

Powering Data Centers with Clean Energy

INL/RPT-24-79307
Revision 0


Gateway for Accelerated
Innovation in Nuclear (GAIN)

A Techno-Economic Case Study of Nuclear and Renewable Energy Dependability

JUNE 2024

Gabriel J. Soto, Botros Hanna, Jakub Toman, Nahuel Guaita, Paul Talbot, Aaron Epiney,
Christopher Shawn Lohse

Idaho National Laboratory



DISCLAIMER

This information was prepared as an account of work sponsored by an agency of the U.S. Government. Neither the U.S. Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness, of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the U.S. Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the U.S. Government or any agency thereof.

Powering Data Centers with Clean Energy

A Techno-Economic Case Study of Nuclear and Renewable Energy Dependability

**Gabriel J. Soto, Botros Hanna, Jakub Toman, Nahuel Guaita, Paul Talbot, Aaron
Epiney, Christopher Shawn Lohse
Idaho National Laboratory**

June 2024

**Idaho National Laboratory
Gateway for Accelerated Innovation in Nuclear (GAIN)
Idaho Falls, Idaho 83415**

<http://www.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

Rising data demands from artificial intelligence (AI) and large language models (LLMs) generating images, videos, and text have prompted increased need for larger and more robust data centers in the United States. Major companies interested in these larger data centers face the choice of linking them to existing regional grids, building stand-alone power supplies onsite, or a combination of both. The request, review, and approval process for new transmission lines to grids in the United States, however, has grown in recent years to times spans rivaling those of new construction for nuclear power plants. Building an islanded power supply for each data center is therefore becoming a prominent option.

In this case study, several technologies are modeled in techno-economic simulations for long-term system costs subject to fixed electricity demand from a singular data center. A 250 MWe data center is assumed with additional 50 MWe for resiliency. Techno-economic simulations are conducted using the Holistic Energy Resource Optimization Network (HERON) software, which is a part of the Framework for Optimization of Resources and Economics (FORCE) tool suite. Technologies considered include solar, wind, lithium-ion batteries, and several types of nuclear reactors: large-scale reactors, small modular reactors, and microreactors. A low- and high-cost estimate for each technology is assumed to develop a range of expected economic performance. Low-cost estimates included several clean energy production tax credits. Different combinations of renewable energy generators with nuclear reactors are considered, ranging from a fully renewable-powered data center to a fully nuclear-powered data center. Historic time series of wind and solar availability from the Texas grid are used to train a reduced order model; this model then generates unique time series with similar characteristics of the training dataset. Multiple scenarios of weather and subsequent operations are simulated for each renewable-nuclear combination to determine total costs throughout the project lifetime.

Fully renewable-powered configurations required large amounts of installed capacity (GW scale) in the simulations to meet the fixed demand of the data center. This is due to some scenarios in the historical dataset which captured low-wind and low-solar days, requiring over-building of these technologies as well as batteries to compensate for the low amounts of electricity generation. Fully nuclear-powered configurations outperformed the fully renewable and mixed renewable-nuclear configurations in terms of cost, with ranges between \$1B and \$10B in 2023 USDs compared to \$40B+ for fully renewable configurations. Of the nuclear technologies, small modular reactors performed better economically than large-scale nuclear models due to lower projected capital costs, and both performed better than the microreactor models. These results demonstrate the applicability of firm, dispatchable electricity resources from baseload generators like nuclear power plants for operating facilities that run at constant power without daily variability.

Page intentionally left blank

CONTENTS

ABSTRACT.....	ii
ACRONYMS.....	viii
1. Introduction.....	1
2. Technology Selection and Economics	2
2.1. Data Centers and Reliability Requirements	2
2.2. Renewable Energy Generators and Storage	3
2.3. Nuclear Power Plants	3
2.3.1. Large Nuclear Power Plant	3
2.3.2. Small Modular Reactors.....	4
2.3.3. Microreactors	4
2.4. Costs and Other Economics	4
2.5. Import and Export Components	5
2.6. Methodology for TEAs	5
3. Results.....	6
3.1. Optimal Costs.....	6
3.2. Optimal Technology Capacities	9
3.2.1. Dispatch Analysis	10
3.2.2. Cost Breakdown.....	12
4. Land Usage and Grid Connection Considerations	15
5. Discussion and Conclusion	17
A.1 ERCOT Training Data	20
A.2 Training Methodology	20
A.2.1 Time Series Analysis Algorithms	20
A.2.2 Segments and Clustering.....	21
A.2.3 Generation of Synthetic Histories	21
A.3 ERCOT Synthetic History Results.....	21
A.4 Bi-Level Optimization	23
A.4.1 Methodology	23
A.4.2 Component and Cashflow Definitions	24
A.4.3 Dispatch Optimization Assumptions.....	25

FIGURES

Figure 1. Range of total costs, in 2023 USD, for 20-yr operation and construction of different nuclear-VRE generation mixes each powering a 300MWe data center. Low- and high-cost estimates were used to create upper and lower bounds.	7
Figure 2. Upper and lower bound costs for 20-yr operation and construction of nuclear-VRE generation mixes as a fraction of the highest and lowest VRE total cost, respectively.	8
Figure 3. Range of total costs. In 2023 USD, for 20-yr operation and construction of different nuclear-VRE generation mixes each powering a 300MWe data center; zoomed in for the 75% and 100% nuclear configurations.	9
Figure 4. Optimal capacities for each of the 11 nuclear-VRE generation mixes. Uses upper bound cost estimates for component costs.	10
Figure 5. Example synthetic history showing a low-wind and low-solar day.	11
Figure 6. Optimal dispatch strategy for the low-wind and low-solar scenario.	12
Figure 7. Cash flow inflows, outflows, and net values for a single realization of the VRE-only configuration for the full project lifetime, using upper bound costs.	13
Figure 8. Cumulative cash flows for full project lifetime as a pie chart for the upper bound costs of VRE-only configuration.	13
Figure 9. Cash flow inflows, outflows, and net values for a single realization of the 50-50 SMR-VRE configuration for the full project lifetime, using upper bound costs.	14
Figure 10. Cumulative cash flows for full project lifetime as a pie chart for the upper bound costs of the 50-50 SMR-VRE configuration.	15
Figure 11. Acres required for installed capacity of each configuration, broken down by technology.	16
Figure 12. Generated synthetic time series or history for wind and solar utilization factors, respectively, as a function of time.	22
Figure 13. Distributions of generated synthetic history for wind and solar utilization factors, respectively, as a function of cluster.	23
Figure 14. Bi-level optimization scheme used in HERON with multiple realizations in the inner levels [diagram originally from Ref. 17].	24

TABLES

Table 1. Macroeconomic and data center cost parameters for the simulations.....	4
Table 2. Summary of costs, both lower and upper bounds, for each technology including references.....	5
Table 3. Summary of nuclear-VRE combinations and the different TEA that were conducted. L+H symbolizes that simulations with low and high costs were separately conducted.	6
Table 4. Assumed number of acres needed per installed MW capacity for different technologies.....	15
Table 5. Parameters of TEA simulations used in HERON for time indexing and dispatch optimization.	24

Page intentionally left blank

ACRONYMS

AI	artificial intelligence
ARMA	auto regressive moving average
CAPEX	capital expenditures
CDF	cumulative distribution function
ERCOT	Electric Reliability Council of Texas
FORCE	Framework for Optimization of Resources and Economics
HERON	Holistic Energy Resource Optimization Network
INL	Idaho National Lab
ITC	investment tax credit
KDD	Knowledge Discovery in Databases
LNPP	large nuclear power plant
LLM	large language models
ML	machine learning
NP	nuclear power
NPP	nuclear power plant
NPV	net present value
O&M	operations and maintenance
PTC	production tax credit
RAVEN	Risk Analysis Virtual Environment software
ROM	reduced order model
SMR	small modular reactor
TEA	techno-economic analysis
TSA	time series analysis
U.S.	United States
USD	United States dollar
VRE	variable renewable energy

Page intentionally left blank

Powering Data Centers with Clean Energy: A Techno-Economic Case Study of Nuclear and Renewable Energy Dependability

1. Introduction

With the advent of large language models (LLMs) and the expansion of more advanced artificial intelligence (AI) and machine learning (ML) algorithms, modern data centers' computational needs—and therefore their energy requirements—have grown exponentially in the United States (U.S.). Data centers will grow in number and size in the coming decade to keep up with requests for auto image generation and other prompted, generative AI. Increases in electricity demands from these data centers will strain regional electricity regulators as future centers, ranging from 200 MWe to the GWe scale, are built in the coming decade [1]. Growth projections of data centers in the U.S. place their percentage of total grid demand between 4.6% and 9.1% of 2030 electricity consumption [2]. Transmission lines to these new data centers can be built but setting up a grid connection for electricity supply is not favorable due to many factors, including cost and length of time necessary to connect these levels of load to the grid, as well as the reliability requirements of these data centers. An islanded operation—one where the power generator is co-located with the data center—is a potential alternative, given that the supply of electricity upholds strict reliability standards. Many data centers require as much as 99.999% uptime for normal operation to avoid damage to data and prevent economic losses for customers [3,4]. Moreover, this growth stage for data center power configuration is a great opportunity to implement a zero-emission, clean energy design to avoid further greenhouse gas emissions from a growing industry. Companies interested in increasing data center production, such as Google, Microsoft, and Amazon, have all made pledges to achieve net-zero emissions or negative emissions between 2030 and 2040 [5–7]. Clean energy solutions to powering data centers will aid these companies in meeting their climate pledges and help the U.S. meet its general net-zero emissions target by 2050 [8].

Requiring zero-emissions clean energy solutions for data centers' electricity demands limits the options of available generation types in the U.S. Localized solutions for hydroelectric or geothermal power could be implemented, but these restrict siting options. More universal solutions include variable renewable energy (VRE) generation via technologies such as photovoltaic solar, offshore and onshore wind turbines, and energy arbitrage via batteries. However, the intermittency of renewable energy generation also poses challenges to the reliability standards of data center electricity demand, even with the use of battery storage. Nuclear power plants (NPPs), with their zero-emission baseload operation, seem better poised to supply the energy needs of these new data centers [9].

Some considerations include costs, reliability, land usage, and construction time. We address the issues of costs and reliability in this report via techno-economic analysis (TEA) of the different nuclear and renewable energy technologies and how effectively they meet data center energy demand using the Framework for Optimization of Resources and Economics (FORCE) [10–12]. TEA has commonly been conducted via the FORCE toolset at Idaho National Lab (INL) for integrated energy systems, where a mixed portfolio of VRE and nuclear power work together to meet regional demand for commodities—e.g., electricity, hydrogen, thermal energy—and sell these commodities to established markets at fluctuating prices. A bi-level optimization is conducted where the inner level optimizes resource utilization or dispatch according to technical constraints and the outer level optimizes component (e.g., power plant, wind farm) sizes. Uncertainty in the demand, electricity prices, and weather conditions for VRE operations is incorporated via reduced order models (ROMs) trained on historical time series. These models generate synthetic time series or histories sharing statistical characteristics of the training time series dataset while still being unique scenarios or realizations. The inner level is expanded to optimize resource utilization over

N unique realizations; this creates a distribution of economic outcomes over all realizations and the outer level can instead optimize the expected value (or some other desired statistic) of economic performance.

TEA for a data center powered by nuclear, VRE, or a combination of both, can be conducted using the FORCE toolset, albeit with some differences as compared to previous studies. Here we do not consider market dynamics, nor do we consider cash inflows from commodity sales. Rather, we focus primarily on capital and recurring operations and maintenance (O&M) costs from the building and operation of the various electricity generating components for comparison studies. The bi-level optimization is better described as a cost minimization problem with electricity demand constraints. The trained model for synthetic time history generation is purely trained on weather conditions for wind and solar availability. The outcome of these TEA simulations therefore answers the question: what is the optimal combination of technologies—different types of nuclear plants and VRE coupled with storage—to reliably meet data center demand at the lowest cost? For this comparison study, it is assumed that the technologies must meet 100% of the data center demand with extra capacity for resilience. Finer resolution reliability studies for the 99.999% uptime requirements and the effects of having multi-unit nuclear reactors for refueling and maintenance outages is a subject of potential future work.

In this report we conduct a comparison TEA study of different clean energy solutions for data center electricity demand. An overview of all technologies considered in the comparison study is presented in Section 2. Results of the TEA and simulations are detailed in Section 3. Some considerations for follow-up studies are proposed in Section 4. A final discussion and conclusion of the current work is summarized in Section 5. Additional information on assumptions and more technical aspects of the simulations are included in the Appendix. In Appendix A.1 and A.2, we give an overview of the selected training dataset for scenario optimization and the training of the synthetic history ROM, with results of the generated time series datasets in A.3. In Appendix A.4 we give an overview of the bi-level optimization strategy conducted via the FORCE toolset as well as some assumptions made.

2. Technology Selection and Economics

A major strength of the FORCE toolset TEA capabilities is the incorporation of uncertainty within the second, inner level via synthesis of unique time series scenarios. These synthetic time series or histories are used to create a distribution of economic performances with imposed technical constraints of the physical systems being modeled in the simulations. In this study, a ROM was trained on datasets from the Electric Reliability Council of Texas (ERCOT), which maintains and regulates a competitive wholesale market of electricity from different generators and consumers. Datasets for multiple years are available for load demand and historic pricing; for the purposes of this study, datasets on solar and wind utilization are used to train a ROM for the subsequent TEA simulations.

Several technologies are considered in this case study as dependable power supply sources for the data center. It is assumed that the data center will not have an outside connection to or from the regional grid. A combination of VRE generators, different types of nuclear plants (existing or advanced designs), and battery storage for electricity, are studied within several TEA simulations. Assumptions and designs are outlined below; more information on how the simulations are carried out in the Holistic Energy Resource Optimization Network (HERON) can be found in Appendix A [13–17]. For all technologies, a lower and upper bound for costs are determined to create a bound of expected techno-economic performance of the system. Costs are based on 2023 USD.

2.1. Data Centers and Reliability Requirements

As demand for cloud and AI services increases across the country, data centers must keep up. Reliability is a key requirement for data centers, as any interruptions or outages will lead to damage of data and

economic losses (approximately \$9,000 per minute) for unplanned outages, according to one study [4]. In this study we consider a data center requiring 300 MWe of capacity. We consider a design where 250 MWe of electrical power is required for operations and an extra 50 MWe is added for resilience and redundancy. Lower and upper bound cost estimates for the data center are provided in Table 1. The project lifetime of 20 years was selected assuming it matches a typical data center’s lifetime. Only capital expenditure (CAPEX) cash flows were considered for the data center [18]. Follow-on studies would benefit from better data on expected cash inflows from specific data centers.

2.2. Renewable Energy Generators and Storage

Several VRE technologies were considered in this case study, including solar and wind power. These ensure a carbon-free energy supply for the data center. Solar generation is assumed to be utility-scale photovoltaic; wind generation is assumed to come from an onshore wind farm. In the outer optimization of HERON, the capacities for these are either optimized or iterated in a parametric sweep; within the inner dispatch optimization, these nameplate capacities are modified based on a time-dependent capacity factor. This capacity factor, for wind and solar respectively, is an hourly synthetic history or time series sampled from the trained ROM, examples of which are shown in Section A.3. In this way, the production of electricity from either of these VRE components is dependent on synthesized weather conditions trained on historic data from Texas. CAPEX and fixed O&M costs were considered for both technologies, the latter being applied as recurring yearly cash flows. The full capacities of these technologies (not modified by capacity factor time series) were used as the drivers for both cash flows.

Energy storage in this case study is modeled using battery storage for electricity. Lithium-ion batteries were selected as the baseline model for this component in the TEA simulations. CAPEX and fixed O&M costs were considered for battery operation, the latter applied as recurring yearly cash flows. An assumed lifetime of 10 years required an additional reconstruction CAPEX cashflow midway through the project lifetime. Additionally, a round-trip efficiency loss of 0.86 was applied to the batteries. Initial storage levels for the batteries were assumed to be 100%, with periodic boundary constraints requiring them to be recharged to full levels by the end of the simulated time segment.

2.3. Nuclear Power Plants

NPPs, with their typical baseload operation, were selected as an additional clean, carbon-free energy supply for the data center. In this case study, nuclear power is considered in combination with VRE and as a stand-alone option for data center supply. Multiple options for NPPs are selected within this case study: a large nuclear power plant design (based on existing technologies), small modular reactors (SMRs), and microreactor designs. Cost estimates for these reactor designs differ based on their technology readiness levels; these ranges are considered by using lower and upper bound estimates and simulating operations at both cost levels or each design and combination with VRE. For each nuclear technology, only the electrical power output was considered, assuming that it would be coupled with a turbine and generator. The conclusions from these TEA simulations are meant to be a demonstration of performance differences between technologies and not a true reliability metric for typical data center uptime requirements (99.999%). Finer resolution studies would be a subject of potential future work to determine an optimal number of units subject to expected outages due to refueling and maintenance.

2.3.1. Large Nuclear Power Plant

Based on available cost numbers, the Vogtle NPP was used as the model for the large NPP (LNPP) technology in this case study. Cost numbers were taken from a 2015 study and converted to 2023 USD

[19]. The net output of the Vogtle plant is 2200 MWe; a scaled-down version would be assumed for this case study at different MWe output levels. CAPEX, fixed O&M, and variable O&M costs were considered in this study, where the variable O&M are applied hourly using power generation as the cash flow driver. A fixed dispatch strategy was assumed for power generation. The 60-year lifetime of the large NPP, being longer than the project lifetime, did not require an additional CAPEX cash flow apart from the first construction.

2.3.2. Small Modular Reactors

Small modular reactors (SMRs) are an advanced reactor design currently in development by several American companies that plan to deploy the technology in the U.S. and Europe in the next decade [20–23]. Since the design has yet to be deployed, the cost estimates for construction and operations are projections and range in values depending on the source. CAPEX and variable O&M costs were taken from a literature review of advanced reactor designs from INL [24]. Here lower and upper bounds were selected for individual analyses. These reactors were assumed to be dispatchable in their electricity production, and so a fixed dispatch strategy was not used. Minimum power operation for these units was assumed to be 30% of their nameplate capacity. Since SMRs are deployed as units, it was assumed that a single unit had a capacity of 150 MWe. Different numbers of units were considered in combination with VRE in this study.

2.3.3. Microreactors

Microreactors are another advanced reactor design being developed at INL for smaller-scale, localized deployment. Cost estimates, similar to those of the SMR, are currently projections based on several studies and have been summarized in a previous literature review. Microreactors are also assumed to be dispatchable in their electricity production (i.e., they produce electricity flexibly). Units of microreactors were assumed to be approximately 3 MWe, so multiple units would need to be deployed together to power the data center even when combined with VRE. The minimum power generation for each unit was also assumed to be 30% of its capacity.

2.4. Costs and Other Economics

Some macroeconomic parameters—such as tax rate, inflation rate, and discount rates—are summarized in Table 1. Tax and inflation rates from a previous INL study were used [25]. The discount rate was taken from the same study and calculated based on weighted average cost of capital for nuclear [25].

Table 1. Macroeconomic and data center cost parameters for the simulations.

Data Center Design Power	Data Center CAPEX, [\$/kWe] [Ref 18]		Project Lifetime	Discount Rate	Tax Rate	Inflation Rate
	<i>Low</i>	<i>High</i>				
300 MW	7,000	12,000	20 years	10.2%	28%	3%

All cost values for the individual technologies considered are summarized in Table 2. These include lower and upper bounds for CAPEX, fixed O&M (where applicable), and variable O&M (where applicable). All costs are given in units of 2023 \$/kW; units for the battery costs are in 2023 \$/kWh. Variable O&M costs are not considered for wind, solar, and battery operation; fixed O&M are not considered for the SMRs and microreactors.

Table 2. Summary of costs, both lower and upper bounds, for each technology including references.

	<i>Low</i>	<i>High</i>	<i>Low</i>	<i>High</i>	<i>Low</i>	<i>High</i>	
	CAPEX [\$/kW]		Fixed O&M [\$/kW-yr]		Variable O&M [\$/kWh]		Reference
Wind	1,025	1,700	20	35	-	-	[19]
Solar	700	1,400	7	14	-	-	[19]
Battery	301*	436*	8	19	-	-	[26]
Large NPP	8,475	13,925	153		0.004	0.005	[19]
SMR	5,280	9,240	-	-	0.020	0.046	[24]
Microreactors	10,560	22,440	-	-	0.092	0.178	[24]

*\$/kWh

To represent a low and high estimate of cost values more accurately, additional adjustment factors were applied based on tax credits from the Inflation Reduction Act. Since these tax credits would reduce costs, the upper bound costs in Table 2 were not modified. The lower bound CAPEX and variable O&M were modified with an investment tax credit (ITC) and a production tax credit (PTC), respectively. An ITC of 40% for clean generators and a PTC of \$33/MWh were assumed as the largest tax credits for a final lower cost estimate.

2.5. Additional Components

While this case study considers a data center and power supply design that is not connected to the grid, the actual simulations do include “import” and “export” entities as a HERON producer and demand component, respectively. These represent an additional influx and outflux of electricity external to the components and are included to improve feasibility of optimization solutions. If, for example, the optimizer or the parametric sweep selects a combination of component capacities that falls short of meeting data center electricity requirements, the import component is available to supply the shortfall. Similarly, in the case of fixed operation, the export component is available as a sink for any additional production. Since these are strictly for feasibility concerns and should not be part of the final power design, usage of either the import or export components is discouraged within the simulations with prohibitively large cost values. Zero utilization of import production was used as a criterion for selecting a feasible combination of technologies that could reliably power the data center.

2.6. Methodology for TEAs

Simulations of reliable data center power supply were conducted for different generation mixes as outlined in Table 3. The generation mix was determined based on a sliding scale of incorporated VRE as a percentage of the full data center demand, assumed to be 300 MWe including redundancies. A total of 11 combinations of nuclear-VRE were considered, each simulated separately with low and high costs. Supply from VRE is assumed to be a combination of solar, wind, and battery.

Table 3. Summary of nuclear-VRE combinations and the different TEA that were conducted. L+H symbolizes that simulations with low and high costs were separately conducted.

Nuclear Technology	Percentage of Supply from VRE (Remainder from Nuclear)				
	0%	25%	50%	75%	100%
LNPP	L + H	L + H	L + H	L + H	L + H
SMR	L + H	—	L + H	—	
Microreactor	L + H	L + H	L + H	L + H	

The 0% VRE column models systems where the entire 300 MWe is supplied by either one LNPP, two SMR units, or 100 microreactors without battery storage. Due to the selected size of SMRs, only a combination of a single 150 MWe unit with the other 50% of data center capacity supplied by VRE is considered. For the LNPP and microreactor options, this case study considers combinations with 25, 50, and 75% of the 300 MWe data center powered by VRE, with the remaining percentage covered by nuclear technology. Additionally, a pair of simulations for low and high costs was conducted with 100% VRE coverage of the data center electricity demand.

For each of the nuclear-VRE combinations in Table 3, the nuclear technology and data center sizes were fixed a priori. The HERON simulations were then posed to determine the optimal size of VRE required to meet the remaining electricity demand from the data center that was *not* already met by the NPPs. For each iteration of VRE sizes, 50 stochastic scenario realizations (as noted in Table 5) were conducted to find the optimal dispatch of electricity for different scenarios of wind and solar availability sampled from the trained synthetic history ROM.

3. Results

Simulations of system operations and economic assessments of the different compositions of nuclear-VRE were conducted using HERON optimization mode. A gradient descent algorithm was chosen to find optimal capacities of VRE components which both meet the data center electricity demand constraint (without need for import or export of electricity) while minimizing costs. Results are summarized in the next subsections.

3.1. Optimal Costs

HERON simulations of the 11 different nuclear-VRE combinations resulted in an expected value of NPV for the specified project lifetime of 20 years. Expected values of NPV for each nuclear-VRE combination were computed in this manner using a low estimate and a high estimate of each component cost. These results are illustrated in Figure 1. Data points on the far left-hand side of the plot represent a range of total costs for a data center powered 100% by VRE. This range is bounded by the two expected values of NPV, one of which uses all the lowest estimated costs for solar, wind, and battery components (including ITC reductions for CAPEX costs) as the lower bound and the highest estimated costs for all VRE components (without any tax credits) as the upper bound. The NPV here is instead labeled as a cost since positive cash flows (from data center operations, for example) are not considered. A gradient descent algorithm was used to find the mix of VRE technologies that led to the lowest total cost throughout the project's lifespan. The lowest VRE-only configuration cost that reliably met the data center demand was approximately \$40B in 2023 USD; the highest cost was larger than \$140B.

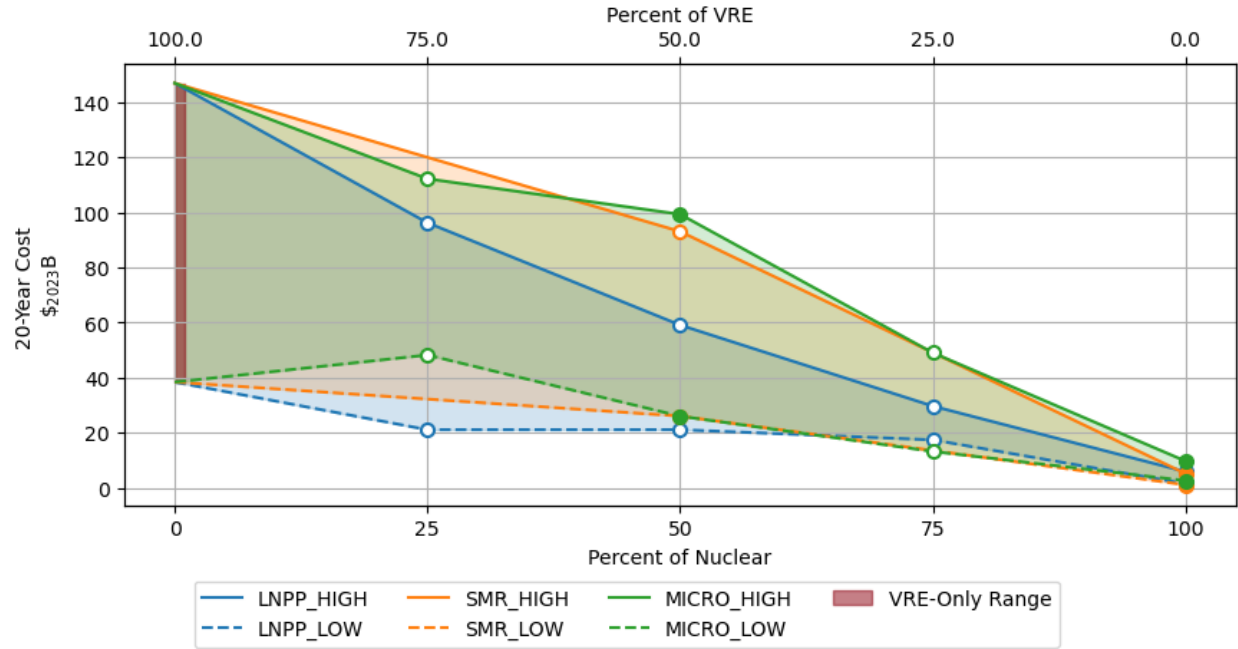


Figure 1. Range of total costs, in 2023 USD, for 20-yr operation and construction of different nuclear-VRE generation mixes each powering a 300MWe data center. Low- and high-cost estimates were used to create upper and lower bounds.

The remaining values of Figure 1 represent total costs for configurations that include nuclear technologies as a method for powering the data center in varying proportions of the required demand. Data points with solid circles represent simulations where the solutions converged with $N=50$ inner samples; others were only able to converge in fewer inner samples due to project time constraints (most in $N=25$, others in $N=47$). Four simulations were conducted for a 25% nuclear, 75% VRE configuration (as a percentage of the total 300 MWe required by the data center). Two of those four simulations model a LNPP, and the others model a microreactor as the nuclear component. There were no simulations for SMR at this proportional level since power output would be less than the nameplate capacity of a single unit. A gradient descent algorithm was again used, this time to determine the optimal mix of VRE technologies that meet 75% of the data center demand (assuming that the nuclear component was sized to meet 25% of the demand). Simulations using a LNPP resulted in lower total costs than the VRE-only cases for both the high-cost and low-cost estimates: about a 30% and almost 50% cost reduction for the upper and lower estimates, respectively. Costs relative to the VRE-only case results are shown in Figure 2. Simulations using microreactor power for 25% of the total demand also resulted in a lower total cost for the upper bound (approximately 25% cost reduction) but a higher cost for the lower bound (approximately 20% cost increase) than the respective VRE configurations.




Figure 2. Upper and lower bound costs for 20-yr operation and construction of nuclear-VRE generation mixes as a fraction of the highest and lowest VRE total cost, respectively.

Simulations using 50% or more nuclear power all resulted in lower total costs than both the upper and lower bound costs for VRE-only configurations as shown in Figure 1. Upper and lower bound costs when modeling 50% nuclear power are still generally prohibitive, ranging in the tens of \$B. When ranking the different nuclear technologies for the 50% nuclear configurations, costs for SMRs and microreactors are higher than those for LNPPs. The respective upper bound SMR and microreactor costs were between 30% and 40% lower than the VRE upper bound costs; the lowest was approximately 30% lower than lower bound VRE costs. LNPP upper bound costs were 60% lower than the highest VRE cost; LNPP lower bound costs were 40% lower than the lowest VRE cost.

Simulations using 75% nuclear power had further cost reductions compared to the previous configurations with higher proportions of VRE. Upper bound costs continued the same trend; lower bound costs for microreactors improved more than LNPPs, however, with a potential factor being higher capacities of VRE required due to uncertainty in the time series realizations. Here, upper bound costs using either nuclear technology range between 60% and 80% lower than the VRE-only upper bound costs; lower bound costs are between 50% and 70% lower than the VRE-only lower bounds.

Nuclear-only simulations, using high and low estimate costs, were conducted for all three nuclear technologies as shown in Figure 1 and in a zoomed in version of the same plot in Figure 3. Since no VRE technologies were considered for this set of simulations and the only source of uncertainty was due to synthetic histories of solar and wind availability, the techno-economic performance of these configurations was calculated with only a single inner realization or scenario. Here, microreactors have the highest ceiling in terms of cost for both the high and low-cost estimates. LNPP and SMR total costs were more similar in both the upper and lower bound results; SMRs here were more cost competitive than the LNPP. Lower bound costs for all nuclear technologies range between \$1B and \$2.5B, with upper bound costs for the LNPP and SMR configurations between \$5B and \$6B.




Figure 3. Range of total costs in 2023 USD, for 20-yr operation and construction of different nuclear-VRE generation mixes each powering a 300MWe data center; zoomed in for the 75% and 100% nuclear configurations.

3.2. Optimal Technology Capacities

Optimal capacities associated with all upper bound total cost results for all 11 nuclear-VRE configurations are shown in Figure 4. Results show that a large amount of VRE capacity was required for most configurations with non-zero VRE contributions, often more than 10-fold the required capacity of the data center. These large amounts of installed VRE—in different proportions of solar, wind, and battery installations—are necessary, mainly due to the constraint of resiliency in power supply to the data center. Power generation from these VRE sources depends on uncertain weather conditions, historical datasets which were used to train a ROM that generates unique instances of these conditions. Capacity planning under these constraints is highly sensitive to outliers in the simulated datasets: a cluster or day with low-wind *and* solar availability will require a large amount of installed wind and solar generation, as well as battery storage, to produce the required 300 MWe. The large amount of built VRE will then be carried throughout the rest of the project even during days when the installed capacity far exceeds requirements.

Some additional features within the optimal capacity results in Figure 4 include the nuclear capacity sizes and the mix of VRE technologies chosen by the optimization. The nuclear capacities required for the nuclear-VRE configurations *do* appear in Figure 4 but appear very small in comparison to the VRE required capacities since they are sized closer to the 300 MWe capacity demanded from the data center. The chosen mix of VRE capacities when coupled with nuclear also varies between nuclear technologies. For example, the 50% LNPP configuration uses a larger proportion of battery storage than wind generation when compared with the 50% SMR configuration and the 50% microreactor configuration. These selections may be due to some outlier events present in the synthetic history profiles: low wind and low solar events, particularly at the start of the day, may require larger capacities of battery storage to make up the remaining demand not already met by nuclear generators. However, because of the periodic boundary constraints, the total installed capacity of wind, solar, and nuclear must also be sufficient to charge the battery capacity back to its starting charge level. This constraint, while ensuring that battery resource levels are conserved

from day to day, may drive required capacities from VRE to large values particularly when an outlier event in the sampled scenarios necessitates those large capacities.




Figure 4. Optimal capacities for each of the 11 nuclear-VRE generation mixes. Uses upper bound cost estimates for component costs.

3.2.1. Dispatch Analysis

One example of how outliers in VRE availability profiles affect capacity selection is shown in a sampled synthetic history in Figure 5 and the associated optimal dispatch in Figure 6. The synthetic history shows hourly values for solar and wind availability for a selected cluster of the year. Solar availability peaks at a factor of 0.5 utilization and is only available for a limited duration of the day. Wind availability peaks at nearly 0.1 utilization of the installed capacity earlier in the morning, with a smaller range of availability than solar for the selected cluster. Combined, these two captured trends from the historical dataset represent a low-solar and even lower-wind scenario, which requires a larger amount of installed capacity to compensate.




Figure 5. Example synthetic history showing a low-wind and low-solar day.

A quick integration of these two curves gives an estimate for a capacity factor of energy production: approximately 2.5 and 0.3 utilization hours for solar and wind, respectively, for the entire day. Considering that the data center requires at least 250 MWe, or 6000 MWh for the day, the amount of pure solar generation that would need to be installed would be 2.4 GWe; for pure wind, 20 GWe. Additional storage would need to be installed to provide the necessary electricity when both solar and wind are not generating sufficiently at the start of the day. Due to the imposed periodic boundary constraints, however, HERON simulations require that the amount of storage at the end of the day matches the beginning of the day to ensure enough electricity is available for the beginning of the next day (for which the VRE capacity factors are not known ahead of time). This requires more generation from the solar and wind components to replenish battery-stored electricity to its starting levels, which can be strenuous during the low-wind and low-solar scenario days. Figure 6 shows an optimal dispatch during the same scenario from Figure 5 using 1.5 GWh of battery storage. The battery is required to power the data center for the first five hours of the day; afterwards, a mixture of solar, wind, and electricity imports is required to recharge the battery while also powering the data center for the remainder of the day. Solar and wind power were not sufficient in this case to be self-reliant. Ultimately, these results show better compatibility between a firm, dispatchable power generation from an all-nuclear configuration and the constant power consumption of data centers.




Figure 6. Optimal dispatch strategy for the low-wind and low-solar scenario.

3.2.2. Cost Breakdown

HERON simulations provide an option to run in a debug mode, which provides illustrative cashflow plots as well as the dispatch plot shown in Figure 6. Cashflow plots were generated for a few of the nuclear-VRE generation mixes. Inflows, outflows, and net cash flow for a full project year, though only for a single realization, are shown in Figure 7 for the VRE-only case using upper bound costs. CAPEX cash flows in the zeroth year dominate the budget sheet, with some fixed O&M costs each year and additional incurred CAPEX costs for batteries due to their component lifetime being 10 years, and therefore requiring reconstruction. In the legend of the figure, the “import” and “export” component cashflows are included in the list of all cashflows but do not show up in the actual graph. These represent the penalty cash flows that arise when the different configurations require excess energy to be imported to or exported from the system (in the case of the former, this would be when the configuration capacities are insufficient to meet the data center demand). For the same configuration, a pie chart in Figure 8 represents the cumulative cash flows. This shows that wind and solar CAPEX cash flows constitute most of the total incurred costs.

Yearly Inflows, Outflows and Net Cash Flow from Project Year 0 to 20




Figure 7. Cash inflows, outflows, and net values for a single realization of the VRE-only configuration for the full project lifetime, using upper bound costs.

Cumulative Inflows vs. Outflows from Project Year 0 to 20




Figure 8. Cumulative cash flows for full project lifetime as a pie chart for the upper bound costs of VRE-only configuration.

Comparatively, cashflow costs for the 50% SMR, 50% VRE configuration (using upper bound costs) in Figure 9 show a much larger proportion of the zeroth year costs dominated by the wind CAPEX cash

flows compared to SMR, battery, and data center cash flows and relevant fixed costs. The cumulative cash flow pie chart in Figure 10 further illustrates how wind CAPEX as well as fixed costs outweigh costs from other components, including SMR construction costs.

Yearly Inflows, Outflows and Net Cash Flow from Project Year 0 to 20

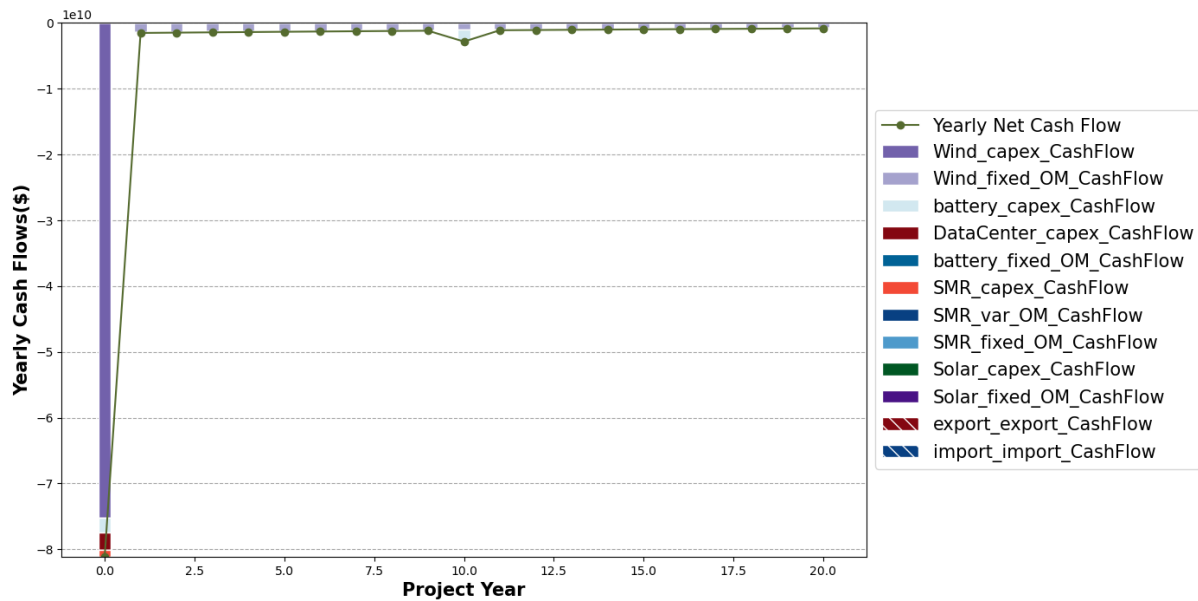


Figure 9. Cash flow inflows, outflows, and net values for a single realization of the 50-50 SMR-VRE configuration for the full project lifetime, using upper bound costs.

Cumulative Inflows vs. Outflows from Project Year 0 to 20

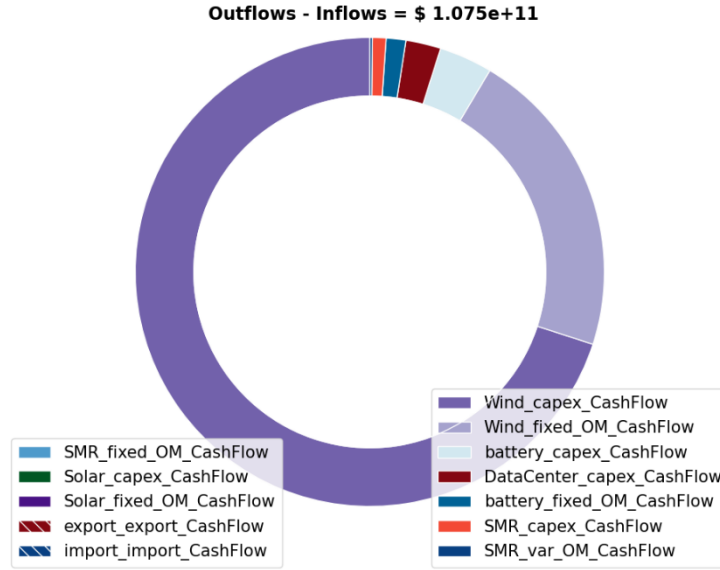


Figure 10. Cumulative cash flows for full project lifetime as a pie chart for the upper bound costs of the 50-50 SMR-VRE configuration.

4. Land Usage and Grid Connection Considerations

Multiple external factors impacting this case study should also be considered when gauging the suitability of different technologies. Land usage is an additional constraint for site selection when considering building VRE as opposed to nuclear plants. GW scale VRE installation of solar or wind could require many acres of land to adequately space apart wind turbines or rows of solar panels. Nuclear plants, on the other hand, require a much smaller footprint to provide the same amount of power generation, with a land use efficiency of 57,000 MWh/year per acre compared to 200 and 3100 MWh/year per acre for solar and wind, respectively [27]. Powering data centers with nuclear technology saves land and has less limited choice for siting locations due to smaller spacing requirements. An estimate of acreage required from each technology configuration was calculated using assumed values of acres/MW from literature and summarized in Table 4.

Table 4. Assumed number of acres needed per installed MW capacity for different technologies.

Technology	Acres/MW	Reference
Wind	2.47105	[28]
Solar	7.9	[29]
LNPP	0.543	[30]
SMR	0.038	[30]
Microreactor	0.06	[31]

These estimates for acres per installed capacity were then used to calculate total amount of acres required for each configuration, using the capacities from Figure 4. A breakdown of acres needed for wind, solar, and each type of nuclear technology are shown in Figure 11. Total acres for all technologies are also shown in a log scale to better demonstrate comparisons between VRE and nuclear requirements. These calculations show that nuclear land usage requirements are between three and five orders of magnitude lower than for VRE. Follow-up studies could use results from this case study to refine land usage requirements and costs for all configurations, then estimate total costs with additional incurred costs in different regions.




Figure 11. Acres required for installed capacity of each configuration, broken down by technology.

At the forefront of the technology selection issue is the question of whether to design an islanded, self-sufficient data center power supply or to connect the data center to the grid for its power supply. Construction times for nuclear plants have been known to range between 3 to 10 years in the United States before the Three Mile Island incident, and up to 15 years afterwards [32]. Construction times in other countries typically range from 5 to 10 years. Though still under development and construction by U.S. companies and manufacturers, modularization of SMRs and microreactors will improve construction learning rates and may further improve construction times [33]. Estimated construction times for SMRs in recent literature reviews have a median of 55 months [34]. In comparison, if considering a grid connection, grid regulators and utilities require entities to file interconnection requests, which include a series of studies and reviews to determine the impact of the new entity on grid requirements and the design and impact of transmission lines to said entity. Recent increases in interconnection requests have created interconnection queues due to limited resources to handle the influx of reviews. In the last decade, time elapsed between filing an interconnection request and beginning commercial operations has increased to 50 months on

average, with an increasing trend [35]. Data for these timelines was assembled for six different regional operators, including California, Texas, and New York. These expected times to commercial operations within different grids rival construction times of NPPs, but without the net-zero carbon emissions guarantee.

5. Discussion and Conclusion

Powering data centers with clean, reliable energy is essential for the net-zero future of the United States and the world. Assuming data centers use a power supply that is disconnected from the grid, comparison studies of different nuclear-VRE configurations used for power generation have shown that nuclear-only options are the least costly designs. The large, constant electricity loads of data centers are better served by firm, non-varying power supply from nuclear, especially when they take 100% of the demand in the power system. Low- and high-cost estimates for individual power generation and storage technologies resulted in ranges of total cost for the nuclear-only options between \$1B and \$10B USD. Low-cost estimates take into consideration some tax credits for clean energy production while high-cost estimates do not. Configurations that incorporate VRE into the energy supply mix increase total cost due to the need for excess capacity of VRE; outlier scenarios of simultaneous low-solar and low-wind availability require more capability than the required demand, including large amounts of battery capacity. VRE-only configurations are likely prohibitively costly for this reason when serving a single, localized data center with no grid connections. Timelines for grid connection request studies, subsequent approvals, and the start of commercial operations have trended upwards in the last few years, rivaling construction times of proposed NPPs with existing designs or advanced modular designs (e.g., the SMRs and microreactors studied here).

Some aspects of the presented case study could be improved in follow-up studies through the application of improved optimization strategies. The gradient descent optimization algorithm can sometimes encounter local minima of the objective function, which is at an increased risk considering the stochastic nature of the proposed problem. Multiple trajectories with different initial conditions can be computed to help mitigate this problem and attempt more exploration of the solution space, but with increased computational resources. An alternate optimization strategy, recently introduced to HERON simulations, is Bayesian optimization, which can more efficiently handle the trade-off between exploration of the solution space and finding the global optimum of the problem. With respect to the inner dispatch optimizations, the initial levels of storage in the battery were assumed to be 100%. In combination with the periodic boundary constraints, this may lead to an over-constrained problem, which would explain some difficulties in finding converging solutions without the need for importing excess electricity from external sources. Future simulations could either relax this constraint or, with a new feature of HERON dispatch optimization, allow the initial battery levels to be a variable within the problem. Alternatively, an expansive Monte Carlo outer level sampling of the domain of component capacities may be used to create a more visual representation of the solution space, including failure regions where the combination of technologies does not meet the required demand of the data center.

Other improvements to the case study rely on more accurate data center cost numbers, including potential estimates of revenue streams to ascertain economic performance in different regions. Expansion of this study to geographical regions other than the Texas market, with more localized historical datasets of wind and solar availability, would lead to more comprehensive analysis of the expected performance of these configurations across the United States. Another new feature in HERON that would improve these case studies is the ability to quantify other sources of uncertainty aside from the market time series scenarios from synthetic history ROMs: uncertainty in cost or cash flows. This uncertainty is represented as user-defined distributions and is simultaneously sampled with the ROM synthetic history generation in each inner realization. Quantification of economic uncertainty can be expanded to component cost numbers for nuclear, solar, wind, and battery. In the future, this can avoid the need to estimate upper and lower bound costs, instead sampling the full probability space of the problem such that the expected value of NPV

considers both market and economic uncertainty. Future studies could also provide a finer resolution reliability study for application of multi-unit nuclear technologies when considering outages due to refueling and how that would affect the data center uptime requirements of 99.999%.

Page intentionally left blank

A. Appendix

Regional Dataset Selection and Synthetic History Training

A.1 ERCOT Training Data

ERCOT has made publicly available datasets on solar and wind utilization dating to 2020 for hourly wind output and starting from 2021 for both solar and wind output. The datasets are for hourly aggregated wind and solar output for the entire ERCOT or Texas grid. Additionally, the datasets provide the total capacity installed for each renewable resource. We define an hourly utilization factor for wind and solar, respectively, as the hourly aggregated output normalized by the total capacity installed:

$$f_t^R = \frac{x_t^R}{C_t^R}$$

where x_t^R is the aggregated output of a resource R (either wind or solar) for time step t , C_t^R is the installed capacity for the resource at the given time step (this may change throughout the year), and f_t^R is the hourly utilization factor for the resource. The utilization factor here is used as an analog for the availability of wind and solar resources in the region.

A.2 Training Methodology

The FORCE toolset uses the Risk Analysis Virtual Environment software (RAVEN, recipient of the 2023 R&D 100 award) as the core engine for optimization, uncertainty quantification, and ROM training [36]. RAVEN has a dedicated time series analysis (TSA) module which hosts a variety of algorithms for characterization, transformation, and generation of time series [11,37]. These algorithms are used to sequentially preprocess and learn characteristics of the time series which are then used to generate unique synthetic time series with the same learned characteristics of the training dataset.

A.2.1 Time Series Analysis Algorithms

The main algorithm used for synthesis of time series is the auto regressive moving average (ARMA) algorithm [38]. The ARMA algorithm models a time series Y_t as a linear combination of p auto regressive terms and q moving average terms as shown below

$$Y_t = \mu + \sum_{i=1}^p \phi_i Y_{t-i} + \sum_{j=1}^q \theta_j \epsilon_{t-j} + \epsilon_t$$

where μ is the expected value of Y_t , ϵ_t is a random variable representing white noise, and ϕ_i and θ_j are coefficients used to fit the signal to the above model, depending on the set of (p, q) selected. The ARMA model captures an autocorrelation of a time series—its relationship with previous steps of the same time series—as well as previous error terms. ARMA models are commonly used in literature for time series prediction [39], but in previous studies at INL they have been used for generating synthetic time series based on the trained coefficients ϕ_i and θ_j [40]. The set (p, q) is typically user-defined, but recent developments in RAVEN have added an auto ARMA algorithm for optimally determining this set of hyperparameters [37].

The ARMA model assumes a stationary signal used for the time series Y_t ; some preprocessing algorithms are therefore typically performed before application of the ARMA algorithm on the time series signals. Detrending of the signal is performed via first fitting a Fourier series (major frequencies of which are determined from a fast Fourier transform) and then subtracting the fitted series from the signal. To

ensure stationarity, a “gaussianizing” step is further applied where the detrended signal is transformed into a stationary signal via its cumulative distribution function (CDF) and the inverse standard normal distribution CDF [40]. An additional preprocessing step is required for solar availability data due to its sharp cutoff during nighttime hours. A zero-filter algorithm is used to mask values of zeros so that subsequent algorithms (e.g., Fourier detrending) only act on the non-zero values. An “out-truncation” filtering algorithm is also used to limit the domain of values to positive numbers.

A.2.2 Segments and Clustering

Time series data from ERCOT is organized by yearly data files with hourly time resolutions specifically for the wind and solar availability (which we convert to utilization factors). In FORCE, the dispatch optimization is currently dictated by the time indexing of the trained synthetic history ROMs. Nominally, ROM training and subsequent dispatch optimizations would be conducted for all data points in the year; computational constraints, however, would render these types of simulations infeasible, particularly when expanding the number of realizations to the hundreds or thousands. One method for easing computation is to divide the yearly data into segments. Here, we implement 365 segments each 24 hours long to isolate and capture daily trends [15]. TSA algorithms can be applied on what is referred to as the “global” yearly signal (pre-division) or on a per segment basis. Fourier detrending is often applied on the global signal for frequencies longer than the segment length, the rest on the segments. The ARMA algorithm, however, is applied for each segment to determine the coefficients ϕ_i and θ_j based on the selected (p, q) hyperparameter set.

An additional method for reducing computation time is to cluster the time series segments into groups of similar characteristics per year. Often, trends and dynamics are repeated throughout the year such that the optimal dispatch of some days would be very similar. To exploit this, dispatch optimization can be conducted on a single representative segment of each cluster rather than on all segments. Training of the synthetic history ROM would still be conducted for all segments. After application of all TSA algorithms, a clustering algorithm—here, we use a Knowledge Discovery in Databases (KDD) classification process with k-means clustering algorithm from scikit-learn [41]—is used to group segments based on similar cluster-able features of each TSA algorithm capable of characterization. An example of a cluster-able feature would be the coefficients ϕ_i and θ_j from the ARMA algorithm. In this report, the number of dispatch optimizations required for a year of each realization is reduced from 365 (per segment) to 20 (per cluster). A multiplicity factor is tallied per cluster to track the number of segments represented by each cluster; this multiplicity is also used to ensure the correct amount of cash flows are represented in the final calculated economic metrics.

A.2.3 Generation of Synthetic Histories

The trained synthetic history ROM can be used for generating unique synthetic time series with the learned characteristics of the historical dataset. If any transformational algorithms are applied in the training process, their inverses are applied in the generation step in reverse order. That is, if in the training step the algorithm order was $A_1 \rightarrow A_2 \rightarrow A_3$, then in the generating steps the algorithms would be applied as $A_3^{-1} \rightarrow A_2^{-1} \rightarrow A_1^{-1}$. Only generating and transforming algorithms are used in the generation step.

A.3 ERCOT Synthetic History Results

The final trained synthetic history ROM was trained on the available ERCOT data from 2021 through 2023. A project life of 20 years is desired for the TEA simulations, so learned characteristics from 2023 were used for the missing years up to and including the year 2040. A generated synthetic history and the

distribution of values is shown in Figure 12. The synthetic history shown in the figure spans 20 years. Time steps up to a year are shown in the figure; for each hourly time step, minimum and maximum bounds are shown to represent the range of values throughout the 20 years. Mean values are also noted in the legend of each plot. On average, the wind utilization factor ranges from 0.13 to 0.57 with a mean of 0.35 of its installed capacity. The absolute maximum utilization or capacity factor is somewhere in the range between 0.6 and 0.8. Solar values are shown after removing zero values using the zero-filter algorithm (here, strictly positive non-zero values are plotted). These utilization factors range, on average, between 0.05 and 0.66 and the mean is 0.37. The solar capacity factors have a lower floor and higher ceiling than the wind availability, although it is worth noting that these values are only available during daylight hours.




Figure 12. Generated synthetic time series or history for wind and solar utilization factors, respectively, as a function of time.

The synthetic history results from the trained ROM are displayed in a different form in Figure 13. Here, values are arranged as a function of cluster per year, with the cluster index on the abscissa. For each cluster, the values for all years are collected and depicted as individual distributions through violin plots. As an example, cluster #0 of the wind utilization has a higher distribution of values, peaking around 0.65 with two horizontal lines bounding the absolute maximum and minimum values of the distribution through the 20-year span. A middle line also marks the expected value of the distribution per cluster.

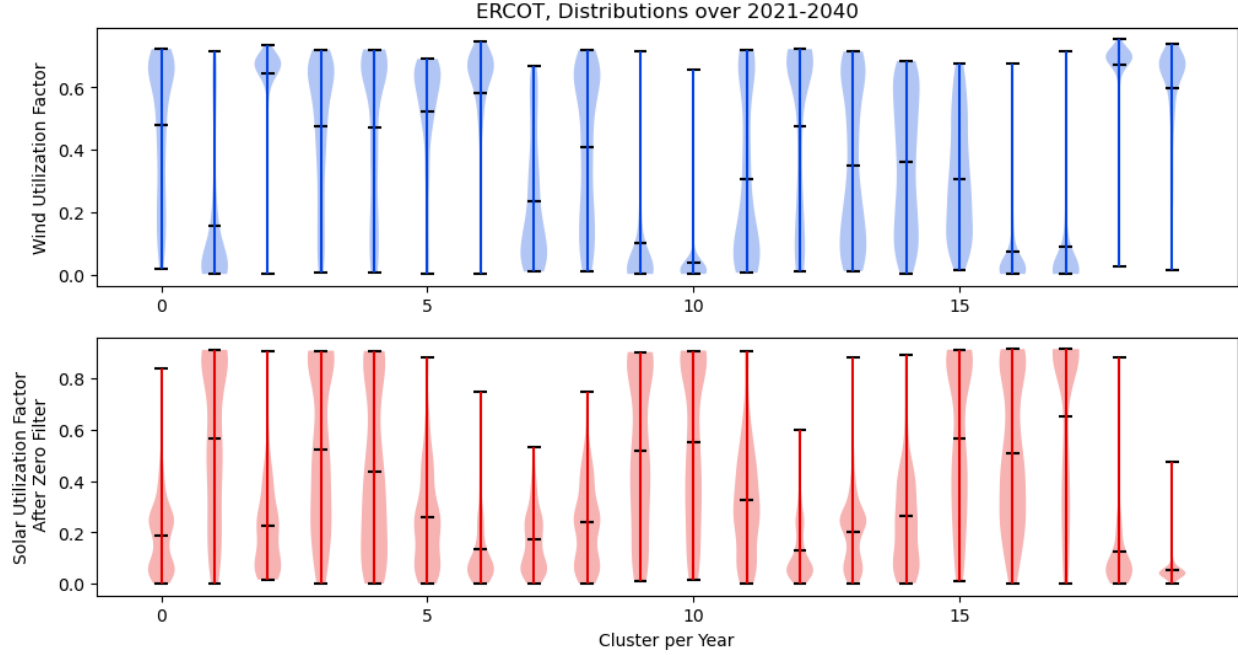


Figure 13. Distributions of generated synthetic history for wind and solar utilization factors, respectively, as a function of cluster.

Each cluster in Figure 13 illustrates similar characteristics of the segments determined by the clustering algorithm. For the wind utilization synthetic history, there are some clusters (namely 2, 18, and 19) where the distribution of values characterizes particularly windy days with high wind utilization. There are some clusters, however, with distributions peaking at low utilization factors—e.g., 10, 16, and 17 for wind and 6, 12, 18, and 19 for solar—which respectively represent low-wind and perhaps cloudy days in the historic data, characteristics of which have been learned and are now evident in the trained ROM. These scenarios, especially when the low-wind and low-solar availability events coincide, present a challenge for a fully renewable energy solution to reliable data center power supply.

A.4 Bi-Level Optimization

The bi-level optimization within the FORCE toolset is conducted using the HERON plugin of RAVEN. HERON handles the creation of RAVEN workflows for the two levels of optimization, which includes sampling of the synthetic history ROM and conducting dispatch optimization to determine optimal, time-dependent resource utilization to improve economic performance subject to technical constraints.

A.4.1 Methodology

Optimization conducted by HERON uses a leader-follower bi-level strategy where the outer level (leader) makes decisions that the inner level (follower) receives and uses to optimize its own decision-making; inner decisions are subsequently communicated to the outer level for further iterations. The outer level in HERON optimizes the sizes or capacities of the components in the modeled system; the inner level simulates N scenarios or realizations of the market or region and conducts dispatch optimization—subject to the capacities determined in the outer level—to determine optimal resource utilization per time step. A more detailed diagram of the optimization strategy is shown in Figure 14. The diagram shows a nested optimization problem where the inner optimization seeks to find the dispatch decisions which maximize the net present value (NPV), the selected economic metric in this case study. The inner level is repeated for

N realizations, each with a different Monte Carlo-sampled synthetic history scenario from the trained ROM. The NPV results from each inner run are collected to create a distribution of economic performance. The outer level then uses the expected value—the selected statistical operator for this case study—of the distribution as the objective metric to maximize with optimal values of component capacities. A gradient descent algorithm in RAVEN is used in the outer level which can handle stochasticity from the inner levels.

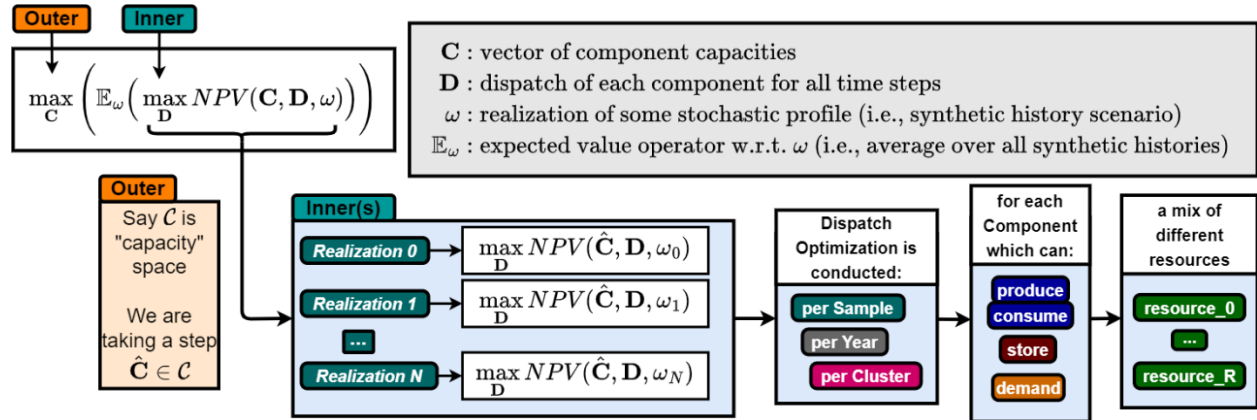


Figure 14. Bi-level optimization scheme used in HERON with multiple realizations in the inner levels [diagram originally from Ref. 17].

Time indexing for the dispatch optimization in the inner levels is determined by the trained synthetic history ROM. For this case study, as described in Section A.2.2, a full year is divided into 24-hour segments; these segments are further grouped into 20 clusters for each year. Dispatch optimization is therefore conducted on a representative segment randomly selected from each cluster: 20 dispatch optimizations per year per realization. These selected parameters are summarized in Table 5.

Table 5. Parameters of TEA simulations used in HERON for time indexing and dispatch optimization.

Number of Inner Realizations	50
Project Lifetime	20 yrs.
Segments per Year	365
Segment Length	24 hrs.
Segment Clusters per Year	20

An additional operational mode for the HERON TEA simulations, not shown in Figure 14, is referred to as a “sweep” mode. This enables a parametric sweep over variables in the outer optimization level (i.e., component capacities). The optimizer algorithm in this case is replaced with a grid sampler that iterates over a list of user-provided capacity values for components.

A.4.2 Component and Cashflow Definitions

Components in HERON are used to help the user represent a physical plant such as a nuclear power plant, wind turbine, electrolyzer, or even subsystems of these. Components are categorized based on how they interact with a set of user-defined resources: they are either a producer, storage, or demand component. Resources are tracked via resource pools in which components can either draw from or add to their supply.

Producer components produce a subset of resources; producer components can also consume a subset of resources to produce another subset subject to a user-defined linear or polynomial transfer function. Storage components can store a resource up to its capacity and then charge or discharge portions of the resource based on optimal decision-making. Demand components only consume resources and are used to model markets or other entities that demand a resource (e.g., electricity, thermal energy, hydrogen).

Cashflows are used to map resource-dependent component activities or capacities to economic values. A universal formula is used to define all cashflows:

$$\hat{F} = \alpha \left(\frac{d}{d'} \right)^\chi$$

where α is the representative price or cost, d is the driver of the cash flow, d' is the reference driver associated with the cost, and χ is a scaling factor representing economies or diseconomies of scale for $0 \leq \chi \leq 1$ and $\chi > 1$ respectively [17]. Three types of cashflows are used in the TEA simulations as well as this case study: one-time, recurring yearly, and recurring hourly. One-time cashflows are applied once per lifetime of the component and represent things like CAPEX. Drivers of CAPEX cashflows are typically the component capacity with some set cost value per unit of capacity. Recurring yearly and hourly cashflows repeat according to their specified timeframes, respectively. Drivers for each are often the component capacity and some determined resource dispatch activity, respectively. Taxes and inflation are optionally applied to the cashflows. A discount rate can also be introduced to determine the present value of future cashflows and is used to determine the final NPV for each inner TEA simulation.

A.4.3 Dispatch Optimization Assumptions

The inner dispatch optimization is defined using Pyomo expressions written in Python with additional technical constraints such as, but not limited to, conservation of all resources, transfer functions for producer components, ramp rates, and maximum and minimum operational bounds [42,43]. A cbc solver was used for the linear programming problem [44]. The objective function within each dispatch optimization (solved for each 24-hr cluster segment) was earlier described as being able to maximize the total NPV of the system: the objective function sums all hourly cashflows since any one-time or yearly cashflows will be constant in the problem. In this way, the inner optimization objective is a subset of the outer objective—the expected value of NPV over all realizations—and therefore more likely represents an optimistic leader-follower relationship where the goals of both levels are synergistic rather than antagonistic. One additional constraint recently added to the HERON Pyomo dispatch formulation is periodic boundary conditions. These specifically enforce constraints on storage levels for storage components. Initial levels of storage are a user input; with the periodic boundary constraints, the final storage levels must be equal to the initial levels to ensure continuity and resource conservation.

Page intentionally left blank

References

1. Rosendo, D., et al. 2019. "A methodology to assess the availability of next-generation data centers." *The Journal of Supercomputing* 75:6361-6385. <https://doi.org/10.1007/s11227-019-02852-3>
2. Aljbour, J., Wilson, T., and Patel, P., 2024. "Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption." *EPRI White Paper* 3002028905. <https://www.epri.com/research/products/3002028905>
3. Zakharov, A. V., Gusev, O. Y., and Cho, G. C., 2018. "Reliability Assessment of Data Centers Power System." *2018 IV International Conference on Information Technologies in Engineering Education* 23-26 October. <https://doi.org/10.1109/INFORINO.2018.8581778>
4. Ponemon Institute Research Report, 2016. "Cost of Data Center Outages." https://www.vertiv.com/globalassets/documents/reports/2016-cost-of-data-center-outages-11-11_51190_1.pdf
5. Google, 2022 "Accelerating Climate Action at Google and Beyond: A Progress Update." <https://www.gstatic.com/gumdrop/sustainability/google-2022-climate-action-progress-update.pdf>
6. Smith, B., 2020. "Microsoft will be carbon negative by 2030." 16 Jan 2020. Accessed 21 June 2024. <https://blogs.microsoft.com/blog/2020/01/16/microsoft-will-be-carbon-negative-by-2030/>
7. Amazon, 2022. "Reaching Net Zero Carbon by 2040: Measuring, Mapping, and Reducing Carbon the Amazonian Way." <https://sustainability.aboutamazon.com/carbon-methodology>
8. United States Department of State and the United States Executive Office of the President, 2021. "The Long-Term Strategy of the United States: Pathways to Net-Zero Greenhouse Gas Emissions by 2050." <https://www.whitehouse.gov/wp-content/uploads/2021/10/us-long-term-strategy.pdf>
9. Last Energy, "Nuclear Energy: A Path to Dependable Carbon-Free Data Centers," *WhitePaper*, 2023. <https://www.lastenergy.com/whitepaper/decarbonizing-data-centers#:~:text=to%20data%20centers,-.Nuclear%20Energy%3A%20A%20Path%20to%20Dependable%20Carbon%2DFree%20Data%20Centers,intensive%20computing%20and%20data%20applications.>
10. Frick, K., Talbot, P. W., McDowell, D. J., *Framework for Optimization of Resources and Economics Ecosystem*. Computer Software. <https://github.com/idaholab/FORCE>. USDOE Office of Nuclear Energy (NE). 9 Mar 2022. <https://doi.org/10.11578/dc.20220511.2>
11. P. Talbot, et al., "2022 FORCE Development Status Update", Idaho National Laboratory Technical Report, 2023, INL/RPT-22- 69275-Rev001 <https://www.osti.gov/biblio/1988136>
12. P. Talbot, et al., "2023 FORCE Development Status Update", Idaho National Laboratory Technical Report, 2023, INL/RPT-23-74915-Rev000 <https://doi.org/10.2172/2279026>
13. Talbot, P. W., et al., *HERON*. Computer Software, <https://github.com/idaholab/HERON>. USDOE Office of Nuclear Energy (NE). 28 Jul 2020. <https://doi.org/10.11578/dc.20200929.2>
14. Epiney, A., et al. 2020. "Economic analysis of a nuclear hybrid energy system in a stochastic environment including wind turbines in an electricity grid." *Applied Energy* 260:114227. [doi:10.1016/j.apenergy.2019.114227.](https://doi.org/10.1016/j.apenergy.2019.114227)
15. McDowell, D. J., et al., "A Technical and Economic Assessment of LWR Flexible Operation for Generation and Demand Balancing to Optimize Plant Revenue," Idaho National Laboratory Technical Report, 2021, INL/EXT-21-65443-Revision-1. <https://doi.org/10.2172/1844211>
16. Talbot, P. W., et al. "HERON as a Tool for LWR Market Interaction in a Deregulated Market". Idaho National Laboratory Technical Report, 2019, INL/EXT-19-56933-Rev000. <https://doi.org/10.2172/1581179>
17. Soto, G. J. and Talbot, P. W., "Economic Parameter Uncertainty Quantification Demonstration", 2024, INL/RPT-24-77193-Rev000. <https://www.osti.gov/biblio/2341277>
18. Zhang, M. "How Much Does It Cost to Build a Data Center" *Dgtl Infra*, 5 November 2023. Accessed: 21 June 2024. <https://dgtlinfra.com/how-much-does-it-cost-to-build-a-data-center/>

19. Bilicic, G. and Scroggins, S. 2023, “Lazard’s LCOE+ (April 2023)”. *Lazard*. Version 16.0.
<https://www.lazard.com/research-insights/2023-levelized-cost-of-energyplus/>
20. TerraPower. “Wyoming” Accessed: 21 June 2024. <https://www.terrapower.com/wyoming/>
21. TerraPower, “TerraPower Begins Construction on Advanced Nuclear Project in Wyoming” 10 Jun 2024. <https://www.terrapower.com/downloads/grounbreaking-press-release.pdf>
22. NuScale, “NuScale Power Joins the Romanian Atomic Forum, Supporting Romania’s Goal of Becoming a Leader in Advanced Nuclear Energy,” 24 May 2023. Accessed: 21 June 2024.
<https://www.nuscalepower.com/en/news/press-releases/2023/nuscale-power-joins-the-romanian-atomic-forum>
23. Reuters, “Romania’s Nuclearelectrica sees preliminary decision on SMR plant in 2025.” 18 Mar 2024. Accessed: 21 June 2024. <https://www.reuters.com/business/energy/romanias-nuclearelectrica-sees-preliminary-decision-smr-plant-2025-2024-03-18/>
24. Abdalla, A., et al., “Literature Review of Advanced Reactor Cost Estimates,” Idaho National Laboratory Technical Report, 2023, INL/RPT-23-72972-Rev000,
<https://doi.org/10.2172/1986466>
25. Guaita, N. and Hansen, J., K., “Analyzing the Inflation Reduction Act and the Bipartisan Infrastructure Law for Their Effects on Nuclear Cost Data,” 2024, INL/RPT-23-72925-Rev000.
<https://www.osti.gov/biblio/2335485/>
26. Mongird, K., et al., “Energy Storage Technology and Cost Characterization Report,” *Pacific Northwest National Laboratory*, 2019, PNNL-28866. <https://www.osti.gov/biblio/1573487>
27. Julie Kozeracki, et al. “Pathways to Commercial Liftoff: Advanced Nuclear”, US Department of Energy, March 2023.
28. Corrie, C., Frank, O., and Laurienti, M., “Land-Based Wind Energy Siting: A Foundational and Technical Resource,” National Renewable Energy Lab, 2021, NREL/TP-5000-78591.
<https://www.osti.gov/biblio/1812706>
29. Ong, S., et al., “Land-Use Requirements for Solar Power Plants in the United States,” National Renewable Energy Lab, 2013, NREL/TP-6A20-56290. <https://doi.org/10.2172/1086349>
30. Idaho National Laboratory, “Advanced Small Modular Reactors.” Accessed: 21 June 2024.
<https://inl.gov/trending-topics/small-modular-reactors/#:~:text=For%20example%2C%20a%20proposed%20920,would%20require%20nearly%20500%20acres.&text=SMRs%20complement%20other%20clean%20energy%20sources%20such%20as%20wind%20and%20solar.>
31. Owusu, D., Holbrook, M. R., and Sabharwall, P., “Regulatory and Licensing Strategy for Microreactor Technology,” Idaho National Laboratory, 2018, INL/EXT-18-51111-Rev000.
<https://doi.org/10.2172/1565916>
32. Lovering, J., Yip, A., and Nordhaus, T., 2016, “Historical construction costs of global nuclear power reactors.” *Energy Policy* 91:371-382. <https://doi.org/10.1016/j.enpol.2016.01.011>
33. Lloyd, C. and Roulstone, A., “A Methodology to Determine SMR Build Schedule and the Impact of Modularization,” *Proceedings of the 2018 26th International Conference on Nuclear Engineering*. <https://doi.org/10.1115/ICONE26-81550>
34. Abdalla, A., et al., “Meta-Analysis of Advanced Nuclear Reactor Cost Estimations,” Idaho National Laboratory, 2024, INL/RPT-24-77048-Rev001. <https://doi.org/10.2172/2371533>
35. Rand., J., et al., “Queued Up: 2024 Edition, Characteristics of Power Plants Seeking Transmission Interconnection As of the End of 2023,” *Lawrence Berkeley National Laboratory*.
<https://emp.lbl.gov/publications/queued-2024-edition-characteristics>
36. RAVEN: Risk Analysis Virtual Environment, Github. Accessed: March 2024.
<https://github.com/idaholab/raven>

37. D. J. McDowell, et al., “2023 FORCE Advanced Synthetic History Development Update,” Idaho National Laboratory Technical Report, 2023, INL/RPT-23-74821-Rev000
38. P. de Jong and J. Penzer, “The ARMA model in state space form,” *Statistics & Probability Letters*, Vol. 70, 2004, <https://doi.org/10.1016/j.spl.2004.08.006>.
39. J. Contreras, “ARIMA models to predict next-day electricity prices,” *IEEE Transactions on Power Systems*, Vol. 18, 2003. <https://doi.org/10.1109/TPWRS.2002.804943>
40. J. Chen and C. Rabiti, “Synthetic wind speed scenarios generation for probabilistic analysis of hybrid energy systems,” *Energy*, Vol. 120, 2017, <https://doi.org/10.1016/j.energy.2016.11.103>
41. Pedregosa, F., et al. 2011. “Scikit-learn: Machine Learning in Python.” *JMLR*, **12** (85): p. 2825-2830. <https://jmlr.csail.mit.edu/papers/v12/pedregosa11a.html>.
42. Bynum, M. L., et al., “Pyomo – Optimization Modeling in Python,” Third Edition. Vol. 67. Springer, 2021.
43. Hart, W. E., et al., “Pyomo: modeling and solving mathematical programs in Python.” *Mathematical Programming Computation*. 2011. 3(3):219-260.
44. Forrest, J., et al., *coin-or.Cbc:Release*, 2023, Zenodo. <https://doi.org/10.5281/zenodo.10041724>