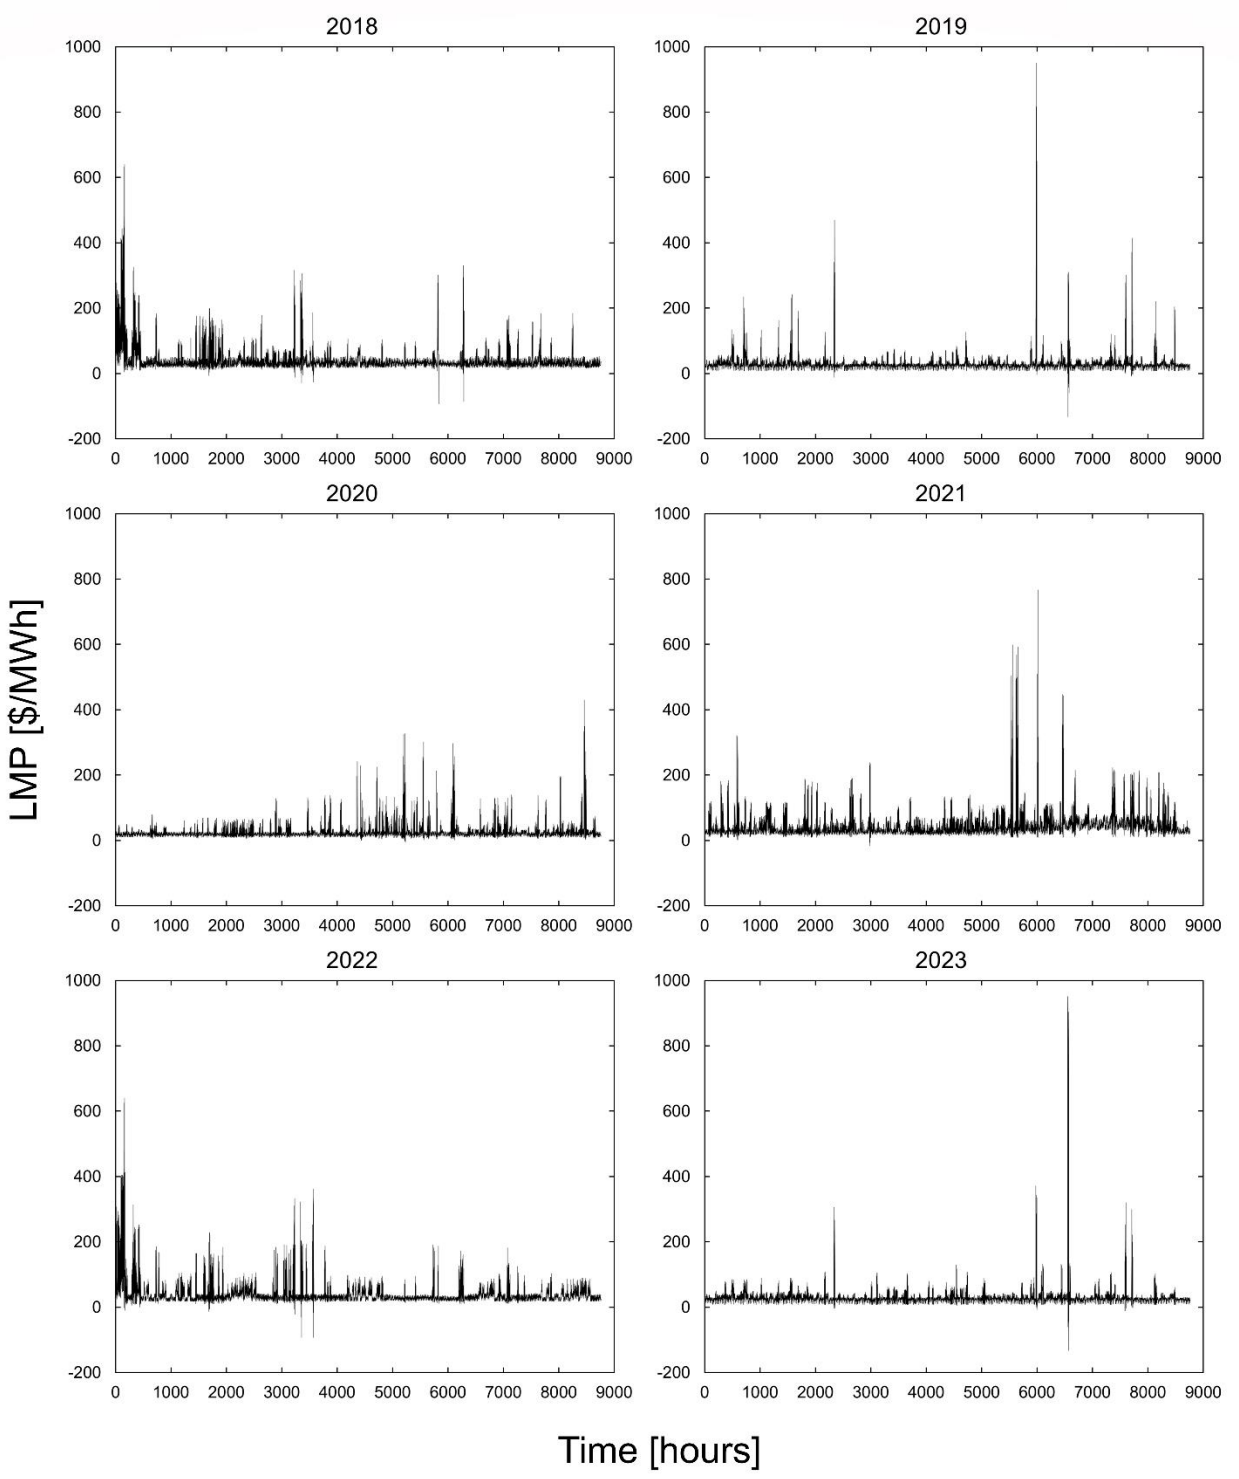

Figure 19. Representative synthetic LMPs for PJM's Brandon Shores node, generated with a segment length of 24 hours. Prices are representative because the values will vary each time the synthetic history object is sampled.

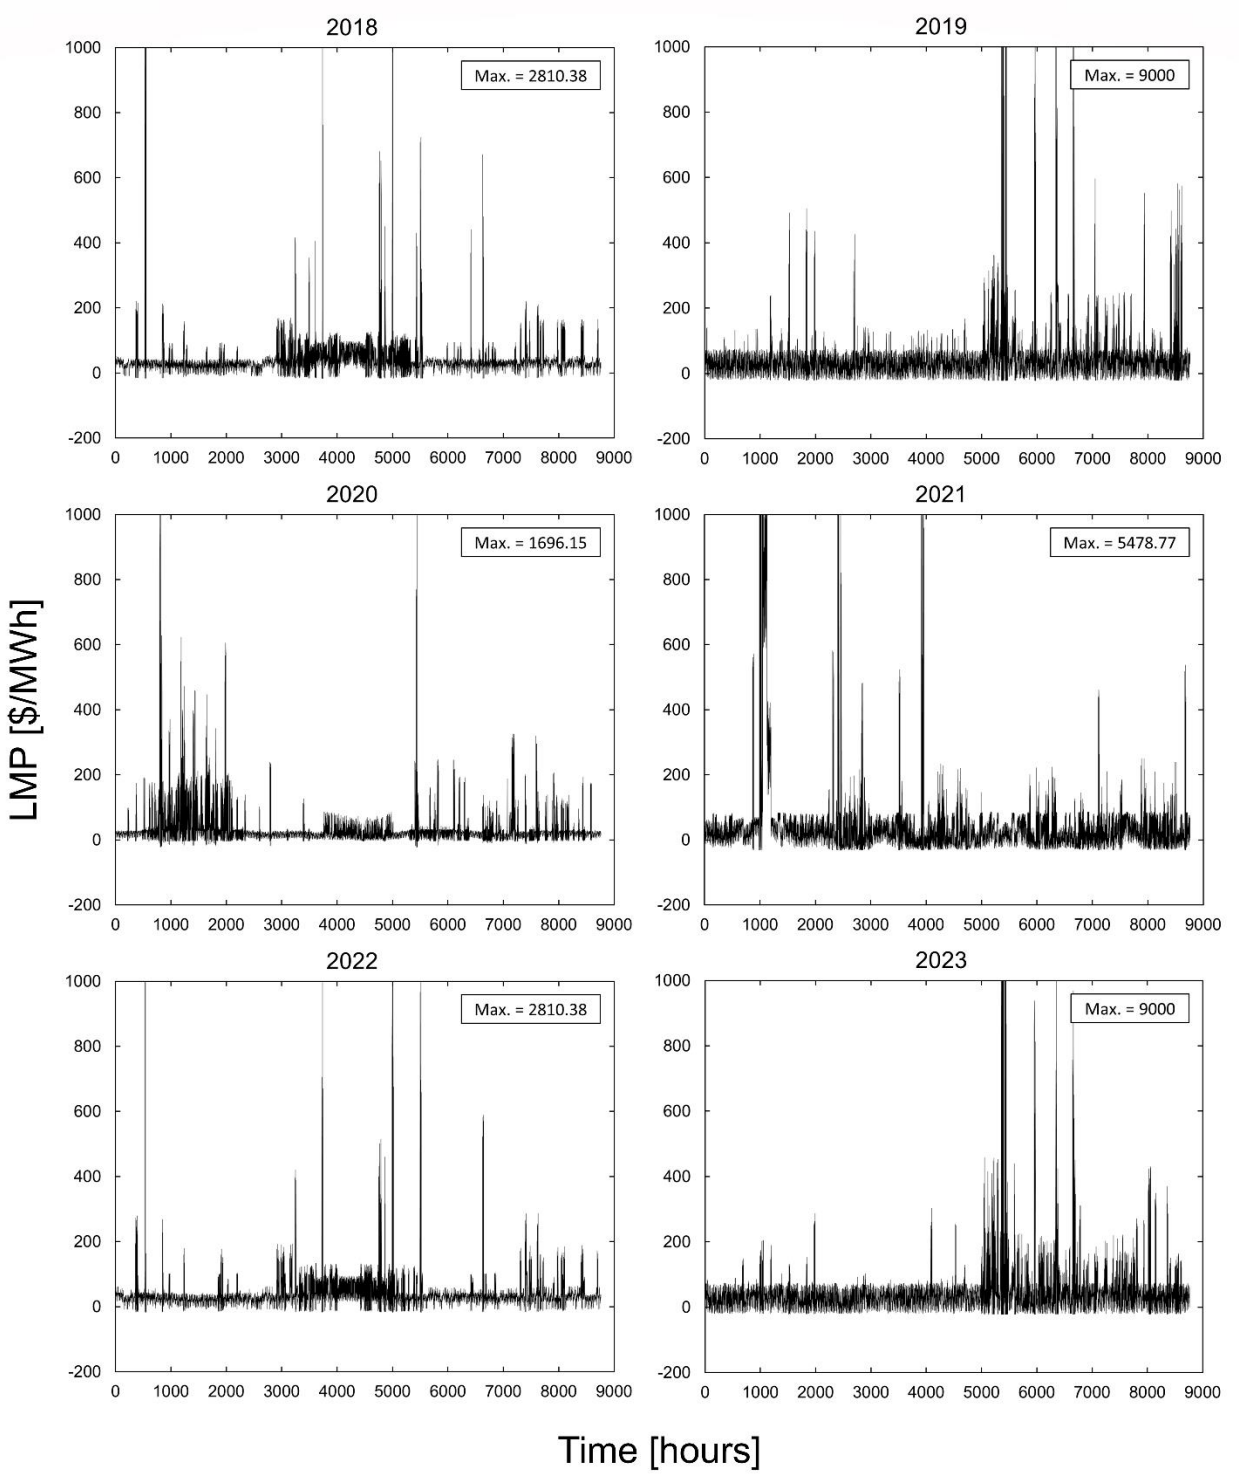

Figure 20. Representative synthetic LMPs for ERCOT's West LZ, generated with a segment length of 24 hours. The values will vary each time the synthetic history object is sampled.

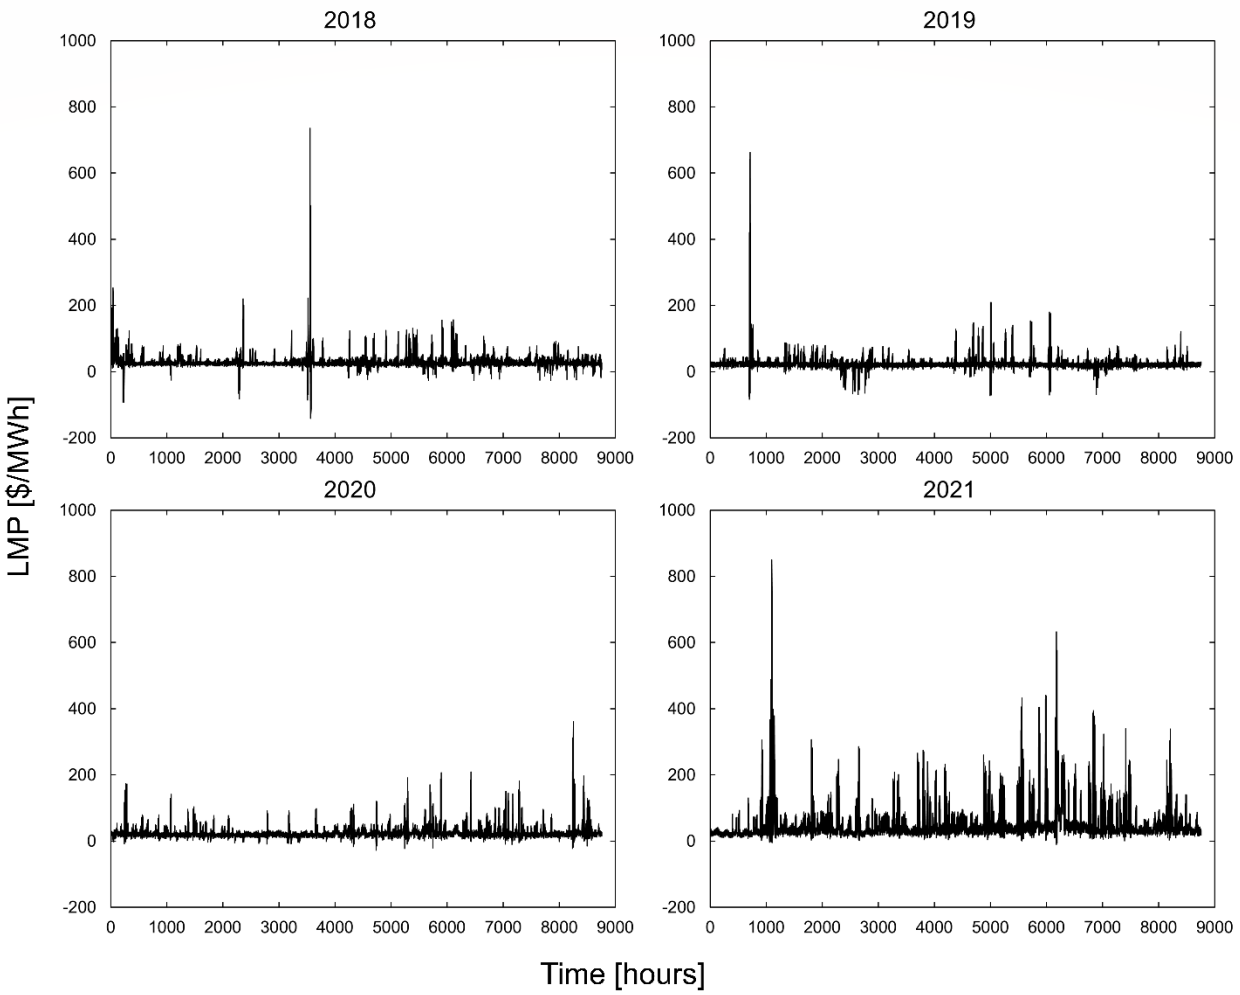

Figure 21. Representative synthetic LMPs for MISO's NSP.HATFIHAT1 node, generated with a segment length of 24 hours. The values will vary each time the synthetic history object is sampled.

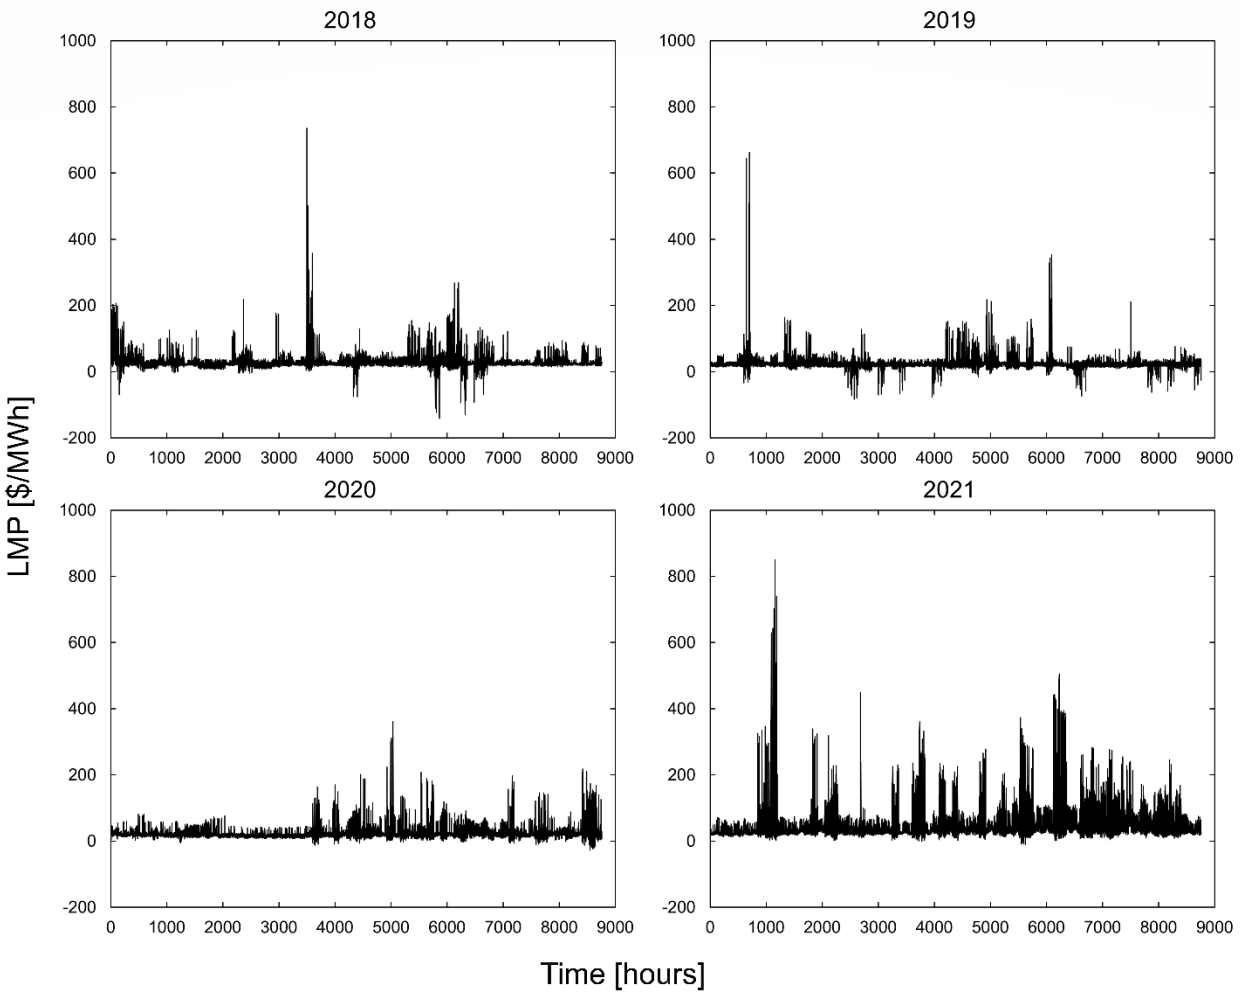

Figure 22. Representative synthetic LMPs for MISO's NSP.HATFIHAT1 node, generated with a segment length of ***120 hours***. The values will vary each time the synthetic history object is sampled.

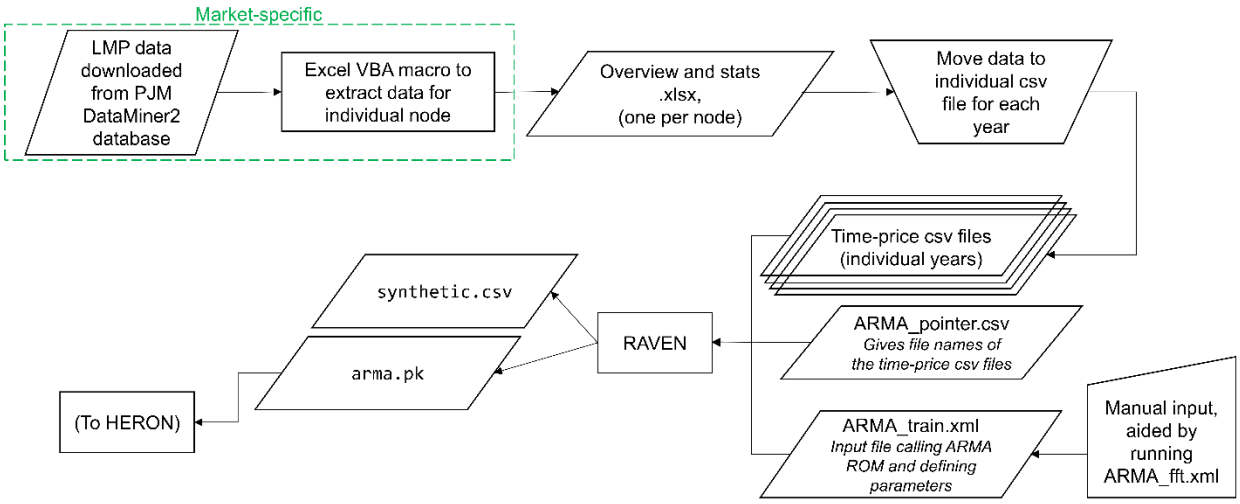

Figure 23. Detailed workflow for synthetic history generation.

5.2 Economic Dispatch and Optimization Methodology

The RAVEN and HERON frameworks were used to conduct an extensive stochastic analysis of TES within the ERCOT, PJM, and MISO markets. In this analysis, by utilizing trained electricity price signals as the primary economic drivers, HERON employed a set of developed cost functions (see Table 11) in order to optimize the costs associated with three TES-coupled HTGR supersets: (1) charge, (2) discharge (including BOP considerations), and (3) TES. By comparing the results, as characterized by net present value (NPV) and statistically robust optimal sizing, we aimed to capture the intricate interplay between electricity prices and the dynamics of these scenarios (including storage operation cycles). Ultimately, this study supports regionally tailored development of TES-coupled AR designs, facilitating a deeper understanding of their interactions within specific markets.

5.2.1 Overview of Methodology and Scenarios

To assess the potential advantages of integrating a TES configuration and the subsequent formulation of charge/discharge superset arrangements, we compared the most economically optimal (i.e., profit-maximized) cases arising from distinct superset mixes, with an emphasis on NPV. Our investigation was specifically directed toward the incremental NPV variations that result from changes in costs and technological factors, including (1) additional capital investments for upgrading existing AR plants into TES-coupled AR systems, (2) thermal-to-electric conversion efficiency, and (3) the minimum generation level for keeping the discharge superset in hot-standby mode. As such, the modeling outcomes presented herein focus on the NPV differential resulting from these variations, not on the absolute mean NPV.

Motivated by the findings introduced in the new BOP design, our study established two baseline scenarios for the NPP evaluation (see Figure 24): (1) The Standard configuration, which was derived from the Xe-100 design detailed in previous analyses and incorporates a steam-heated TES with a thermal-to-electric conversion efficiency of 42% [3]; and (2) a new BOP design suggested for the fully decoupled configuration, characterized by a thermal-to-electric conversion efficiency of 43%, coupled with an intermediate salt heat-transfer loop—an approach parallel to that employed by TerraPower [27]. In this context, our objective was twofold: (1) separate out the novel discharge design's impact on enhanced efficiency and (2) assess such an innovative design's economic competitiveness across diverse market scenarios.

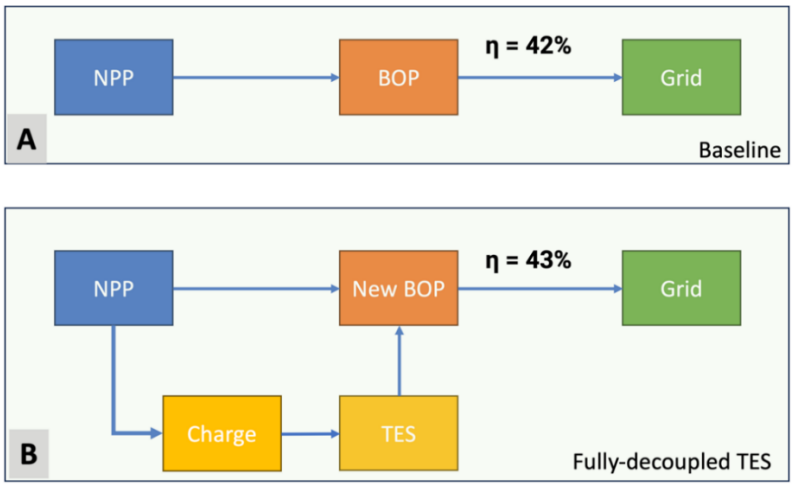

Figure 24. Schematic representation of the HERON model configurations: baseline scenario (top) and fully decoupled TES scenario (bottom).

The HERON optimization procedure combines two distinct stages: capacity optimization and dispatch optimization. In the capacity optimization phase, various configurations for component capacities are generated, with each configuration held constant throughout the subsequent dispatch optimization. The primary objective of the dispatch optimization is to identify an economically efficient portfolio that minimizes costs (or maximizes profits). For the cases examined within this study, the components subject to optimization include the charging, storage, and discharging supersets ("New BOP" in Figure 24B). These components were optimized based on derived cost functions (as outlined in Table 11), which serve as economic drivers. These cost functions, in conjunction with electricity price signals, were provided for the final NPV calculation. In the modeling framework employed, certain components such as the reactor design ("NPP" in Figure 24A and Figure 24B) and grid infrastructure remain fixed due to their established status (assuming they are given). Conversely, components integral to the coupling—namely, the charging, storage, and discharging supersets—possess a range of values throughout the optimization process. Table 15 gives an overview of the necessary inputs to the HERON simulation, along with its resulting outputs.

For computational tractability, the optimization procedure is principally divided into two key stages within the presented framework for enhanced efficiency: (1) the sweep mode, wherein a combination of optimal capacities for each superset is evaluated (these serve as the input to the subsequent detailed optimization in the opt mode); and (2) the opt mode, which employs a gradient descent algorithm to converge upon an optimal set of capacities. For example, in each computational simulation, the discharge superset spans 210 to 300 MW, incrementing by 10 MWth. Subsequently, the span extends from 300 to 460 MW, with 20 MW increments (see

Table 16). The resulting set of optimal capacities is then used to define bounds (i.e., <opt_bounds>) in the opt mode. This approach is strategically employed to minimize the initial condition dependency that may arise from utilizing the gradient descent optimization [28].

Table 15. HERON model inputs and outputs.

Key inputs	Key outputs
Optimization variable resolution; see • Table 16 for further details	• Optimized generation mix • Total system cost (NPV) • Hourly dispatch profile of each superset

** Perfect foresight of future hourly electricity prices during each dispatch horizon (i.e., segment length) is assumed.*

Table 16. Optimization variable resolution for the HERON analyses.

Variable	Range [increment]
NPP	203 MW_{th} (fixed)
Charge	178–203 MW_{th}
Discharge	190–460 MW_{th}
TES	0–1300 MWh_{th}

Table 17. Modeled scenarios for different markets, technological assumptions, and simulation settings.

Scenario	Market	Thermal-to-electric efficiency [%]	Minimum generation level [%]	Electricity price volatility	Segment length [hour]
Base	ERCOT	42	0	Reference	24, 120, 146, 219
	PJM	42	0	Reference	24, 120, 146, 219
	MISO	42	0	Reference	24, 120, 146, 219
New BOP without TES	ERCOT	43	0	Reference	24, 120, 146, 219
	PJM	43	0	Reference	24, 120, 146, 219
	MISO	43	0	Reference	24, 120, 146, 219
Fully decoupled system	ERCOT	43	0	Reference	24, 120, 146, 219
	PJM	43	0	Reference	24, 120, 146, 219
	MISO	43	0	Reference	24, 120, 146, 219
Minimum generation level	ERCOT	43	10, 15, 25, 40	Reference	120
Electricity price volatility	ERCOT	43	0	Low[1*]	120
	PJM	43	0	Low[1*]	120
	MISO	43	0	Low[1*]	120

[1] *Electricity price profiles with reduced volatility*

***NOTE:** *The capacities of the charge and discharge supersets remain constant at 203 MWth, aligning with the thermal output of the reactor.*

The modeling setting consists of 43 cases resulting from five possible scenarios: the standard BOP configuration, the new BOP design within the discharge superset (without TES), the fully decoupled TES system, varying the minimum generation level constraints, and different levels of input price signal volatility. In Table 17, these scenarios are denoted as Base, New BOP without TES, fully decoupled system, Minimum generation level, and Electricity price volatility. To highlight the temporal factors' various sensitivities to system economics and sizing, our analysis further accounts for four different segment lengths (or dispatch time horizons) when considering a subset of a full year. For the fully decoupled system, minimum generation level, and electricity price volatility scenarios, TES is modeled with a round-trip efficiency (RTE) of 90%, and incorporates a periodic condition, thus ensuring a fixed storage level (i.e., 75% TES capacity) at both the start and end of each segment length. Importantly, the model restricts simultaneous charging and discharging operations. In all scenarios, a project lifetime of 30 years is assumed for the capital recovery period across all superset investments, along with a discount rate of 8%. In our modeling approach, it is important to note that market intercomparison does not reflect policy considerations, market expectations, or transmission constraints. Instead, these factors are regarded a posteriori, as represented by the electricity price signals used in HERON runs.

5.2.2 Sensitivity of HERON Modeling to Temporal Factors

Given the intricate nature of capacity and dispatch optimization in the context of HERON simulations, previous research has often employed input time series (e.g., demand profiles and price data) trained on a 24-hour segment length. Typically, these inputs were subsequently grouped into 3, 20, or 30 clusters to represent a full-year modeling period [29],[8],[30]. Empirical studies have demonstrated that breaking down segment lengths, such as grouping 8760-hour segments of original time series into

representative days or weeks, does not lead to substantial differences in total system costs. This observation holds true when a periodic condition is applied, maintaining a consistent storage level at the beginning and end of each segment length [31],[32]. However, restricting our HERON analysis to 24-hour segment lengths may introduce the potential for noteworthy disparities in TES operations. [29] This consideration becomes especially significant in markets where high electricity prices (or energy arbitrage opportunities) are infrequent throughout the week, necessitating strategic TES scheduling.

We also observed a noticeable increase in the standard deviation (or electricity price volatility) of our trained price signals with extended segment lengths and when the cluster number constant is fixed at 20 during ARMA model training. As depicted in Figure 25, which illustrates the mean electricity prices and their corresponding standard deviations for each segment length and market, the increasing trend is evident when comparing against the input—namely, the original electricity price data for ARMA models. This outcome can be attributed to an insufficient number of clusters when training an ARMA model with an extended segment length.

Considering the proposed TES system's preference for exploiting periods of frequent oscillations in high electricity prices so as to efficiently flex stored thermal energy into the grid, the deployment potential of the TES-coupled system may be higher in markets characterized by elevated electricity price volatility. This notion prompted us to examine HERON's sensitivity to price volatility concerning NPV, optimal superset sizing, and TES operational cycles. Per this objective, we derived test ARMA models with reduced volatility, based on utilizing a higher number of clusters (i.e., 73) for the sensitivity analysis (see Table 18).

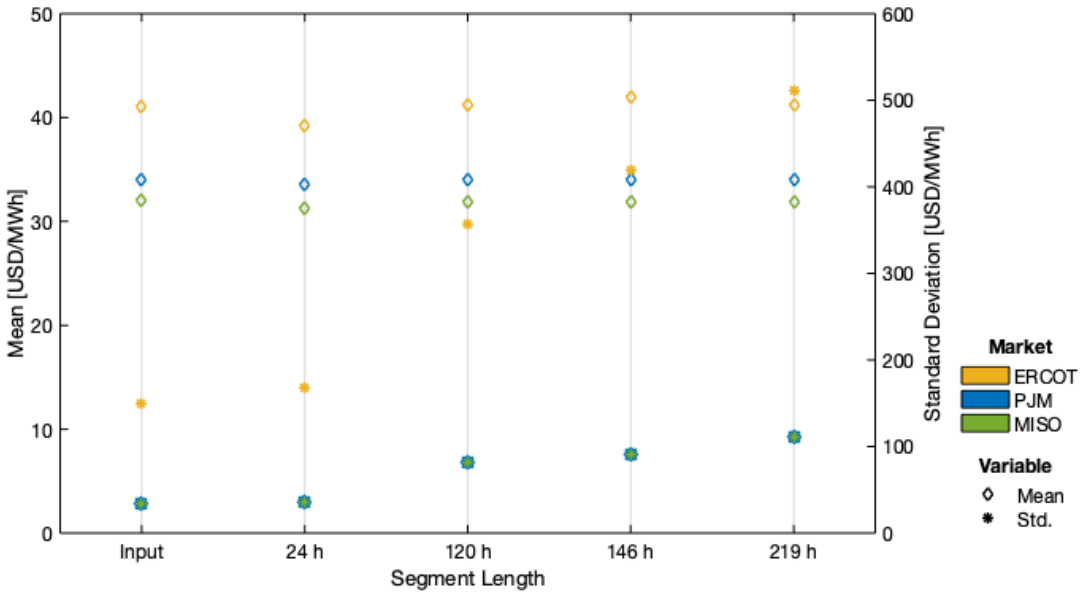

Figure 25. Real-time market price statistics across the chosen markets and segment lengths.

Table 18. Comparing the mean and standard deviation for the reference ARMA model and low-volatile ARMA model.

Market	Electricity price volatility	Segment length [hour]	Cluster number	Mean [USD/MWh]	Standard deviation [USD/MWh]
ERCOT	Reference	120	20	41.28	357.20
	Low	120	73	41.84	198.69
PJM	Reference	120	20	34.01	81.76
	Low	120	73	34.08	42.88
MISO	Reference	120	20	31.91	81.72
	Low	120	73	32.02	42.83

5.2.3 Sensitivity of HERON Modeling to Technical Parameters

The operational modes of TES-coupled systems are expected to exhibit more irregular behavior during energy arbitrage operations. This irregularity may result in frequent startups, shutdowns, and prolonged periods without dispatch. Consequently, maintaining an equilibrium temperature for each superset becomes critical to prevent cooling down. This ensures that thermal power is rapidly dispatched from the TES in response to electricity price signals. Our previous analysis suggested that diverting approximately 1% of the primary cycle's nominal mass flow will effectively meet the heat requirements of the TES systems during standby periods [3]. To characterize an appropriate operational strategy for the TES system, we incorporated five discrete minimum power generation levels, ranging from 0% to 40%, from the TES BOP in HERON simulations (see Table 17). A comprehensive set of charge, discharge, and electricity production modes was analyzed to understand the interplay between operational margins and system economics.

5.3 HERON Results

5.3.1 General Comparison

Figure 26 and Table 19 summarize trends and key parameters across three distinct markets, based on modeled scenarios. Each row is contrasted with the corresponding Base scenario (featuring a 42% thermal-to-electric conversion ratio) in order to highlight the potential economic advantages derived from adopting a fully decoupled system. Our analysis centered on reconfiguring the existing reactor plant design into fully decoupled systems. Thus, the figure presents the extra capital investments dedicated to this transition, clarifying the absence of corresponding costs in the minus NPV column for the Base scenario (located on the right-hand side of Figure 26).

Error bars are included to account for errors stemming from varying the segment lengths. The obtained NPVs within the PJM and MISO markets remain consistent across different segment lengths; however, the ERCOT market exhibits a 0.5–4.3% discrepancy in terms of error variance. Table 19 includes the Delta NPV column, which captures changes in NPV relative to the Base scenario. This shift in NPV reflects the additional economic advantages stemming from electric sales.

In our findings, the fully decoupled system demonstrated the most economic promise in the ERCOT market, followed by the PJM market and then the MISO market. This aligns well with our initial expectations, as the relative electricity prices integrated into the model mirror the historical price patterns across ERCOT, PJM, and MISO. Notably, ERCOT has exhibited historically higher electricity prices than PJM and MISO. This explains why New BOP, designed to be optimized according to its cost function,

finds its application exclusively within the ERCOT market by offsetting additional capital investments. In contrast, the enhanced efficiency does not produce improved NPVs in PJM and MISO, primarily due to the relatively modest electricity price levels (see Figure 25 and

Table 20). This implies that the fully decoupled system may be preferable to the standard BOP configuration, particularly in markets characterized by high electricity prices. Note that the variations in NPV changes within the same market are mainly attributed to the unique ARMA sampling. This is because, although the price signals used in both the New BOP and Base scenarios were sourced from the same ARMA file (for each segment length), they are not entirely identical. If the ARMA sampling process is configured to generate identical electricity signals for both scenarios, any differences in NPV would then reflect extra capital investments, such as the higher costs associated with the enhanced BOP in the New BOP scenario.

Interestingly, the average optimal storage capacities for each market closely correlate with the electricity price volatility. Specifically, the ERCOT market demonstrates frequent oscillations between high and low electricity prices, prompting the necessity for larger storage capacities to maximize revenue generated during times of high electricity prices. The optimal storage capacity for the PJM market slightly exceeds that of MISO, due to its marginally larger electricity volatility (see Figure 25), though this distinction may be indiscernible in Figure 26. Throughout all the scenarios explored, the sizing of the charge and discharge supersets generally converges around the reactor's thermal output of 203 MWth, producing a consistent discharge-to-charge ratio of 2.03, on average. While the optimal storage capacity fluctuates in accordance with the respective market prices, its contribution to the mean NPV becomes less evident when viewed from the perspective of the capital recovery period, which spans 30 years.

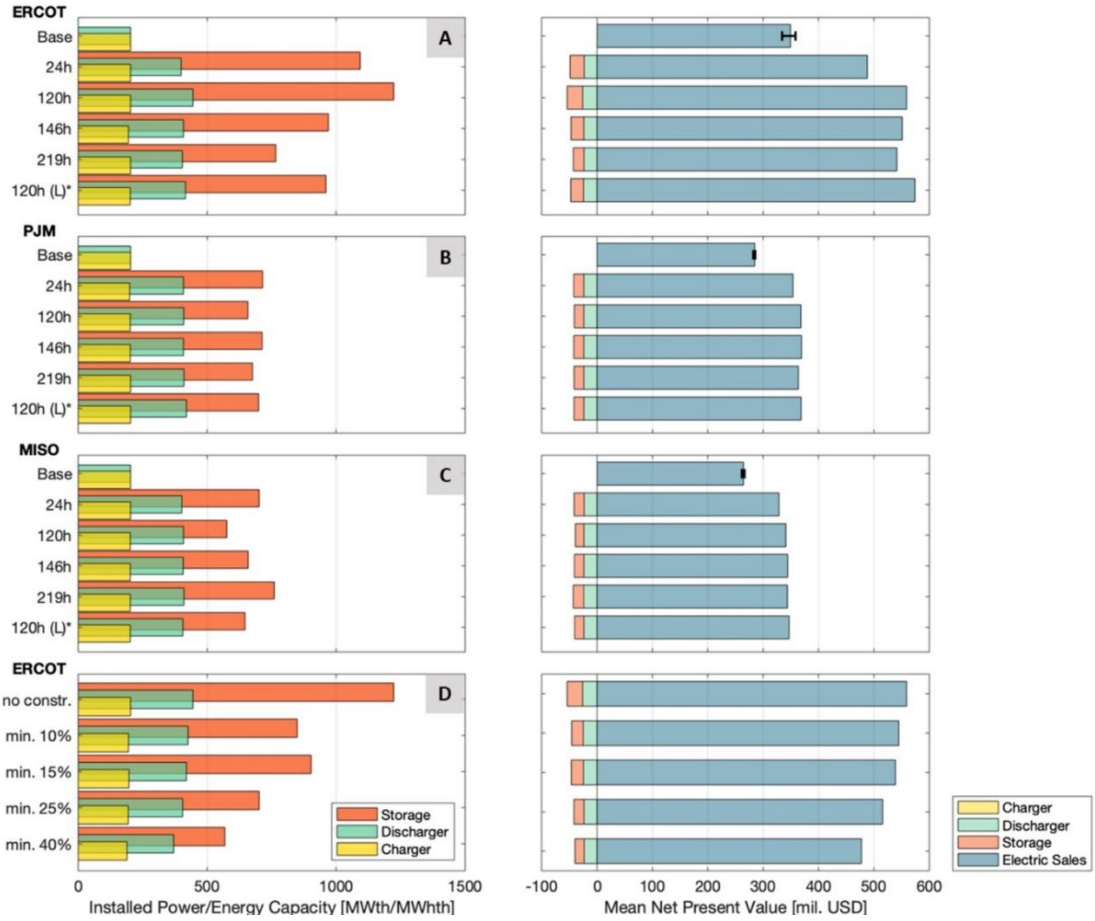

Figure 26. HERON-optimized cases for the HTGR-TES coupling superset capacity (left) and NPV (right).

Table 19. HERON-optimized parameters for the fully decoupled system and electricity price volatility scenarios.

Market	Segment Length [hour]	Delta NPV[1] [mil. USD]/[% difference]		Delta NPV/ MWh_{th} [mil. USD]	Discharge-to-Charge Ratio	Charge Duration [hour]	Discharge Duration [hour]
ERCOT	24	105.11	31.4	0.10	1.96	5.38	5.59
	120	147.04	41.1	0.12	2.19	6.03	5.06
	120 (L)*	168.67	47.1	0.18	2.06	4.75	4.49
	146	145.54	40.6	0.15	2.09	4.97	4.56
	219	151.97	43.9	0.20	1.99	3.78	3.81
PJM	24	30.27	10.8	0.04	2.05	3.58	3.42
	120	41.58	14.6	0.06	2.02	3.26	3.18
	120 (L)*	42.30	14.8	0.06	2.01	3.42	3.40
	146	41.36	14.5	0.06	2.02	3.52	3.47
	219	36.58	12.8	0.05	2.02	3.33	3.26
MISO	24	26.04	10.0	0.04	1.98	3.46	3.53
	120	36.70	13.8	0.06	2.02	2.85	2.80
	120 (L)*	41.06	15.5	0.06	2.01	3.20	3.18
	146	37.74	14.2	0.06	2.02	3.26	3.21
	219	34.64	13.0	0.05	2.04	3.77	3.64

[1] Delta NPV: $\Delta NPV = NPV_{case} - NPV_{Base}$

NOTE: *The discharge duration is calculated by dividing the TES capacity by the maximum discharge capacity. Due to heat losses from the TES system, which result in 90% RTE, the heat transferred from the reactor is prioritized for dispatch during periods of high electricity demand. Thus, in our shared BOP configuration, in which the TES does not have its own dedicated BOP system, the maximum discharge capacity of TES (thermal output) cannot exceed the difference between the BOP thermal input capacity and the reactor thermal output.*

Table 20. Impact of enhanced BOP efficiency on the NPV, relative to the Base scenario.

Market	Segment Length [hour]	NPV Change [%][1]
ERCOT	24	0.35
	120	2.23
	146	0.90
	219	3.01
PJM	24	-8.18
	120	-0.64
	146	-0.61
	219	-0.64
MISO	24	-0.62
	120	-0.26
	146	-0.32
	219	-0.30

[1] *The NPV change values are shown as the differences in NPV between cases with a 42% thermal-to-electric efficiency (Base scenario) and cases with 43% efficiency (New BOP [without TES])*

5.3.2 Impact of Temporal Factors

In the framework of our modeling, wherein future electricity price data for the associated segment length are accurately forecasted, the optimal operation of TES is directly linked to the value of the stored thermal energy. This value is, in turn, shaped by two primary factors: segment length (reflecting the quantity of future data influencing current operational decisions) and price volatility level. Therefore, our emphasis lies in understanding the two-fold impact of electricity prices on TES operations, by conducting a detailed investigation of the price points as charging/discharging events occur. This analysis spans varying segment lengths ranging from 24 to 219 hours, as well as different levels of volatility (see Figure 27). It is important to note that we confined the x-axis of ERCOT within the range of -200–1000 to effectively capture the evolving trends associated with the price points of charge and discharge. This adjustment is supported by the fact that 99.9% of charge events and 99.3% of discharge events in our test scenarios for ERCOT occurred below the price point of $1000/MWh.

We first explored the impacts of augmented volatility as the segment length increases. As discussed in Section 5.2.2, to align this increased volatility with the original input data, the number of clusters grouped into a subset representing a full year's data can be adjusted, as shown in Table 18. Figure 27 illustrates that higher volatility results in more sporadic charge and discharge patterns across the price spectrum, particularly within the $0–200/MWh range across all markets. In cases where the standard deviation is decreased (indicated as [L]* in the third row of Figure 27), this variability is mitigated, yet the depth of charge and discharge still remains distinct when compared to cases based on a 24-hour segment length. This indicates that the effect of price volatility can be examined independently of the effect of segment length (or information quantity) during optimization. Counterintuitively, increased price volatility does not necessarily yield more energy arbitrage opportunities. Rather, reduced volatility led to a 4.5% increase in discharge events in ERCOT, whereas the impact was not statistically significant for PJM and MISO, as evidenced in Table 21. When these findings were translated into electric sales revenue, only marginal increments of 2.7%, 0.1%, and 1.7% were observed for ERCOT, PJM, and MISO, respectively.

The regional differences in the effect of volatility on discharge frequency and revenue can be explained by price patterns. In ERCOT, a dominant daily peak occurs in the evening hours during the summer, characterized by increased volatility. This leads to a more concentrated discharge pattern, with higher amplitude. However, when the volatility is reduced, the threshold prices that trigger discharge decrease, resulting in a higher frequency of discharge events (see Table 21). Additionally, decreased volatility reduces the instances of negative prices. Within specific time frames, the system must still discharge to meet a periodic condition or respond to prolonged negative electricity pricing for certain clusters. In contrast, the daily peaks of MISO and PJM are more evenly distributed throughout the year, with two peaks occurring in the morning and evening. Consequently, in these regions, higher volatility does not result in an increased number of discharge events. Furthermore, even though increased price volatility leads to more high-value selling opportunities, this added advantage is offset by the rise in costs incurred during periods of negative prices, as depicted in Figure 27. Instead, volatility's impact is clearly seen to be more pronounced on the optimal sizing of TES, and this impact varies across different markets—the optimal TES size decreased by 21% in ERCOT but increased by 5% and 12% in PJM and MISO, respectively, when reduced price signals were used, as shown in Figure 26.

Our analysis also reveals that as the segment length expands, occurrence of active discharge (-50% to -100%) within the price range of $0–400/MWh becomes prevalent, generating added electricity sales revenue thanks to the longer look-ahead price information. The revenue rises rapidly as the segment length increases from 24 to 120 hours, but then gradually approaches a plateau at a segment length of 146 hours. However, our observations indicate that the advantages afforded by longer look-ahead price information are fully realized when combined with high electricity prices featuring "favorable" volatility, as observed in ERCOT. For cases in which electricity prices remain relatively stable throughout the year and do not exhibit significant peaks, the impact of longer look-ahead information remains limited (see

Figure 26). It is worth noting that, as previously stated, the decline in electric sales revenue in ERCOT after extending the 120-hour segment length to 146 and 219 hours is primarily attributed to the duration of negative price periods being extended due to the increased price volatility.

Table 21. Sensitivity of optimal storage operation to price volatility.

		Number of events per year	
Market	Volatility	Charge	Discharge
ERCOT	Reference	1826	1849
	Low	1918	1934
PJM	Reference	1966	2016
	Low	1976	2009
MISO	Reference	1958	2037
	Low	1984	2039

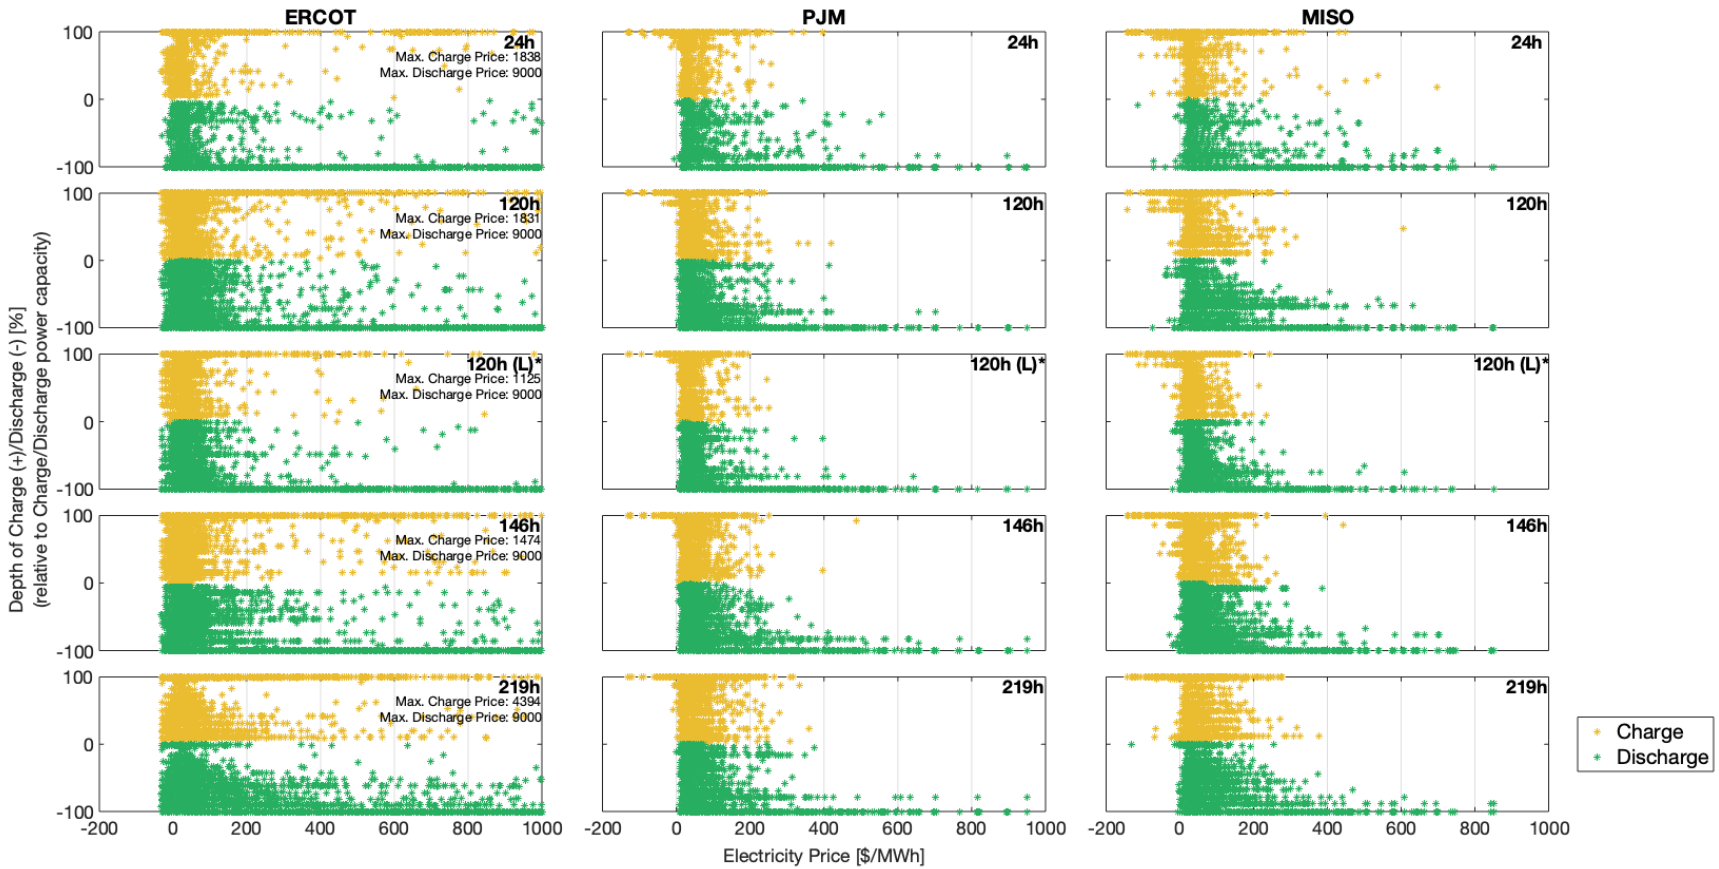

Figure 27. Scatter plots depicting the relation between electricity prices and the net changes in TES level: the realized revenue from electricity sales for each period is determined by multiplying the discharge amplitude (MWh) by the corresponding price point ($/MWh). Note that the plots correspond to scenarios involving a fully decoupled system and electricity price volatility (cases denoted as [L]*).

5.3.3 Impact of Technical Parameters (Feedback from Dymola)

The optimal duration of a single charge or discharge event depends on the perceived value (based on available price information) of stored thermal energy and their respective maximum power capacities. To gain deeper insights into the factors influencing an optimal TES duration, the distribution of charge and discharge operation modes was investigated in light of three scenarios: a fully decoupled system, electricity price volatility, and minimum generation level (see Figure 28). The results show that the majority of TES modes correspond to an optimal duration of 1 hour, particularly in cases with a 24-hour segment length. In general, discharge events tend to have shorter durations than charge events. This trend is particularly pronounced in the case of PJM and MISO, for which moderate peaks are evenly distributed throughout the modeled period of 30 years. This is because TES optimization responds to price signals, which tend to be more volatile than the aggregated demand used in macro-energy optimization. As a result, the "carryover value" from TES operation predominantly applies to shorter-duration storage applications (1 or 2 hours), even when the optimization horizon extends beyond a single week (e.g., 219 hours). While one might anticipate a higher frequency of mid-to-long-duration modes over 219 hours, it remains unlikely that TES operators would base operational decisions on price information extending beyond this timeframe.

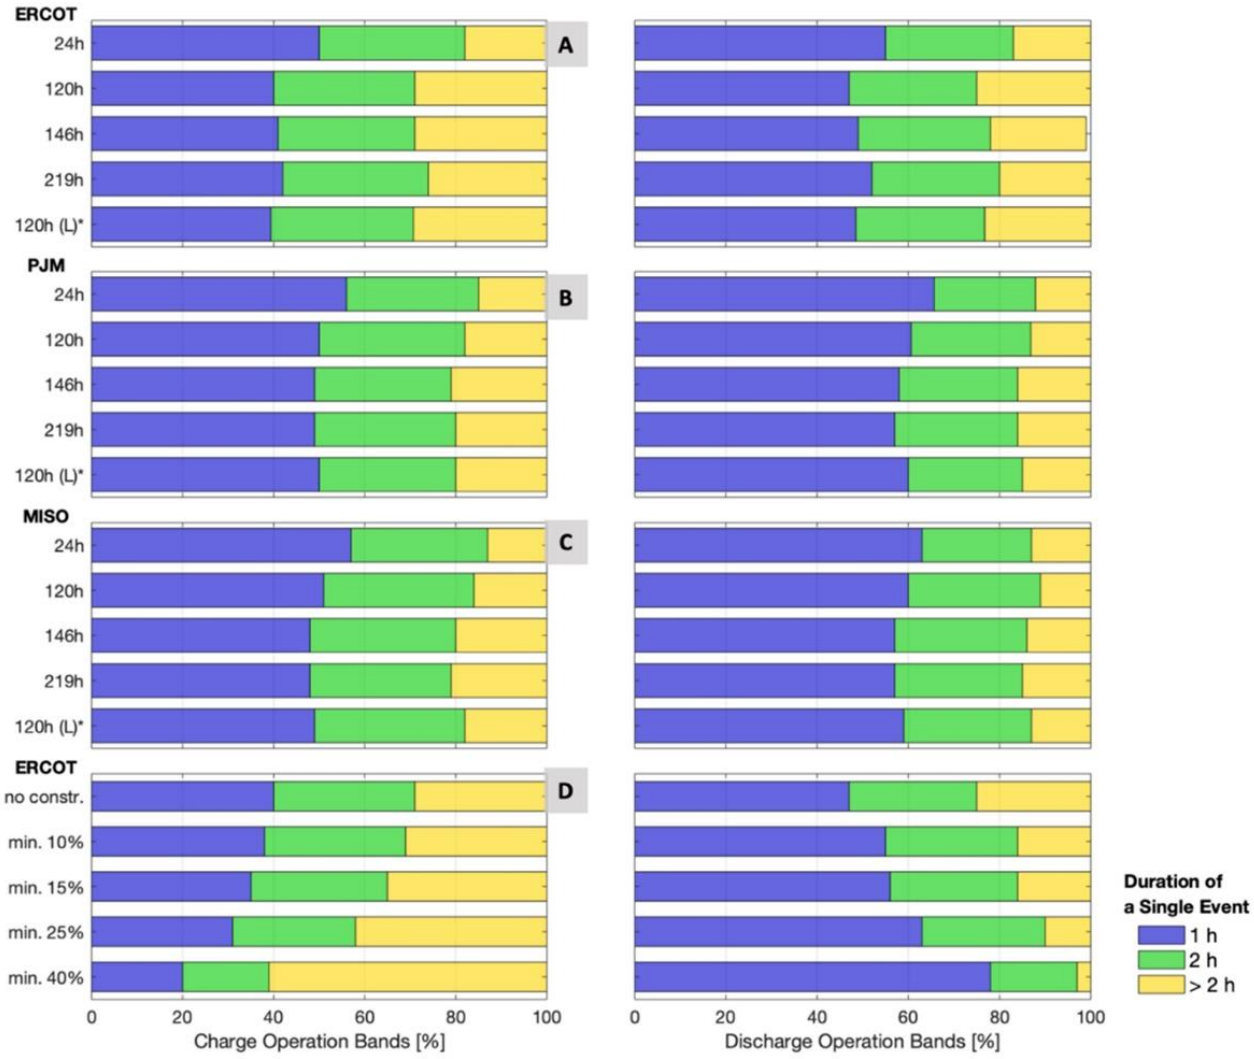

Figure 28. Contribution of different charge/discharge operation modes to the TES state-of-charge profile.

Figure 28 also illustrates how the charge and discharge modes respond to variation in the minimum generation levels. Recall that this scenario was explored specifically for the ERCOT market, under a segment length of 120 hours. The introduced constraint of minimum generation level demands continuous diversion of thermal energy into the BOP, causing the TES to take longer to build up to the required level before initiating discharge. This decreases the number of mid-duration discharge operations that occur, as these arbitrage opportunities are covered by the electricity generated from the BOP, due to the imposed constraint. Enforcing a minimum generation level directly affects the sizing of the TES, consequently reducing the requirement for BOP sizing (or discharge superset capacity in Figure 26).

Table 22. Impacts of a minimum generation level constraint on optimal sizing in ERCOT.

Minimum Load [% of BOP capacity]	Capacity of Charge, Storage, Discharge (MW_{th}, MWh_{th}, MW_{th})			Mean NPV [mil. USD]	Delta NPV1 [mil. USD]/[% difference]		Delta NPV1/ MWh_{th} [mil. USD]
0	203	1223	445	504.95	147.04	41.1	0.12
10	195	849	426	498.97	141.07	39.4	0.17
15	197	903	420	492.23	134.33	37.5	0.15
25	194	702	405	474.00	116.09	32.4	0.17
40	189	569	370	437.91	80.00	22.4	0.14

Due to the saved capital costs for the TES and BOP at up to a 15% minimum generation level, there is only a marginal decrease of 1% and 3%, respectively, in the total NPV. However, beyond the 15% level, the decrease in NPV becomes more significant, reaching 6%, and 13% for the respective 25% and 40% cases, in comparison to the unconstrained case (see Figure 26D and Table 22). We note the potential for regional variations in TES and discharge sizing, as well as electricity sales revenue, when different minimum generation level constraints are imposed. Therefore, additional sensitivity studies for the PJM and MISO markets should be carried out to thoroughly characterize the effects of such constraints.

6. TRANSIENT STATE AND OFF-DESIGN DYNAMIC MODELICA MODELS

Transient models were developed to showcase the system's ability to follow the desired dispatch profiles, thus showcasing its dynamic behavior. These models also provide valuable insights into the system's off-design performance predictions, as well as into the operational limitations imposed by certain physical phenomena. Dynamic modeling enables predications that offer a comprehensive evaluation of the system's behavior under various real-world conditions.

The primary aspect of the performed work was to develop an integrated model of the newly proposed, fully decoupled configuration. As the configuration and system operation significantly differ from the previously developed Modelica models, some of them required major modifications or had to be developed as new models.

6.1 Subsystem Models

The Modelica BOP model was built in accordance with the process diagram used for designing the thermodynamic model (see Figure 29). The model was built from preexisting or newly developed components in the Transform and NHES libraries, which are contained in the open-source HYBRID repository. The cycle was configured to involve a preheater and an evaporator that features both, a drum and forced circulation. Such an approach enables the use of single-phase (i.e., boiling only) heat exchangers, significantly improving the likelihood and speed of model convergence during the simulation. As a result of its configuration specifics, the BOP model contains heat input heat exchangers—something not the case for previously developed BOP models. An open feedwater heater,

acting also as a deaerator, was implemented as well. As the model considers an actual vessel with a liquid level, it more closely reflects actual systems than to previous BOP configurations, which only considered a mixing node. Another feature is partial admission control for the HPT and low-pressure turbine (LPT) so as to maintain most of the cycle parameters when under a partial load. The partial admission affords controlled extraction for the feedwater heating steam. If uncontrolled extraction was considered, the feedwater temperature would drop during a partial load. (The effect of partial admission is more readily apparent when compared side-by-side with the performance achieved without it.)

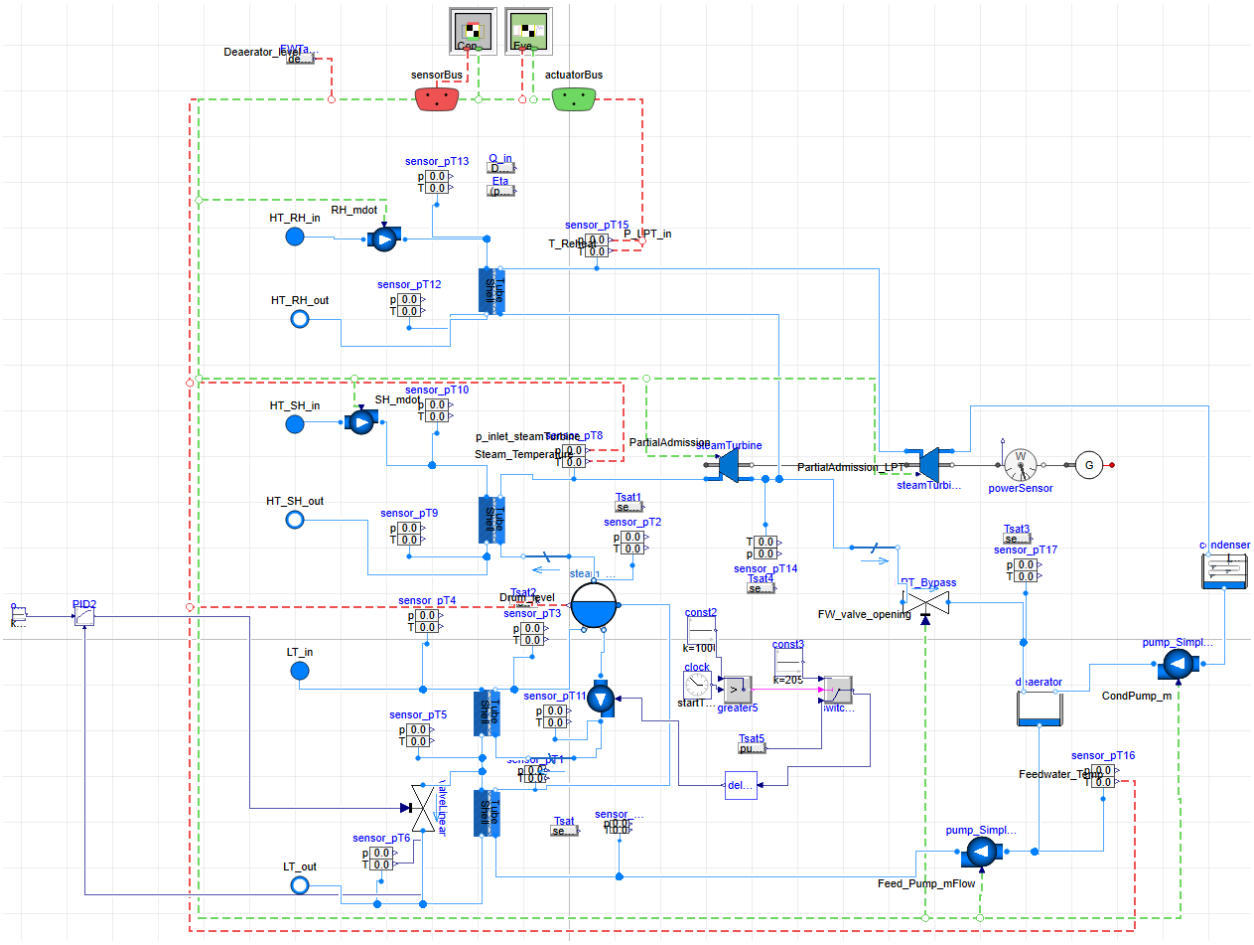

Figure 29. Modelica model for the reheat-cycle BOP.

Figure 30 shows the BOP control system, which runs the feedwater pump in order to maintain a constant liquid level in the drum (matching the flowrate to the incoming heat from the LT-TES molten salt). The evaporator circulation pump is controlled to maintain a reasonable circulation ratio. Partial admission is controlled so as to maintain the nominal turbine inlet pressure. Superheated steam is controlled by pumps that manage the flowrate from the HT-TES, with separate pumps for the superheater and reheater. The extraction valve leading into the deaerator is controlled based on its outlet temperature. The condensate pump is controlled to maintain a constant liquid level in the deaerator. An additional controller is implemented on the preheater bypass so as to limit fluctuation of the LT-TES return temperature, in turn limiting oscillation of the reactor inlet temperature.

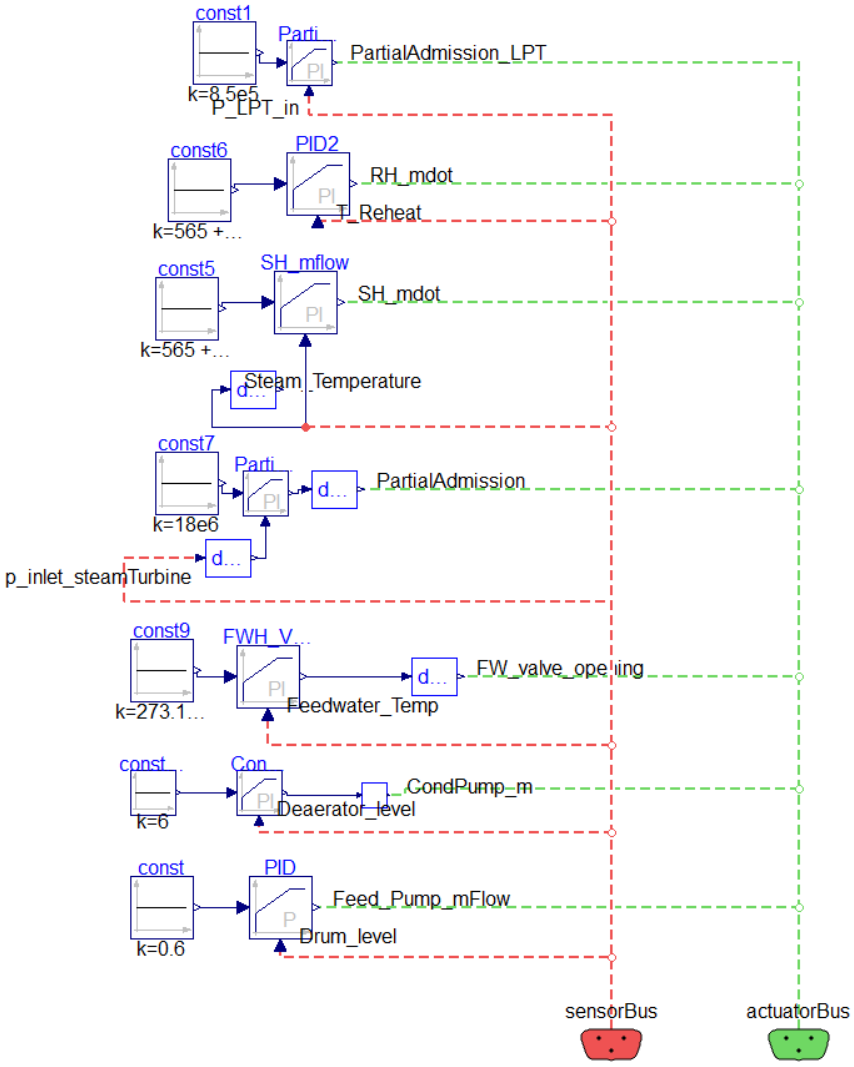

Figure 30. Modelica model for the reheat-cycle BOP control system.

The storage tanks and reactor model were modified to operate without the integrated heat exchangers, allowing for helium to be directly used at the reactor outlet and for salt to be routed to the BOP, as shown in Figure 31.

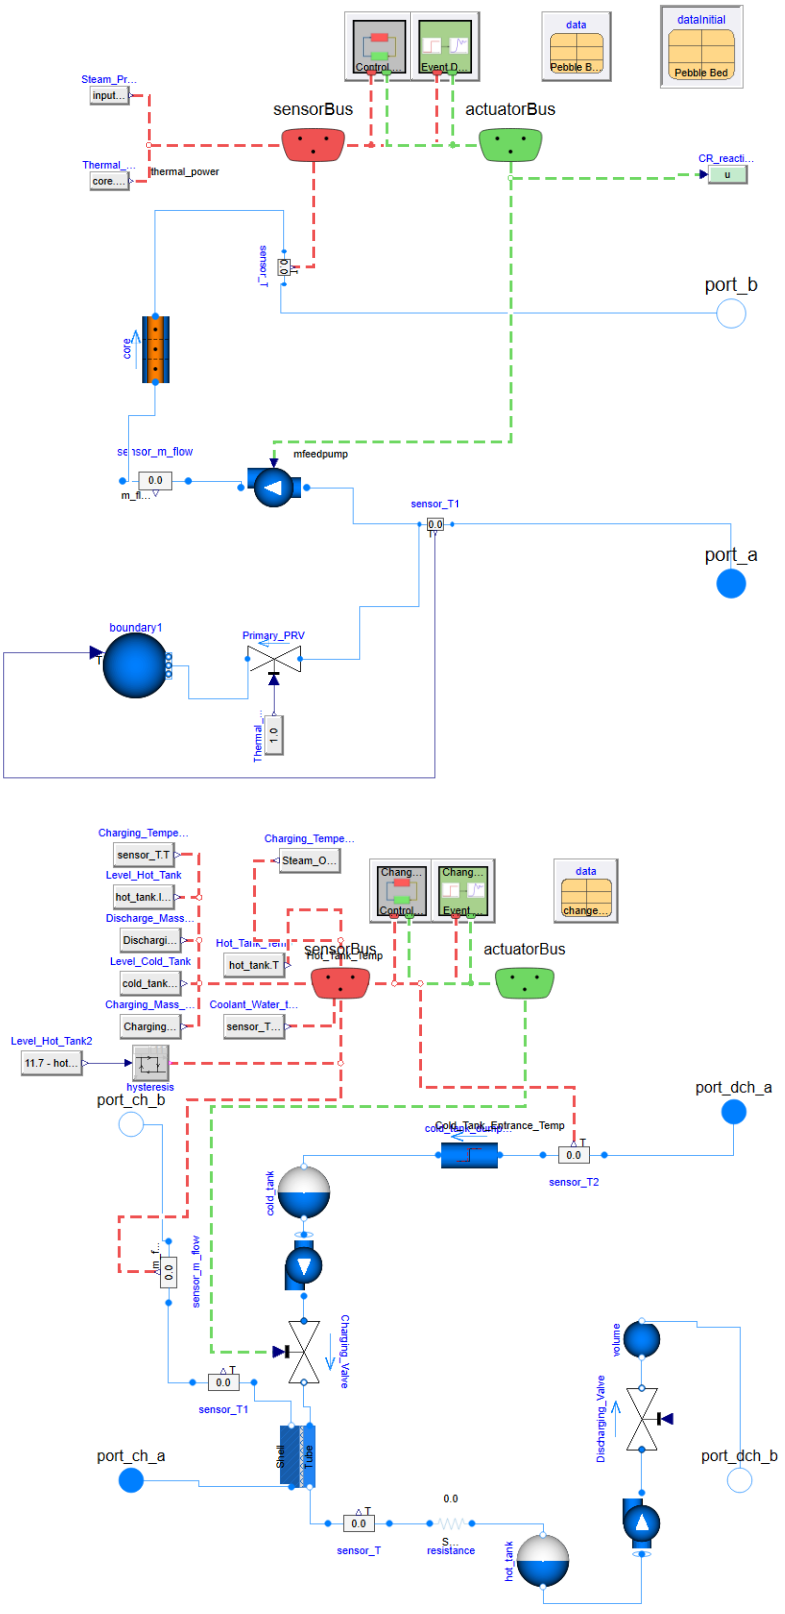

Figure 31. Modified Modelica models (created based on the legacy models) for the HTGR and storage tanks.

6.2 Overall System and Control Strategy

Figure 32 shows the full model representing the use case that consists of the previously described subsystems and upper-level control systems. As for the control strategy, the reactor is controlled for nominal (constant) thermal power, and charging of the TES is identical to the original models in order to satisfy the prescribed hot tank temperatures. The LT-TES discharge rate is then controlled according to the desired power output; the BOP controllers follow based on this thermal input, and the electrical power is eventually adjusted. The HT-TES discharge rate also follows the system, meeting the desired temperatures for the superheated steam.

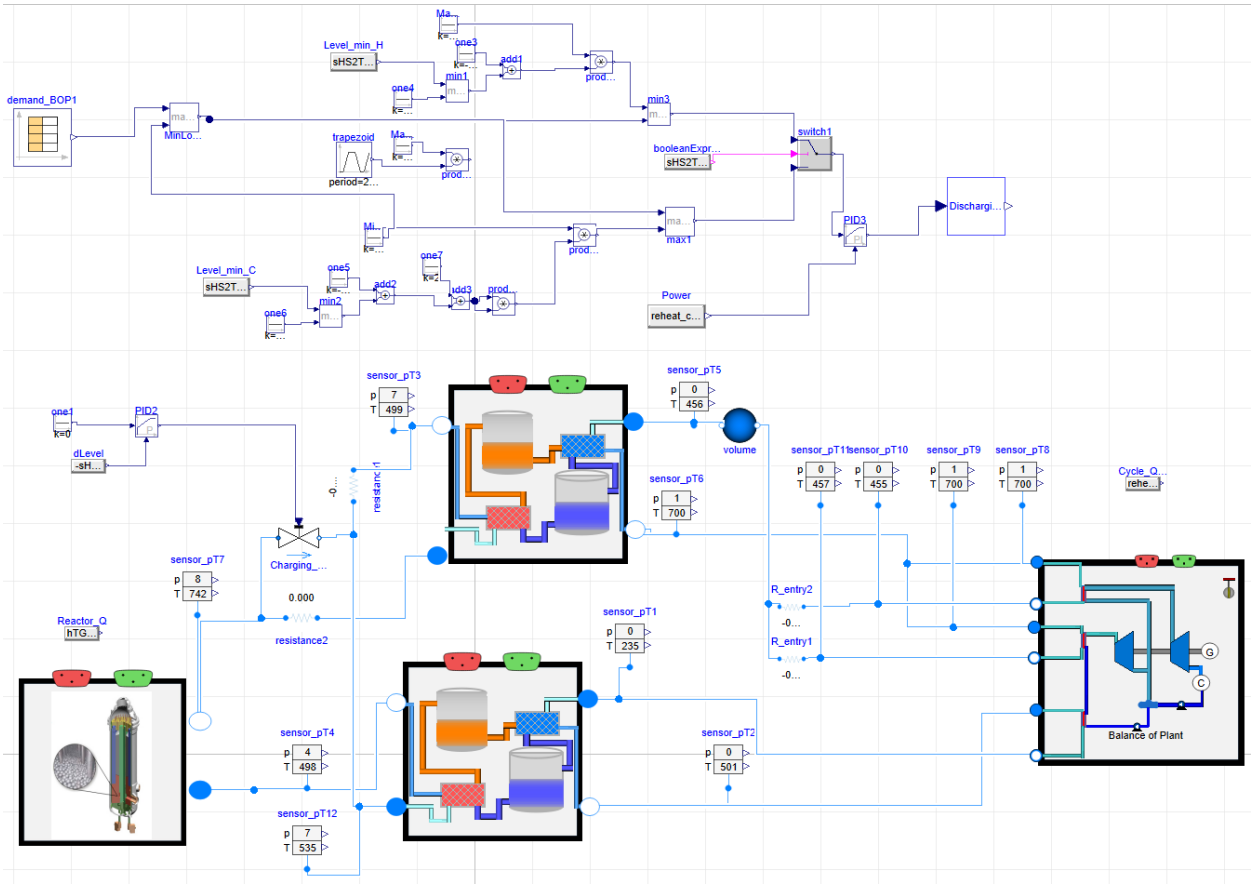

Figure 32. Overall Modelica model for the single-phase heat transfer, two-TES system.

During operation, several potential situations necessitate the implementation of fail-safes. Regardless of the degree of design system balancing of the heat distribution between the HT-TES and LT-TES, the non-steady nature of operation will always cause certain mismatches to occur. The system is thermodynamically designed such that it is the HT-TES that can become overcharged, not the LT-TES. A bypass valve on the helium stream affords a means of thermal balancing. This valve is controlled by the difference between the HT-TES and LT-TES (relative) levels. The overall TES system can also become empty. A near-empty hot tank prompts an override of the desired power level, reducing it until the level stops decreasing and becomes steady. Until the desired power is decreased, the system is thus driven into steady-state operation. Similarly, when the TES is nearly full, the power level is appropriately increased such that the hot tanks will not overflow. Thus, in the top part of the figure, next to the input of the desired power, represented by the timetable “Demand_BOP,” a fail-safe mechanism is implemented for increasing or decreasing the power to prevent emptying or overflowing of the tanks.

Even though actual power systems are typically further limited by the ramp rates of specific components (e.g., MW/min for the turbine and K/min for any thermally stressed component such as a heat exchanger or turbine), these parameters are vendor- and design-specific, and were not included as limiters in the control system.

6.3 Shakedown Tests, Operational Limits, and Fail-Safes

The system model was subjected to testing in order to find the operational limits and ensure its reliable operation. The ramping speed was set to be 20 minutes, providing a 20-minute interval between reaching each steady setpoint. The control system was tuned so that the actual power would closely follow the setpoint, as seen in the shakedown test represented by Figure 33.

The limitations on the power levels at which a system can be operated are affected by implementing partial admission. With partial admission on both the HPT and LPT, the system can theoretically be operated down to zero load. When partial admission is implemented only on the HPT—or no partial admission is implemented at all—the lower limit of the power level was found to be about 15% of full load. Below this level, the current model cannot be operated using the proposed control system, mainly because the deaerator pressure would reach vacuum conditions. Going below this point would require either controlled extraction or multiple extraction points, with steam being routed based on the actual operation regime.

The second impact of partial admission is improved partial load efficiency. As the pressure upstream of the HPT—and eventually the LPT—remains constant, so does the cycle efficiency, as demonstrated via a shakedown test (see Figure 33). In all cases, the demand profile is well met. The cycle efficiency of a system with partial admission control for both the HPT and LPT has the least fluctuations, the efficiency during the steady state operation remains near-constant, and the efficiency is always higher than in the other cases, apart from the aggressive transient peaks seen during sudden changes in setpoint. The reason the efficiencies do not entirely match is that, to better show the difference here, the maximum power setpoint is a few percent below the system's maximum power. The extent of the efficiency decrease is highest when no partial admission is implemented.

To illustrate the impact of partial admission on the system parameters, Figure 34 plots the HPT and LPT inlet pressures for this shakedown test, which involved three different cases pertaining to partial admission. Without partial admission, the HPT inlet pressure decreased to 50 bar at low power levels, whereas partial admission would have otherwise kept it near-constant. In both cases, the extraction pressure for feedwater heating fluctuated at similar values when partial admission was not implemented in the LPT.

Note that model initialization is included in all the reported figures stemming from the model development conducted under the scope of the present work, and does not reflect physical behavior but rather numerical initialization of the system. The starting values were selected so that the set of equations will successfully converge, but they typically do not represent an actual operating state. Either sufficient time must be allowed to get beyond these initialization transients, or, in later analyses, the model can be initialized from a previously computed final operation point.

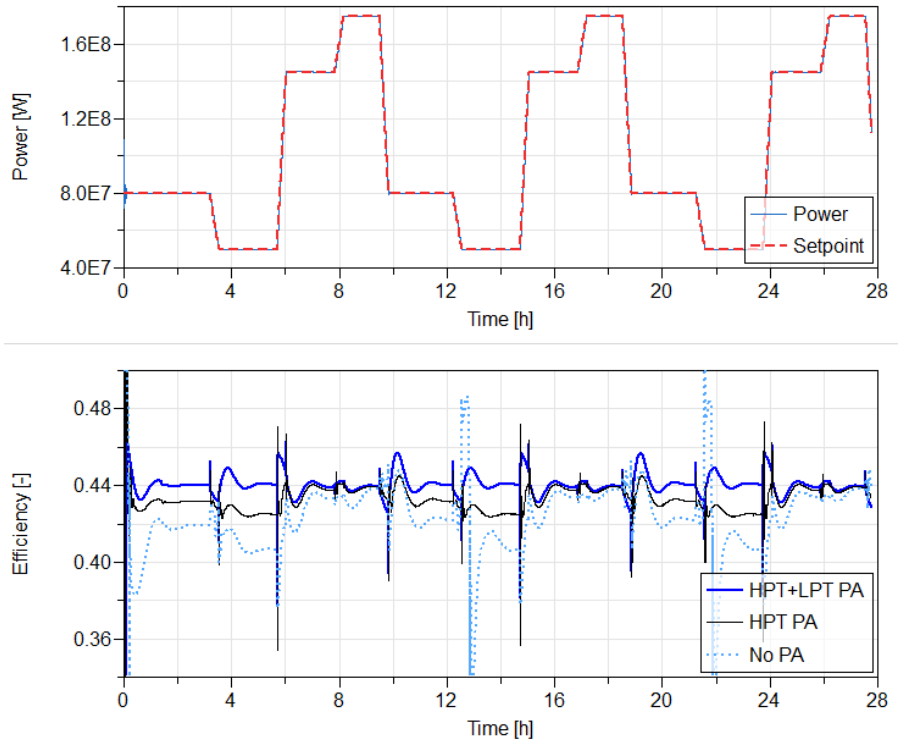

Figure 33. Shakedown test showing, for the single-phase heat transfer, two-TES, reheat-cycle systems, the following of the prescribed power output (left), as well as the impact of partial admission on cycle efficiency (right).

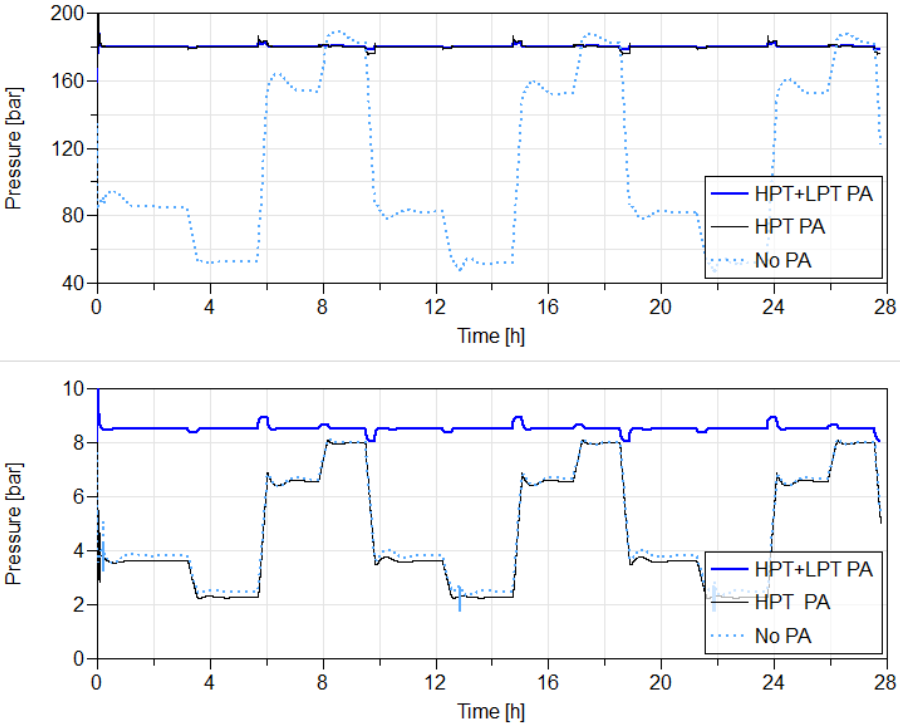

Figure 34. Bowl (inlet) pressures of the HPT (top) and LPT (bottom) during the shakedown test, for different explored scenarios of partial admission implementation.

The system requires that certain measures be implemented to ensure safe, continuous operation when its limits are reached. In evaluating IES, two situations commonly arise due to limited (finite) storage volume: a fully discharged system (empty hot tank) or a fully charged one (empty cold tank). If either of these occurs, the power must be decreased/increased so as to eventually attain a balance with the reactor power. In the control system, this was implemented in a manner similar to that employed for the tank models developed in prior work. When the limiting level is approaching, the required power production is overridden via a multiplier that increases its intensity as the zero level approaches. Application of this fail-safe was demonstrated on tank levels, actual power production, and the desired production signal (see Figure 35).

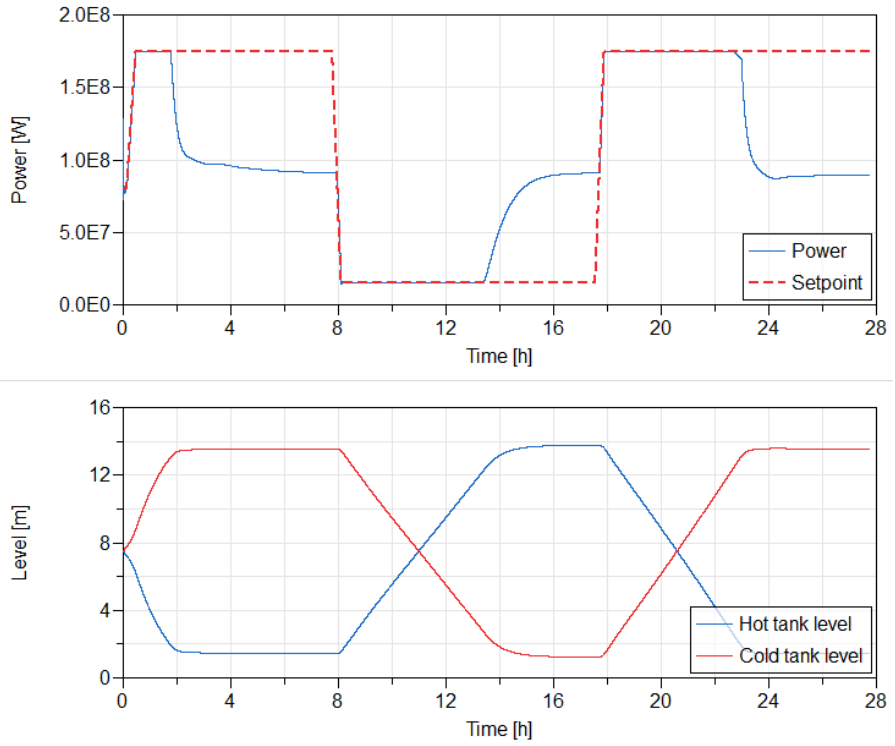

Figure 35. The tank level limits being reached, and the system responding accordingly by adjusting the power output.

As mentioned, the two-TES configuration requires a heat distribution between the HT-TES and LT-TES. The mismatch between them is managed by controlling a helium bypass valve for the HT-TES. In addition, by regulating the charging operations of each TES, it is possible to keep the tanks balanced, as shown in Figure 36.

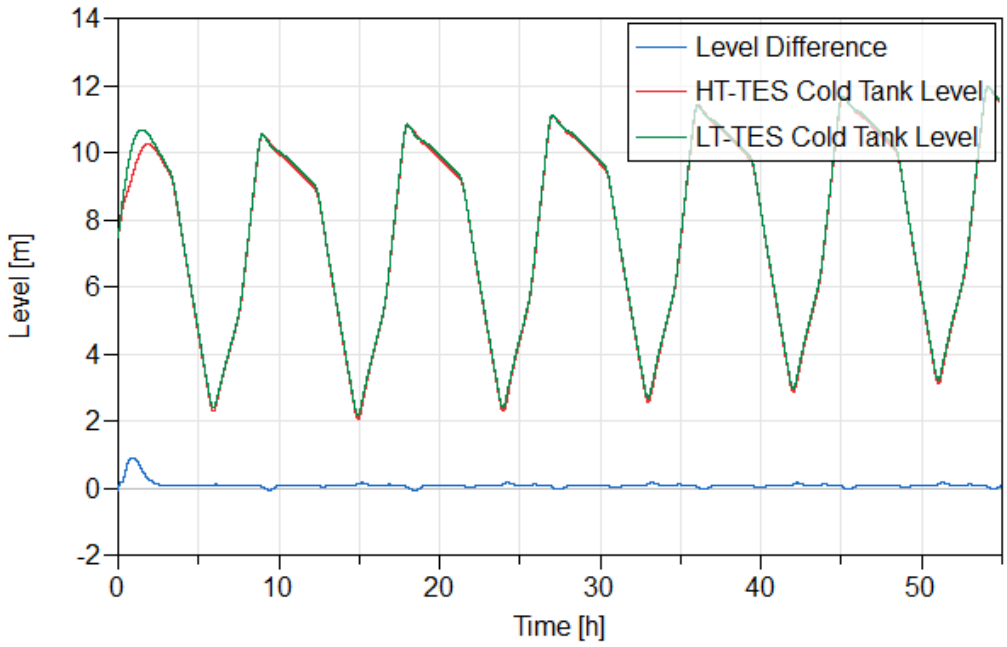

Figure 36. Level variation between the HT-TES and LT-TES during the shakedown test.

6.4 Application/Operation of the New Decoupled System Models using Dispatch Profiles

The dispatch profiles from the HERON optimization are given in hourly intervals, in accordance with the electricity market (e.g., day-ahead market), often with large changes on an hour-to-hour basis. Imposing such sharp changes as input may cause numerical issues in the convergence of the developed models, or lead to transients too fast for the adopted control systems to manage. Ramping for the purpose of softening these changes was implemented to ensure proper operation of the models.

Figure 37 shows the resulting agreement between the desired and the actual power profiles—along with the corresponding efficiency and close-up details—in light of the overall selected demonstration time period. The presented case is a selected cluster from a 120-hour-segment-length optimized case for the MISO market, in which rather aggressive ramping was seen across multiple power levels. In general, the power output matches the requirements well, showing certain transients and the impact of a minimum power load (here, 8.5% of the nominal output). Note that even coal plants typically have a minimal operating load of 30%, with some exceptions being around 15% [33]. Steady operation at below 10%, even for a plant purposefully built to afford maximal flexibility, is highly unlikely. The sharp peaks in efficiency arised from minor parameters oscillation during the transition into or out of very low output within the implemented control system. Real system limitations, the control system, and the shape of transient behavior would result from the specific detailed design and the vendor parameters driven by thermal stresses, not from this thermodynamic-level analysis.

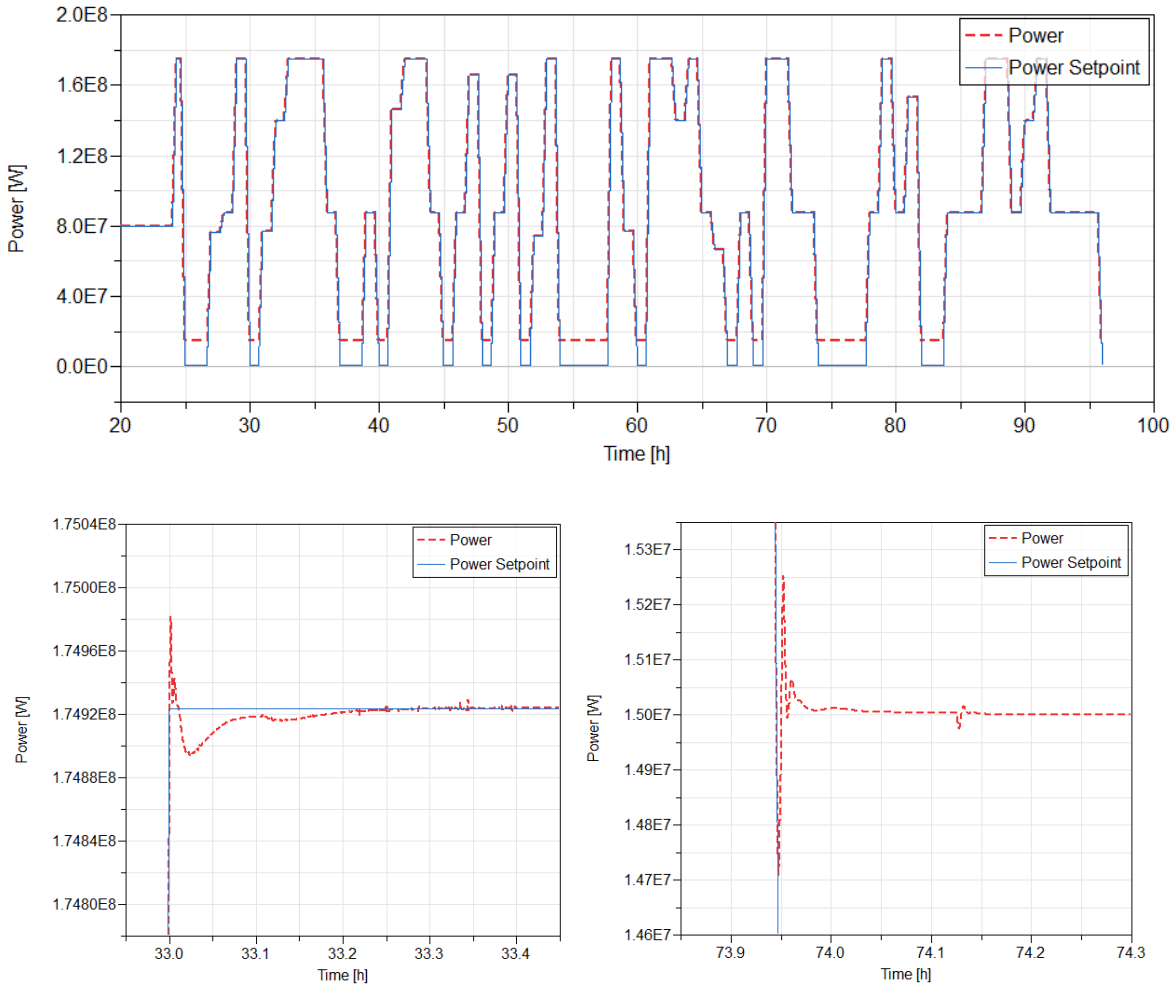

Figure 37. Two-TES, reheat-cycle system when operated to follow the HERON optimized dispatch profile.

Figure 38 similarly shows the operational efficiency reported from the shakedown test, along with the controlled thermal input into the cycle. As the ramps are faster and some of the transitions a little rough, spikes are seen in the efficiency. Still, the decrease in average efficiency as a result of the partial load is relatively low, thanks to partial admission having been implemented.

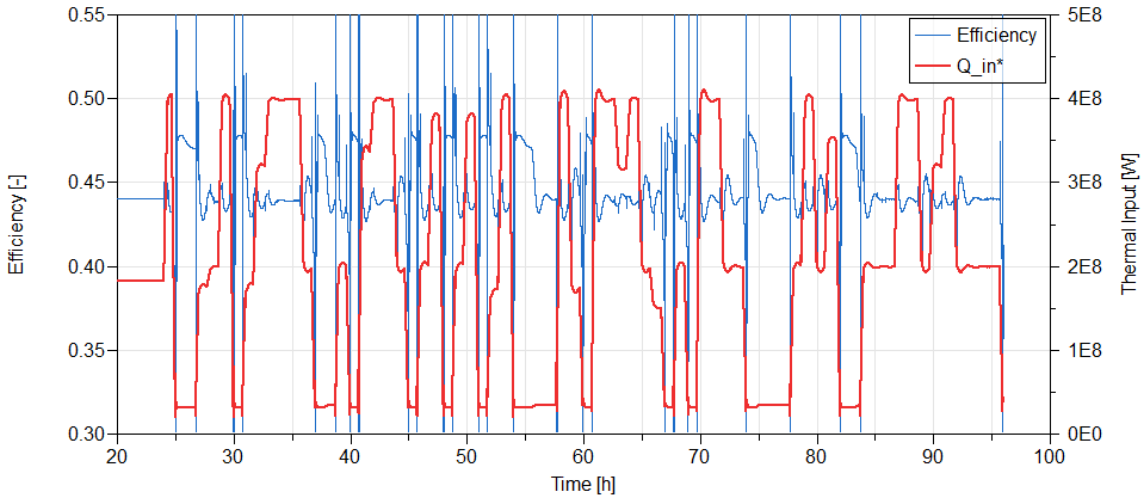

Figure 38. Efficiency based on the immediate values of produced power and the thermal input into the cycle.

Figure 39 shows the variations in reactor thermal power for this test run. The maximum variation that occurs after reaching a steady-state operation (from initialization) remains within 0.1% of the nominal reactor power. The impact on reactor control elements, control rods (as represented by reactivity), and the mass flow rate through the circulation compressor—with the current system factoring in most of variations, based on changes in circulation flow—are shown as well.

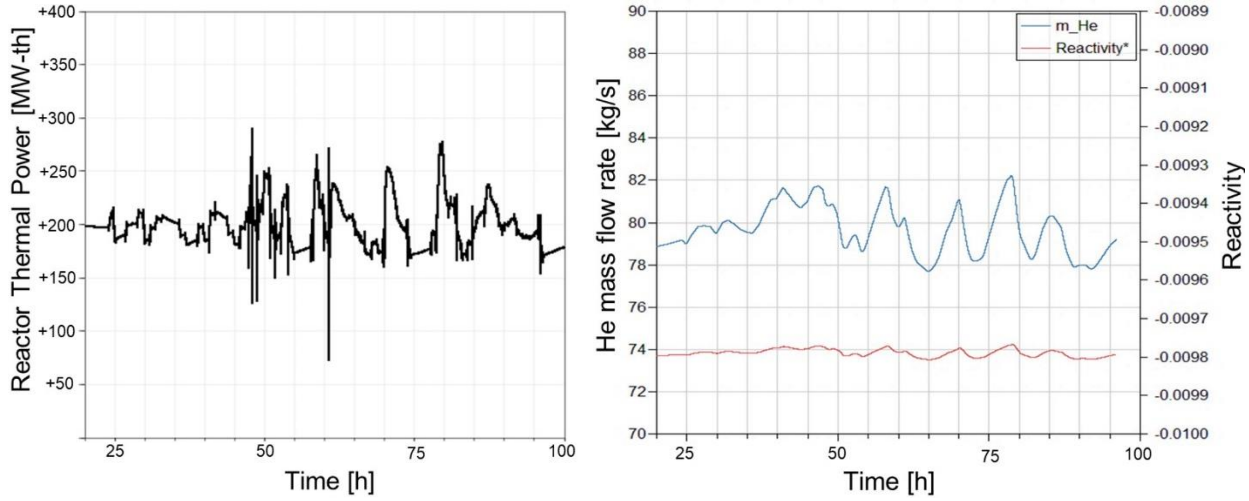

Figure 39. Fluctuation in the reactor power and reactor control elements.

Fluctuations in reactor inlet temperature are the main cause of any system impact on the controlled reactor power. Such fluctuations are due to LT-TES hot tank (charging) temperature control. Small variations in the cold LT-TES tank (as return salt from the BOP changes the temperature, per the operational regime) and incoming helium into the LT-TES CHX are propagated to the temperature of the helium returning to the reactor. These temperatures are represented in Figure 40. Decreasing the strength of this controller causes fluctuations in the hot tank inlet temperature (see the magenta line in Figure 40). Its impact on the power cycle is minimal, but it dampens the impact on the return helium temperature and, in turn, the reactor power. Note that this impact can again be fine-tuned using more detailed and complex controllers.

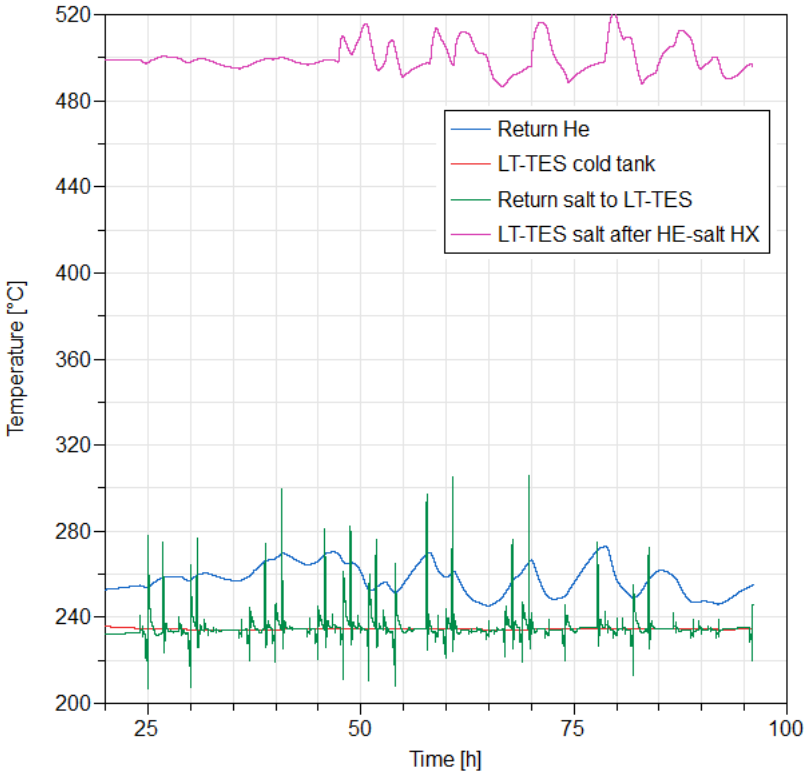

Figure 40. Salt and helium temperature of the LT-TES system.

7. CONCLUSIONS

Single-phase-to-single-phase heat transfer systems for coupling a nuclear HTGR to TES were investigated to provide a complex analysis of TES options for flexible-operation NPPs. To highlight the potential of the newly proposed systems, the hot storage fluid was maintained as close to the helium temperature as possible. As the exergy losses from storage are minimized, both systems provide theoretically nominal efficiency during discharge. The first of the proposed configurations was based on the Xe-100, with 42.2% efficiency, while the turbine efficiency was artificially increased to compensate for simplified feedwater heating in the models. This configuration can be considered partially decoupled, generating a HDR of up to about 33% from the reactor to the TES, and it uses a single set of HT-TES system (745°C) that operates based on $MgCl_2$-KCl salt. The second configuration fully decouples the reactor heat from the conversion system by employing an intermediate molten-salt system that acts not only as a TES storage fluid but also as a heat transfer intermediate loop. In this manner, it theoretically provides an unlimited HDR. The downside of the complexity involved is that no readily available molten salt can cover a temperature span ranging from less than 250°C to nearly 750°C; thus, the TES had to be split into high- and low-temperature subsystems. The HT-TES uses $MgCl_2$-KCl, whereas the LT-TES uses HITEC salt. This approach leads to limitations as well as opportunities on the discharging side, in our case leaning toward a single reheat steam cycle. As a result, an improved cycle efficiency of 43.1% has already been achieved with realistic turbine efficiency and a single feedwater heater, and this percentage could perhaps be pushed significantly higher if a full feedwater heating system is considered.

The sizing of the components that was needed in addition to the baseline nuclear plant was performed for the fully decoupled case. Special attention was paid to selecting suitable materials, especially for very-high-temperature and high-pressure components. The resulting design considers Incoloy 800 for the highest temperature components, stainless steel for the intermediate-temperature components, and carbon steel for the low-temperature (~250°C) components. After the sizing, a component costing analysis was

performed, followed by development of costing functions for supersets pertaining to discharging, storage, and charging operations (the charging being fixed by only a single point, as it always occurs under nominal conditions).

Using the cost functions and the dispatch optimizing tool, HERON, the cost-optimal interaction between price signals and optimal superset capacities, and various performance attributes was derived. The includes the optimal discharge-to-charge ratio, charge/discharge duration, and distribution of state-of-charge. This assessment was performed within a stochastic framework from the trained price signal ARMA profiles. Our findings revealed significant regional variations in NPV—variations primarily driven by electric sales revenue. The distinct volatility and concentrated price peak patterns within ERCOT led to the highest TES sizing, in contrast to the PJM and MISO markets. However, the sizes of both charging and discharging displayed relatively consistent trends across these markets. In a sensitivity analysis, we observed that a minimum generation constraint significantly influences both TES sizing and operation, including reducing the need for TES, lowering electricity sales revenue, and producing notable changes in the charge and discharge operation patterns. When the TES size is cost-optimal, it is generally operated for short duration storage (<3 h) across all the markets. With the lower cost and higher RTE, the role of TES is likely to expand to mid-duration storage (3–6 h). Thus, further research is required to characterize optimal TES operation while considering TES costs and RTE.

The fully decoupled system was developed in the Modelica environment. Herein, partial admission of the turbine was explored to minimize the efficiency drop during partial-load operation. Additionally, operation limits and the need for several fail-safe mechanisms were identified. These fail-safes were then implemented (specifically, control of the level of mismatch between the HT- and LT-TES, and new tank overcharge/over-discharge protection). The impact of power ramping and load following of the dispatch profile from HERON was demonstrated. And lastly, it was shown that full decoupling and the implementation of suitable control systems produces a minimal impact on reactor parameters, even during an aggressive load profile.

8. FUTURE WORK

In a broader view of future work, this research provides valuable insights into the feasibility of TES systems that utilize a two-tank molten-salt design; however, further investigation is needed to assess the long-term performance and explore alternative HT-TES options such as alternative molten salts and packed beds with gaseous heat transfer fluids. In the field of fully decoupled systems, as well as with gas-solid packed bed systems, multiple alternative conversion systems can be analyzed, such as the steam supercritical cycle or application of the Brayton cycle (perhaps also in conjunction with Brayton cycle pumped TES).

By providing a baseline for nuclear-TES coupled reactors for electricity production, we can now explore another level of application. Integrated flexible electricity and zero-carbon high-temperature heat supply systems would utilize the solutions developed here. The application value could even be extended by added thermal extraction at suitable points. With known heat demand profiles and intermittent production from renewables, interesting economic dispatch optimization cases could be set up.

However, it is important to also keep the systems within technically feasible parameters. Exploring and implementing real limitations such as component ramp rates (MW/min), actual minimal loads, thermal losses, and auxiliary system power consumption, or investigations into predictions of high-renewable-penetration grids would properly position the systems for use in future real-world application scenarios.

9. REFERENCES

[1] Saeed, R. M., Shigrekar, A., Frick, K. L., and Bragg-Sitton, S. M. (2022). "Phenomenon Identification and Ranking Table Analysis for Thermal Energy Storage Technologies Integration with Advanced Nuclear Reactors." INL/EXT-21-65459 Rev. 000, Idaho National Laboratory, Idaho Falls, ID. https://doi.org/10.2172/1875124.

[2] Saeed, R. M., Frick, K. L., Shigrekar, A., Mikkelson, D., and Bragg-Sitton, S. (2022). "Mapping thermal energy storage technologies with advanced nuclear reactors." *Energy Conversion and Management*, 267 (Sept.), 115872. https://doi.org/10.1016/j.enconman.2022.115872.

[3] Saeed, R. M., et al. (2022). "Multilevel Analysis, Design, and Modeling of Coupling Advanced Nuclear Reactors and Thermal Energy Storage in an Integrated Energy System." INL/RPT-22-69214 Rev. 000, Idaho National Laboratory, Idaho Falls, ID. https://doi.org/10.2172/1890160.

[4] Saeed, R. M., et al. (2024). " Design and optimization of flexible decoupled high-temperature gas-cooled reactor plants with thermal energy storage"/ *Energy Conversion and Management*, Volume 321. https://doi.org/10.1016/j.enconman.2024.119098.

[5] Weller, T., Jockenhöfer, H., Johnson, M., and Bauer, D. (2020). "Design and build of a first laboratory CHEST system." In *2nd International Workshop on Carnot Batteries*, Stuttgart, Germany. https://www.researchgate.net/publication/344634969_Design_and_build_of_a_first_laboratory_CHEST_system.

[6] Knobloch, K., et al. (2022). "A partially underground rock bed thermal energy storage with a novel air flow configuration." *Applied Energy*, 315 (2022): 118931. https://doi.org/10.1016/j.apenergy.2022.118931.

[7] Hyme. (n.d.). "New Molten Salt Storage for the Energy Transition." CVR 42822027, *Hyme Energy APS*. https://www.hyme.energy/technology.

[8] Saeed, R. M., Shigrekar, A., and Toman, J. (2023). "Synthetic Electricity Market Data Generation and HERON Use Case Setup of Advanced Nuclear Reactors Coupled with Thermal Energy Storage Systems." INL/RPT-23-70945 Rev. 000, Idaho National Laboratory, Idaho Falls, ID. https://doi.org/10.2172/1960133.

[9] Nomoto, H. "Development in materials for ultra-supercritical (USC) and advanced ultra-supercritical (A-USC) steam turbines." *Advances in Steam Turbines for Modern Power Plants*. Woodhead Publishing, 2017. 263-278. https://doi.org/10.1016/B978-0-08-100314-5.00012-9.

[10] Kelly, B. D. (2010). "Advanced Thermal Storage for Central Receivers with Supercritical Coolants." DOE/GO18149, *Abengoa Solar, Inc*. https://www.osti.gov/servlets/purl/981926.

[11] Turchi, C. S. and Heath, G. A. (2013). "Molten salt power tower cost model for the system advisor model (SAM)." NREL/TP-5500-57625, National Renewable Energy Laboratory (NREL), Golden, CO. https://www.nrel.gov/docs/fy13osti/57625.pdf.

[12] Turton, R., Bailie, R. C., Whiting, W. B., and Shaeiwitz, J. A. (2008). Analysis, Synthesis and Design of Chemical Processes." *Pearson Education*. https://scholar.google.com/scholar?q=Analysis,+Synthesis+and+Design+of+Chemical+Processes+Pearson+Education&hl=en&as_sdt=0&as_vis=1&oi=scholart.

[13] Maxwell, C. (2023). "Cost Indices." Towering Skills, updated Jun. 06, 2023. https://toweringskills.com/financial-analysis/cost-indices/.

[14] MEPS. (2023). “Stainless Steel 316 Composite Prices.” ©2023 *MEPS International Ltd*. https://mepsinternational.com/gb/en/products/stainless-steel-316-composite-prices?#product-unit-measurement-conversion-form.

[15] Relative Cost Ratio, Steel Tank Institute/Steel Plate Fabricators Association. https://www.steeltank.com/Portals/0/Pressure%20Vessels/SSWseminarOct2012/Relative%20Cost%204%2015%202012.pdf.

[16] Guanyu Tube. (n.d.). “Stainless Steel Relative Cost Data.” *Guanyu Tube*. https://tubingchina.com/Stainless-Steel-Relative-Cost-Data.htm.

[17] Incoloy Alloy 800. *Aesteiron Steels, LLP*. https://www.aesteiron.com/incoloy-800.html#price.

[18] Price list © 2001-2019. by Steel Tubes India, last updated Jun. 01, 2019. https://www.stindia.com/inconel-alloy-600-625-718-price-per-kg-india.html.

[19] Incoloy 800 supplier, © 2001-2019. by *Steel Tubes India*. https://www.stindia.com/incoloy-alloy-800-800h-800ht-supplier.html#price.

[20] Bonk, A., et al. (2022). “Thermal Energy Storage using Solar Salt at 620°C: How a reactive gas atmosphere mitigates corrosion of structural materials.” *Proc. in SolarPaces International Conference 2022*, Albuquerque, NM. https://elib.dlr.de/190729/.

[21] SAM 201 report. Data from *Abengoa* 2010. https://sam.nrel.gov/concentrating-solar-power/csp-cost-data.html.

[22] NREL. (2017). “SAM 2017.9.5.” National Renewable Energy Laboratory. https://sam.nrel.gov/download/version-2017-9-5.html.

[23] SAM 2019. https://www.nrel.gov/docs/fy19osti/72856.pdf.

[24] SAM 2022. https://sam.nrel.gov/download/version-2021-12-02.html.

[25] González-Roubaud, E., Pérez-Osorio, D., and Prieto, C. (2017). “Review of commercial thermal energy storage in concentrated solar power plants: Steam vs. molten salts.” *Renewable and Sustainable Energy Reviews*, 80 (2017): pp. 133-148. https://doi.org/10.1016/j.rser.2017.05.084.

[26] Gen3 CSP. https://www.energy.gov/eere/solar/generation-3-concentrating-solar-power-systems-gen3-csp.

[27] Forsberg, C. (2022). “Separating nuclear reactors from the power block with heat storage to improve economics with dispatchable heat and electricity.” *Nuclear Technology* 208(4) (2022): 688–710. https://doi.org/10.1080/00295450.2021.1947121.

[28] Epiney, Aaron, Rabiti, C. Talbot, P. and Alfonsi, A. (2020). “Economic analysis of a nuclear hybrid energy system in a stochastic environment including wind turbines in an electricity grid.” *Applied Energy* 260, (2020): 114227. https://doi.org/10.1016/j.apenergy.2019.114227

[29] Cho, So-Bin, Epiney, A. S., Talbot, P. W., Sun, X., and Allen, T. (2023). “Study of Storage Requirements and Costs for Shaping Renewables and Nuclear Energy: FY23 Summary Report.” No. INL/RPT-23-72642-Rev000, Idaho National Laboratory, Idaho Falls, ID. https://www.osti.gov/biblio/1984047.

[30] Talbot, Paul W., et al. (2020). “Evaluation of Hybrid FPOG Applications in Regulated and Deregulated Markets Using HERON.” No. INL/EXT-20-60968, Idaho National Laboratory (INL), Idaho Falls, ID; Electric Power Research Institute (EPRI), Palo Alto, CA. https://doi.org/10.2172/1755894.

[31] Frew, B. A. and Jacobson, M. Z. (2016). “Temporal and spatial tradeoffs in power system modeling with assumptions about storage: An application of the POWER model.” *Energy* 117(1) (2016): 198-213. https://doi.org/10.1016/j.energy.2016.10.074.

[32] Li, Binghui, Talbot, P. W., McDowell, D. J., and Hansen, J. K. (2021). “Release a public version of HERON (HERON 2.0) with improved algorithms for the treatment of energy storage.” No. INL/EXT-21-65473-Rev000, Idaho National Laboratory, Idaho Falls, ID. https://doi.org/10.2172/1838610.

[33] Fiebrandt, M., Röder, J., Wagner, H-J. (2022). “Minimum loads of coal-fired power plants and the potential suitability for energy storage using the example of Germany.” *Int J Energy Res*., 46(4), 4975- 4993. https://doi.org/10.1002/er.7490.

Page intentionally left blank

APPENDIX A

Implementation Analysis Using Modelica's Transient Physical Modeling for HTGR

This appendix describes the transient model design for the HTGR system with steam-heated TES. It reports on its control methodology and implementation, and a justification is made for the steam diversion ratio limits, based on the initial design. This appendix also presents the model output based on the economic dispatch profile generated by HERON for a charge capacity limit of 50%, and off-design performance is evaluated as well. All models presented here are available in the HYBRID repository of Modelica models, as found on GitHub.

A-1. Understanding the Need for Transient Physical Analysis

Per the economic optimization performed in HERON, a steam diversion ratio approaching 100% appears optimal. The process modeling in Aspen HYSYS® was performed for design states under the assumption of 0% and 50% steam diversion ratios, based on engineering judgement, since coal plants, which are typically more flexible than nuclear systems, have a technical minimal operating load at about 30%, with some exceptions being around 15% [33]. Furthermore, a physics-informed justification for this decision is required. In addition, the design models used in Aspen HYSYS® only capture full charge and discharge states, while the economic dispatch allows for (and indeed finds) optimum plant dispatch strategies based on states that fall between these two extremes. In these states, the turbines operate under partial loads, and generating these off-design performance data requires a more flexible modeling approach. As for the implementation of such a design, a robust control system methodology must be designed to account for following load demand profiles while remaining within safe limits of operation.

The HYBRID collection of IES models written in the Modelica language affords this flexibility, providing off design performance analysis and the ability to investigate transient effects while following the desired economic dispatch.

A-2. Control Design and Conceptual Operation

The control system discussed in the following section was developed based on system power dispatch control. Given a power setpoint, three operating regimes may occur. The state control decision logic behind these three modes is seen in Figure A-1. When the power requirement is lower than the nominal reactor electrical output, the system will operate in a charging mode. Turbine power is decreased, and an excess of generated steam is directed to the TES while keeping the reactor power constant. This charging mode can range from a minimal steam mass flow rate to the TES in order to maintain hot standby operation, to a specified maximum of 50%. When the power requirement exceeds the nominal reactor electrical output, the system will operate in discharging mode and meet the power requirement by dispatching energy from the TES. Again, this discharge mode can range from a minimum output power from the independent BOP corresponding to hot standby, to a maximum output BOP capacity dictated by the economic optimum found by HERON. Steady operation occurs when the electrical demand has a heat rate identical to the reactor thermal output and normally corresponds to the nominal main plant electrical output. Note that steady operation occurs much more frequently during economically driven operation than during load-following operation, given the efficiency penalty for passing energy through storage.

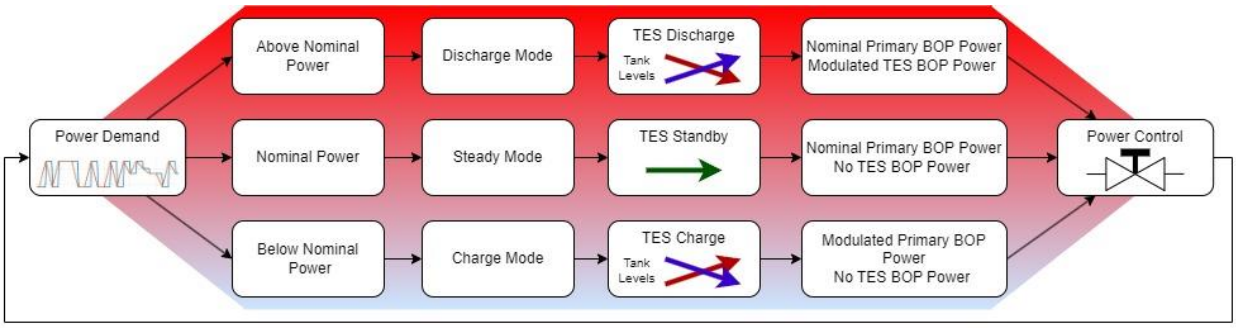

Figure A-1. State control logic behind power control.

Some terms must be introduced before discussing the IES deployment scenario. The thermal integration point for heating the TES is the charging heat exchanger (C-IHX). The flow depositing heat across the C-IHX is referred to as "bypass" while hot, and "condensate" when flowing out of the C-IHX. The mass flow in the C-IHX on the TES side is the "charging" flow. The thermal dispatch integration point is the discharging heat exchanger (D-IHX). The mass flow in the D-IHX on the TES side is the "discharging" flow, and the steam flow on the other side is the "TES feed" and "TES steam."

During TES standby mode, the TES control systems would operate minimal flow on both the charging and discharging sides. Modelica is unable to consistently handle the physics of zero flow, and thus these minimal flows are set to 1.93 kg/s (around 2.5% of the nominal steam flowrate) for the C-IHX and 2.13 kg/s for the D-IHX (producing 1.7 MWe in BOP). There are nine controlled elements within the full IES, as seen in Figure A-2. The flow rate in the HTGR primary is controlled by a helium circulator, with the output pressure ratio of this circulator being controlled based on the desired thermal power of the reactor. The primary side controls the primary coolant pressure via mechanisms outside the consideration of IES operation. The pressurizer operates via nominal methods in all operating paradigms. Control rods on the primary side will also operate according to the nominal method of controlling output coolant temperature.

Within the primary BOP, the controllable mechanisms include the turbine bypass valve (TBV), turbine control valve (TCV), feedwater heating valve (FHV), and feedwater control pump (FWCP). The independent TES BOP will also have a separate TCV and FWCP set denoted with a "D-" (standing for discharge) before the name (i.e., the D-TCV and D-FWCP), and this set must be controlled. The TES model used is sized so that constantly operating pumps can produce the flow necessary to meet 100% of the TES charge/discharge demand. Internally, a charging control valve (CCV) and a discharging control valve (DCV) moderate the TES fluid flows.

To meet demand, control for the NPP-integrated BOP must be able to operate during all three operational modes. The overarching system demand is the primary input for the control systems. Table A-1 summarizes the BOP control methods. Turbine output is the primary operational setpoint and will be controlled via the TBV. Note that this TBV is independent of a pressure relief valve for ensuring that over-pressurization does not occur. The TCV operates to maintain the steam pressure at the outlet of the steam generator. The FHV position meets the feedwater temperature during all operational modes. To maintain the steam generator output temperature, the FWCP changes its speed to increase or decrease the pumping power. The fact that the components combine in a non-linear system enables the overall system conditions to be met. The goal of the FHV, TCV, and FWCP subsystem is to provide a constant heat sink to the primary side in order to avoid impacting reactor conditions whenever possible.

The TES BOP system is relatively simpler. Feedwater is sourced directly from the TES BOP condenser and pumped through the D-FWCP, which operates to maintain steam pressure. The D-TCV opens to meet the electrical demand (i.e., the system demand less the nominal BOP production) in conjunction with a second valve, D-FCV, which effectively operates to meet large changes in demand, while the D-TCV makes smaller adjustments to the power output.

Within the TES, the CCV changes position such that the charging flow meets a reference hot tank temperature. Similarly, the DCV changes position to meet the reference cold tank temperature. Both valves have two additional mechanisms that aid in modal operation. The first is an anticipatory signal that forces open the valves to let some mass through at the initiation of charging/discharging. This counteracts the fact that, according to the valve logic, simply not opening is a mathematically acceptable way to meet the reference temperatures. The second mechanism is a safety mechanism that forces the valves to close when there is a lack of source media (meaning that the CCV looks at the cold tank level and the DCV looks at the hot tank level). This is done by multiplying the opening value by a fraction if the level becomes too low.

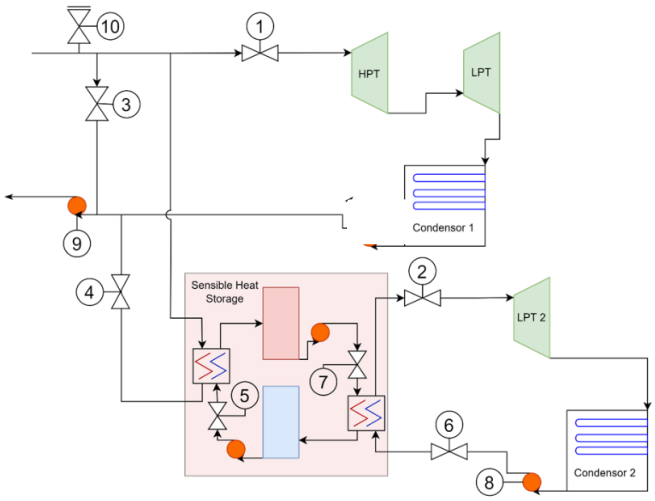

Figure A-2. Block diagram of the control system implemented within the Modelica model.

Table A-1. Control methods summary.

Label	Name	Controlling	Setpoint
1	Turbine Control Valve (TCV)	Steam Inlet Pressure (HPT)	165 bar
2	Discharge Turbine Control Valve (D-TCV)	Power (LPT 2)	If demand power > Q_nominal = demand power – Q_nominal else is 0.
3	Feedwater Heating Valve (FHV)	Feedwater Temperature	193°C
4	Turbine Bypass Valve (TBV)	Power (HPT and LPT)	If demand power > Q_nominal = Q_nominal else is demand power.
5	Charging Control Valve (CCV)	SHS Hot Tank Temperature	400°C
6	Discharging Feedwater Control Valve (D-FCV)	Power (LPT 2)	If demand power > Q_nominal = demand power – Q_nominal else is 0.
7	Discharging Control Valve (DCV)	SHS Cold Tank Entry Temperature	240°C
8	Discharging Feedwater Control Pump (D-FWCP)	Steam Inlet Pressure (LPT 2)	105 bar
9	Feedwater Control Pump (FWCP)	Steam generator output temperature	565°C
10	Pressure Relief Valve	Pressure Overloads	Release at 180 bar

A-3. Limits of Operation and off Design Performance

Off design performance refers to operating the steam turbines in the model at any mass flow rate other than their nominal design setpoint. For the NPP BOP this occurs whenever a fraction of the steam is diverted to charge the TES. We define the steam diversion ratio as:

$$\text{Steam Diversion Ratio} = \frac{\text{Steam Mass Flow Rate in C-IHX}}{\text{Steam Generator Mass Flow Rate}} \tag{1}$$

The main NPP-BOP efficiency is defined as:

$$\text{Main BOP Efficiency} = \frac{(\text{NPP-BOP Electrical Power}) - (\text{Pumps Power})}{\dot{m}_{\text{TCV}} \cdot \left(h_{\text{NPP-BOP, in}} - h_{\text{NPP-BOP, out}} - h_{\text{FWP, out}} + h_{\text{FWP, in}}\right)} \tag{2}$$

Figure A-3 captures how the efficiency of the main BOP varies for steam diversion ratios ranging from 2.52% to 86%. It is seen that, at up to 60% diversion ratio, the efficiency consistently drops with the reduction in steam mass flow rate. As mentioned above, the HERON results favor high steam diversion ratios and maximizing the energy storage size within the given timescale. When implementing this in Modelica for the given design, it can be observed that the minimum exit temperature of the condensate from the C-IHX is 240°C, as limited by the TES cold tank temperature. Without additional plant infrastructure to utilize this condensate, or indeed rejecting its heat at steam diversion ratios greater than 60%, we cannot maintain the feedwater temperature at 193°C when the FHV is fully closed. This uncontrolled steady-state temperature rise is seen in Figure A-3. The loss of feedwater temperature control is undesirable and should be avoided—from a safety perspective—by implementing a robust operating margin on temperature control. Designs that reject this heat significantly reduce their cycle efficiency but benefit from the ability to maximize steam diversion. In this research, we deemed such designs to be outside the scope of the present study. As a result, we elected to cap the steam diversion rate at 50%, aligning with the initial assumption and extending to the transient analysis for the selected case from the HERON solution.

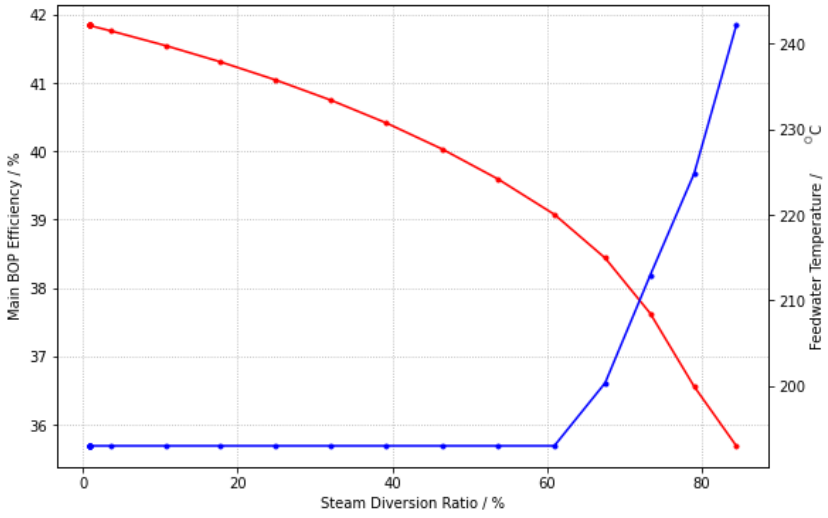

Figure A-3. Off-design efficiency of the NPP BOP (left-hand axis) and controlled feedwater temperature (right-hand axis) under varying steam diversion ratios.

The TES BOP's turbine operates over a greater mass-flow-rate range than the main BOP turbine. It must be able to cope with power levels ranging from 1.7 MWe (power at hot standby flow) to 39.84 MWe. This full range is considered for the sake of simplicity of analysis, knowing that the actual design might not be able to go below a 15% load. The TES-BOP efficiency is defined as:

$$\text{TES-BOP Efficiency} = \frac{(\text{TES-BOP Electrical Power}) - (\text{Pump Power})}{\dot{m}_{\text{D-TCV}} \cdot (h_{\text{TES-BOP, in}} - h_{\text{TES-BOP, out}})} \quad (3)$$

Over the range of operation, this efficiency varies greatly, as seen in Figure A-4.

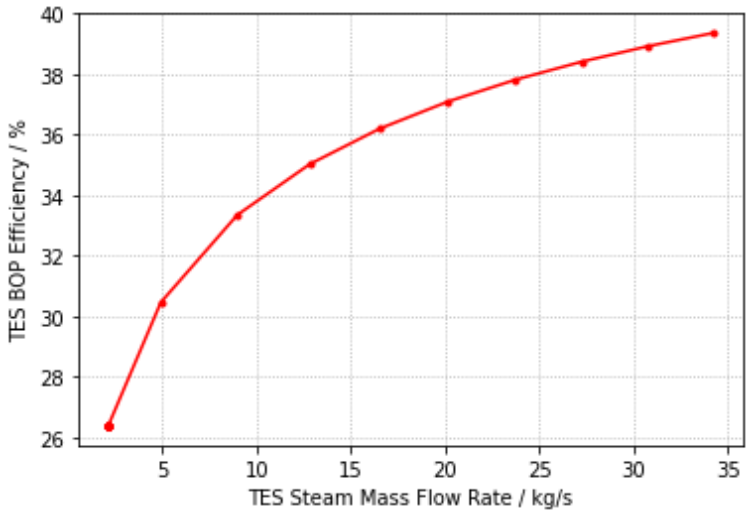

Figure A-4. Off-design efficiency of the TES-BOP for varying steam mass flow rates in the D-TCV, from the minimum mass flow rate of 2.13 kg/s up to the nominal mass flow rate at 34.13 kg/s, corresponding to 39.84 MWe.

This leads to adjustments that may be accounted for in future HERON runs, including a minimum power rate from the TES BOP, which was set at 1.7 MWe in this case, and changing the thermal-to-electric conversion efficiency in BOPs, based on the power rate. A minimum power rate on the TES BOP stems from two sources. The first is a modeling limitation: transient models tend to be computationally intractable when systems alter their operation modes too frequently. By requiring a minimum power rate, the entire TES BOP can always generate at least some power, and no zero-flow errors arise. The second is an as-of-yet poorly characterized but well-understood need to maintain at least hot standby conditions within a turbogenerator system. Thermal and rotational ramping limits are neglected in turbogenerator modeling, using only the Stodola equation for off-design steady state, and a simplifying assumption is used to maintain minimum power generation. The former limit of avoiding zero-flow errors within our transient simulations imposes a minimum bypass flow of 1.93 kg/s from the power cycle to the C-IHX—and consequently a minimum steam diversion ratio of 2.52%. This also would need to be applied to HERON in future iterations.

A-4. Heron Dispatch Analysis

The optimal dispatch that HERON found for the 50% steam diversion limited case served as an economic dispatch profile to demonstrate the linkage between the HERON and HYBRID tools. Power level changes were scheduled to be 10-minute ramps, so a change from hour x to hour x+1 occurs from hour x:55:00 to hour x+1:05:00. Figure A-5 shows the difference in dispatch profile when these flattened-demand shorter ramps are imposed on the system.

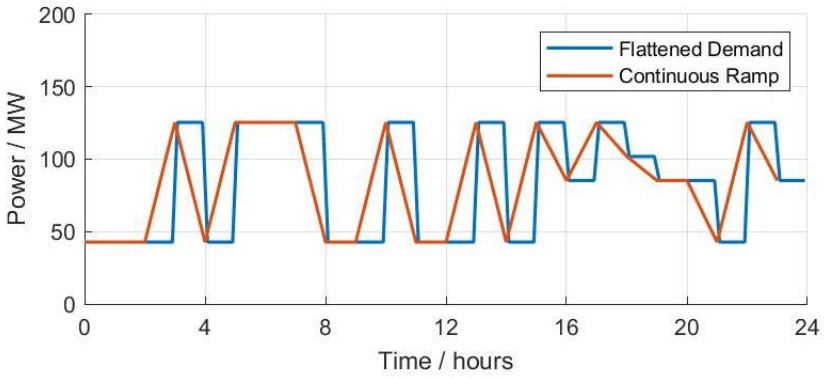

.

Figure A-5. HERON dispatch demand, showing continuous ramp (linear connection between hourly setpoints) demand and 10-minute inter-hour ramping demand.

Figure A-6 shows the IES electricity demand and production throughout the simulation. This value is generally matched, apart from a minor degree of overshoot during down transients and a minimal amount of oscillation at steady state. Tuning of the control system required significant integral action to achieve the desired ramping, leading to the oscillations during the most extreme 10-minute ramps. Figure A-7 shows the TES BOP demand profile, which is dispatched to provide additional power over the nominal BOP power. The variation in efficiencies of both BOPs is dependent on their required power dispatch, as seen in Figure A-8. Figure A-9 shows the impact on reactor power as a result of this mode of operation. The reactor power changes on short timescales, due to feedback from changing conditions in the steam generator, but system controls via the helium circulator speed and control rods maintain +/-0.2% operating conditions.

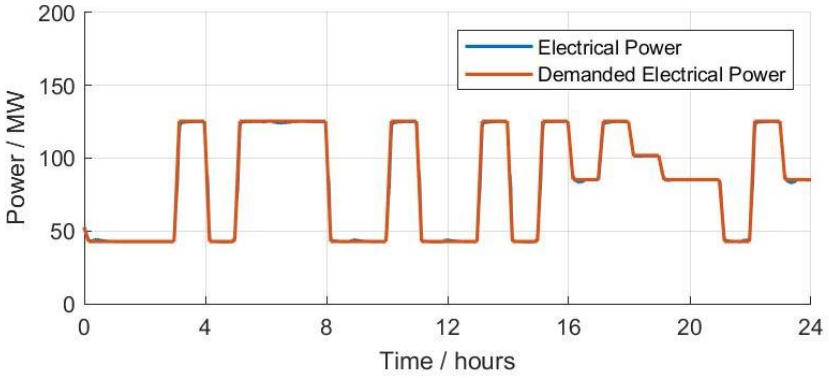

Figure A-6. Electricity production vs. demand throughout a 24-hour dispatch test.

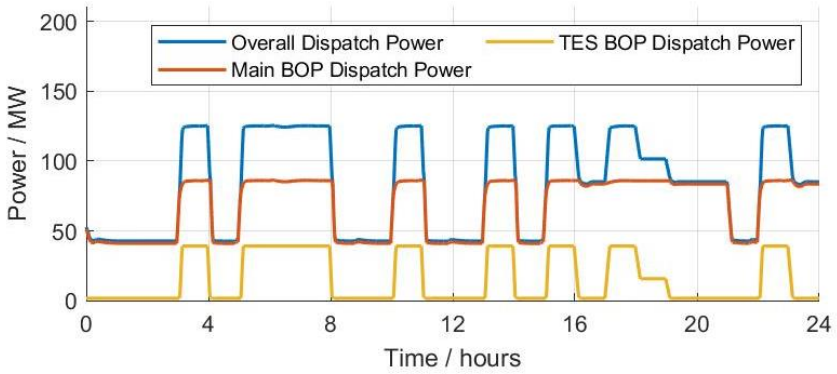

Figure A-7. BOP dispatch powers throughout a 24-hour dispatch test.

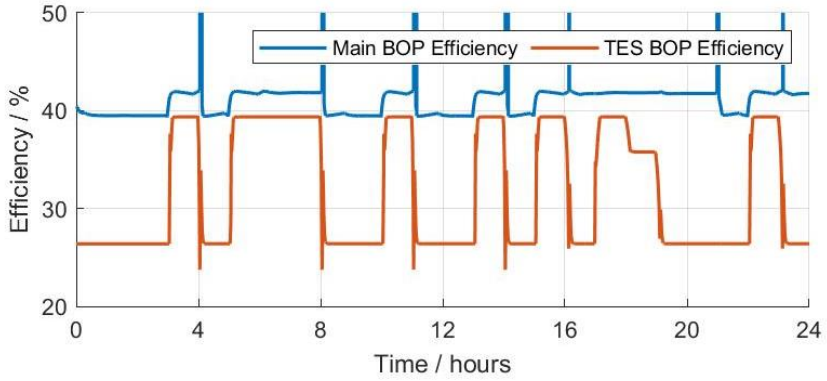

Figure A-8. BOP efficiencies (with definitions taken from Eqs. (1) and (3)) throughout a 24-hour dispatch test. The transient peaks have been cut off to improve clarity.

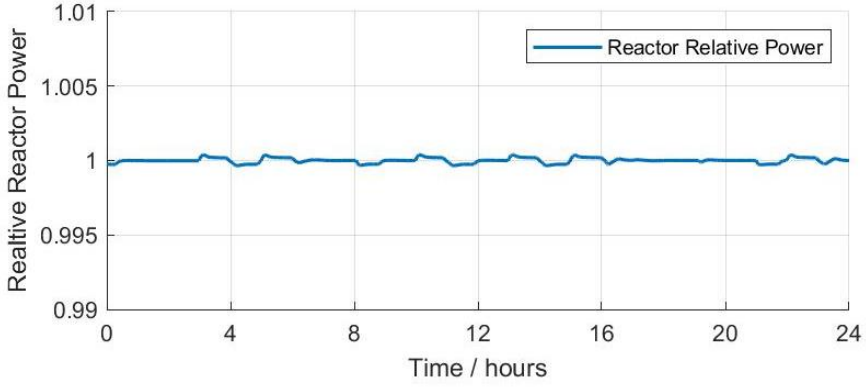

Figure A-9. Relative reactor power throughout the dispatch test. The power is maintained within +/-0.5% of nominal power.

Perhaps of more concern is this state switching's impact on the feedwater temperature shown in Figure A-10. While the FHV acts rapidly to maintain a fixed feedwater temperature, the impact of large mass flow rate swings in condensate returning from the C-IHX, along with the interplay of control actuation on the TBV and FWCP, necessitate a difficult balancing act to keep the feedwater temperature constant. In addition, as was discussed earlier, at high steam diversion ratios the valve operates close to its

fully closed position, limiting its ability to optimally counteract temperature overshoots. Future designs should consider a method of steam rejection from the condensate so as to not only aid in meeting the steam diversion limits but also to enable improved feedwater temperature control.

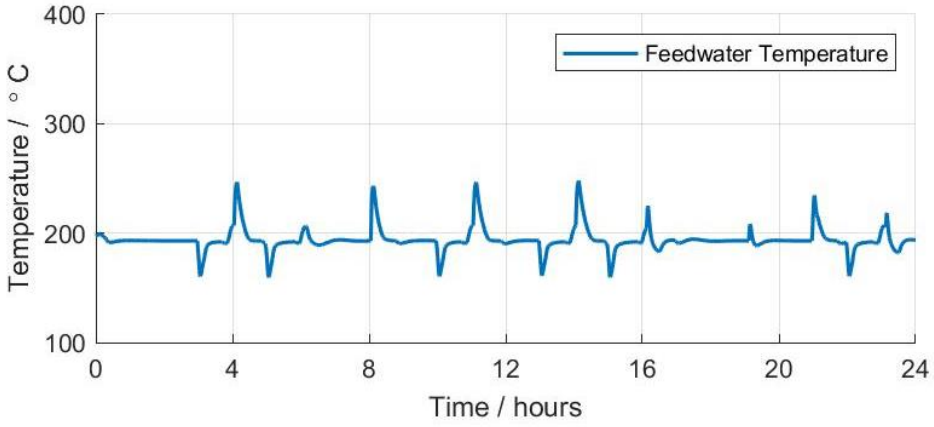

Figure A-10. Feedwater temperature throughout the 24-hour dispatch profile. Significant temperature spikes are observed upon state switching, especially when transitioning from low to high steam diversion ratios.

Figure A-11 shows the tank levels throughout this simulation. Overall, the system discharges an amount of heat similar to what it gains during charging. The initial HERON simulation estimated that these values should be within 10% of each other. The tank temperatures are shown in Figure A-12. The complex control methodology of the TES is seen to work well by maintaining tank temperatures within appropriate ranges yet allowing sufficient dispatch to meet the desired demand. The hot tank is the more volatile of the two, due to its coupling with the C-IHX, which experiences more changeable conditions than the simpler TES BOP on the D-IHX side.

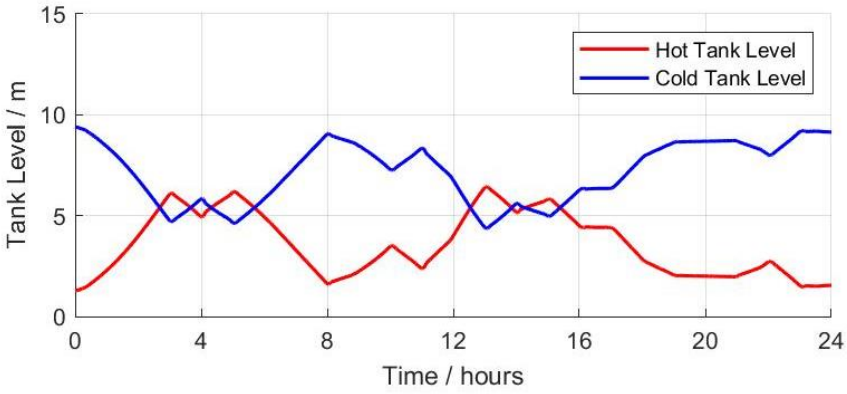

Figure A-11. Hot and cold tank levels throughout the 24-hour dispatch period.

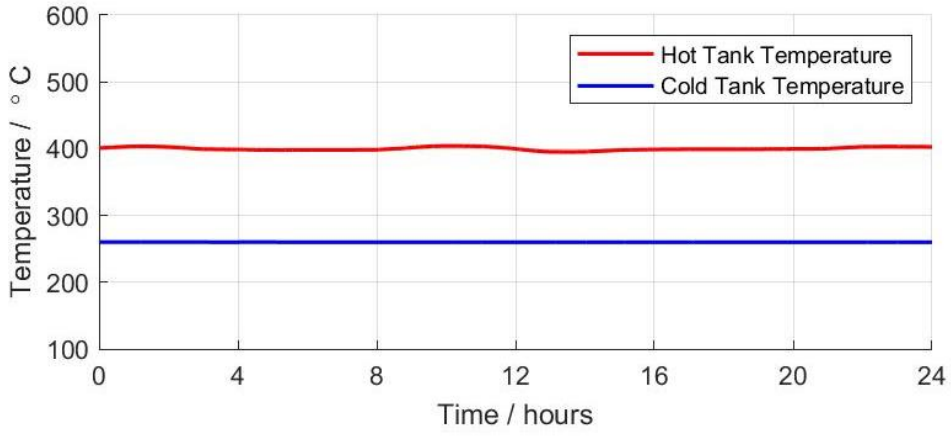

Figure A-12. The TES tank temperatures are kept constant throughout the 24-hour dispatch period.

Page intentionally left blank

APPENDIX B

Contents of Input File for Identification of Fourier Modes

This RAVEN input file was used to find Fourier modes that could be employed in the later input file utilized in the synthetic history generation process, as described in Section 5.1.1.

```
<Simulation verbosity="debug">
 <RunInfo>
  <WorkingDir>.</WorkingDir>
  <Sequence>read,process</Sequence>
 </RunInfo>

 <Files>
  <Input name="inFile">FFT_pointer.csv</Input>
 </Files>

 <Models>
  <PostProcessor name="fft" subType="FastFourierTransform">
   <target>price</target>
  </PostProcessor>
 </Models>

 <Steps>
  <IOStep name="read">
   <Input class="Files" type="">inFile</Input>
   <Output class="DataObjects" type="HistorySet">inData</Output>
  </IOStep>
  <PostProcess name="process" pauseAtEnd="False">
   <Input class="DataObjects" type="HistorySet">inData</Input>
   <Model class="Models" type="PostProcessor">fft</Model>
   <Output class="DataObjects" type="HistorySet">fft_out</Output>
   <Output class="OutStreams" type="Print">fft_final</Output>
  </PostProcess>
 </Steps>

 <OutStreams>
  <Print name="fft_final">
   <type>csv</type>
   <source>fft_out</source>
  </Print>
 </OutStreams>

 <DataObjects>
  <HistorySet name="inData">
   <Input>scaling</Input>
   <Output>price, hour</Output>
   <options>
    <pivotParameter>hour</pivotParameter>
   </options>
  </HistorySet>
  <HistorySet name="fft_out">
   <Output>price_fft_period,price_fft_amplitude</Output>
   <options>
    <pivotParameter>price_fft_frequency</pivotParameter>
   </options>
  </HistorySet>
 </DataObjects>

</Simulation>
```

Page intentionally left blank

APPENDIX C

Contents of a Synthetic History Generation Input File

The following input file is for synthetic history generation with a 24-hour segment length, as described in Section 5.1.1.

```
<?xml version="1.0" ?>
<Simulation verbosity="debug">
 <RunInfo>

<WorkingDir>Output_2018_2021_to2045_presvInpCDF_0_1_F1095_F168_Updt_fouryearrepeat</Wo
rkingDir>
  <Sequence>load, train, meta, serialize, sample, plotting</Sequence>
 </RunInfo>

 <Files>
  <Input name="input">../ARMA_pointer_fouryearrepeat.csv</Input>
  <Input name="pk">arma.pk</Input>
 </Files>

 <Steps>
  <IOStep name="load">
   <Input class="Files" type="">input</Input>
   <Output class="DataObjects" type="HistorySet">input</Output>
    <Output class="OutStreams" type="Print">test1</Output>
  </IOStep>
  <RomTrainer name="train">
   <Input class="DataObjects" type="HistorySet">input</Input>
   <Output class="Models" type="ROM">arma</Output>
  </RomTrainer>
  <IOStep name="meta">
   <Input class="Models" type="ROM">arma</Input>
   <Output class="DataObjects" type="DataSet">meta</Output>
   <Output class="OutStreams" type="Print">romMeta</Output>
  </IOStep>
  <IOStep name="serialize">
   <Input class="Models" type="ROM">arma</Input>
   <Output class="Files" type="">pk</Output>
  </IOStep>
  <MultiRun name="sample">
   <Input class="DataObjects" type="PointSet">placeholder</Input>
   <Model class="Models" type="ROM">arma</Model>
   <Sampler class="Samplers" type="MonteCarlo">mc</Sampler>
    <Output class="OutStreams" type="Print">test2</Output>
   <Output class="DataObjects" type="DataSet">synthetic</Output>
   <Output class="OutStreams" type="Print">synthetic</Output>
  </MultiRun>
  <IOStep name="plotting">
   <Input class="DataObjects" type="HistorySet">input</Input>
   <Input class="DataObjects" type="HistorySet">synthetic</Input>
   <Output class="OutStreams" type="Plot">plotters</Output>
  </IOStep>
 </Steps>

 <DataObjects>
  <PointSet name="placeholder">
   <Input>scaling</Input>
   <Output>OutputPlaceHolder</Output>
  </PointSet>
  <HistorySet name="input">
   <Input>scaling,Year</Input>
```

```
  <Output>price,hour</Output>
  <options>
   <pivotParameter>hour</pivotParameter>
  </options>
 </HistorySet>
 <DataSet name="synthetic">
  <Input>scaling</Input>
  <Output>price</Output>
  <Index var="hour">price</Index>
   <Index var="Year">price</Index>
 </DataSet>
 <DataSet name="meta"/>
</DataObjects>

<Models>
 <ROM name="arma" subType="ARMA">
  <pivotParameter>hour</pivotParameter>
  <Features>scaling</Features>
  <Target>price,hour</Target>
  <P>0</P>
  <Q>1</Q>
  <Fourier>8760,4380,2190,1095,168,24,12</Fourier>
  <Segment grouping="interpolate">
   <macroParameter>Year</macroParameter>
   <subspace pivotLength="24" shift="zero">hour</subspace>
      <evalMode>full</evalMode>
      <Classifier class="Models" type="PostProcessor">classifier</Classifier>
  </Segment>
  <reseedCopies>True</reseedCopies>
  <preserveInputCDF>True</preserveInputCDF>
  <seed>42</seed>
 </ROM>
  <PostProcessor name="classifier" subType="DataMining">
      <KDD labelFeature="labels" lib="SciKitLearn">
      <Features>price</Features>
      <SKLtype>cluster|KMeans</SKLtype>
      <n_clusters>20</n_clusters>
      </KDD>
  </PostProcessor>
</Models>

<OutStreams>
 <Print name="romMeta">
  <type>csv</type>
  <source>meta</source>
 </Print>
  <Print name="test1">
  <type>csv</type>
  <source>input</source>
 </Print>
  <Print name="test2">
  <type>csv</type>
  <source>placeholder</source>
 </Print>
 <Print name="synthetic">
  <type>csv</type>
  <source>synthetic</source>
 </Print>
  <Plot name="plotters" subType="SyntheticCloud">
      <training>input</training>
      <samples>synthetic</samples>
      <macroParam>Year</macroParam>
      <microParam>hour</microParam>
```

```
      <variables>price</variables>
    </Plot>
  </OutStreams>

  <Samplers>
   <MonteCarlo name="mc">
    <samplerInit>
     <limit>1</limit>
     <initialSeed>42</initialSeed>
    </samplerInit>
    <constant name="scaling">1</constant>
   </MonteCarlo>
  </Samplers>

</Simulation>
```