MIT CEEPR
Center for Energy and Environmental Policy Research

Working Paper Series

Steady Heat, Flexible Power: Flexibility by Proxy in Nuclear-Thermal Energy Storage Systems

Nicolas Trosino, John E. Parsons, Ruaridh Macdonald, Charles Forsberg, and Yannick Perez

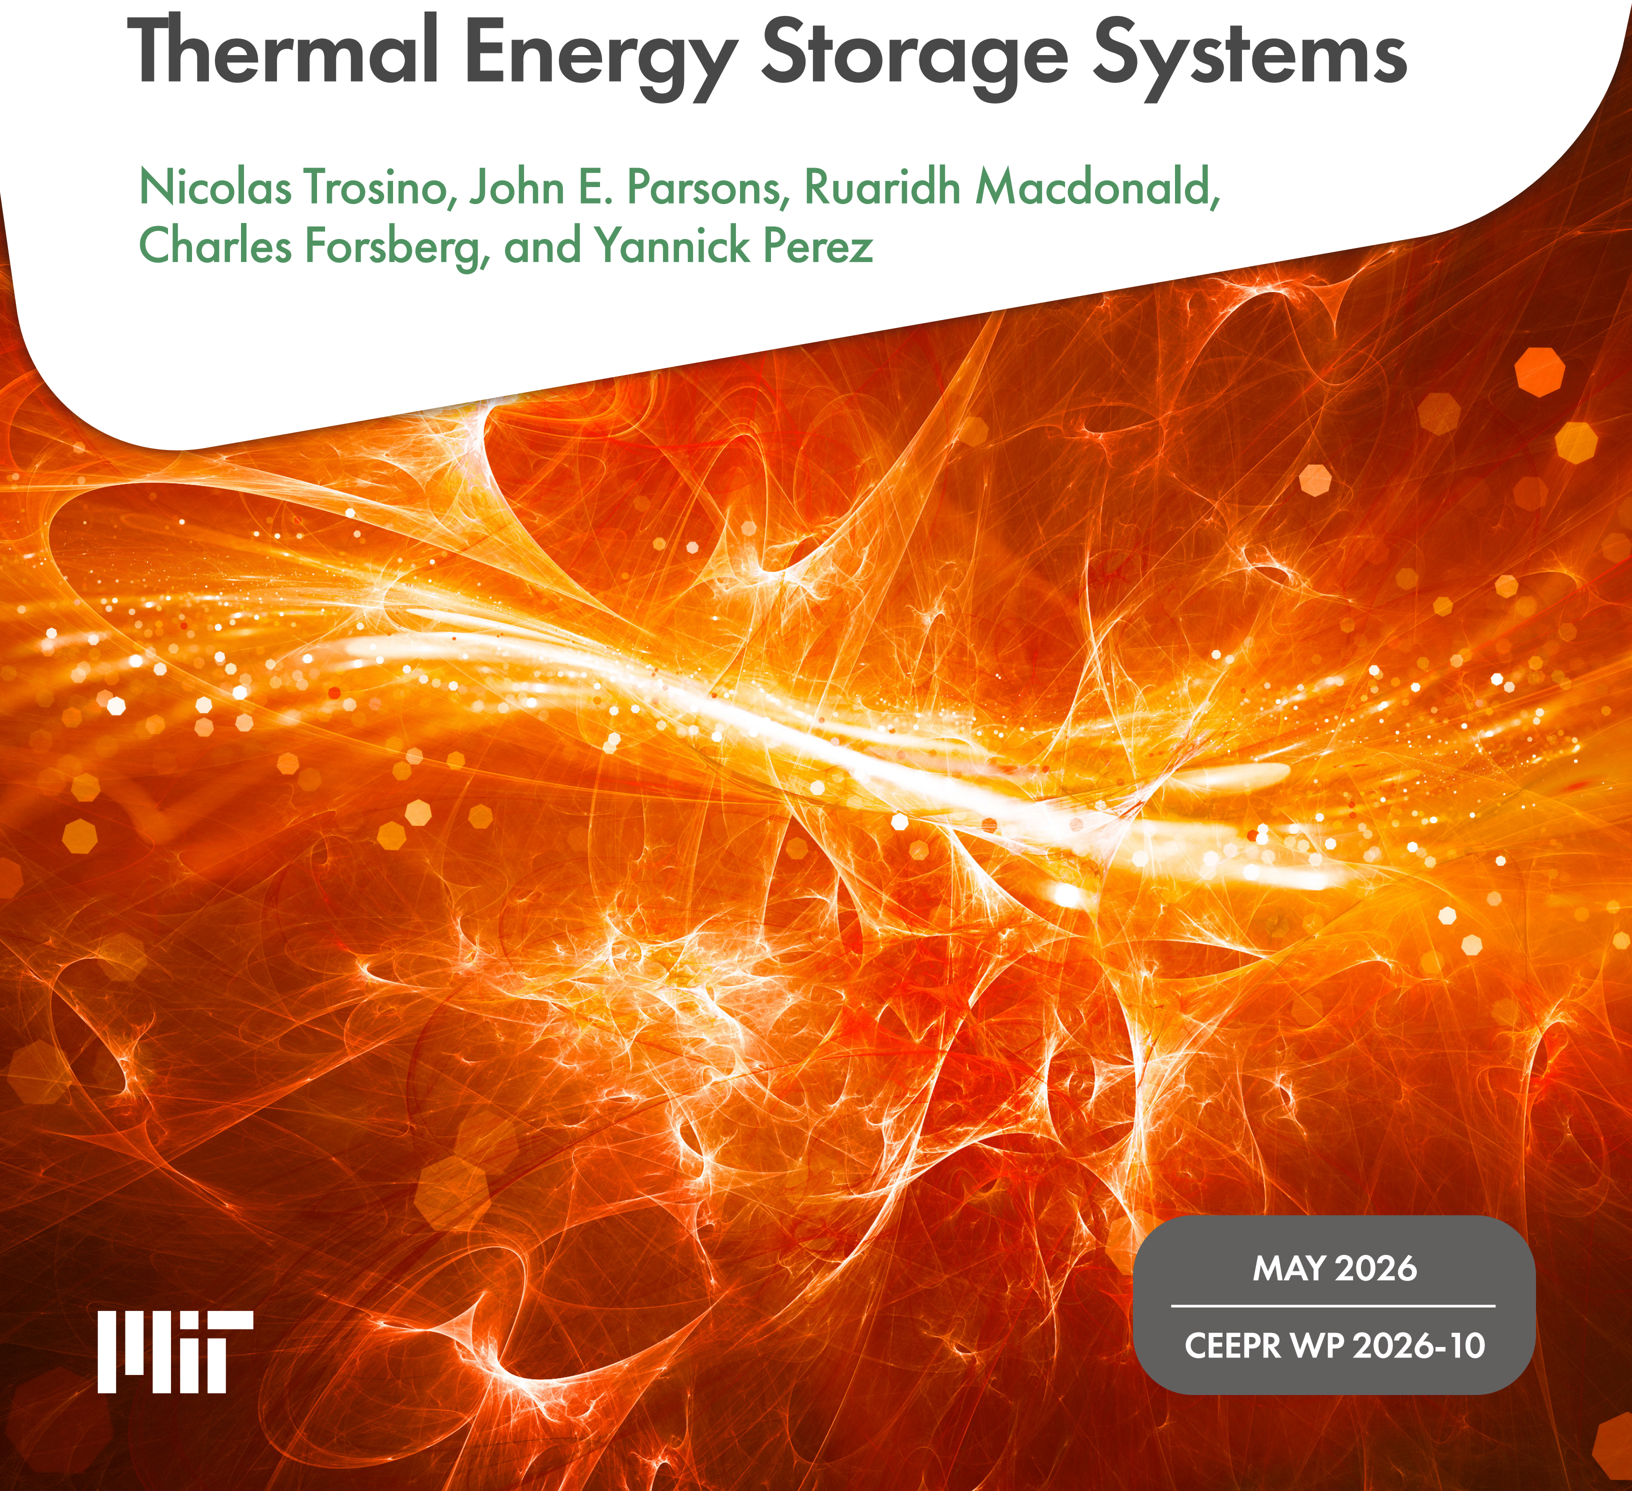

MAY 2026

CEEPR WP 2026-10

Working Paper Series.

Steady heat, Flexible Power: Flexibility by Proxy in Nuclear-Thermal Energy Storage Systems

Working paper

Authors

Nicolas Trosino*, Fulbright PhD Candidate at CentraleSupélec, MIT and HEXANA

John E. Parsons, MIT Sloan School of Management and MIT Center for Energy and Environmental Policy Research

Ruaridh Macdonald, Energy Systems Research Lead at the MIT Energy Initiative

Yannick Perez, Full Professor in Energy and Mobility Economics at CentraleSupélec

Charles Forsberg, Principal Research Scientist at MIT Nuclear Science and Engineering

*: corresponding author, nicolas.trosino@centralesupelec.fr

Abstract

We use a capacity expansion and dispatch model to study the value of integrating a Thermal Energy Storage (TES) system into a fusion power plant. Storage enables the plant to run the fusion reactor at very high capacity factors, while generation of electricity fluctuates significantly. TES enables penetration of fusion even at relatively high capital costs. Our model optimally sizes both the power and the energy capacity of the TES, and we find that it chooses very long durations. The TES has a lower turnover than grid batteries, but a higher utilization. The TES captures a higher spread on its turnovers than grid batteries, and its operating profitability is more concentrated in a fewer number of hours.

Keywords: Thermal Energy Storage, Capacity Expansion and Dispatch Modeling, Electric System Flexibility, Next-generation Nuclear Energy, Fusion Energy, Cost Minimization.

1. Introduction

Thermal Energy Storage (TES) is a promising technology to reconcile the economic constraints of capital-intensive low-carbon thermal generators such as nuclear power plants with the operational flexibility required in deeply decarbonized systems. The increasing penetration of variable renewable energy drives a growing spread in the marginal value of generation across hours through a day and throughout the year. On the one hand, when renewables are plentiful and net load is small or even zero, the marginal value of generation is zero. On the other hand, when there is a drought of renewable energy and net load threatens to exceed firm capacity, the marginal value of generation spikes. Integrating Thermal Energy Storage enables baseload production of heat[1], even during hours when the marginal value of electricity is low, which can then be converted to electricity when the marginal value of generation is high. This creates a new operational paradigm in which nuclear assets can behave simultaneously as firm baseload providers to maximize rate-of-return and as flexible resources to match system needs.

Understanding the system-level value of such a configuration requires moving beyond component-level techno-economics. We embed nuclear - Thermal Energy Storage couplings within a full-scale Capacity Expansion and Dispatch Model (CEM) in GenX[2]. GenX identifies the cost-minimizing portfolio of generation, energy storage, and transmission investments needed to meet a pre-defined system demand, while adhering to various technological and physical grid operation constraints, resource availability limits, and other imposed environmental, and policy constraints. Here, we co-optimize investment and hourly operation across a portfolio of competing low-carbon technologies.

This framework allows us to quantify not only the optimal sizing of generators, storage volumes, and turbine capacities, but also the underlying drivers of their value: avoided ramps, reduced volumes of non-served energy, enhanced integration of renewables, and improved utilization of expensive nuclear capital. Our results show that Thermal Energy Storage fundamentally reshapes the economic frontier for next-generation nuclear power, dramatically increasing the breakeven cost at which investment in fusion plants is cost-effective for the grid from $8{,}000\ \$/\mathrm{kW}_e$ up to $19{,}000\ \$/\mathrm{kW}_e$. Incorporating TES and therefore deploying fusion compresses the spread in the marginal value of generation across hours of the day and through seasons of the year. Our model optimizes the size of the TES energy reservoir as well as the size of the turbine generating electricity. The optimized TES includes a very large energy reservoir, giving the TES the ability to deliver power for a long duration of 10 hours. We analyze the operation of the TES and contrast it with the operation of other storage technologies.

The remaining sections of the paper are organized as follows: Section 2 provides the literature review on thermal energy storage, Section 3 introduces our modeling framework, Section 4 presents our results, Section 5 discusses our three primary findings, and the final section concludes.

---

[1] Low-carbon heat generating technologies include fusion, fission, geothermal and concentrated solar plants. Our modeling methodology remains valid for alternative sources of heat.

[2] The MIT Energy Initiative and Princeton University's Zero-carbon Energy systems Research and Optimization (ZERO) Lab have developed an open-source tool for investment planning in the power sector, offering improved decision support capabilities for a changing electricity landscape: github.com/GenXProject/GenX.jl

2. Literature Review

This research builds on existing capacity expansion dispatch models (CEM) such as GenX. GenX uses a set of technologies defined by costs and technological properties, and outputs the optimal capacities and dispatch of the system, given load, resource availability and fuel price time series. Optimal can be defined from the perspective of an asset owner (profits' maximization) or from the system planner (cost minimization). Boiteux (1960, 1964) key result states that profit's maximization of all stakeholders is equivalent to a system cost minimization. In this article, we use system's cost minimization as an effective proxy to abstract from any market design in stakeholders' decisions.

However, CEM can include specific policies such as decarbonization targets (references). In some cases, to account for existing assets or limited land availability, we also add minimum or maximum capacities. A key indicator of the competitive space for a given technology is the optimal generation capacity (power in MW_e). When considering storage assets, energy reservoir capacity (in MWh_e) is a second dimension and adds complexity in the merit-order (Antweiler & Müsgens, 2024). When considering next-generation assets such as fusion reactors, an alternative is to document the ceiling capital cost (i.e., the maximum capital cost to reach a given penetration), instead of the optimal capacity. This enables design-to-cost approaches for project developers.

Until now, CEM literature relied on various technological options ranging from fossil thermal generators to variable renewable generators. Two studies focused on the US decarbonized grid in 2050 (Schwartz et al., 2023; Armstrong et al., 2024) introduced fusion power plants as an option. Both used a capital-intensive pulsating[3] tokamak-style fusion reactor, coupled to a thermal energy storage, for which costs are uncertain. The former (Schwartz et al., 2023) used the "dual" approach, documenting ceiling cost, whereas the latter (Armstrong et al., 2024) surveyed the optimal capacities found in several scenarios combining capital cost assumptions and decarbonization scenarios.

Because the pulsating reactor is evaluated in a model with hourly operational granularity, including TES in the design is not a mere option to even out the thermal production of the core. Conservatively, we only consider the role of the TES as a proper storage asset. It charges on one side with the heat production of the reactor and discharges on the other side with the turbine supply. A key specificity of TES is its asymmetry: the charge capacity is equal to the reactor capacity, whereas the discharge capacity equals the turbine capacity. These asymmetrical storage assets have got limited attention in the literature.

First, Schwartz et al. (2023) modeled a simplified Thermal Energy Storage as an energy reservoir with capital costs proportional to its energy capacity. It appeared to contribute significantly to the optimal mix in some US regions, including New-England, which is the focus of this paper. Beyond that first approach, Saeed et al. documented the need for finer cost models, including charge and discharge capacities, architecture, and integration in systems. This leaves a gap to tackle on the value of TES asymmetry: our research here will look for the competitive space of Thermal Energy Storage to foster the penetration of next-generation capital-intensive nuclear assets.

---

[3] Many proposed designs of fusion reactors operate for some period of time, shut down to remove fusion waste products from the fusion reaction chamber and then restart. In our analysis we use a typical heat generating cycle of 20 minutes of heat production followed by 1 minute of dwell time between pulses. We consider the hourly averaged

Until now, Thermal Energy Storage literature has been focusing on technological modeling (Computational fluids dynamics, architectural optimization regarding thermal transfers) and little is published on the penetration of Thermal Energy Storage in the broader energy system.

Prior literature in a similar vein includes, first, Saeed et al. (2022) and (2024), exploring various TES-nuclear coupling architectures, and Schmalensee et al. (2015), documenting TES-concentrated solar couplings. We also include in this category other studies conducted on Thermal Energy Storage - Nuclear Power Plant (hereafter TES-NPP) architecture or applications. Finally, Novotný et al. (2025) is a recent and very sound compilation of many possible architectures, focused on the process flow modeling of such couplings.

In a different approach based on custom versions of GenX, Schwartz et al. (2023) and Armstrong et al. (2024) investigated the value of Fusion power plants – coupled to TES – in a decarbonized future and analyzed the results across a variety of scenarios.

Expanding Saeed et al. (2022, 2024) assumptions and turning upside down the focus of past studies conducted with GenX, we specifically document TES value in fostering "upstream" plant penetration and modeling assumptions impacts in generation expansion planning models. "Upstream" plant refers to the thermal source supplying the TES with thermal energy to store, without any specific attention to the technology effectively implemented in that black box.

Existing literature on Thermal Energy Storage - Nuclear Power Plant couplings

First, we identify in the literature two possible Thermal Energy Storage architectures: Figure 1 presents the "direct" architecture, whereas Figure 2 illustrate the "indirect" alternative. In Appendix A, we provide a comprehensive set of mathematical equations to include any of these in a capacity expansion dispatch model. The choice of one architecture or the other is out of this paper perimeter and discussed by Novotný et al. (2025).

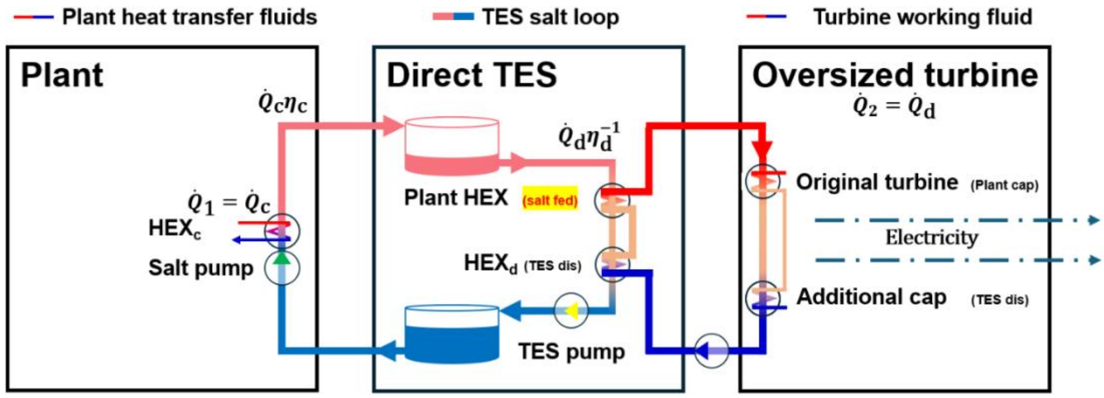

Figure 1: Direct TES architecture. The plant provides thermal energy to the salt loop ($\dot{Q}_1$, feeding HEX$_c$) which independently supplies the turbine loop ($\dot{Q}_2$, through the salt-fed Plant HEX and HEX$_d$). When the TES charges, the plant provides more energy than the salt loop discharges to the turbine ($\dot{Q}_1\eta_c - Q_2\eta_d^{-1} > 0$). Therefore, excess salt is transferred from the cold tank (light blue) to the hot tank (light red). Conversely, when the TES discharges, the TES pump transfers salt from the hot tank to the cold tank ($\dot{Q}_1\eta_c - Q_2\eta_d^{-1} < 0$), resulting in a boosted generation at the turbine. The red slice in the hot tank can be seen as the thermal energy available, and the blue slice as the storage availability. This architecture requires a salt loop and a salt-fed power block (no additional cost if the salt loop is included by-design with the plant, HEX$_c$ and power block technological change cost otherwise). It also requires an additional heater HEX$_d$, a pump, and an oversized turbine. Additional capacity is the TES net discharging capacity. This design presents a molten-salt two-tank design, but alternatives exist, like heterogenous fluid 1 → storage media → fluid 2 or isocline.

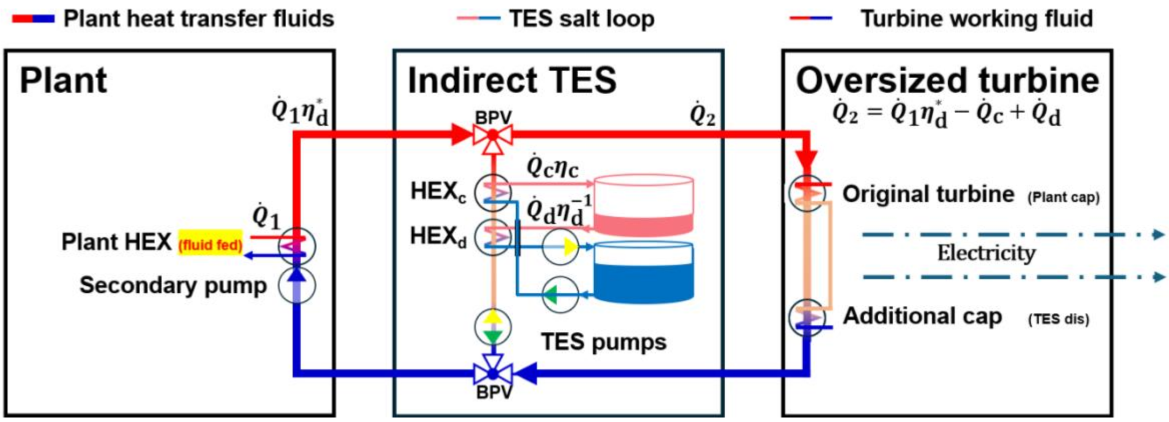

Figure 2: Indirect TES architecture. The plant provides thermal energy to the turbine through its heater "Plant HEX" (thick red line is hot working fluid). After the turbine, the cooled down fluid returns to the plant through the cold branch (thick blue line is exhaust gas/liquid). When the TES charges, some thermal energy is diverted from the hot branch (top BPV), providing energy to the TES salt loop through its heat exchanger (HEX$_c$) and rejected (bottom BPV). The green pumps are active. It fills the hot tank with hot salt (light red) and empties the cold tank (light blue). Conversely, when the TES discharges, cold fluid is diverted from the cold branch (bottom BPV), heated (through HEX$_d$, fed with hot salt), and mixed with hot fluid at top BPV. The yellow pumps are active. It transfers hot salt from the hot tank to HEX$_d$ to the cold tank. The red slice in the hot tank can be seen as the thermal energy available, and the blue slice as the storage availability. This architecture requires at least two TES heat exchangers, three pumps, and an oversized turbine. Additional capacity is the TES discharging capacity. Check valves and piping details are not represented.

For each of these architectures, the capacity expansion dispatch models can optimize TES-nuclear investment across three dimensions.

An upstream **plant and Thermal Energy Storage charging capacity**. It can be anything in the technological options available if it produces thermal energy. We illustrate our example with a tokamak-style fusion power plant, but more mature alternatives could be Sodium-cooled Fast Reactor such as the one being built by TerraPower in the United States or being developed by HEXANA in France. Both include molten-salt thermal energy storage in their designs.

A downstream **turbine and Thermal Energy Storage discharging capacity**. Many turbine technologies are available or under development. To remain conservative, we illustrate our example with the most mature Rankine steam cycle, whereas Brayton air and Brayton supercritical CO_2 are promising alternatives.

The **thermal energy storage** itself, defined as an energy reservoir. We represent the reservoir with two molten-salt hot-and-cold tanks, but alternatives exist. What is specific to TES compared to traditional batteries or most pumped-hydro storage is that it stores thermal energy before conversion and is by essence asymmetrical with an energy inflow (plant) not necessarily coming from the grid.

Our emphasis is on the optimal investment in charge/storage/discharge TES capacities, whereas the previous literature has a strong plant-scale technical focus: Denholm et al. (2012) and Alameri and King (2013) respectively explored the TES-NPP couplings and the reaction of a TES-NPP core to unplanned transients. The first emphasized that TES-NPP literature was sparse as of 2012, not because there is a conceptual challenge, but rather because the design space is large. Moreover, as nuclear assets are capital intensive, long to build and subject to sought safety licensing, it is not as simple as learning-by-doing or trial & error. Hence, future research may leverage on more common couplings between thermal energy storage and concentrated solar plants, multiple of which were tested in the past decade (Medrano et al. (2010), Gomez (2011), Herrman and Kearney (2002)).

The performance of multiple Thermal Energy Storage - Nuclear Power Plant couplings media was documented by Green et al. (2013) and Edwards et al. (2016) through exergy efficiency, and Forsberg et al. (2018) explored technological options, emphasizing the value of low-cost TES coupled to NPP in a grid with important renewable penetration. Particularly, Forsberg et al. (2018) expects TES-NPP to be more profitable than TES coupled to Concentrated Solar Plants ("CSP", at comparable TES capacities) because of economies of scale and steady-state year-round nuclear production. Frick et al. (2018) simulated small modular reactor operations with a TES and concluded that they were i) able to mitigate the power swings during periods of high wind and solar production and ii) enhancing the capacity factor and reducing the stress on plant components.

Furthermore, Keith et al. (2023) and (2025) computed the economic dispatch of a TES connected to the steam loop of a Nuclear Power Plant, which supplies the grid and a TES-desalination plant coupling. They demonstrated the importance of the relative economic value between electricity and other nuclear-produced commodities. This research paves the way for further TES-NPP co-generation studies.

In 2021, Knighton et al. summarized the energy storage options for integration with nuclear power. Thermal energy storage being identified as a promising candidate according to its lower leveraged cost of stored energy (LCOS), compared to batteries or dihydrogen (power-to-H_2-to-power).

Finally, Saeed et al. (2022) developed models for various steady-state TES-nuclear coupling architectures within ASPEN HYSIS®. Three architectures are explored in the report: i) "Standalone NPP-TES coupled with a secondary power generation cycle"; ii) "Directly coupled NPP-TES system"; and iii) "Integrated NPP-TES coupled with an oversized primary turbine". These models were then optimized with three types of Nuclear Power Plants: a) Advanced Light Water Reactor (A-LWR); b) High-Temperature Gas-cooled Reactor (HTGR) and c) Liquid-Metal Fast Reactors (LMFRs). Saeed et al. (2022) then developed tailored cost functions for these models, using Charge, Storage and Discharge supersets of components (see Figure 3).

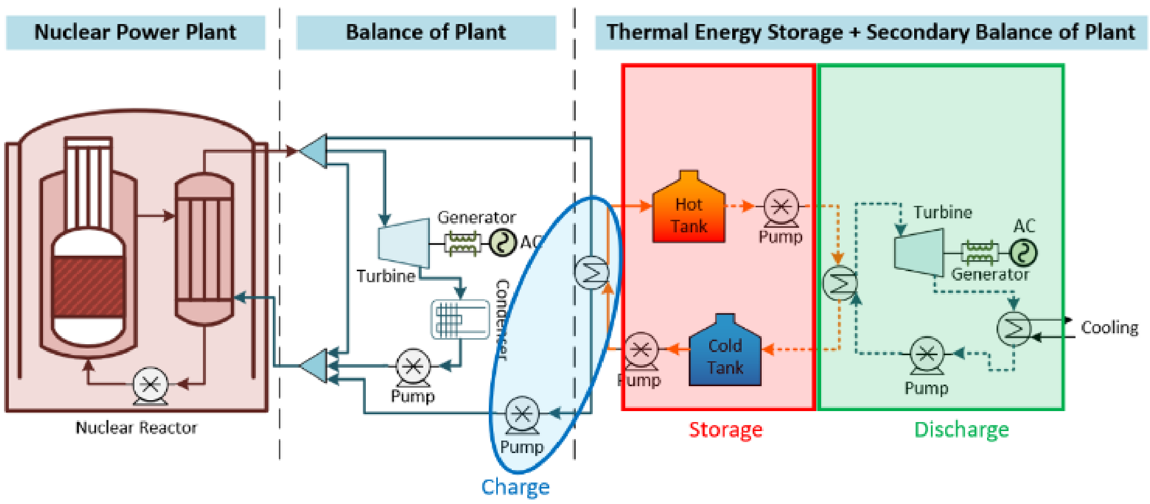

Figure 3: Component groupings to derive the superset cost functions for the first coupling option. Figure 25 in Saeed et al. (2022), p. 37

In line with this article perimeter, Saeed et al. (2024) conducts capacity expansion and dispatch optimization using HERON and price time series from RAVEN. Whereas the report presents capacity and dispatch optimization results, HERON-RAVEN coupling is still under investigation at

MIT Energy Initiative: significant concern can be raised on the coherence between the synthetic time series produced by RAVEN and the historical time series[4].

Table 1 summarizes the literature review, comparing our approach to three past articles. It emphasizes the detail given to Thermal Energy Storage modeling and contribution analysis, whereas existing literature focused on the generation of electricity or on the technology itself.

Table 1: Summary of literature review: Where does this paper's approach stands compared to existing literature?

	Model design	Strengths	Gaps
Saeed et al. (2022, 2024)	Techno-economic assessment of TES-Nuclear Power Plant couplings within HERON-RAVEN framework.	Explores a variety of coupling architectures with current and future fission technologies. Provides a sound review of cost functions for each component of the TES system. Considers asymmetrical systems. Uses multiple synthetical time series to reduce the advantage from perfect foresight.	Optimization is conducted on synthetic time series with increased volatility, not representative of real-world data. TES direct architecture is coupled after the steam cycle, which is suboptimal for most next-generation nuclear systems. No comparison against competing storage technologies
Armstrong et al. (2024, rev 0)	Techno-economic assessment of TES-Fusion Power Plant couplings using GenX, in a decarbonized New-England, in 2050. Primal approach using FPP cost scenarios.	Surveys the penetration of next-generation fusion power in various cost and decarbonization scenarios. Optimizes the mix against 20 independent sets of assumptions for the load, resource availability, and fuel prices. These time series are extrapolations of historical datasets, ensuring tractability in volatility and statistical moments. This provides results robust to various system conditions.	Do not consider non-fusion nuclear designs. Results are highly dependent on local conditions (New-England). Methodological caveat in the introduction of Thermal Energy Storage: mathematical incentive to not include TES capacities. Indirect TES architecture (see Figure 2) is suboptimal when coupled to a

[4] Short, there is a dilemma between conserving statistical momentums and autocorrelations and producing synthetic series that are different enough from source data. It deserves dedicated research out of this article perimeter.

		Limited geographic scale helps unpacking the value drivers of operations. Includes competition with a set of technologies ranging from fossil-fired plants to intermittent renewable energy sources. Limited size and complexity of the system enable replication of results with limited computing resources.	plant with a built-in salt loop. Fusion power plant is assumed to be perfectly flexible, which limits intrinsically the value of the TES system. Turbine capital costs are too expensive compared to literature. TES is forced to be symmetrical, while this clearly increases capital costs.
Schwartz et al. (2023)	Techno-economic assessment of TES-Fusion Power Plant couplings using GenX. Dual approach documenting the target Fusion investment cost required to reach a given penetration.	Large-scale modeling of the US grid in 2050, incorporating multi-zone effects. Simplified model of Thermal Energy Storage as a simple reservoir to understand the fundamental drivers of its value to the grid. Documents the influence of TES on the break-even capital cost of Fusion plants.	Low cost of the TES overestimates the optimal capacity of TES. No losses in the TES, creating an advantage over other technologies. The charge/discharge capacities are free and do not accurately document the role of TES as an asymmetrical storage asset, nor the influence of direct or indirect architectures.
Our choices	This study employs GenX as a linear capacity-expansion and hourly dispatch framework to evaluate the techno-economic performance and system value of a fusion–thermal energy storage coupling in a decarbonized New England power	Explores in greater detail Thermal Energy Storage value with an architecture found in commercial projects and designs (e.g., Natrium and HEXANA for fission, CFS for fusion). Direct TES (see Figure 1) is able to decouple the heat production and the electricity generation.	Subject to fusion-specific constraints. Hence, the results do not depict the role of TES coupled with fission designs. Results are highly dependent on local conditions (New-England). Fusion power plant is assumed to be

	system, in 2050. We assume perfect foresight. Setting is the same as in Armstrong et al. (2024). In particular, we keep the approach robust to 20 scenarios and the limited size of the system.	Compares the system value of TES to alternative storage technologies with quantitative indicators. Provides two open-source TES models to replicate the study with other modeling frameworks. Strengths of Armstrong et al. (2024). Contrasts with the common set of single year scenarios and proves instrumental in evaluating storage value.	perfectly flexible, which limits intrinsically the value of the TES system. Perfect foresight enables storage assets to operate without uncertainty. The pure cost-minimization approach is somehow comparable to an energy-only market design, whereas many systems have implemented energy-capacity designs.

3. Our Model

3.1 The GenX model

We use the GenX model described in Armstrong et al. 2024, which identifies the cost-minimizing portfolio of generation and storage investments for the six states of New England, without any market design assumption. A distinctive feature of that model is its parameterization with 20 scenario years for hourly load and wind and solar resource availability.

Figure 4 shows the load-duration curves for these 20 scenario years. The model identifies a single portfolio of investments to minimize the expected cost across all 20 scenario years, i.e., this is a version of robust optimization.

Generating technologies available in the model include utility-scale PV, commercial rooftop PV, residential PV, onshore wind, fixed offshore wind, floating offshore wind, legacy run-of-river hydro, natural gas combined-cycle, and natural gas combined-cycle outfitted with carbon capture and sequestration. Storage technologies include legacy pumped hydro-facilities as well as Li-ion batteries. Notably, since New England imports a significant amount of power from Quebec's large hydro system, the GenX model also includes an hourly load profile in Quebec, together with an hourly water inflow profile, and the model optimizes the exports of power to New England while respecting specified reservoir system constraints such as minflow and seasonal maximum capacities.

The investment, operating and maintenance costs of each technology are provided in Appendix .

The solution space can be constrained with respect to exogenous policies or limits, such as decarbonization targets or support mechanisms for new technologies. As for carbon policies, one option is to constrain the average carbon intensity of each kWh_e consumed. In this paper, the carbon limit is $12\ \mathrm{gCO_2/kWh}_e$, a 95 % reduction relative to 1990. Then, to introduce support mechanisms, some technologies may have minimum capacities. Finally, land availability or other limits are implemented through maximum capacity values.

These constraints will effectively cut parts of the solution space to comply with policies, and the dual variable bear information about how costly these constraints are. For the carbon constraint, the dual variable can be interpreted as the marginal cost of one ton of carbon. Finally, a binding minimum capacity constraint will introduce a subvention payment, whereas a maximum capacity limit will translate into a Ricardian rent.

Each generating asset's carbon intensities and min/max capacities are specified in Appendix .

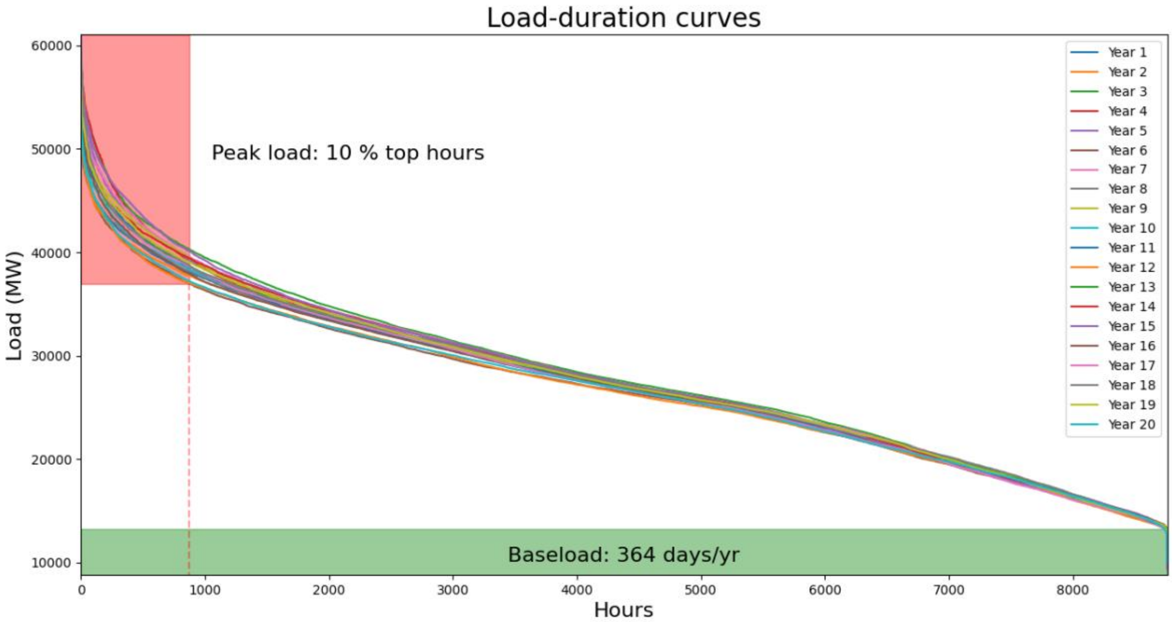

Figure 4: Twenty load-duration curves for NE-ISO, extrapolated from historical data using 2050 demand assumptions, used in Armstrong et al. (2024), see their section B.2.2.

3.2 The Fusion Thermal Plant

Alongside the other generation technologies, Armstrong et al. 2024 introduce a tokamak-type fusion power plant. The representation of a thermal power plant in GenX involves only a few essential parameters, such as thermal efficiency and operational constraints, and does not require elaboration of the many details particular to this or that fusion plant. Nevertheless, it is useful to fix ideas, especially when it comes to benchmarking the parameter values chosen to represent the fusion technology. Therefore, we briefly describe a few selected features of a generic, imagined deuterium-tritium magnetic confinement fusion power plant as detailed in Armstrong et al., 2024. However, in this paper's linear framework, the optimizer can invest in a decimal number of plants. Hence, there won't be any discrete effect with capacities indicated hereafter.

The base case fusion power capacity of $1{,}000$ megawatts (MW_{th}) is a standard metric in the academic literature on fusion. In addition to the energy released by the fusion reactions, heat is also generated by reactions in the blanket, energy is injected into the plasma, and energy is used to pump the molten salt blanket, which brings the base case plant's total thermal power capacity to $1{,}095\ \mathrm{MW}_{th}$. Turbine efficiency depends on the chosen power cycle. Here we assume an efficiency of $40\ \%$, in which case the plant's gross electric power capacity is $438\ \mathrm{MW}_e$. Some 111 MW_e of power is needed for station load within the plant to cool the magnets, drive the fusion reactor, pump the molten blanket, and supply electric power for the tritium processing system and other plant operations. Thus, the net electric capacity is $327\ \mathrm{MW}_e$, which is the standard metric in the electric power industry.

For our modeling, it will be important to understand how the station load varies as the thermal power is ramped down. We assume $\mathrm{FixedSL} = 10\ \mathrm{MW}_e$ is the fixed load and $\mathrm{VarSL} = 0.083\ \mathrm{MW}_e/\mathrm{MW}_{th}$ is the variable load. The Fusion Power Plant can balance its self-consumption by collecting part of its electricity generation or, without any efficiency prejudice, by consuming electricity from the grid ($\mathrm{FPP\_imports}_{t\in T}$). The latter option effectively increases the system load. Both options can be used separately or as a combination. Overall, the plant's heat-to-electricity efficiency is $\eta_{\mathrm{FPP}} = 29.4\ \%$.

The Fusion Power Plant cost is the annualized sum of multiple components' costs, assuming a 6 % discount rate. Economic lifetimes are 2 years for the vessel and 40 years for the other parts. While this paper keeps the reference costs used by Armstrong et al., 2024 ($8{,}500\ \$/\mathrm{kW}_e$ net of a benchmark Fusion Power Plant), one major difference is that we notionally separate hereinafter the turbine capacity sized to match the reactor capacity and the extra turbine capacity to discharge the reservoir (see 3.3 The Thermal Energy Storage). This reference cost only includes the reactor-matching turbine capacity. The marginal cost of producing one unit of energy is noted $\mathrm{MC}_{\mathrm{plant}}$.

3.3 The Thermal Energy Storage

The Thermal Energy Storage is the key component between heat production and electricity generation. We revisit and revise the Fusion Power Plant – Thermal Energy Storage - Turbine coupling developed for Armstrong et al., 2024.

There is a charging interface on one side, whose capacity (in MW_{th}) is equivalent to the thermal capacity of the Fusion Power Plant described in the previous section. On the other side, there is a discharging interface to supply the turbine with thermal energy. Its capacity (in MW_{th}) equals the turbine inlet capacity. Finally, in the middle, there is an energy reservoir whose capacity is in MWh_{th}.

Our paper architecture is presented in Figure 1, whereas Armstrong et al., 2024 used an indirect architecture (see Figure 2).

We model losses arising in the Thermal Energy Storage at three points of operation: charging, storing, and discharging. The charge efficiency (η_{c}) reflects losses as thermal energy is deposited from the reactor into the reservoir. The discharge efficiency (η_{d}) reflects losses as thermal energy is sent from the reservoir to the turbine. The self-discharge (sd) reflects the rate of leakage of thermal energy out of the reservoir.

Because our architecture is direct with a plant already featuring a salt loop, we assume $\eta_{\mathrm{c}} = \eta_{\mathrm{d}} = 1$. This is because all salt loop losses are already accounted for in the Fusion Power Plant model. We only consider self-discharge, with $\mathrm{sd} = 2\ \%/\mathrm{day}$.

3.4 The Turbine model

In many capacity expansion and dispatch models, generation technologies are defined by their net power capacities. In this model we track both the fusion technology's gross power capacity of the oversized turbine as well as the self-consumption of the reactor. This can be important to a model with TES because the operating mode that produces the peak injection of power to the grid involves turning down the reactor's thermal generation to reduce the fusion reactor's significant self-consumption and using heat from the TES to drive the turbine. Conversely, in nominal reactor operations, part of the turbines will be used to balance the station load of the plant, resulting in less power being injected on the grid. This is another detail in which the model presented here departs from the model in Armstrong et al. 2024.

We define the “net generation” of the plant by simply balancing the power at the power station node between gross generation, grid consumption on one side, and self-consumption and net generation on the other side. This change enables Thermal Energy Storage to reach its maximum discharge capacity even when the fusion reactor is not operated at maximum capacity. Additionally, it simplifies the accounting of investment costs and makes this model closer to other generic thermal plants used in GenX.

Table 2: Summary of the Fusion Power Plant - Thermal Energy Storage coupling’s features

Benchmark plant without TES		
Reactor thermal output	MW_{th}	1,095.0
Gross electric output, baseload	MW_{e}	438.0
Self-consumption, permanent	MW_{e}	10.0
Self-consumption, fixed when in use[5]	MW_{e}	10.0
Self-consumption, variable parameter	$\mathrm{MW}_{e}/\mathrm{MW}_{th}$	0.083
Self-consumption, variable max value	MW_{e}	90.9
Self-consumption, baseload	MW_{e}	110.9
Net electric output, baseload	MW_{e}	327.1
Turbine baseload, gross	MW_{e}	438.0
TES		
Reservoir energy capacity	MWh_{th}	29,818.2
Reservoir energy capacity, equivalent	MWh_{e}	11,927.3
Turbine boost	MW_{e}	737.9
Turbine total, gross	MW_{e}	1,175.9

A key feature of this coupling is its ability to discharge faster than it can charge. This is what we call “boost ratio”:

$$\text{Boost ratio} = \frac{\text{Dis\_cap}}{\text{Cha\_cap}}$$

Also, define Thermal Energy Storage net discharge capacity and net discharge as:

- $\text{Net\_dis\_cap} = \text{Dis\_cap} - \eta_{\mathrm{c}}\eta_{\mathrm{d}}\text{Cha\_cap} < \text{Dis\_cap}$
- $\text{net\_dis}_{t\in T} = \frac{1}{\eta_{\mathrm{d}}}\text{dis}_t - \eta_{\mathrm{c}}\text{cha}_t \in [-\text{Cha\_cap}, \text{Dis\_cap}]$

A remarkable feature of this TES architecture is that net discharge is not limited to the net discharge capacity, but to the total discharge capacity. This is an operational mode that we name “boosted discharge”, in which the Fusion plant station load becomes a flexible load. We define:

[5] Here, linear commitment in fact transforms this fixed self-consumption into a variable self-consumption. This results in an additional 0.009 $\mathrm{MW}_{e}/\mathrm{MW}_{th}$. The total variable self-consumption becomes 0.092 $\mathrm{MW}_{e}/\mathrm{MW}_{th}$, which is a slightly higher rate than in Armstrong et al., 2024.

$$\text{boosted\_dis}_{t \in T} = \text{net\_dis}_t - \text{Net\_dis\_cap}$$

Figure 5 illustrates the operating space of the reactor - thermal energy storage coupling.

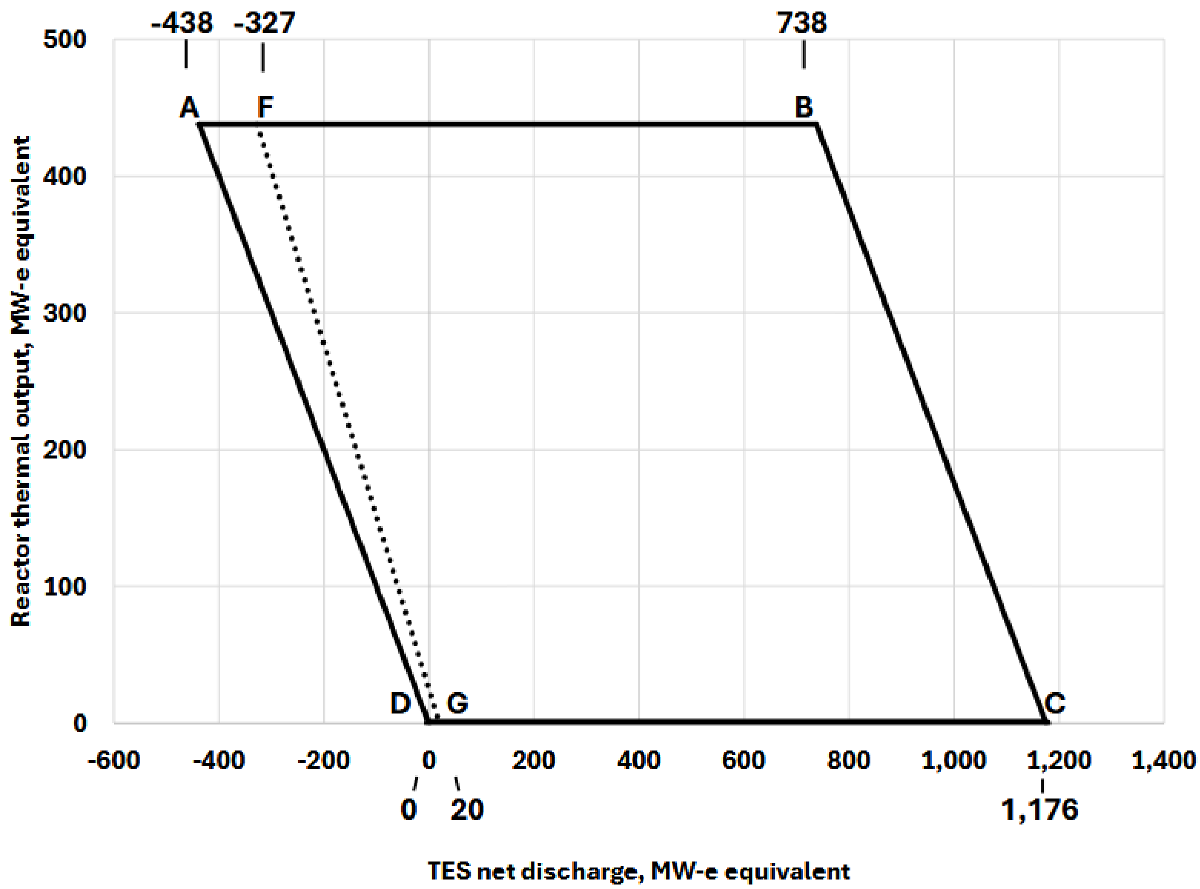

Figure 5: Feasible operating space for reactor thermal output and TES net discharge. Both variables are thermal megawatts, but the values have been converted to electric equivalent by multiplying by the 40 % conversion efficiency. The parallelogram ABCD defines the complete space. As operation moves within the space, the net electricity to the grid varies. On the line segment FG, net electricity to the grid is exactly zero. To the right, net electricity to the grid is positive. To the left, net electricity to the grid is negative—i.e., the plant is importing electricity to satisfy its self-consumption.

Figure 5 shows the feasible operating space for reactor thermal output and TES net discharge, before considering the state-of-charge. The vertical axis is the reactor thermal output, however the units have been re-denominated into MW-e equivalents by multiplying by the 40 % conversion factor. The reactor's thermal output can vary between 0 and the reactor's maximum capacity of 438 MW-e. The horizontal axis is the net discharge from the TES, also re-denominated in MW-e equivalent. Note the TES net discharge can be negative or positive. When it is negative, and therefore charging, it is because the reactor's thermal output sent to the TES is greater than the thermal output sent from the TES to the turbine. When it is positive, and therefore discharging, it is because the reactor's thermal output sent to the TES is less than the thermal output sent from the TES to the turbine.

The greatest rate of TES charging occurs at point A, where it is taking all of the reactor's thermal output while delivering nothing to the turbine. This rate of charge is 438 MW-e. Along line segment AB, the greatest rate of TES discharging occurs at point B, where it is supplementing all of the thermal energy received from the reactor with thermal energy from the TES up to the turbine's capacity. This rate of discharge is 738 MW-e.

Net electricity delivered to the grid is a function of the two operating parameters mapped in the parallelogram ABCD. This adds a third dimension to the example not graphed in Figure 5, but which we can appreciate. For example, at point B, the turbine is operating at its maximum gross capacity of 1,176 MW-e, and has self-consumption of 111 MW-e, so that the net electricity delivered to the grid is 1,065 MW-e. Although at point B the turbine is at maximum gross capacity, this is not the maximum net capacity to the grid. Along line segment BC[6], the reactor thermal output is declining and therefore self-consumption is declining, and at the same time discharge from the TES is displacing the reactor thermal output to maintain the turbine at its maximum capacity. At point C, net discharge from the TES is 1,176 MW-e so that the turbine is still operating at its maximum gross capacity of 1,176 MW-e, but self-consumption has declined to 10 MW-e, so that the net electricity delivered to the grid is 1,166 MW-e.

At point D, both the reactor thermal output and the TES net discharge are zero. However, because there are a permanent part and a fixed part of self-consumption, the net electricity to the grid is -20 MW-e. Line segment FG is defined by the property that net electricity to the grid is zero. At point F, the TES is charging at 327 MW-e, and the reactor thermal output is at capacity of 438 MW-e, so that the turbine is generating the 111 MW-e difference which covers the self-consumption. At point G, the reactor thermal output is zero, so the TES net discharge must supply 10 MW-e to the turbine in order to exactly cover the permanent part of the self-consumption.

In the following, we briefly describe how storage assets are modeled and define a set of indicators to describe investments and operations.

3.5 Storage Technology Modeling and Evaluation

Modeling storage assets: investments, operations, and accounting

GenX optimizes the investments in the portfolio of generation and storage assets, including the dimensions of the fusion plant's TES system. Most storage assets are usually modeled as symmetric in charging and discharge capacities, so that optimization is made across the two dimensions of power and storage. The TES system has three dimensions that can be optimized:

- charging capacity, $\mathrm{Cha\_cap}$, defined as the maximum thermal energy leaving the fusion reactor; the thermal energy arriving in the TES will be lower due to efficiency losses (η_{c}); note that this capacity is equivalent to the fusion reactor capacity;
- discharging capacity, $\mathrm{Dis\_cap}$, defined as the maximum thermal energy received by the turbine; the thermal energy departing from the TES will be higher due to efficiency losses (η_{d}); note that this capacity is equivalent to the fusion plant's turbine capacity;
- reservoir capacity, $\mathrm{Sto\_cap}$.

GenX also optimizes the dispatch of the fusion reactor, utilization of the TES, and utilization of the turbine. With $T \in \mathbb{N}$ the number of timesteps and $N \in \mathbb{N}$ the number of years considered, operations are described with three time series:

- $\mathrm{cha}_{t\in T}$, the charging time series, seen from the fusion reactor;
- $\mathrm{dis}_{t\in T}$, the discharging time series, seen from the turbine inlet;
- $\mathrm{sto}_{t\in T}$, the state-of-charge time series, seen from the reservoir.

---

[6] This is where $\mathrm{boosted\_dis}_{t\in T} > 0$.

The storage balance equation links together the three variables:

$$\text{sto}_{t\in T} - (1-\text{sd})\cdot \text{sto}_{t-1} = \eta_{\text{c}}\text{cha}_t - \frac{1}{\eta_{\text{d}}}\text{dis}_t = -\text{net\_dis}_t$$

3.5.1 Evaluating storage operations

Given the results of an optimization, it will be useful to analyze properties of the optimal investment and dispatch across different storage technologies. To that end, we define some useful metrics.

- Charge duration t_{c}: filling the reservoir from $0\ \%$ to $100\ \%$ using all the charging capacity and without any discharge;
- Discharge duration t_{d}: emptying the reservoir from $100\ \%$ to $0\ \%$ using all the discharging capacity and without any energy inflow.

For the TES, it makes sense to define a second discharge duration making the converse assumption about energy inflow:

- Discharge duration "on" $t_{\text{d,on}}$: emptying the reservoir from $100\ \%$ to $0\ \%$ using all the capacities and with an energy inflow equal to the upstream plant production.

With these durations we can quantify the maximum number of turnovers in a year by assuming the storage system cycled continuously:

$$\text{MaxTurnovers} = \frac{8760}{t_{\text{c}} + t_{\text{d}}}$$

Using the average annual discharge of the TES, we define the actual turnovers[7] in a given year as:

$$\text{AnnDischarge} = \frac{\sum_{t=0}^{T} \text{net\_dis}_t}{N}$$

$$\text{AnnTurnovers} = \frac{\text{AnnDischarge}}{\text{Sto\_cap}}$$

and translate this into a storage capacity factor:

$$\text{CF}_{\text{sto}} = \frac{\text{AnnTurnovers}}{\text{MaxTurnovers}}$$

To compare the operations' dynamics between technologies, we track the duration spent in the reservoir for each unit of energy charged. This is done using a first-in, first-out approach.

3.5.2 A Granular Analysis of Cost Savings

The hourly dispatch of the TES is chosen to contribute to the cost minimization. We can use the hourly values $(\lambda_t)_{t\in T}$ of the dual variable on the energy balance constraint to gain some insight into when and how the dispatch most contributes to reducing cost. GenX is a cost minimization model and makes no assumption about the market for electricity, so technically we cannot speak of prices and profit. Nevertheless, it is well understood that under certain assumptions a cost

[7] If the Fusion Plants were operated at full capacity, and all of the heat was immediately transferred to the turbine (i.e., bypassing the Thermal Energy Storage), just as if the TES were never used, AnnDischarge and AnnTurnover would be equal to 0.

minimization model can be reformulated as a competitive, profit maximization model. The main debate is whether or not the necessary assumptions hold in a specific context. Here, we are not concerned with the realism of the assumptions. We merely want to utilize the terminology appropriate for a competitive, profit-maximizing owner of the TES because we think that terminology is familiar to all readers. The same ideas could be exposited in terminology strictly appropriate to cost minimization, but in a more cumbersome and less familiar way. So, going forward, we are going to speak of the dual variables to the energy balance equation $\lambda_{t\in T}$ as if they are prices in the electricity market, and we are going to describe the Fusion plant as earning revenue and capturing operating profit from storing and releasing energy in its thermal energy storage.

In addition, although the TES is integrated together with the fusion reactor and the turbine into a single facility operated as one, we are going to create an accounting fiction of two separate entities. One entity is the reactor together with a turbine sized to operate at baseload capacity. The other entity is the TES together with the extra turbine capacity above baseload. This fiction will allow us to evaluate the profitability of the TES itself as opposed to the profitability of a standalone reactor without storage.

When the TES wants to accumulate stored energy, it will purchase the heat from the reactor, and it will have to pay the reactor's opportunity cost. In hours when the reactor would have sent the heat to the turbine and generated electricity to the grid, the opportunity cost is the reactor's foregone sale of electricity. This will be true for hours when the price of grid electricity is above the reactor's marginal operating cost $\mathrm{MC_{plant}}$. In hours when the price of grid electricity is below the reactor's marginal operating cost, the opportunity cost is the reactor's marginal operating cost.

We now track the operating profit from the dispatch of the TES. We maintain an accounting of the cost incurred to accumulate the state-of-charge in any hour:

- $\mathrm{unitCost}_{t>0} = \frac{\mathrm{sto}_{t-1}\cdot\mathrm{unitCost}_{t-1}+\mathrm{cha}_t\cdot\mathrm{MCc}_t}{(1-\mathrm{sd})\cdot\mathrm{sto}_{t-1}+\mathrm{cha}_t} \geq 0$, calculated at the end of timestep t;
- $\mathrm{unitCost}_{t=0} = \frac{1}{\mathrm{sto}_0}(\mathrm{sto}_{8759}\cdot\mathrm{unitCost}_{8759} + \mathrm{cha}_0\cdot\mathrm{MCc}_0 - \mathrm{dis}_0\cdot\mathrm{unitCost}_{8759})$

Discharging captures a spread defined as,

- $\mathrm{spread}_{t\in T} = \lambda_t - \mathrm{unitCost}_t \in [0, \lambda_t]$
- $\mathrm{OP}_{t\in T} = \mathrm{spread}_t \cdot \mathrm{net\_dis}_t$

For Thermal Energy Storage, at times when $\mathrm{boosted\_dis}_t > 0$, positive opportunity costs will appear. This is because the nuclear-TES system has turbines for the nuclear plant output (flowing through the TES) and additional turbines to discharge the reservoir during peaks. Because the Thermal Energy Storage has a "direct" architecture, the TES can use all the turbine capacity. If the plant located upstream of the TES has significant variable self-consumption, the optimal dispatch will occasionally shut down some plants (just as if it was flexible demand).

Incidentally, when $\mathrm{boosted\_dis}_t > 0$, TES will use additional turbine capacity, and the plant station load will decrease. During these specific timesteps, the TES must compensate for the missed opportunities of the plant ($\mathrm{MissedOpp}_{\mathrm{FPP},t\in T}$ in MW_e). Simultaneously, the offline plants have avoided station load ($\mathrm{AvoidedSL}_{\mathrm{FPP},t\in T}$ in MW_e). Hence, we define:

$$\mathrm{MissedOpp}_{\mathrm{FPP},t\in T} = (\mathrm{Cha\_cap} - \mathrm{cha}_t)\cdot\eta_{\mathrm{FPP}}$$

$$\mathrm{AvoidedSL}_{\mathrm{FPP},t\in T} = (\mathrm{Cha\_cap} - \mathrm{cha}_t)\cdot\mathrm{VarSL} + (\mathrm{FPP\_units} - \mathrm{FPP\_com_t})\cdot\mathrm{FixedSL}$$

$$\text{BoostedDisCost}_{t\in T} = \begin{cases} \mathbf{1}_{\text{boosted\_dis}_t>0}(t)\cdot\left(\text{MissedOpp}_{\text{FPP},t}\cdot\lambda_t - \text{AvoidedSL}_{\text{FPP},t}\cdot \text{MC}_{\text{plant}}\right) & (\text{TES}) \\ 0 & \text{for other technologies} \end{cases}$$

Conversely, another situation arises when the marginal cost of electricity goes below the marginal cost of the fusion plant. There is an advantage collecting electricity to balance the station load, while charging the TES with the thermal energy produced by the plant. The TES is effectively charged with electricity from the grid, converted to heat through the fusion plant. In such cases, the cost of consuming electricity should be associated with the TES.

$$\text{GridChaCost}_{t\in T} = \begin{cases} \mathbf{1}_{\text{cha}_t>0}(t)\cdot \text{FPP\_imports}_t\cdot\lambda_t & (\text{TES}) \\ 0 & \text{for other technologies} \end{cases}$$

Finally, we define FPP opportunity costs that TES must compensate:

$$\text{OppCost}_{t\in T} = \text{BoostedDisCost}_t + \text{GridChaCost}_t$$

For each technology, the storage operations yield operational profits (noted AnnOP for the annual average). It is calculated as:

$$\text{AnnOP} = \frac{1}{N}\sum_{t=0}^{T}\left[\text{OP}_t - \text{OppCost}_t\right]$$

Considering AnnCosts the annual costs faced by the storage asset (CAPEX and OPEX), we define net profitability (noted AnnNP for the annual average) as a simple subtraction:

$$\text{AnnNP} = \text{AnnOP} - \text{AnnCosts}$$

Note: for a greenfield cost-optimization (no existing capacities) and no limit to additional units of capacity, $\text{AnnNP} = 0$ when the capacity is optimal. When $\text{AnnNP} > 0$, it implies a Ricardian rent (more capacity would be better). When $\text{AnnNP} < 0$, it supposes external revenues such as subsidies to break even (minimum required capacity).

Intuitively, storage assets extract profits off high-lambda times following low-lambda times. This means that operational profits will naturally concentrate on the discharging hours. However, high-lambdas vary, and some discharging hours may carry more operational profits than others. A legitimate question is whether this concentration is important, or limited. This is pivotal to evaluating the influence of market designs on storage value.

To study the concentration of profits in some hours, we leverage on a common tool in inequalities economics, the Lorenz curve. To build this graph, hours of operation are sorted with respect to the operational profit captured. The Lorenz curve is the plot of the cumulative operational profit in these sorted hours. The closer the Lorenz curve to the first bisection, the less the storage relies on specific high-profit hours to achieve cost reductions. Conversely, the closer the Lorenz curve to the bottom right corner, the more economic value is reliant on extreme events.

As an addition to our study on the concentration of profits in specific hours, we document how profits are concentrated in specific years characterized by specific conditions (e.g., low renewables availability, high price of fuel). To that end, we create a graph that we name "profitability threshold graph." In this graph and for each year, we plot the average operational profit earned from a MWh_e discharged (y-axis) regarding the count of turnovers operated (x-axis). On the same plot, we superimpose the target average operational profit that would enable Thermal Energy Storage to break even (the "profitability threshold"). Being above the line indicates that the TES is profitable that year, whereas being under the line indicates net losses. On average, the barycenter must be on the profitability threshold.

Finally, to compare the technologies, we provide a summarized view, with each year's operational profits expressed as a percentage of the annualized costs, sorted by increasing operational profits.

4. Results

The GenX model has dual values reflecting the hourly marginal value of electricity, which can be tied back to the hourly marginal value of stored thermal energy. The cost-minimizing dispatch of thermal energy storage charges the thermal reservoir when the marginal value of thermal energy is low and discharges when the marginal value is high. It will be convenient to describe the storage as if it were capturing an "operating profit" by purchasing and selling thermal energy, although technically the model says nothing about a market in thermal energy.

Hence, to enhance legibility, we define a notional operational profit although there is no market structure in our cost optimization. Specific to storage, we note "operating profit" the cost benefit achieved by charging a storage at a given time step of lower lambda and discharging it later when lambda is higher. More generally, we commonly use "lambda" in the following to describe the marginal cost of electricity. Under specific assumptions, this notion is closely related to the price of electricity in an energy-only market.

In addition to simply documenting the investment or lack of investment in TES, we use the model outputs to provide a detailed analysis of how Thermal Energy Storage is operated in different contexts. For example, we document the frequency of cycling—how often does it cycle on a daily or longer basis—and document the profitability of different cycles and events—is it earning the same average level of profits off many events, or is a large proportion of profits earned on infrequent events? Is it earning profits from arbitraging system lambdas within a usual range, or from supplying electricity in rare scarcity events? What is the influence of Thermal Energy Storage on the profile of the lambda duration curve?

Short, our results demonstrate the value of Thermal Energy Storage to foster the penetration of expensive advanced designs of thermal generators, combining baseload reactor operations (maximizing return on investment) and load-following turbine operations (maximizing operational profit). It also effectively limits the volatility of the system lambda.

4.1 Capacity Optimization without Thermal Energy Storage

This section presents Capacity Expansion and Dispatch results for the reference case. Main assumptions are $8{,}500\ \$/\mathrm{kW}_e$ capital cost for the benchmark Fusion Power Plant and no Thermal Energy Storage can be built.

Table 3 compiles the optimal capacities and dispatch, whereas Table 4 document storage assets' capacities for the two available technologies.

Table 3: Optimal capacities and dispatch without TES. Fusion generation is displayed net of imports. Storage charge/discharge is presenting battery and pumped hydro operations.

Optimal Capacities and Dispatch Without TES			
	Capacity	Generation	Capacity

	(GW_e)	(TWh_e)	Factor
Natural gas	8.9	2.21	3 %
Natural gas w/ CCS	23.1	50.83	25 %
Solar PV	22.0	27.59	14 %
Onshore wind	9.0	30.27	38 %
Offshore wind (fixed)	27.3	114.79	48 %
Run-of-river hydro	2.0	7.78	45 %
Fusion	0.0	0.00	
Pumped hydro	1.8	2.61	
Battery	18.6	25.48	
Total	112.7	261.56	
Net imports	2.3	13.15	
Total generation		233.48	
Storage charge		-33.32	
Storage discharge		28.09	
Net Energy supplied		241.39	

Table 4: Optimal energy and power capacities of storage assets without TES.

	Energy	Power		Duration	
	Capacity	Charge	Discharge	Charge	Discharge
	(GWh_e)	(GW_e)	(GW_e)	(h)	(h)
Battery	56.6	18.6	18.6	3.3	2.8[8]
Pumped hydro	22.0	1.8	1.8	13.3	10.8
Fusion TES			Not available		

The Fusion Power Plant investment cost is too high to be economically viable. Its optimal capacity is $0\ GW_e$. The total capacity installed (generation + storage) is $112.7\ GW_e$, which is equivalent to $185\ \%$ of the maximum load ($61.0\ GW_e$).

[8] It is worth noting that the optimal duration for battery is in the lower range of what can be found in the literature for BESS applications. This is even more true when TES is available (see Table 6).

The lambda duration curve of the system is presented in Figure 6. The median system lambda is $60.57\ \$/\mathrm{MWh}_e$ and there are $53\ \mathrm{GWh}_e$ of non-supplied energy over the 20 years.

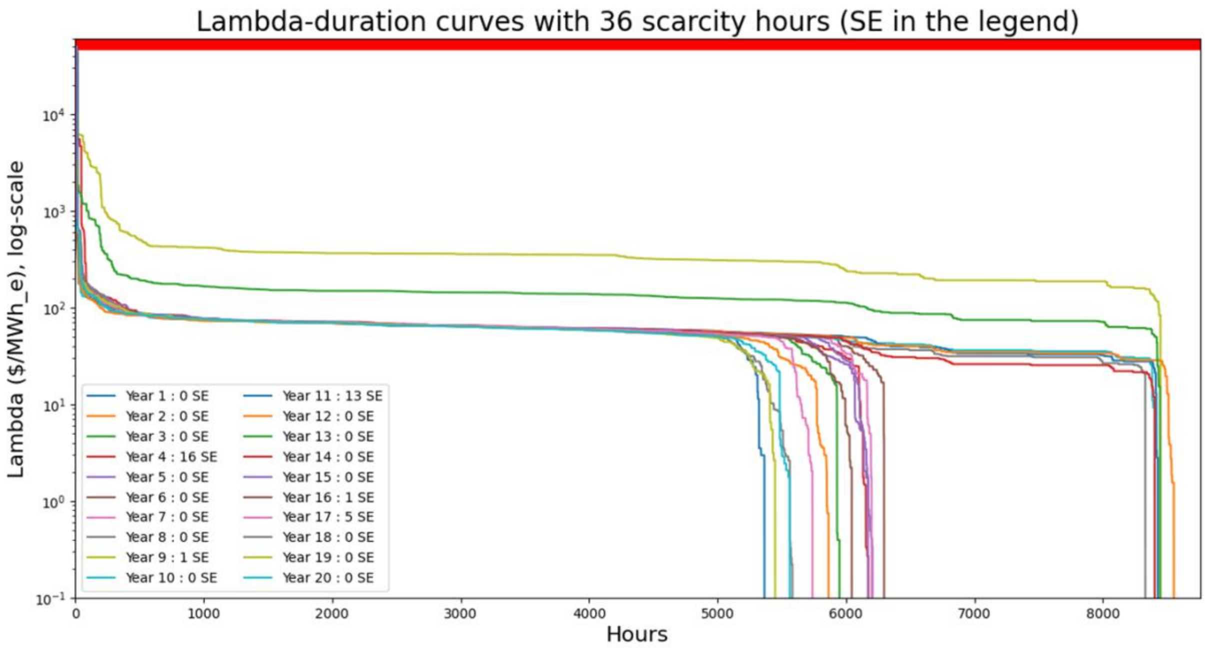

Figure 6: Lambda duration curves without TES, each line is an individual year in the sample.

4.2 Capacity Optimization with Thermal Energy Storage

We now add Thermal Energy Storage to the design of the Fusion Power Plant. This section presents Capacity Expansion and Dispatch results in the same setting as in the previous section.

Table 5 compiles the optimal capacities and dispatch and compares figures with/without TES, whereas Table 6 document storage assets' capacities for the three available technologies.

Table 5: Optimal capacities and dispatch with TES. Delta columns compare figures with TES to figures without TES presented in Table 3. Fusion generation is displayed net of imports. Storage charge/discharge is presenting battery and pumped hydro operations.

Optimal Capacities and Dispatch With TES

	Capacity (GW_e)		Generation (TWh_e)		Capacity
	With TES	Delta	With TES	Delta	Factor
Natural gas	9.8	+10 %	2.83	+28 %	3 %
Natural gas w/ CCS	19.4	-16 %	46.11	-9 %	27 %
Solar PV	22.0	+0 %	27.70	+0 %	14 %
Onshore wind	9.0	+0 %	30.47	+1 %	39 %
Offshore wind (fixed)	23.3	-15 %	98.37	-14 %	48 %
Run-of-river hydro	2.0	+0 %	7.78	+0 %	44 %
Fusion	7.7		17.46	+0 %	26 %
Pumped hydro	1.8	+0 %	1.97	-25 %	
Battery	12.5	-33 %	13.55	-47 %	
Total	107.5	-5 %	246.24	-6 %	
Net imports	2.3	+0 %	13.59	+3 %	
Total generation			230.72	-1 %	
Storage charge			-18.44	-45 %	
Storage discharge			15.52	-45 %	
Net Energy			241.39		

Table 6: Optimal energy and power capacities of the storage assets with TES.

	Energy Capacity (GWh_e)	Power Charge (GW_e)	Power Discharge (GW_e)	Duration Charge (h)	Duration Discharge (h)
Battery	28.6	12.5	12.5	2.5	2.1
Pumped hydro	22.0	1.8	1.8	13.3	10.8
Fusion TES*	78.3	2.9	7.7	27.2	16.2

* For the Fusion TES, the discharge duration is calculated assuming the reactor is supplying thermal energy at full power.

Introducing Thermal Energy Storage raised the breakeven cost for the Fusion Power Plant – TES coupling. It now has 7.7 GW_e of optimal capacity (7.2 % of the total capacity). The total capacity installed (generation + storage) is 107.5 GW_e, which is equivalent to 176 % of the maximum load (61.0 GW_e).

The lambda duration curve of the system is presented in Figure 7. The median system lambda is $59.61\ \$/\mathrm{MWh}_e$ and there are $35\ \mathrm{GWh}_e$ of non-supplied energy over the 20 years.

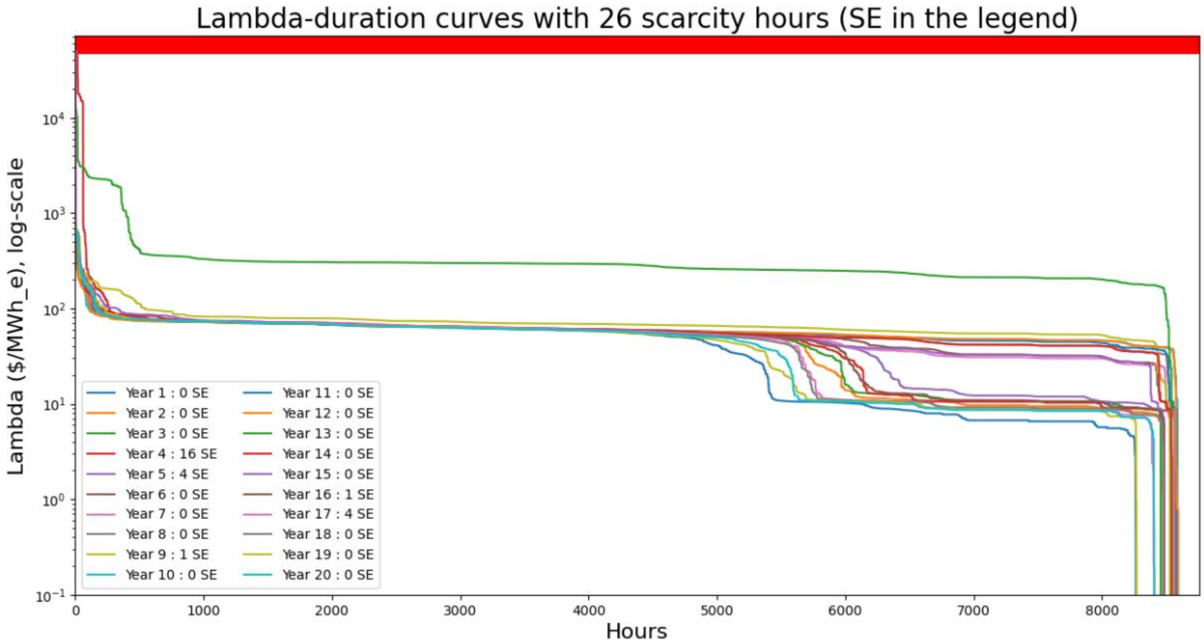

Figure 7: Lambda duration curves with TES, each line is an individual year in the sample.

4.3 System-scale influence of Thermal Energy Storage

This section provides a system-scale summary of the influence of TE. First, Table 7 presents the influence of Fusion Power Plants and TES on the system lambdas. Finally, Table 8 documents the difference between the leftmost parts of Figure 6 and Figure 7, during scarcity hours when some demand cannot be served.

Insignificant or non-defined deltas are simplified.

Table 7: Statistical description of the system lambda distribution with and without Thermal Energy Storage.

		Without TES ($\$/\mathrm{MWh}_e$)	With TES ($\$/\mathrm{MWh}_e$)	Delta
Mean		85.12	82.67	-3 %
StdDev		792.37	733.48	-7 %
Percentiles				
	0%	0.00	0.00	
	1%	0.00	0.00	
	5%	0.00	8.43	
	10%	0.00	9.70	
	25%	26.58	40.76	+53 %
	50%	60.57	59.61	-2 %
	75%	71.88	70.05	-3 %

As Table 7 documents, the consequences of the introduction of Thermal Energy Storage is also visible on the lambda-duration curves (compare Figure 6 and Figure 7). Thousands of hours with very low lambda ($\lambda < 0.1\ \$/\mathrm{MWh}_e$) vanish when Thermal Energy Storage is available. This drives down lambda volatility. This is because when there is less storage ($+123\ \%$ with TES compared to the base case), excess renewable generation capacity must be curtailed when net demand is null, setting the marginal cost of production to $0\ \$/\mathrm{MWh}_e$.

Table 8: Non-served energy and scarcity hours with and without Thermal Energy Storage.

	Without TES	With TES	Delta
Volume (GWh_e)	53	35	-34 %
Average depth (GW_e)	1.46	1.34	-8 %
Median depth (GW_e)	1.90	0.86	-55 %
Max depth (GW_e)	3.52	5.67	+61 %
Scarcity hours	36	26	-28 %
VoLL ($\$/\mathrm{MWh}_e$)	50,000	50,000	
Cost ($)	2,650,000	1,750,000	-34 %

4.4: Analysis of TES operations

TES appear to be a key enabler of capital-intensive next-generation thermal plants. This section unpacks the main driver of this effect at the scale of the coupling itself. It explicates the effective role of TES within the plant-TES-turbine value chain and compares it with alternative storage technologies.

4.4.1 Thermal Energy Storage turnovers

The optimal coupling found has a large energy reservoir ($78.3\ \mathrm{GWh}_e$). It is combining an equivalent $2.9\ \mathrm{GW}_e$ of Fusion capacity with almost twice as much ($7.7\ \mathrm{GW}_e$) of Thermal Energy Storage discharge capacity. When the Fusion Thermal Plant is operated at full capacity, the Thermal Energy Storage can discharge up to $7.7\ \mathrm{GW}_e$ for 16.15 hours. When operated alone, this discharge duration decreases to 10.14 hours. Its charge and discharge durations are longer than most alternatives found in storage literature. As documented in Table 9 the Thermal Energy Storage can operate a maximum of 202 turnovers per year.

Table 9: Storage assets turnover durations and available annual turnovers.

	Turnover Duration (hours)	Max Annual Turnovers
Battery	4.6	1,913
Pumped hydro	24.1	363
Fusion TES*	43.4	202

* For the Fusion TES, the discharge duration is calculated assuming the reactor is supplying thermal energy at full power.

It appears from Table 6 that Thermal Energy Storage is effectively an energy reservoir, while the battery behaves more as a power reservoir. Table 10 indicates that both assets are discharging comparable quantities of energy. The former operates with fewer, longer cycles than the latter. This has a consequence on its operational profits structure.

In the Table 10 below, we try to distinguish the operational profits made from storing energy in the Thermal Energy Storage from those made from selling "flowing-through" energy (at date of production). This result is where the opportunity cost is involved, to account for the transactions between the TES and the upstream plant. At the end, it turns out that separating these two revenue streams is difficult. In our estimation of opportunity costs, TES appear to earn a positive net profit. However, the average annual net profit of the plant+TES, is close to $0\ \$/\mathrm{yr}$, because the net loss of the plant alone balances TES net profit.

Table 10: Optimal dispatch of storage assets

	Annual Discharge	Annual Turnovers		Operational Profit	Net Profit	Mean Spread	Median Spread	Median Duration
	$(\mathrm{GWh}_e/\mathrm{yr})$	$(\#/\mathrm{year})$	(% of max)	($M/year)	($M/year)	$(\$/\mathrm{MWh}_e)$	$(\$/\mathrm{MWh}_e)$	(h)
Battery	13,552	474	25%	963	0	44.99	4.42	8
Pumped hydro*	1,967	90	25%	283	197	63.52	21.46	36
Fusion TES**	12,368	158	78%	788	34	87.64	9.61	17

* Legacy pumped hydro facilities have a positive net profit, because they do not face capital costs

** Thermal Energy Storage has a positive net profit, balanced by Fusion net losses.

The median MWh_e stored in the TES is generated 17 hours after the reactor produced the heat, whereas the battery has a median duration (grid balance) of less than half of that duration. Building on longer storage duration, the TES also capture higher median spreads. During peak/scarcity events lasting for more than 3 hours, TES value is dramatic, when battery discharge duration is insufficient to cover the event in its entirety.

Overall, the Thermal Energy Storage is operating $78\ \%$ of the turnovers it could perform if it were operating continuously. This is more than twice the utilization rate of alternatives like batteries and Pumped storage hydropower ($25\ \%$).

4.4.2 Fusion Thermal Production and Turbine Generation are decoupled

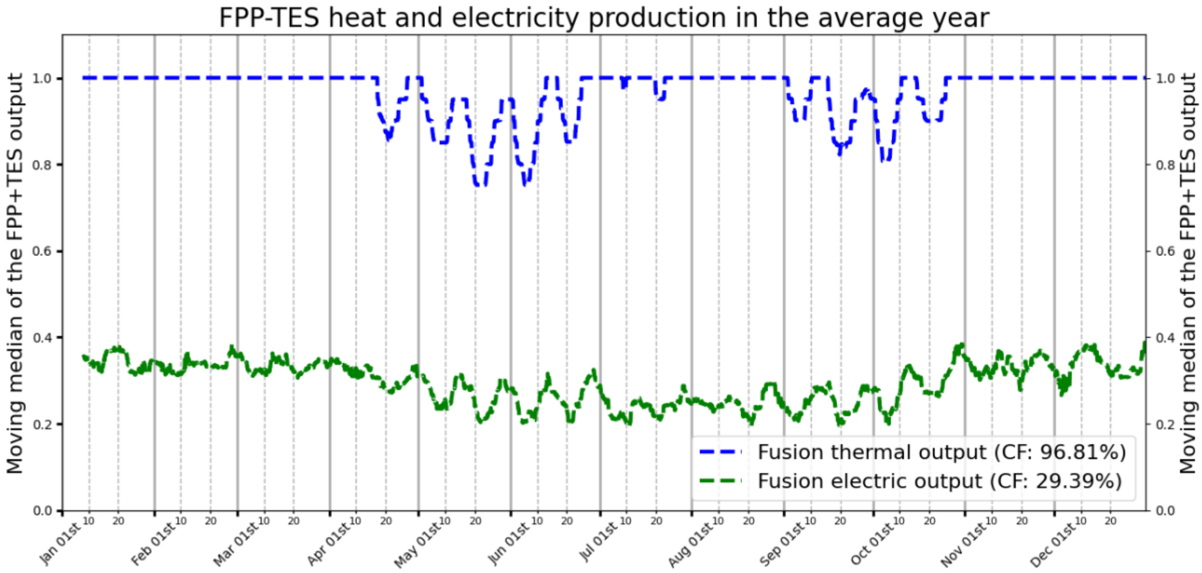

Figure 8: Heat and electricity production of the FPP+TES system on a calendar x-axis. The median reactor output of each week of the average year is the blue dashed line. Same logic applies for the turbine output (green dashed line). The reactor is basically operated in baseload (97 % capacity factor), whereas the turbine acts as a load-following asset (29 % capacity factor).

Figure 8 illustrates that the three independent dimensions of Thermal Energy Storage are optimized to decouple the reactor operations from the generation. The TES effectively allows almost-continuous reactor operations at nameplate capacity, whereas the TES discharges when the generation yields higher profits.

In other words, adding Thermal Energy Storage to a capital-intensive thermal plant significantly alters the set of hours in which its energy is delivered to the system. Without storage, the plant delivers its energy to the grid when the average marginal value of energy is X. Whereas with Thermal Energy Storage, the plant shifts the profile of hours when its energy is delivered, raising the average marginal value to Y. Thus, adding TES can justify investment in a plant even when the average cost of energy produced is Z, with $X < Z < Y$.

4.4.3 Thermal Energy Storage operations

To go beyond this simple explanation, Table 11 unpacks the Thermal Energy Storage discharge operations using the reference points set in Figure 5. Segment BC and floating points are hidden as they contribute to less than one percent of the total operational profits and hours.

Table 11: Analysis of the operational profits made, segmented by operational mode as identified in Figure 5.

Net thermal discharge		**Share of operational profits**		**Share of hours**	
MW-e	%	AB	CD	AB	CD
< 542.30	< 90%	< 2.8%		< 18.1%	
602.56	100%	44.7%		20.0%	
662.82	110%		0%		0%
723.07	120%		0%		0%
783.33	130%		0%		0%
843.58	140%		0%		0%
903.84	150%		0%		0%

964.10	160%	51.3%	1.0%

To illustrate how the respective shares of operational profits and hours can be so different, we use the same x-axis as in Figure 5. It appears on Figure 9 that operations occur at a variety of price levels, whereas operational profits are heavily localized into scarcity hours as illustrated in

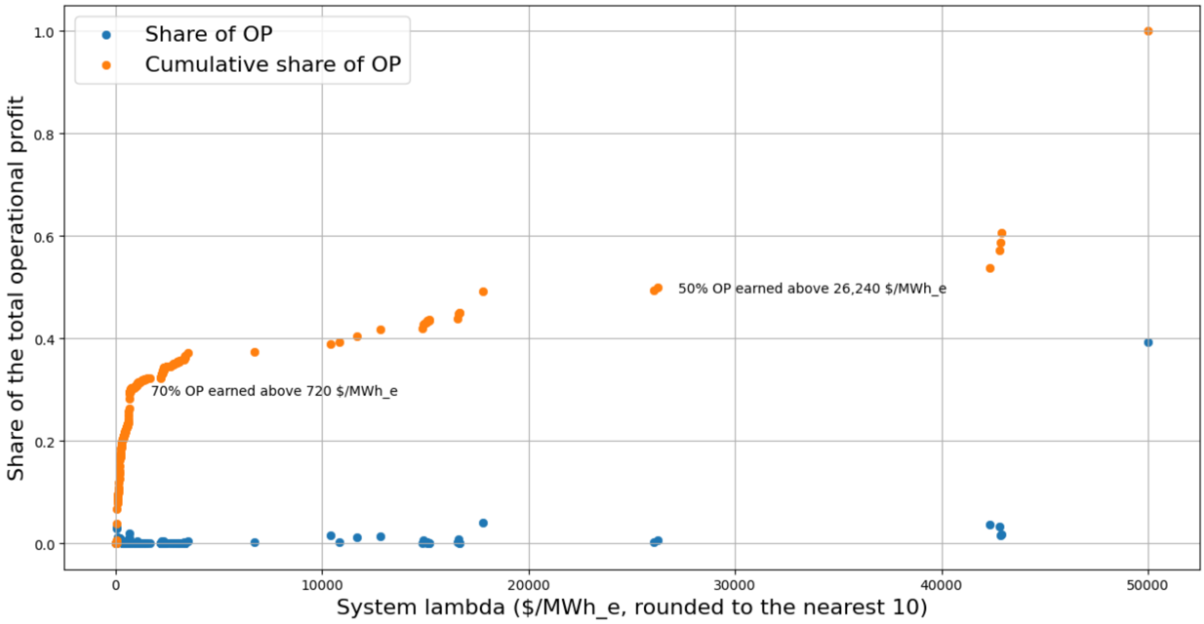

Figure 10.

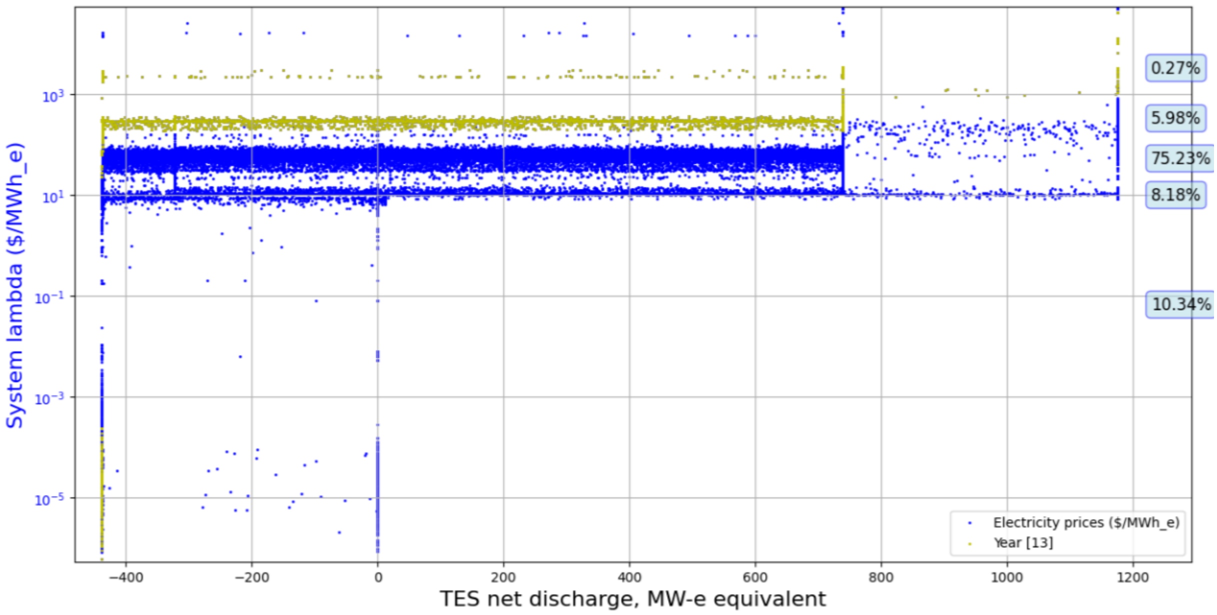

Figure 9: TES net charge/discharge operations contextualized with system lambda ($/MWhe).

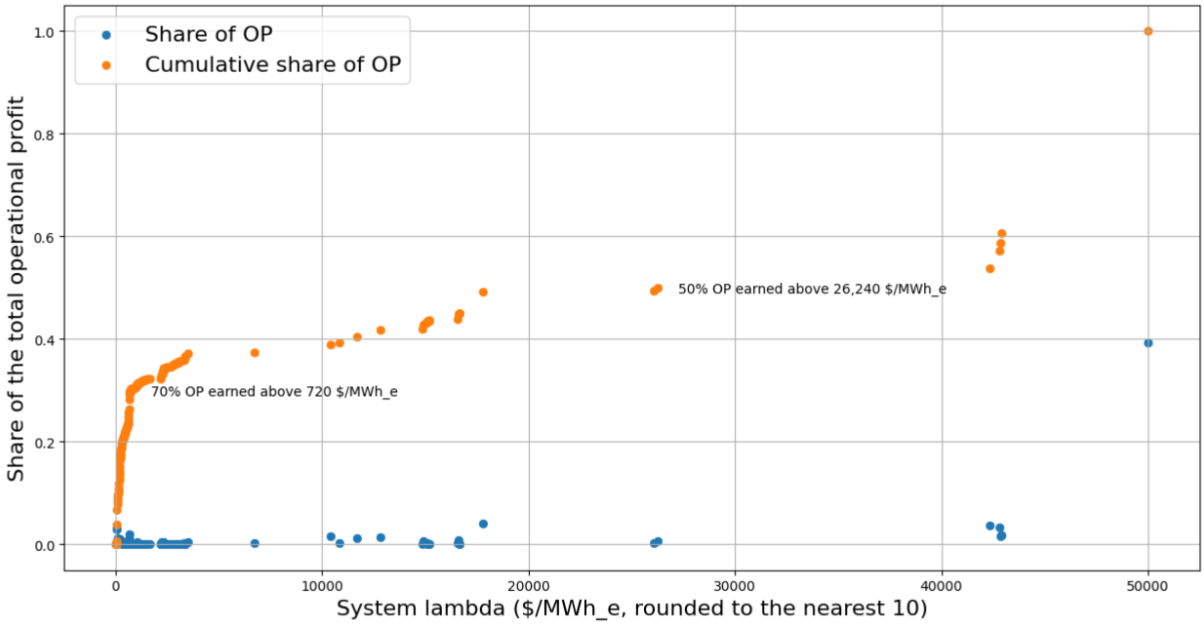

Figure 10: Cumulative share of operational profits earned at various system lambda levels.

Figure 9 indicates that a total of $90\ \%$ *of charge/discharge operations occur in clearly visible lambda bands. Year by year, the exact position of the band will depend on the lambda distribution. When a year has lambdas significantly higher than the others in the dataset, it is easy to identify it (see yellow marks identifying Y13 operations). However, Table 11 indicates that operational profits are concentrated in two vertical segments at* $net\_dis_t = Net\_dis\_cap$ *and* $net\_dis_t = Dis\_cap$*. This is also confirmed in*

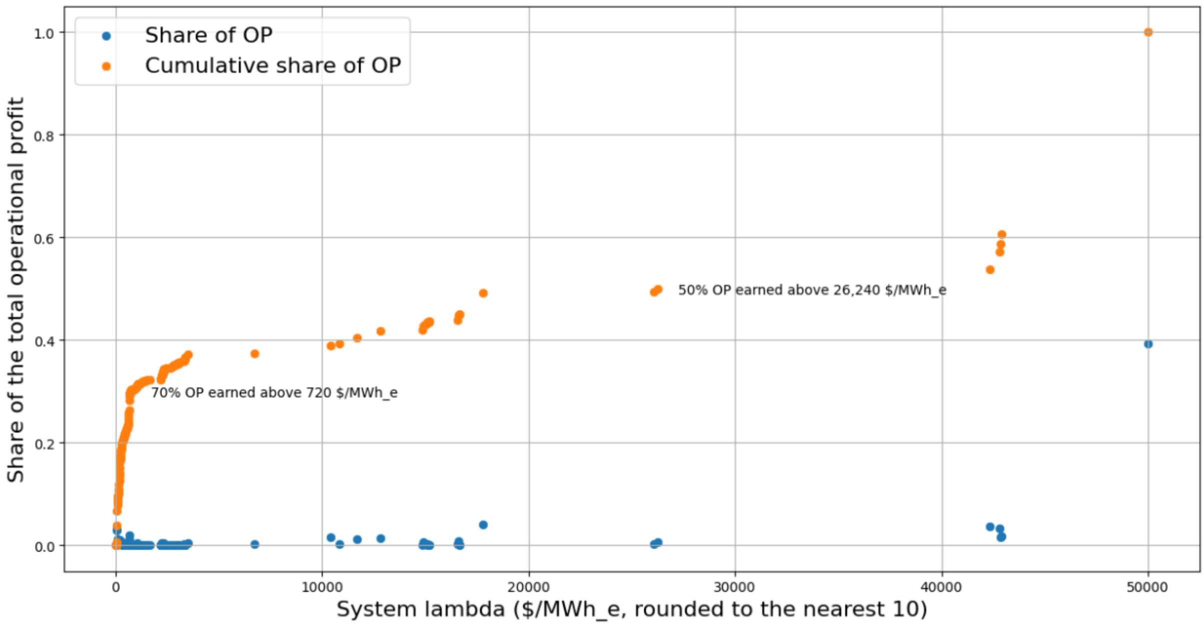

Figure 10, with 70 % of total operational profits earned when lambda is above 720 $\$/\mathrm{MWh}_e$ and 50 % above 26,240 $\$/\mathrm{MWh}_e$ alone.

Overall, we demonstrate that adding TES to a Fusion Power Plant is an effective way to minimize the system costs. Moreover, adding TES and turbines without oversizing the reactor enable boosted generation above the nameplate capacity of the reactor, focused on higher lambda hours. In some cases, when the variable recirculating needs of the reactor are using too much turbine capacity, some reactors can be shut down to reduce load and the TES then operates at even higher discharge rates.

4.5 Concentration of storage profits in hours and years

Understanding the concentration of operational profits is pivotal for several reasons. It documents how sensitive profitability is to main assumptions such as available technologies, or real-world price caps. In our model, we chose to not use any market structure *per se* to ensure tractability of our observations beyond any local specificities or regulatory paradigms.

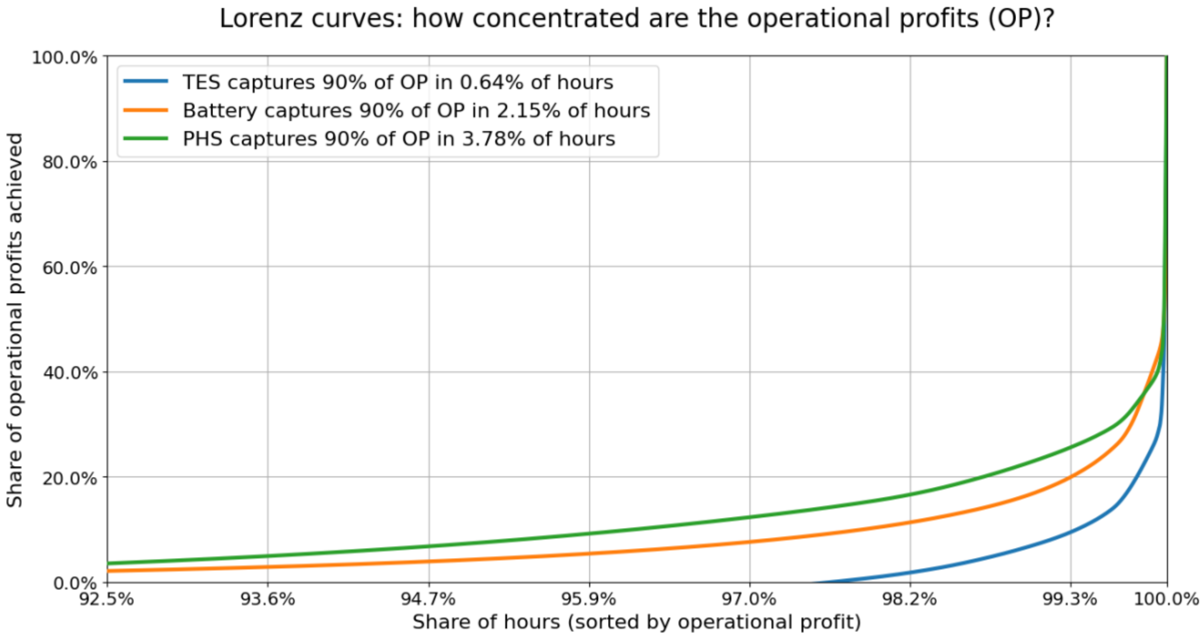

Figure 11: Lorenz curve of storage assets' conceptual operational profits. It documents if the remunerative cost reductions are concentrated in a few hours (curve in the bottom right corner, as pictured here) or, conversely, evenly spread across the year (closer to the first bisection).

In a cost minimization approach, the storage assets capture most of their profits off a few hours in the dataset (peak lambda hours and scarcity events). We demonstrate that TES profits are more concentrated than those of batteries and PHS:

Table 12: Statistical description of the concentration of operational profits.

Duration to capture … of operational profits	99 %	90 %	80 %	50 %
Battery	10.20 %	2.15 %	0.68 %	0.02 %
Pumped hydro	11.20 %	3.78 %	1.30 %	0.02 %
Fusion TES	2.04 %	0.64 %	0.21 %	0.02 %

For all tested storage specifications, half of the annual operational profits are made during extreme events. The modeling shows that TES is valuable, in that it lowers system cost. An overwhelmingly large amount of this cost reduction is concentrated in a small number of hours with an extremely high system lambda. Hence, should a price cap be applied in a market-based approach, it appears that TES – yet more generally storage assets – capacity accreditation is a major issue for capacity market designers to ensure the inclusion of the optimal capacity of this technology.

Moving beyond concentration of operational profits in a single year, we use 20 independent years to ensure robustness of our results. Whereas Thermal Energy Storage profits are concentrated in fewer hours than batteries and Pumped storage hydropower, Figure 12 illustrates the fact that

Thermal Energy Storage value is carried in a number of years similar to alternatives. Thermal Energy Storage profits are concentrated in 4 years out of 20, while battery's ones are concentrated in 6 out of 20. While Pumped storage hydropower were already built in our system, profits are even more concentrated for that technology (2 out of 20).

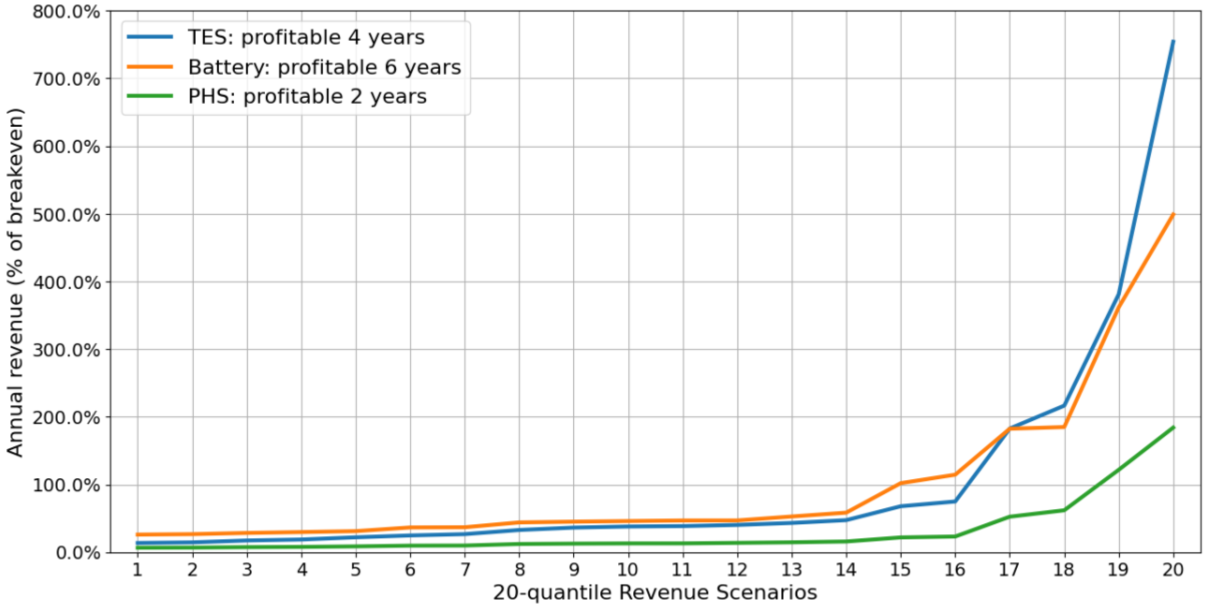

Figure 12: 20-quantile revenue scenarios for storage technologies. A year at 100 % indicates that annual operational profits balance the fixed costs of the asset. Being above means being profitable, whereas falling below translates into net losses.

The optimal capacities of storage assets, robust to 20 annual samples, appear profitable only one-in-a-few years. Battery pays back its costs in $25 - 30$ % of the sample, TES in 20 % and Pumped storage hydropower is profitable only 10 % of the simulated years. This is a decisive demonstration of the importance of multi-year scenarios, given the uncertainties on demand growth, resource availability or fuel prices. Ultimately, TES appears not significantly riskier than alternatives, given that ability to catch peaks is crucial for all storage technologies profitability (50 % of operational profits).

We illustrate in greater detail this concentration of operational profitability in a few years for TES, for which the four profitable years are generating significant profits, whereas two of the six profitable years of batteries are almost at-the-money.

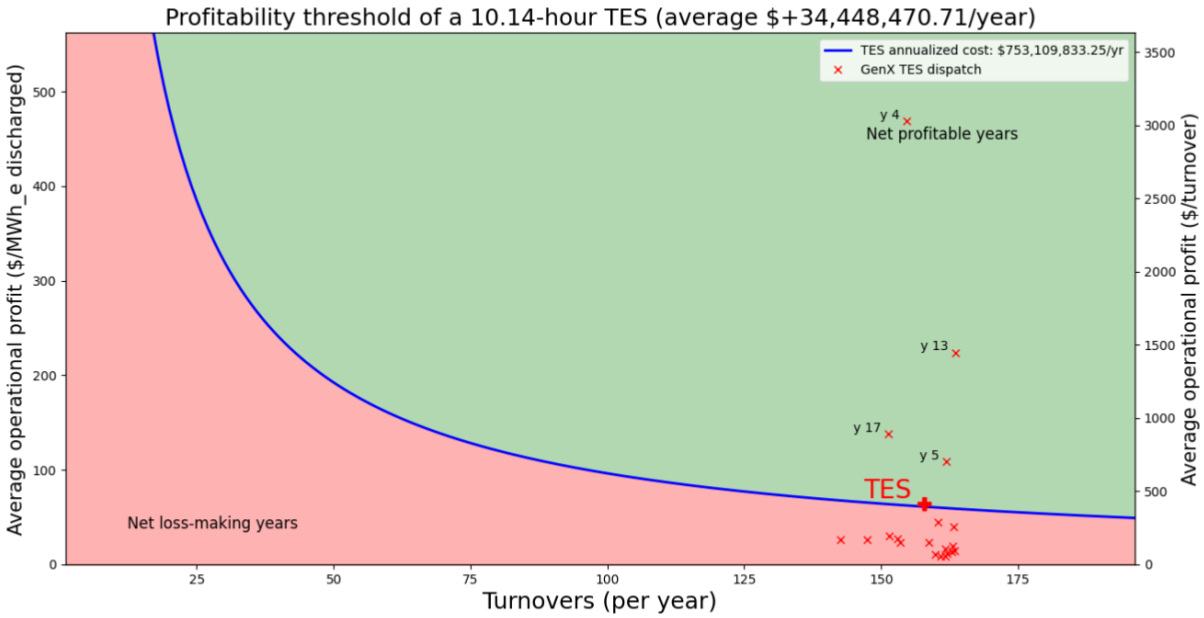

Figure 13: Profitability threshold graph for TES. The blue line depicts the operational profit per unit (OP/u) to break-even (y-axis) given a number of turnovers in the year (x-axis). Green and red filled areas are respectively in-the-money and out-of-the-money. Small red marks are the years taken individually. The thick red mark is the barrycenter, or the average turnovers and average (OP/u). Being in the green area means "net profitable", whereas in the red area indicates losses.

Figure 13 identifies the four years from which TES yields its value (in the green area). We note that three of the four present scarcity events (4, 5 and 17), whereas the last (13) doesn't. The lambda-duration curve of the year 13 has the distinctive feature of being higher in median and average compared to other years in the sample (poor resource availability). There is no significant advantage of years with scarcities compared to year 13 without, as year 13 is the second most profitable year of the sample, meaning two years with scarcities are less profitable.

5. Discussion

Thermal Energy Storage creates a competitive space for next-generation capital-intensive nuclear technologies

We demonstrate that introducing Thermal Energy Storage can effectively enable next-generation nuclear technologies to penetrate the optimal mix. This also means that it reduces the system's total cost. This effect is reflected by significant reductions visible in the statistical distribution of lambda (compare Figure 6, Figure 7, and see Table 7, and Table 8).

Capacity of fusion increase dramatically from $0\ \mathrm{GW}_e$ to $7.7\ \mathrm{GW}_e$ when adding Thermal Energy Storage to the reference case ($8{,}500\ \$/\mathrm{kW}_e$). The consequence of doing so is clearly visible on lambda values: the average cost of generating $1\ \mathrm{MWh}_e$ decreases by $3\ \%$ with TES, whereas the median decreases by $2\ \%$. This illustrates, in fact, two distinct effects: first, the lambda duration curve becomes significantly less volatile ($-7\ \%$). The thousands of hours of low lambda visible in Figure 6,are traits of renewable energy surplus and disappear almost totally when TES is available (see Figure 7).

Second, there is a pivotal effect during peak and scarcity hours. These extreme events are not visible in the lambda-duration curves but remain significant for system operators. Scarcity events, times when all customers cannot be supplied and energy is scarce, are key indicators of the reliability of the grid. Without TES, there are $53\ \mathrm{GWh}_e$ of non-supplied energy in 36 hours, whereas with TES this volume shrinks to $35\ \mathrm{GWh}_e$ in 26 hours. This $34\ \%$ decrease in volume and, more importantly, the 10 avoided hours of scarcity indicate that TES has a very positive influence on grid reliability.

As Figure 14 illustrates, the dramatic value of Thermal Energy Storage can be observed for all tested capital costs of the benchmark Fusion Power Plant, up to $19{,}000\ \$/\mathrm{kW}_e$.

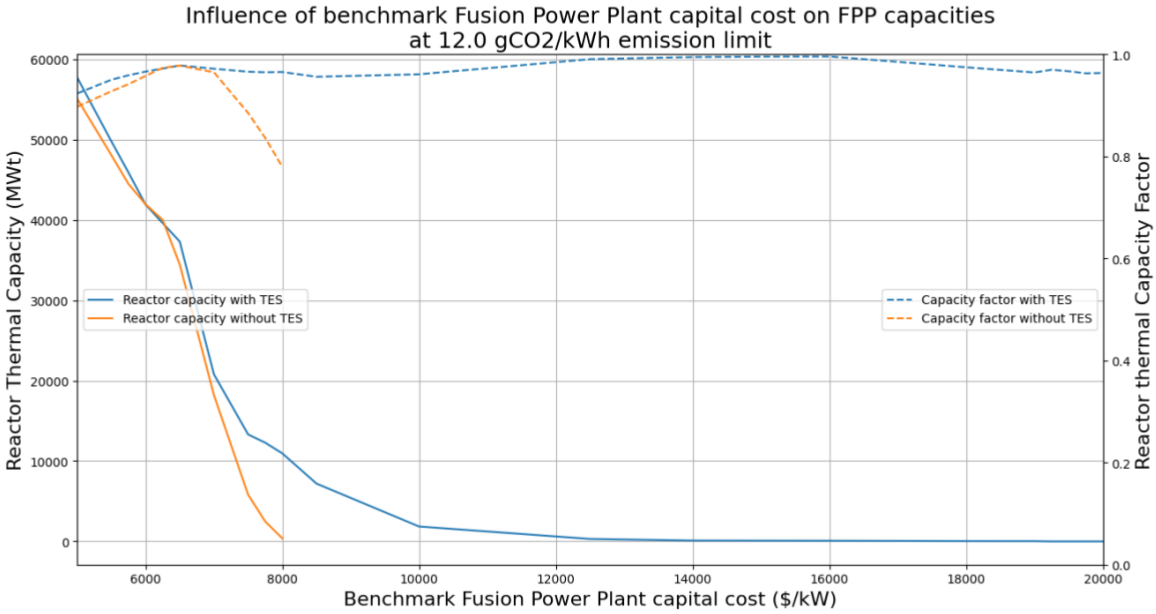

Figure 14: Evolution of the penetration of next-generation nuclear asset with and without Thermal Energy Storage. First, it appears that TES consistently increase the optimal turbine capacity. More importantly, the capital cost limit above which the optimal capacity is zero is pushed far to the right with TES. Second, for each tested capital cost, the reactor appear operated at a higher capacity factor when there is Thermal Energy Storage coupled to the nuclear reactor.

We demonstrate that this impact has far reaching consequences. It effectively raises the breakeven cost of the thermal plant from $8{,}000\ \$/\mathrm{kW}_e$ to more than $19{,}000\ \$/\mathrm{kW}_e$, as pictured in Figure 14. The former is optimistic, whereas the latter is credible, even for recent First-of-a-Kind projects. Moreover, Figure 14 illustrates another positive effect of coupling the reactor to Thermal Energy Storage: the reactor capacity factor is consistently increased compared to cases without. In fact, it depicts that Thermal Energy Storage preserves the reactor from modulations during lambda depressions. This is a key enabler of next-generation assets that are not allowed to modulate, either for technological reasons or financial reasons.

Finally, we discuss how Thermal Energy Storage competes with battery storage. First, Figure 15 illustrates how TES dimensions evolve with benchmark Fusion Power Plant capacity costs, whereas Figure 16 depicts how battery storage capacities are consistently reduced when Thermal Energy Storage is available. Last, Figure 17 documents the TES durations evolution: the TES charge duration increase significantly with rising capital costs of the benchmark Fusion Power Plant, whereas its discharging duration remains somehow constant.

Moreover, battery's discharge duration without TES appears somehow constant across the cost scenarios tested. With TES however, the battery's discharge capacity appears to increase faster than the storage capacity: the discharge duration will present a "belly" with lower values between 8,000 $\$/\mathrm{kW}_e$ and 10,000 $\$/\mathrm{kW}_e$.

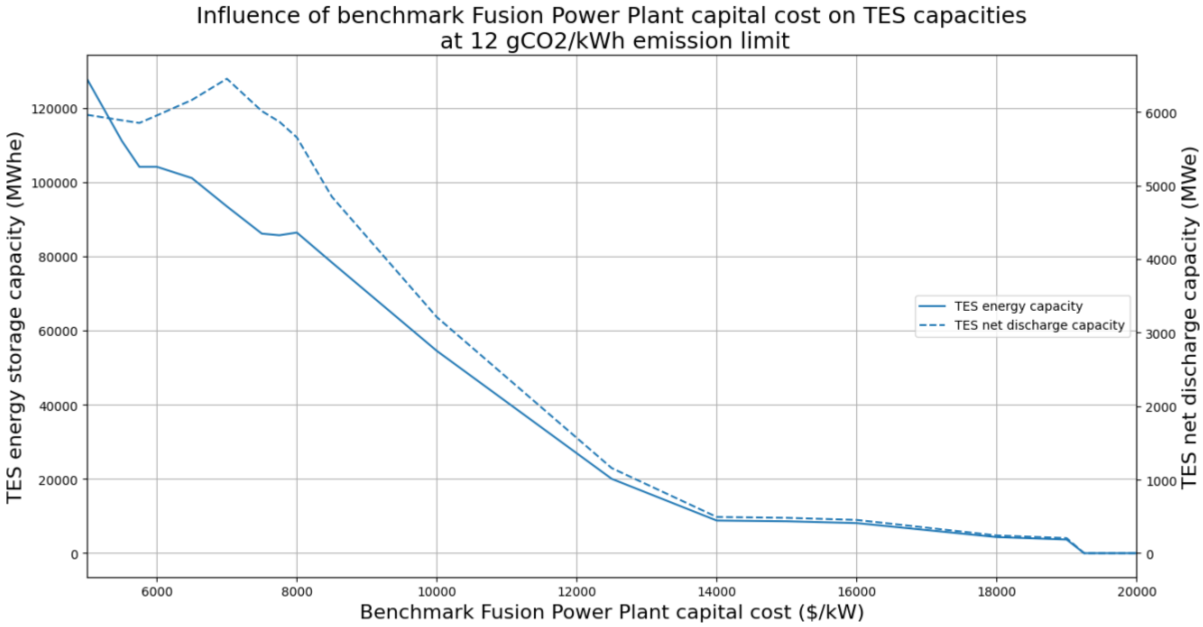

Figure 15: Evolution of TES capacities when varying the capital costs for the benchmark Fusion Power Plant.

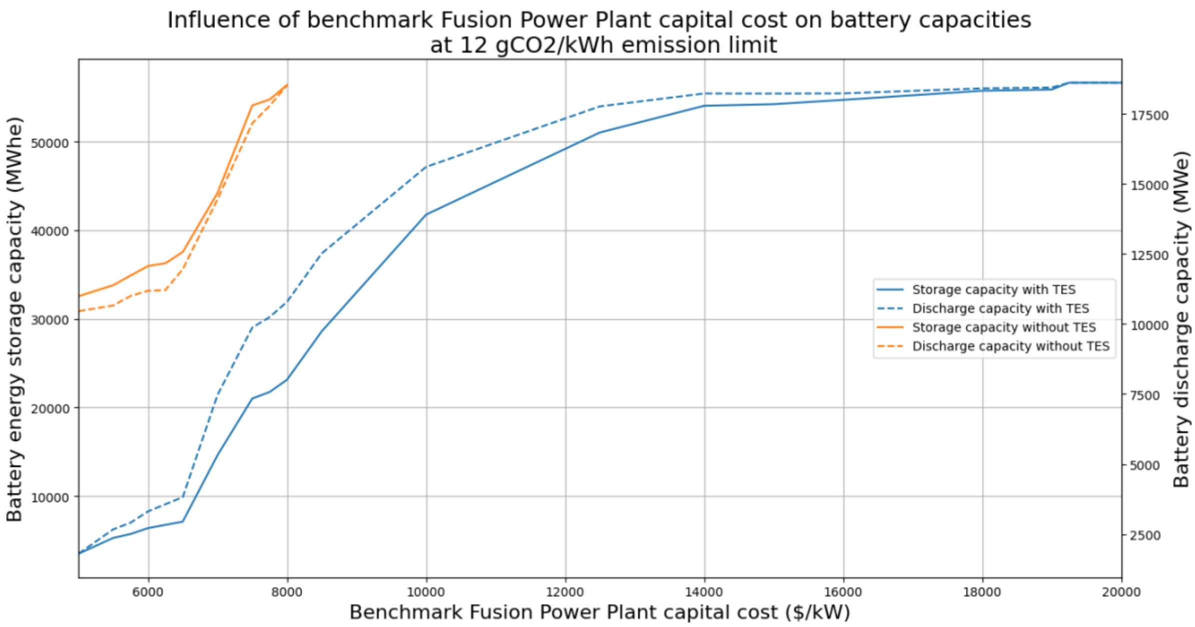

Figure 16: Evolution of battery capacities with and without Thermal Energy Storage when varying the capital costs for the benchmark Fusion Power Plant.

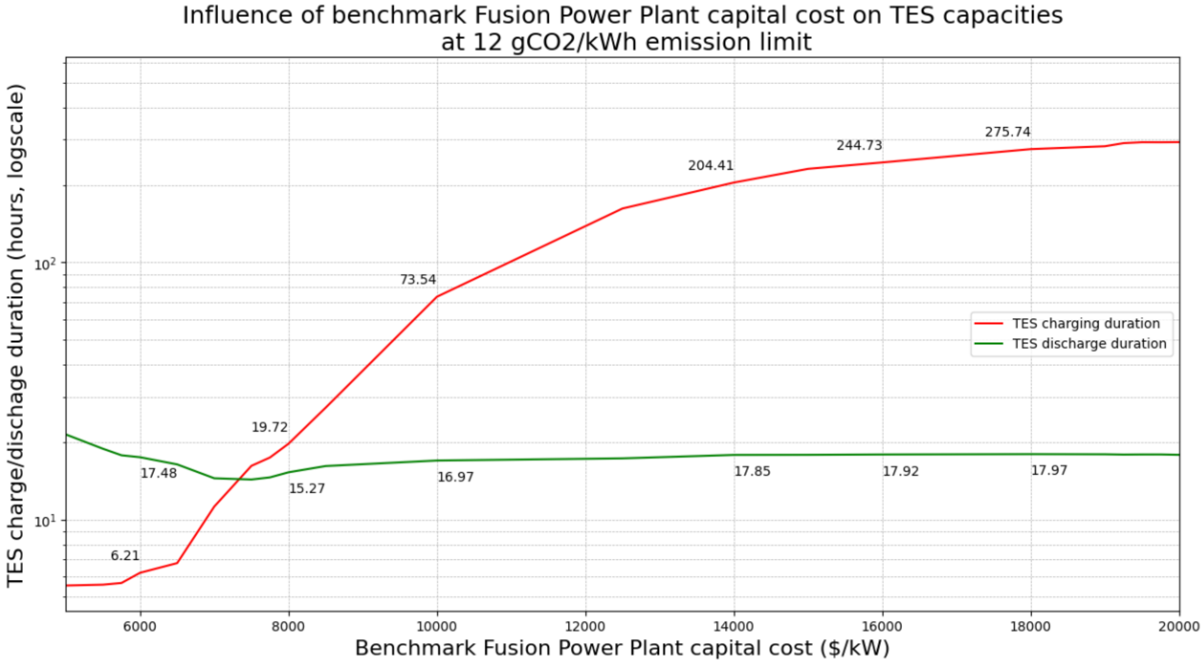

Figure 17: Evolution of TES charging/discharging durations when varying the capital costs for the benchmark Fusion Power Plant.

This first methodological limit of this paper is the assumption of linearity for assets' commitment. The optimizer can set the number of operating plants of each technology without any integer constraint. We justify this choice because of the algebraic complexity of such constraints in a large model simulated each hour for twenty years, so for 175,200 timesteps. Moreover, and as the core matter of this paper, the theoretical next-generation nuclear asset modeled in this paper is assumed to be perfectly flexible and linear. This results in an unequivocal underestimation of Thermal Energy Storage role in smoothing the plants' operations. This deserves specific additional research, especially in the case of the integration of realistic flexible nuclear plants in capacity expansion dispatch models. For example, to implement this paper's model with next-generation pressurized-water reactors, the implementation of xenon transients is studied by Jenkins et al. (2018), and the one of limited maneuverability because of pellet-cladding interaction studied by Trosino et al. (2026) may be considered.

Shortly, our main goal was to demonstrate the system value of decoupling reactor and turbine operations using Thermal Energy Storage. In fact, Thermal Energy Storage may enable additional cost reductions.

Thermal Energy Storage turns inflexible technologies into firm dispatchable plants to coexist with renewables

More than just illustrating the penetration or not of next-generation capital-intensive assets, our results document deeper plant-scale effects. With Thermal Energy Storage we demonstrate that heat production and electricity generation are not necessarily occurring simultaneously. It goes beyond the next hour and effectively decouples the reactor from the turbine.

This feature is essential for the system to be economically viable and to cover its costs. With TES, the reactor can produce more thermal energy, almost all year long (97 % of capacity factor) because it can delay the conversion to electricity to more remunerative hours when lambda is higher. It will generate power synergistically with the intermittency of other assets but keep an almost constant heat production.

Figure 8 is particularly eloquent on that matter: the reactor production shows a limited seasonality, whereas the turbine is used only $29\ \%$ of the time. In the reference case without Thermal Energy Storage, the two lines would be superimposed, and the plant would not be able to leverage the storage to optimize its operations. The capital-intensive reactor would only operate when it is economically marginal or inframarginal, which is not true for a significant part of the years simulated.

Figure 14 depicts how this effect is independent from capital cost assumptions: the capacity factor of the reactor is consistently close to $100\ \%$ when coupled to a Thermal Energy Storage. It is interesting to note the "bell" effect on the leftmost part of the graph: the reactor capacity factor increases and then decreases. This is because when nuclear capital costs are low enough, the economic incentive to operate at full power becomes less stringent. However, when nuclear capital costs are rising, the reactor can leverage on Thermal Energy Storage to maximize its rate of return on investment.

Beyond this plant-scale effect, we compare Thermal Energy Storage dimensions to other technologies of storage. Pumped hydroelectricity capacities being fixed, batteries are the only effective and scalable alternative to TES. A striking result is the tenfold ratio between TES and batteries' durations.

It implies quite different operational patterns for both technologies; the former being used as an energy reservoir with fewer turnovers available per year, and the latter as a power reservoir able to cycle two thousand times a year. The optimal dispatch illustrates that TES is using a significantly higher share of its available turnovers than batteries, $78\ \%$ for Thermal Energy Storage versus $25\ \%$ for batteries (up to $31\ \%$ in the reference case without TES).

However, these results do not account for a significant role that Thermal Energy Storage can play in real-world systems. While batteries do not provide any inertia other than synthetical, TES turbines are massive rotating pieces of steel. They would make a major contribution to grid stability, no matter if they produce at nameplate capacity or not. Figure 8 clearly shows that the turbine shaft is basically rotating all year round, providing this potentially remunerative ancillary service at no cost. It is a completely uncharted perspective here and a conservative limit of this paper, which deserves additional research.

Thermal Energy Storage has a hedge value against extreme events

The last category of this paper's results is centered on the conceptual operational profits. In our first best model (or cost-optimal mix), the hours in which the storage assets achieve the remunerative cost reductions are two times concentrated. First, Figure 11 depicts the high concentration of operational profits in some hours of each year: the scarcity or highest peak hours. It means that storage economic viability is highly reliant on extreme lambda values, up to $50{,}000\ \$/\mathrm{MWh}_e$.

However, such values are not socially acceptable out of cost minimization. Under specific assumptions, cost minimization is somehow equivalent to energy-only market designs. In these, lambda would be an indicative cost of electricity paid by customers. In fact, real market designs may dramatically lower the maximum value (for example $12{,}000\ \$/\mathrm{MWh}_e$ or even lower), obliterating the viability of peak assets. To achieve a second best, it is necessary to compensate for the operational losses of generators that are in the optimal mix but not profitable anymore with sole energy trading.

To that end, other mechanisms, such as capacity markets, are created. It appears that most of the operational profits of storage assets would be trimmed, requiring significant revenues to come from alternative markets. In the case of capacity markets, storage capacity accreditation becomes a major issue to support the inclusion of sufficient capacities in the system. It is outside the scope of this paper but deserves thorough consideration in future work.

Second, Figure 13 illustrates that Thermal Energy Storage viability is not evenly spread throughout the years but concentrated in 4 out of 20. This is also true for other storage technologies; batteries' operational profits being concentrated in 6 years and pumped hydroelectricity in only 2. This puts to the fore that storage assets' cost reductions occur in years with scarcities or with low renewable availability (year 13 in our sample). The single optimal mix presented in this paper is robust to twenty possible load and resource availability profiles: single-year results would differ significantly. This result is consistent with the recent reports by the Independent System Operator of New-England (ISO-NE): certain assets that are essential for reliability will be operated once in a few years, depending on weather conditions (ISO-NE EPCET report, 2024). They need either to leverage on extreme price levels, or to be compensated for their available capacity no matter what their effective production is. This is the core principle of capacity markets, though outside this paper's perimeter.

This last element is pivotal for future work. It sets a new standard for future capacity expansion dispatch model simulations: there is a need for longer simulations to realistically estimate the cost reduction achievable through storage. Otherwise, if the inputs were inducing too many scarcities, it would systematically overestimate the value and the optimal capacities of storage. Conversely, with "easier" years featuring fewer peak hours or scarcity events, coping with the latter would be pricier or impossible, ultimately causing reliability issues or blackouts.

The second methodological limit of this paper is the assumption of perfect foresight. It enables the model to accurately predict the peaks and to size the capacities accordingly. In the last paragraph, we demonstrated that Thermal Energy Storage would be very valuable in the case of scarcity events not anticipated more than a day before the loss of reliability. While one-in-ten-year events cannot be planned with great precision, it appears legitimate to assume that grid operators can predict accurately the need for a demand-side response a day or more before a blackout, at least in non-accidental scenarios. This intuition could be confirmed with imperfect foresight simulations such as rolling horizons in future work.

Hence, we demonstrate the critical *hedging* value of Thermal Energy Storage during scarcity events and the need for additional research in capacity accreditation of asymmetrical storage systems.

6. Conclusion

In this paper, we demonstrated the key value of Thermal Energy Storage (TES) as a technology able to **reconcile the economic constraints of capital-intensive low-carbon thermal generators such as nuclear power plants with the operational flexibility required in deeply decarbonized systems**. TES capacities can fully decouple the reactor operations from the turbine generation. The consequence of this result is twofold: (i) it creates a competitive space for advanced low-carbon technologies that are expensive to build, actively reducing the impact of market volatility on effective operations; (ii) as variable renewable energy reshapes residual load patterns, it enhances the flexibility of assets with operations constrained to baseload output, hence reducing the cost of operations and maintenance or shutdowns.

Using GenX, a Capacity Expansion and Dispatch Model, we demonstrate that integrating TES fundamentally changes the system-scale dynamics, reducing the total system cost and reducing the lambda volatility. However, we inform that the value of TES is highly reliant on specific events and/or specific years. Some extreme events can carry most of the profits to make the TES break even. We demonstrate that the profit concentration for TES is in the common range, between batteries and pumped hydro storage.

Applying our results into real world systems requires tackling a critical issue for system resilience. How to remunerate thermal energy storage in capacity markets, given that discharge capacity is not always available? This question will be central to enable the effective integration of nuclear-TES couplings in decarbonized systems. Additional research is also required on the value of inertia carried by TES turbines in ancillary services and frequency regulation.

By examining the interplay between system conditions, technology costs, and operational constraints, we demonstrated that nuclear-TES hybrids can play a pivotal role in future decarbonized electricity mixes. Their ability to deliver firm, low-carbon energy while dynamically responding to net-load variability positions them as a strategic complement to renewables and as a long-duration storage solution. This work provides a quantitative foundation for assessing when, where, and why such hybrid systems become economically optimal, and what this implies for nuclear deployment strategies in the coming decades.

Acknowledgments

Nicolas Trosino gratefully acknowledges financial support for this research by the Fulbright Program, which is sponsored by the U.S. Department of State and the Franco-American Commission – Fulbright France. Its contents are solely the responsibility of the author and do not necessarily represent the official views of the Fulbright Program, the Government of the United States, or the Franco-American Commission.

Drs. Parsons' and Forsberg’s contributions were made possible by a grant from the DOE Office of Nuclear Energy's Nuclear Energy University Program (22-26882) and from the MIT Energy Initiative’s Future Energy Systems Center.

References

Alameri S. A. & King J. C. (2013). “A coupled nuclear reactor thermal energy storage system for enhanced load following operation.” *International Nuclear Atlantic Conference*, Recife, PE, Brazil.

Antweiler, W., & Müsgens, F. (2025). “The new merit order: The viability of energy-only electricity markets with only intermittent renewable energy sources and grid-scale storage.” *Energy Economics*, *145*, 108439.

Armstrong R. C. et al. (2024). “The role of fusion energy in a decarbonized electricity system.” Report by MIT Energy Initiative

Boiteux M (1960). “Peak-load pricing.” *The Journal of Business*. Vol. 33, No. 2, pp. 157-179.

Boiteux M (1964). "The choice of plant and equipment for the production of electric energy." In J. Nelson (Ed.), *Marginal Cost Pricing in Practice*. Englewood Cliffs, N.J.: Prentice-Hall.

Denholm, P., King, J. C., Kutcher, C. F., & Wilson, P. P. (2012). "Decarbonizing the electric sector: Combining renewable and nuclear energy using thermal storage." *Energy Policy*, *44*, 301-311.

Edwards, J., Bindra, H., & Sabharwall, P. (2016). "Exergy analysis of thermal energy storage options with nuclear power plants." *Annals of Nuclear Energy*, *96*, 104-111.

Forsberg, C., & Sabharwall, P. (2018). "Heat Storage Options for Sodium, Salt and Helium Cooled Reactors to Enable Variable Electricity to the Grid and Heat to Industry with Base-Load Reactor Operations (No. INL/EXT--18-51329-Rev000)." Idaho National Laboratory (INL), Idaho Falls, ID (United States).

Frick, K., Misenheimer, C. T., Doster, J. M., Terry, S. D., & Bragg-Sitton, S. (2018). "Thermal energy storage configurations for small modular reactor load shedding." *Nuclear Technology*, *202*(1), 53-70.

Gomez, J. C. (2011). "High-temperature phase change materials (PCM) candidates for thermal energy storage (TES) applications (No. NREL/TP-5500-51446)." National Renewable Energy Laboratory (NREL), Golden, CO (United States).

Green, M., Sabharwall, P., Mckellar, M. G., Yoon, S.-J., Tech, G., Abel, C., Petrovic, B., & Curtis, D. (2013). "Nuclear Hybrid Energy System: Molten Salt Energy Storage." Idaho National Laboratory (INL), Idaho Falls, ID (United States).

Herrmann, U., & Kearney, D. W. (2002). "Survey of thermal energy storage for parabolic trough power plants." *J. Sol. Energy Eng.*, *124*(2), 145-152.

Jenkins, J. D., Zhou, Z., Ponciroli, R., Vilim, R. B., Ganda, F., de Sisternes, F., & Botterud, A. (2018). "The benefits of nuclear flexibility in power system operations with renewable energy." *Applied energy*, *222*, 872-884.

Junge, C., Mallapragada, D., & Schmalensee, R. (2022). "Energy storage investment and operation in efficient electric power systems." *The Energy Journal*, *43*(6), 1-24.

Keith E, Lindley B, Wagner M J. "Thermodynamic and Technoeconomic Analysis of Desalination with Thermal Energy Storage." ANS Winter Meeting 2023, Nov 12-15, 2023. Washington, DC

Keith, E., Lindley, B., & Wagner, M. J. (2025). "Optimized dispatch and component sizing for a nuclear-multi-effect distillation integrated energy system using thermal energy storage." *Optimization and Engineering*, *26*(2), 913-938.

Knighton, L. T., Shigrekar, A., Wendt, D. S., Frick, K. L., Boardman, R. D., Elgowainy, A. A., ... & Reddi, K. R. (2021). "Energy arbitrage: Comparison of options for use with LWR nuclear power plants (No. INL/EXT-21-62939-Rev000)." Idaho National Laboratory (INL), Idaho Falls, ID (United States).

Medrano, M., Gil, A., Martorell, I., Potau, X., & Cabeza, L. F. (2010). "State of the art on high-temperature thermal energy storage for power generation. Part 2—Case studies." *Renewable and Sustainable Energy Reviews*, *14*(1), 56-72.

Novotný, V., Kim, J., Cho, S. B., & Saeed, R. M. (2025). “Nuclear—thermal energy storage configurations for industrial combined heat and power supply—conceptual and thermodynamic study with high temperature gas-cooled reactor.” *Energy*, *322*, 135544.

Parzen, M., Fioriti, D., & Kiprakis, A. (2023). “The value of competing energy storage in decarbonized power systems.” *arXiv preprint arXiv:2305.09795*.

Saeed, R. M., Shigrekar, A., Mikkelson, D. M., George Rigby, A. C., Yang Hui Otani, C. M., Garrouste, M., ... & Bragg-Sitton, S. M. (2022). “Multilevel analysis, design, and modeling of coupling advanced nuclear reactors and thermal energy storage in an integrated energy system (No. INL/RPT-22-69214-Rev000).” Idaho National Laboratory (INL), Idaho Falls, ID (United States).

Saeed, R. M., Novotny, V., Cho, S. B., Shigrekar, A., Otani, C., Toman, J., ... & Mikkelson, D. M. (2023). “A Multidisciplinary Approach to Integrated Energy Systems: Advanced Nuclear Plants with Thermal Storage for Dynamic and Flexible Operation in Diverse Markets (No. INL/RPT--23-74725-Rev000).” Idaho National Laboratory (INL), Idaho Falls, ID (United States).

Saeed, R. M., Novotný, V., Cho, S. B., Shigrekar, A., Otani, C., Toman, J., & Mikkelson, D. (2024). “Design and optimization of flexible decoupled high-temperature gas-cooled reactor plants with thermal energy storage.” *Energy Conversion and Management*, *321*, 119098.

Schmalensee R. et al. (2015). “The Future of Solar Energy.” Report by MIT Energy Initiative

Schwartz, J. A., Ricks, W., Kolemen, E., & Jenkins, J. D. (2023). “The value of fusion energy to a decarbonized United States electric grid.” *Joule*, *7*(4), 675-699.

Turvey R (1968). “Optimal Pricing and Investment in Electricity Supply: An Essay in Applied Welfare Economics.” Cambridge, MA: MIT Press.

Appendix A

Mathematical model of direct/indirect Thermal Energy Storage

With $T \in \mathbb{N}$ the number of timesteps, we note the thermal flows (normalized to their maximum) $\dot{Q}_c \in [0;1]^T$ extracted by HEX$_c$ (before losses), $\dot{Q}_d \in [0;1]^T$ injected by HEX$_d$ (after losses). Finally, $\dot{Q}_1 \in [0;1]^T$ is the plant's primary thermal flow and $\dot{Q}_2 \in [0;1]^T$ feeds the turbines.

We use multiple exogeneous values, set by the modeler:

- Energy capacity capital cost:	$\mathrm{CAPEX_e}$	given in $\$/\mathrm{MWh_{th}}$
- HEX$_c$ (plant > salt) capital cost:	$\mathrm{CAPEX_c}$	given in $\$/\mathrm{MW_{th}}$ input[9]
- HEX$_d$ (salt > fluid)[10] capital cost:	$\mathrm{CAPEX_d}$	given in $\$/\mathrm{MW_{th}}$ turbined[11]
- Original HEX[12] (plant > fluid) capital cost:	$\mathrm{CAPEX_d^*}$	given in $\$/\mathrm{MW_{th}}$ turbined[11]
- TES turbine + BoP[13] capital cost:	$\mathrm{CAPEX_T}$	given in $\$/\mathrm{MW_{th}}$ turbined[11]
- Original turbine + BoP capital cost:	$\mathrm{CAPEX_T^*}$	given in $\$/\mathrm{MW_{th}}$ turbined[11]
- Original plant nameplate capacity:	$\mathrm{Plant\_cap}^*$	given in $\mathrm{MW_{th}}$
- Original plant thermal output timeseries:	$\mathrm{Plant}^*_{t\in T}$	given in $\mathrm{MW_{th}}$
- Plant>SL charging efficiency ratio[14]:	$\eta_c(\dot{Q}_c) \in [0;1]$	higher the better
- SL>Turbine discharging efficiency ratio[12,14]:	$\eta_d(\dot{Q}_d) \in [0;1]$	higher the better
- Plant>Turbine efficiency ratio[14]:	$\eta_d^*(\dot{Q}_1) \in [0;1]$	higher the better
- New turbine efficiency ratio[12,14]:	$\eta_T(\dot{Q}_2) \in [0;1]$	higher the better
- Original turbine efficiency ratio[14]:	$\eta_T^*(\dot{Q}_2) \in [0;1]$	higher the better
- TES self-discharge rate:	$\mathrm{sd} \in [0;1]$	lower the better

Hereafter, we introduce a handful of notations:

- The plant+TES thermal output timeseries:	$\mathrm{Plant}_{t\in T}$	given in $\mathrm{MW_{th}}$
- The turbine thermal input timeseries:	$\mathrm{Turbine}_{t\in T}$	given in $\mathrm{MW_{th}}$
- Notional[15] net discharge of the TES:	$\mathrm{net\_dis}_{t\in T}$	given in $\mathrm{MW_{th}}$
- Power block change cost:	$\Delta\mathrm{CAPEX_{SH}}$	lower the better[16]
- Power block change efficiencies[14]:	$\Delta\eta_d$ and $\Delta\eta_T$	higher the better[17]

For the sake of generality, we also include optional variables:

- Storage capacity limits:	$[\mathrm{min\_sto\_cap}\ ;\ \mathrm{max\_sto\_cap}] \in \mathbb{R}_*^{+2}$
- Charging capacity limits:	$[\mathrm{min\_cha\_cap}\ ;\ \mathrm{max\_cha\_cap}] \in \mathbb{R}_*^{+2}$

---

[9] The HEX input goes up to nameplate capacity, whereas output to salt loop is input minus losses (η_c).

[10] Annual Technology Baseline for concentrated solar plants (CSP) (salt/water): $\mathrm{CAPEX_c} + \mathrm{CAPEX_d} = 903\ \mathrm{k\$/MW_{th}}$ and $\eta_d\eta_T = 41.2\ \%$ (NLR, 2024), 7 % direct contingencies and 20 % indirect costs included.

[11] 1 $\mathrm{MW_{th}}$ "turbined" refers to 1 $\mathrm{MW_{th}}$ entering the turbine (thus exclusive of turbine efficiency)

[12] The World Nuclear Association published in 2020 a nuclear CAPEX breakdown with a conventional island + BoP share of 33 %. This yields $\mathrm{CAPEX_d^*} + \mathrm{CAPEX_T^*} \approx 564\ \mathrm{k\$/MW_{th}}$ and $\eta_d^*\eta_T^* = 36.2\ \%$ on average for Taishan 1-2 and Olkiluoto 3 (EPR, water>steam), 7 % direct contingencies and 20 % indirect costs included.

[13] BoP stands for "Balance-of-Plant" and includes all auxiliary components.

[14] Efficiency ratio is a function of the thermal flow through the exchanger, introducing a non-linearity.

[15] We note $\mathrm{net\_dis}_{t\in T} = \mathrm{dis}_t - \mathrm{cha_t}$ seen from the turbine. Positive = net discharge, negative = net charge.

[16] $\Delta\mathrm{CAPEX_{SH}} = \eta_c\eta_d\eta_T(\mathrm{CAPEX_d} + \mathrm{CAPEX_T}) - \eta_d^*\eta_T^*(\mathrm{CAPEX_d^*} + \mathrm{CAPEX_T^*})$. The sign is not relevant as changing the technology may increase efficiency ($\eta_c\eta_d\eta_T > \eta_d^*\eta_T^*$), thus the generation capacity even without TES.

[17] $\Delta\eta_d = \eta_d/\eta_d^*$ and $\Delta\eta_T = \eta_T/\eta_T^*$. It is above 1 (resp. below) if the new technology is more (resp. less) efficient than the original one. See footnotes **Error! Bookmark not defined.** and 12.

- Discharging capacity limits: $[\text{min\_dis\_cap} ; \text{max\_dis\_cap}] \in \mathbb{R}_*^{+2}$

We assume that the storage asset can operate on its full range, from $0\ \%$ to $100\ \%$. Most storage technologies require to set operational limits and reduce that range (e.g., from $30\ \%$ to $80\ \%$, arbitrary values). To keep it simple, we do not consider these limits because they could be included in the energy capacity capital costs, through a simple multiplicator.

For the optimization *per se,* **three capacity variables** should be instantiated, whatever the architecture, with their possible range:

- $\text{Sto\_cap} \in [\text{min\_sto\_cap};\ \text{max\_sto\_cap}]$ is the TES energy capacity: the maximum energy that can be stored in the reservoir.
- $\text{Cha\_cap} \in [\text{min\_cha\_cap};\ \text{max\_cha\_cap}]$ is the TES charging capacity: the maximum energy that can be transferred **from the plant** to the salt loop in a single timestep.
- $\text{Dis\_cap} \in [\text{min\_dis\_cap};\ \text{max\_dis\_cap}]$: is the TES discharging capacity: the maximum energy that can be transferred from the salt loop **to the working fluid** in a single timestep.

Moreover, it is necessary to instantiate **three lists of operations variables**, whatever the architecture. Note that in capacity expansion and dispatch models, because the variables' ranges are endogenous, boundaries will usually be set through constraints, not in the static definition below. We align with this practice in this paper, to enhance the clarity of this open-source model.

- $\text{sto}_{t\in T}$, the quantity of heat stored in the TES at time t
- $\text{cha}_{t\in T}$, the quantity of heat charged at time t, seen from the plant (before HEX_c losses)
- $\text{dis}_{t\in T}$, the quantity of heat discharged at time t, seen from the turbine input

Then, **three pairs of constraints** should be declared to restrict the assets' operating range:

1. $\text{sto}_{t\in T} \geq 0$ and $\text{sto}_{t\in T} \leq \text{Sto\_cap}$
2. $\text{cha}_{t\in T} \geq 0$ and $\text{cha}_{t\in T} \leq \text{Cha\_cap}$
3. $\text{dis}_{t\in T} \geq 0$ and $\text{dis}_{t\in T} \leq \text{Dis\_cap}$

Finally, the **storage balance coupling TES variables** can be introduced as a constraint[18]:

$$\text{sto}_{t\in T} - (1 - \text{sd}) \cdot \text{sto}_{t-1} = \text{cha}_t \cdot \eta_\text{c} - \text{dis}_t \cdot \frac{1}{\eta_\text{d}} \tag{1}$$

Other equations such as the TES part in the objective function, or the plant-TES-turbine capacities equivalences, are architecture-dependent and are defined in the sections below.

Please note that for the sake of simplicity in our system-scale model, the leakage is assumed to be a static ratio ($\text{sd} = 2\ \%/\text{day}$ for example, relative to the energy loaded in the TES) whereas in reality, the modeling of the thermal losses requires at least temperature calculations that cannot be implemented for each TES in a model at the scale of the electric system. Disciplinary literature on thermal transfers addresses, specifically, this issue.

Theoretically, we rely on the frameworks developed by Boiteux (1960, 1964) and Turvey (1968). The mathematical implications of Boiteux-Turvey style models for storage have been documented in the literature by Junge et al. (2022). The TES-NPP model presented above is written in plain

[18] If there are heat consumers (such as industries) connected to the TES or an interest in releasing heat through a dedicated heat rejection system (usually of negative value), storage balance would include it.

mathematical equations, but in a way that can easily be translated into any programming method. The code implementation in GenX (Julia with JuMP toolbox) is accessible in open source[19].

Generally speaking, the cost minimization objective function can be written as:

$$J(\mathbf{cap}, \mathbf{power}_{t\in T}) = \mathbf{CAPEX} \cdot \mathbf{cap} + \mathbf{OPEX} \cdot \sum_{t\in T} \mathbf{power}_t$$

with:

- $\mathbf{CAPEX}$, the vector of investment costs for all available technologies
- $\mathbf{OPEX}$, the vector of O&M costs for all available technologies
- $\mathbf{cap}$, the vector with the capacities of each technology
- $\mathbf{power}_t$, the vector with the production of each technology at a timestep t

The problem can be written as:

$$\text{minimize } J(\mathbf{cap}, \mathbf{power}_{t\in T}), \text{ subject to} \begin{cases} \boldsymbol{g}(\mathbf{cap}, \mathbf{power}_{t\in T}) \leq \mathbf{0} \\ \boldsymbol{h}(\mathbf{cap}, \mathbf{power}_{t\in T}) = \mathbf{0} \end{cases}$$

where:

- $\boldsymbol{g}$ is a vector of inequality constraints (for example min capacities, max emissions...)
- $\boldsymbol{h}$ is a vector of equality constraints (for example system balance, fixed capacities...)

This classical problem of energy economics is often solved using Karush-Kuhn-Tucker theorem and Lagrangian/Hamiltonian optimization. This approach has the distinctive feature of explicating $\lambda_{t\in T}$, the dual variable of the load balancing constraint $P = Q$. Under specific conditions, this long-term marginal cost includes simultaneously the short-term marginal cost and the long-term investment signal.

Box 1 (Additional turbine capacity)

A significant remark is that whenever the self-discharge rate is positive ($sd > 0$), or the TES interfaces imperfect ($\eta_c\eta_d < 1$), the energy stored into the TES has a higher cost per unit compared to the one directly produced by the upstream plant. Assuming perfect foresight and linearity[20], it comes from the models presented in Boiteux-Turvey style models that the TES will be discharged only after the upstream plant is operated at its nameplate capacity. This has been documented in the literature by Junge et al. (2022). Hence, discharging the TES when the plant is operating (and supplying the turbine) necessarily requires building an additional turbine capacity (equal to the discharging capacity of the TES). This is because "without TES" the turbine's maximum thermal input would be equal to the plant capacity (minus the losses). This consequence has been already explored by Saeed et al. (2022), in which two options are identified: building a dedicated turbine for the TES (see their option 1) and oversizing the existing turbine (see their options 2 and 3). We emphasize that modelers should be careful in setting up the turbine capacity constraints, as it should accommodate both the plant nameplate capacity and the TES discharging capacity (this is architecture dependent, see below) or allow for dedicated turbines to be built.

[19] Link to be specified.

[20] With uncertainty or non-linearities (such as start-up costs or ramps) the marginal value of a unit of stored energy ($\mu_{t\in T}$, dual variable of equation 1) can in fact fall below the marginal cost of energy from the plant. In such cases, the storage may discharge when the plant is not operating at nameplate capacity.

Box 2 (TES Boost service)

With an additional turbine capacity available (see Box 1), the TES can "boost" the output of the plant during peak hours (and as a consequence, of the turbine). This is because the TES can charge during hours of energy surplus (see literature on the Californian "duck curve") enabling the upstream plant to produce at baseload capacity even when the market price is below the marginal cost of production, before standing by or shutting down. In short, it allows longer operations by reducing the need for modulation and hence it increases the asset's profitability. This is because it can discharge energy later when the price signal is high enough. For capital intensive assets like nuclear plants, this is a promising feature. It is also sought after for dispatchable CSP plants.

For example, when $\mathrm{TES}_{\text{boost ratio}} = 110\ \%$ (arbitrary value) it means that adding the TES can increase the thermal input of the turbine by $10\ \%$ during peak hours (relative to the original plant capacity), without increasing the system cost by $10\ \%$ of the plant costs. This is a key feature of TES, especially when retrofitting an existing system or when considering capital intensive thermal sources such as nuclear reactors.

Box 3 (TES asymmetry)

We provided an adaptable linear formulation for storage assets with possible asymmetries (charging capacity ≠ discharging capacity). However, the TES charging and discharging interfaces may be, in some cases, symmetrical (as in batteries for example). In this paper, we choose to notionally separate these interfaces to allow asymmetrical TES modeling (just as in Saeed et al. (2022) and mentioned Parzen et al. (2023)), without prejudice of being able to model systems that are, in fact, symmetrical (for engineering or costing reasons). This symmetry assumption has been made, for example, by Armstrong et al. (2024) to simplify the model by removing the optimization on charging capacity. But it comes with the cost of either overestimating TES costs or limiting the discharge potential during peak hours. In all cases, this results in an underestimation of TES value to decarbonizing grids.

In the two following sections, we describe simple models for the two architectures and their associated investment cost. We also define their "TES boost ratio" (that explicate how much the TES can provide additional power during peak hours) and their charge/discharge/roundtrip durations. To enable comparisons with the broader storage literature, we refer to the "canonical" (discharge) duration as the number of hours that the TES can operate at its design capacity using only the heat from the storage system (Schmalensee, 2015). Both architectures pros and cons will be commented on in the Discussion section.

Direct TES architecture

The direct architecture we define can be seen as the addition of a component within the value chain of the power plant. On the upstream side, the TES charges with the thermal energy provided by the plant ($\dot{Q}_1 = \mathrm{cha}_t$). On the downstream side, it discharges to supply the turbine with hot working fluid ($\dot{Q}_2 = \mathrm{dis}_t$). The TES net discharge is the difference between these thermal flows (minus the losses and the self-discharge, see equation 1).

In the direct TES architecture, the TES is most of the time simultaneously charging and discharging, with positive flows on both sides of the TES. The TES can be seen as a pool of molten salt inside the salt loop, much like an exorheic lake with an upstream river pouring new water and a downstream river draining the lake. Here, the turbine acts like a dam regulating the drain flow from the lake to the downstream river.

Note that in direct TES architecture <u>with an existing salt loop</u>, we recommend using $\eta_c = 1$ and $\mathrm{CAPEX_c} = 0$. This is because the charging heat exchanger already exists (no added costs and losses). Moreover, $\Delta\eta_d = \Delta\eta_T = 1$ and $\Delta\mathrm{CAPEX_{SH}} = 0$ because technologies remain unchanged. This is a key advantage for direct when the salt loop already exists, see the Discussion.

It is the architecture used in Schwartz et al. (2023) and documented by Edwards et al. (2016) and by Saeed et al. (2022) (see their option 2, though with structural differences commented above). Its distinctive feature is that this system stores thermal energy in a salt loop standing between the plant and the turbine, instead of storing energy in a side system, such as in pumped hydro storage, batteries, or flywheels. Such a system is presented in Figure 1.

Direct TES architecture requires at least two investments to add the TES itself:

- Storage capacity (hot and cold tanks)
- Discharging capacity (HEX_d + additional turbine capacity)

If there is no salt loop in the design, the salt loop also requires two additional investments:

- Salt loop charging capacity (HEX_c)
- Power cycle technological change (Plant HEX and turbine)

The thermal balance between the upstream plant and the turbine is:

$$\mathrm{Plant}_t = \mathrm{cha}_t \tag{2}$$

$$\mathrm{dis}_t = \mathrm{Turbine}_t \tag{3}$$

Note that cha_t and dis_t are already linked through the TES storage balance (see equation 1). It appears clearly that the plant operations are algebraically uncoupled from the turbine operations, as long as the TES is not full (in such case $\mathrm{Turbine}_t = \eta_c\eta_d\mathrm{Plant}_t$).

The objective function J (in a cost-minimization problem) should be expanded to include the cost elements of the direct TES system:

$$J {+}{=}\ \mathrm{CAPEX_e} \cdot \mathrm{Sto\_cap} + (\mathrm{CAPEX_d} + \mathrm{CAPEX_T}) \cdot \left(\underbrace{\mathrm{Dis\_cap} - \overbrace{\eta_c\eta_d\mathrm{Plant\_cap}^*}^{\text{Turbine cap for the plant}}}_{\text{TES net discharge capacity}} \right) \tag{4}$$

If there is no salt loop in the design to re-use, the objective function J should also be expanded to include the charging heat exchangers sized to accommodate the nameplate capacity of the plant (plant fluid > salt, $\mathrm{CAPEX_c}$), and the cost of replacing the original power block (plant fluid > original working fluid + turbine, $\mathrm{CAPEX_d^*} + \mathrm{CAPEX_T^*}$) with the salt-fed heater (salt > working fluid + turbine, $\mathrm{CAPEX_d} + \mathrm{CAPEX_T}$). Both exchangers' capacities are defined with respect to $\mathrm{Plant\_cap}^*$ considered here as exogenous (and neglecting salt loop inline losses and piping costs)[21].

$$J {+}{=} \left(\underbrace{\mathrm{CAPEX_c}}_{\text{Salt loop}} + \underbrace{\Delta\mathrm{CAPEX_{SH}}}_{\text{Technology change cost difference}} \right) \cdot \mathrm{Plant\_cap}^* \tag{5}$$

[21] Capacities are relative to $\mathrm{Plant\_cap}^*$ because additional discharge capacity is already accounted in (4). Note that cases were $\mathrm{CAPEX_c} + \Delta\mathrm{CAPEX_{SH}} < 0$ may exist, indicating that it would theoretically be more cost-effective to change the technology and introduce the salt loop even without TES. See footnote **Error! Bookmark not defined.** for some reasons. However, other reasons such as past standardization may hamper such structural changes.

For the direct TES, the "boost" service (see Box 2) can be accurately pictured by:

$$\mathrm{TES}_{\text{boost ratio}_{\text{direct}}} = \frac{\overbrace{\eta_{\mathrm{d}}\eta_{\mathrm{T}}\mathrm{Dis\_cap}}^{\text{Decoupled discharge capacity}}}{\underbrace{\eta_{\mathrm{d}}^{*}\eta_{\mathrm{T}}^{*}\mathrm{Plant\_cap}^{*}}_{\text{Original discharge capacity}}} \geq 1 \tag{6}$$

Then the TES will operate as a cycling asset: filling itself with energy (not necessarily entirely) and releasing it (not necessarily until totally empty). In addition to the traditional duration (full discharge at full power), we can define the turnover duration of a complete cycle ($\mathrm{TES}_{\text{turnover duration}_{\text{direct}}}$) as the sum of the duration required to charge Sto_cap into the TES ($\mathrm{TES}_{\text{charge duration}_{\text{direct}}}$). We also define the duration to discharge Sto_cap from the TES when the plant is active ($\mathrm{TES}_{\text{discharge duration}_{\text{direct}_{\text{on}}}}$):

$$\mathrm{TES}_{\text{charge duration}_{\text{direct}}} = \frac{\mathrm{Sto\_cap}}{\mathrm{Cha\_cap}}\frac{1}{\eta_{\mathrm{c}}} \tag{7}$$

$$\mathrm{TES}_{\text{discharge duration}_{\text{direct}_{\text{on}}}} = \frac{\mathrm{Sto\_cap}}{(\mathrm{Dis\_cap} - \eta_{\mathrm{c}}\eta_{\mathrm{d}}\mathrm{Plant\_cap}^{*}) \cdot \frac{1}{\eta_{\mathrm{d}}}} \tag{8}$$

$$\mathrm{TES}_{\text{duration}_{\text{direct}}} = \frac{\mathrm{Sto\_cap}}{\mathrm{Dis\_cap} \cdot \frac{1}{\eta_{\mathrm{d}}}} \tag{9}$$

$$\mathrm{TES}_{\text{turnover duration}_{\text{direct}}} = \mathrm{TES}_{\text{charge duration}_{\text{direct}}} + \mathrm{TES}_{\text{discharge duration}_{\text{direct}_{\text{on}}}} \tag{10}$$

Comments:

- $\mathrm{TES}_{\text{duration}_{\text{direct}}}$ is the canonical definition of duration: it is a valuable indicator for the plant-owner as it informs on the duration of plant events that could be hedged by TES autonomous production, such as unplanned outages or other nuclear "scram" events. However, it is not representative of the normal operational discharge duration.
- $\mathrm{TES}_{\text{discharge duration}_{\text{direct}_{\text{on}}}}$ is a primary indicator for both the system-planner and the plant-owner as it informs on the duration of system-scale events that could be covered by TES through boosted production such as *dunkelflauten* (droughts) or peak price events.
- $\mathrm{TES}_{\text{charge duration}_{\text{direct}}}$ could be a secondary indicator for the plant-owner as it informs on the ability of TES to absorb the baseload production during sudden price drops (without modulation), or the ability to smoothen the downward modulation during grid-scale events requiring the islanding of the power plant (such as the loss of a transmission line).
- $\mathrm{TES}_{\text{turnover duration}_{\text{direct}}}$ is a useful indicator, it informs the analyst on the number of turnover available in a year of normal operations ($\mathrm{TES}_{\text{max turnover}_{\text{direct}}}$) and can be compared to the optimal dispatch cycling statistics (substituting the load factor for a powerplant):

$$\mathrm{TES}_{\text{max turnover}_{\text{direct}}} = \frac{\text{8760 hours in a year}}{\mathrm{TES}_{\text{turnover duration}_{\text{direct}}} \text{ for a full cycle}} \tag{11}$$

Indirect TES architecture

The indirect architecture we define is a remarkably simple adaptation of most storage assets models found in literature. It is the one used in Armstrong et al. (2024) and documented by Saeed et al. (2022) (see their options 1 and 3) and Schmalensee et al. (2015). Its distinctive feature is that it stores thermal energy through a heat exchanger between the plant turbine-loop heater and the turbine, instead of storing electricity from the grid. Such a system is presented in Figure 2.

Whenever the indirect TES adds losses compared to the main power block ($\eta_c\eta_d < 1$), the charge and discharge operations will be mutually exclusive. When $\eta_c\eta_d = 1$ (perfect TES), the simultaneous charge and discharge operations are numerically possible. The indirect TES can be seen as a pumped-hydro storage reservoir but storing thermal energy instead of gravitational energy.

Indirect TES architecture requires three investments in all NPP cases (pumps neglected):

- Storage capacity (hot and cold tanks)
- Charging capacity (HEX_c)
- Discharging capacity (HEX_d and additional turbine capacity)

The thermal balance between heat production and heat consumption becomes:

$$\underbrace{\mathrm{Plant}_t + \mathrm{dis}_t}_{\text{heat production}} = \underbrace{\mathrm{Turbine}_t + \mathrm{cha}_t}_{\text{heat consumption}} \tag{12}$$

And the objective function J (in a cost-minimization problem) should be expanded to include the cost elements of the indirect TES system:

$$J += \mathrm{CAPEX_e} \cdot \mathrm{Sto\_cap} + \mathrm{CAPEX_c} \cdot \mathrm{Cha\_cap} + (\mathrm{CAPEX_d} + \mathrm{CAPEX_T}) \cdot \mathrm{Dis\_cap} \tag{13}$$

In the indirect TES architecture, whatever the working fluid (in the turbine loop), the TES can charge by extracting thermal energy from the hot branch (red in Figure 2). Conversely, it can be discharged by supplying the hot branch with heated-up working fluid. For example, as Saeed et al. (2022) documented, it is possible to connect this indirect TES to the secondary steam loop of a nuclear power plant. This choice is commented on in the discussion section.

For the indirect TES, the "boost" service (see Box 2) can be accurately pictured by:

$$\mathrm{TES}_{\text{boost ratio}_{\text{indirect}}} = 1 + \frac{\eta_d\eta_T\mathrm{Dis\_cap}}{\eta_d^*\eta_T^*\mathrm{Plant\_cap}^*} \geq 1 \tag{14}$$

For indirect TES using a single oversized turbine, there is no technological change at the turbine, hence $\eta_T = \eta_T^* \Rightarrow \Delta\eta_T = 1$, whereas indirect TES with dedicated turbine may present $\eta_T \neq \eta_T^*$.

The TES will operate as a cycling asset: filling itself with energy (not necessarily entirely) and releasing it (not necessarily until totally empty). In addition to the traditional duration (full discharge at full power), we can define the turnover duration of a complete cycle ($\mathrm{TES}_{\text{turnover duration}_{\text{indirect}}}$) as the sum of the duration required to charge Sto_cap into the TES ($\mathrm{TES}_{\text{charge duration}_{\text{indirect}}}$), with losses, departing from an empty state:

$$\mathrm{TES}_{\text{charge duration}_{\text{indirect}}} = \frac{\mathrm{Sto\_cap}}{\mathrm{Cha\_cap}}\frac{1}{\eta_c} \tag{15}$$

$$\mathrm{TES}_{\text{duration}_{\text{indirect}}} = \frac{\mathrm{Sto\_cap}}{\mathrm{Dis\_cap} \cdot \frac{1}{\eta_d}} \tag{16}$$

$$\mathrm{TES}_{\text{turnover duration}_{\text{indirect}}} = \mathrm{TES}_{\text{charge duration}_{\text{indirect}}} + \mathrm{TES}_{\text{duration}_{\text{indirect}}} \tag{17}$$

Comments:

- $\mathrm{TES}_{\mathrm{duration}_{\mathrm{indirect}}}$ is a primary indicator for the system-planner as it informs on the duration of events that could be covered by TES such as *dunkelflauten* (droughts) or peak prices.
- $\mathrm{TES}_{\mathrm{charge\ duration}_{\mathrm{indirect}}}$ could be a secondary indicator for the plant-owner as it informs on the ability of TES to absorb a part of the production during sudden price drops (limiting modulation), or the ability to smoothen the downward modulation during grid-scale events requiring the islanding of the power plant (such as the loss of a transmission line).
- $\mathrm{TES}_{\mathrm{turnover\ duration}_{\mathrm{indirect}}}$ is a useful indicator because it informs the analyst on the number of turnover available in a year ($\mathrm{TES}_{\mathrm{max\ turnover}_{\mathrm{indirect}}}$) and can be compared to the optimal dispatch cycling statistics (substituting the load factor for a powerplant):

$$\mathrm{TES}_{\mathrm{max\ turnover}_{\mathrm{indirect}}} = \frac{8760 \text{ hours in a year}}{\mathrm{TES}_{\mathrm{turnover\ duration}_{\mathrm{indirect}}} \text{ for a full cycle}} \tag{18}$$

Methodological sanity test

To ensure that both architectures are coherent with each other and that none of these introduce a significant advantage, we encourage our readers to compare the results with the following assumptions:

- $\eta_{\mathrm{c}} = 1$ (no losses to charge the salt loop)
- $\mathrm{CAPEX_c} = \Delta\mathrm{CAPEX_{SH}} = 0$ (no capital cost for the salt loop / salt loop already installed)[22]

Both TES architectures should produce identical results when comparing the optimal TES energy capacity. The charging/discharging capacities will differ, as the direct TES includes enough capacity to transfer the baseload thermal flow to the turbine. However, the additional capacities $\mathrm{Dis\_cap} - \mathrm{Cha\_cap}$ for direct and $\mathrm{Dis\_cap}$ for indirect should be comparable.

Numerical example: break-even point indirect/direct

Let us assume, in all cases:

- $\mathrm{Plant\_cap}^* = 1000\ \mathrm{MW_{th}}$ (arbitrary) without salt-loop by design
- Indirect and direct TES have the same dimensions
- $T = 40 \cdot 8{,}760 = 350{,}400$ hours (40 years, the economic lifetime of the plant)
- $\mathrm{Plant}^*_{t\in T} = 90\ \% \cdot \mathrm{Plant\_cap}^*$ (equivalent to an average 90 % load factor)
- $\mathrm{CAPEX_e} = 45{,}000\ \$/\mathrm{MWh_{th}}$ (Armstrong et al. 2024)
- $\mathrm{CAPEX_c} = 375{,}000\ \$/\mathrm{MWh_{th}}$ (Armstrong et al. 2024)
- $\mathrm{CAPEX_d} + \mathrm{CAPEX_T^*} = 903{,}000\ \$/\mathrm{MW_{th}}$ (see footnote 10, assuming symmetric cost)
- $\mathrm{CAPEX_d^*} + \mathrm{CAPEX_T^*} = 564{,}000\ \$/\mathrm{MW_{th}}$ (see footnote 12)
- $\eta_{\mathrm{c}} = 95\ \%$ (arbitrary, assumed constant, pessimistic value)
- $\eta_{\mathrm{d}}^*\eta_{\mathrm{T}}^* = 36.2\ \%$ (assumed constant, see footnote 12)
- $\mathrm{sd} = 0\ \%$ (neglected for the sake of simplicity)

Note that we don't account for turbine technological change in these examples ($\Delta\eta_{\mathrm{T}} = 1$).

Then, let us construct three hypothetical cases:

Table 13: Numerical examples of various break-even points.

[22] This is because a plant-scale salt loop changes the losses and this has a cost. See L_{plant} in section "**Error! Reference source not found.**" below.

Case	A (reference)	B	C
$\text{price}_{t\in T}$	190 $\$/\text{MWh}_{\text{th}}$	$\mathbf{100}$ $\$/\mathbf{MWh}_{\mathbf{th}}$	$\mathbf{100}$ $\$/\mathbf{MWh}_{\mathbf{th}}$
$\eta_{\text{d}}\eta_{\text{T}}^{*}$	39.00 %	39.00 %	$\mathbf{41.20}$ %
Yields			
$\Delta\text{CAPEX}_{\text{SH}}$	130,393.5 $\$/\text{MW}_{\text{th}}$	130,393.5 $\$/\text{MW}_{\text{th}}$	149,266.2 $\$/\text{MW}_{\text{th}}$
$\eta_{\text{c}}\Delta\eta_{\text{d}}\Delta\eta_{\text{T}} - 1$	+2.3 %	+2.3 %	+8.1 %
L_{plant}	-373.2 M$	-196.4 M$	-643.1 M$
With a break-even point at (% of nameplate capacity)			
$\frac{\mathbf{Cha\_cap}_{\mathbf{\Delta=0}}}{\mathbf{Plant\_cap}^{*}}$	$\mathbf{35.2}$ %	$\mathbf{82.4}$ %	$\mathbf{-31.7}$ %

Interpretation:

A. Reference: When the charging cap $\text{Cha\_cap} < 352.5\ \text{MW}_{\text{th}}$, the indirect TES will be cheaper than the direct TES, whereas for $\text{Cha\_cap} > 352.5\ \text{MW}_{\text{th}}$ the direct architecture is more valuable. The break-even point is at 35.2 % of the plant nameplate capacity.

B. In a system with lower electricity prices, the break-even points increase to 82.4 % of the nameplate capacity. The indirect TES will be less expensive than direct architecture up to that $3{,}295.6\ \text{MW}_{\text{th}}$ limit. This is because the lower prices reduce the profitability of investing in a system increasing the system efficiency by "only" 2.3 %.

C. In the same system as B., if the salt-fed power cycle boosts the plant efficiency by 8.1 %, the direct TES architecture is always less expensive. Note that $\Delta\text{CAPEX}_{\text{SH}}$ is relatively higher, because the efficiency gain requires to invest in additional turbine capacity.

Appendix B

CAPEX and OPEX of available technologies

Table 14: CAPEX and OPEX of available technologies. Italic lines are technologies with fixed existing capacities without capacity expansion possible.

Technology	Investment cost		Fixed O&M cost		Variable O&M cost	Heat rate
	$\mathrm{k\$/MW_e/yr}$	$\mathrm{k\$/MWh_e/yr}$	$\mathrm{k\$/MW_e/yr}$	$\mathrm{k\$/MWh_e/yr}$	$\mathrm{\$/MWh_e}$	$\mathrm{MMBTU/MWh_e}$
Hydro (Quebec)	*319.96*	-	*31.00*	-	-	-
Natural gas (combined cycle)	83.01	-	24.00	-	1.61	6.20
NGCC+CCS	121.72	-	39.00	-	3.32	7.01
Solar PV	47.75	-	13.00	-	-	-
Onshore wind	69.14	-	23.00	-	-	-
Offshore wind (fixed)	196.69	-	71.00	-	-	-
Offshore wind (floating)	318.15	-	61.00	-	-	-
Battery	26.51	15.48	7.00	3.50	-	-
Run-of-river	*295.46*	-	*18.70*	-	-	-
Pumped hydro	*454.52*	-	*47.00*	-	-	-
Fusion	564.92	-	-	-	2.00	8.53
Thermal energy storage[23]	96.87	2.99	-	-	-	-

[23] Thermal Energy Storage's capital cost includes the storage capital cost and the extra discharging capacity capital cost. The latter includes a discharging heat exchanger (12.46 $\mathrm{k\$/MW_e/yr}$) and additional turbine capacity (84.41 $\mathrm{k\$/MW_e/yr}$). The charging heat exchanger cost is included in the upstream fusion plant.

Appendix C

Carbon intensity and min/max capacities of available technologies

Technology	Carbon intensity		Min/Max capacities	
	gCO_2/kWh_e	$t_{CO_2}/MMBTU$	GW_e min	GW_e max
Hydro (Quebec)	-	-	41.11	41.11
Natural gas	352.41	0.057	-	-
Natural gas (CCS)	37.76	0.005	-	-
Solar PV	-	-	-	22.00
Onshore wind	-	-	-	9.00
Offshore wind (fixed)	-	-	5.95	37.00
Offshore wind (floating)	-	-	-	275.00
Battery	-	-	-	-
Run-of-river	-	-	1.96	1.96
Pumped hydro	-	-	1.83	1.83
Fusion	-	-	-	-

Contact.

MIT CEEPR Working Paper Series
is published by the MIT Center for Energy and Environmental Policy Research from submissions by affiliated researchers.
For inquiries and/or for permission to reproduce material in this working paper, please contact:

General inquiries: ceepr@mit.edu
Media inquiries: ceepr-media@mit.edu

MIT CEEPR
Center for Energy and
Environmental Policy Research

MIT Center for Energy and
Environmental Policy Research
Massachusetts Institute of Technology
77 Massachusetts Avenue, E19-411
Cambridge, MA 02139-4307
USA

ceepr.mit.edu

MASSACHUSETTS INSTITUTE OF TECHNOLOGY