INL/EXT-17-43441
Revision 1

Status Report on Modeling and Simulation Capabilities for Nuclear-Renewable Hybrid Energy Systems

C. Rabiti, A. Epiney, P. Talbot, J. S. Kim, S. Bragg-Sitton, A. Alfonsi (INL)
A. Yigitoglu, S. Greenwood, S. M. Cetiner (ORNL)
F. Ganda, G. Maronati (ANL)

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

September 2017

INL/EXT-17-43441
Revision 1

Status Report on Modeling and Simulation Capabilities for Nuclear-Renewable Hybrid Energy Systems

C. Rabiti, A. Epiney, P. Talbot, J. S. Kim, S. Bragg-Sitton, A. Alfonsi (INL)
A. Yigitoglu, S. Greenwood, S. M. Cetiner (ORNL)
F. Ganda, G. Maronati (ANL)

September 2017

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

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

EXECUTIVE SUMMARY

This report summarizes the current status of the simulation framework developed for the economic assessment of Nuclear-Renewable Hybrid Energy Systems (N-R HES). There are four main cornerstones of the simulation framework: generation of stochastic time series, a probabilistic analysis and optimization set of algorithms implemented in RAVEN, models for representation of the physical behavior of N-R HES, and a new RAVEN module that maps physical performance into economic performance.

The capability to generate stochastic time series is used to reproduce the stochastic behavior of variable renewable electricity sources, electricity demand and electricity prices. This is used to create, from limited databases of time histories (wind speed, electricity demand etc.), an unlimited number of representative time histories that are never exactly the same as, but are statistically similar to, the time histories in the databases.

The unlimited source of time histories is used to perform statistical analysis and optimization of the economic performance of N-R HES. This type of analysis is driven by the RAVEN code, which has been modified to perform stochastic optimization and processing of large data sets.

Several models for simulating the physical response of the N-R HES have been developed for different levels of accuracy (in increasing accuracy order): copper plate model (the system has no inertia), linear surrogate (simplified representation of physical systems), and detailed Modelica-based representation of each subsystem of the N-R HES.

The physical behavior of the system is recorded by a new module of the RAVEN code, which, from physical quantities such as system component size, fuel consumption, and electricity production, generates economic performance metrics including Net Present Value (NPV), Internal Rate of Return (IRR), and Levelized Cost Of Electricity (LCOE). The whole system is built in a scalable fashion such that the size of the system components and subsystems can be accounted from both a physical point of view and economic point of view.

The report presents several realistic test cases that were evaluated to assess the proper functioning of the framework, initiate the exploration of the N-R HES performance and to analyze the characteristics with respect to conventional systems. This first attempt to perform this level of complex analysis will require time to digest all the information generated both on the behavior of the N-R HES and on how to tune and best use the tools that have been developed.

Computational challenges that reflect the complexity of the problem are highlighted. These identified challenges are already being addressed via ongoing work to improve RAVEN scalability and to port Modelica into a High-Performance Computing Cluster (i.e. Titan at ORNL).

The conclusion of this work is positive with respect to the capability and the high degree of flexibility of the framework. The amount of information generated even running the simple test cases reported here highlights the complex behavior of the N-R HES. Further investigation will aid in system design and future deployment and operation.

CONTENTS

FIGURES

ACRONYMS

ARMA	Auto Regressive Moving Average
BOP	Balance of Plant
CS	Control System center
CAPEX	CAPital EXpenditure
CCGT	Combined Cycles Gas Turbines
CPI	Consumer Price Inflation
CSV	Comma-Separated Values
CT	Combustion Turbines
EG	Electrical Grid
EM	Energy Manifold
EPC	Engineering, Procurement and Construction
ES	Energy Storage
HES	Hybrid Energy Systems
HPC	High-Performance Computing
HTSE	High-Temperature Steam Electrolysis
IP	Industrial Process
IRIS	International Reactor Innovative and Secure
IRR	Internal Rate of Return
LCOE	Levelized Cost Of Electricity
LWR	Light Water Reactor
MACRS	Modified Accelerated Cost Recovery System
MSBR	Molten Salt Breeder Reactor
N-R HES	Nuclear-Renewable Hybrid Energy Systems
NHES	Nuclear HES
NPV	Net Present Value
O&M	Operations and Maintenance
ORNL	Oak Ridge National Laboratory
PHS	Primary Heat System
PWR	Pressurized Water Reactor
RAVEN	Risk Analysis Virtual ENvironment
RNSD	Reactor and Nuclear Systems Division
SES	Secondary Energy Source
SOEC	Solid Oxide Electrolysis Cells
SPSA	Simultaneous Perturbation Stochastic Approximation
SY	Switch Yard
XML	eXtensible Markup Language

Status Report on Modeling and Simulation Capabilities for Nuclear-Renewable Hybrid Energy Systems

1. INTRODUCTION

This report summarizes the current status of the modeling and simulation capabilities developed for the economic assessment of Nuclear-Renewable Hybrid Energy Systems (N-R HES). The increasing penetration of variable renewables is altering the profile of the net demand, with which the other generators on the grid have to cope. N-R HES analyses are being conducted to determine the potential feasibility of mitigating the resultant volatility in the net electricity demand by adding industrial processes that utilize either thermal or electrical energy as stabilizing loads. This coordination of energy generators and users is proposed to mitigate the increase in electricity cost and cost volatility through the production of a saleable commodity. Overall, the financial performance of a system that is comprised of peaking units (i.e. gas turbine), baseload supply (i.e. nuclear power plant), and an industrial process (e.g. hydrogen plant) should be optimized under the constraint of satisfying an electricity demand profile with a certain level of variable renewable (wind) penetration. The optimization should entail both the sizing of the components/subsystems that comprise the system and the optimal dispatch strategy (output at any given moment in time from the different subsystems). Some of the capabilities here described have been reported separately in [1, 2, 3]. The purpose of this report is to provide an update on the improvement and extension of those capabilities and to illustrate their integrated application in the economic assessment of N-R HES.

First, minor improvements have been implemented in the Modelica physical representation of N-R HES components/subsystems. These improvements have primarily focused on enhancing the modularity of the systems, upgrading the libraries used, and improving the computational performance in terms of speed and reliability. The computational time is currently in the range of one to two hours for a week of real time simulation. A set of tests using random dispatch time histories have been employed to verify the minimization of spurious failure which could impair the optimization algorithm. As will be clarified later in the report, the speed up of each single Modelica calculation is key to extending the real simulation time to more than one week. For these reasons, an initial study was carried out to port the N-R HES Modelica model to the Titan supercomputer located at Oak Ridge National Laboratory (ORNL). The current status of this effort and the Modelica model enhancements are described in section 2 of the report.

The growth in the amount of Modelica subsystem models and the newly added capability to scale the subsystems required extension of the available economic data; hence, the economic dataset has been extended accordingly. The status of the data available and the extension of that data are reported in section 3 of this report.

The infrastructure of the Gitlab repository hosting the software developed for hybrid system modeling has also been improved. The regression system is used to ensure that independent software developments will not interfere with one another. Every time a modification is introduced in the software repository, a series of tests is run automatically to ensure that no existing capability is altered unwittingly. Up to this point coverage of the regression system has been limited to Modelica-based software. The recent enhancements now extend the regression system to also cover Python-based code. Python-based code is used to interface the Modelica models with the Risk Analysis Virtual ENvironement (RAVEN) [4] package while performing all the manipulation to convert the physical output of the system model in Modelica into financial performance.

Several improvements have been implemented in the optimization scheme described in [1], leading to a multi-level approach where two complete independent algorithms can be used for optimization of the capacity of each component (or subsystem) and dispatch (level of utilization of each subsystem). These

improvements capture the different mathematical aspects of dispatch optimization versus the optimal capacity sizing for the system. A summary of those improvements is reported in section 5.

A set of increasingly challenging tests have been performed to demonstrate the capability of the simulation framework and to characterize the behavior of N-R HES. The first tests are simple parametric sweeps (Cases 1 and 2) where the economic performance of the systems is monitored as a function of N-R HES component sizes, and as a function of the hydrogen production price (industrial process used in the reference N-R HES configuration chosen). In this case the system model is replaced by a marginal cost dispatch model. A third test (Case 3) has been performed to optimize profit using a linear surrogate of the subsystems. The fourth test reinstates the sweep on the subsystem sizes but uses the Modelica representation of the system and the dispatch is optimized using RAVEN rather than being based on marginal cost. This test is followed by a case (Case 5) in which the sweep over the subsystem sizes is replaced by an optimization search for the optimal size. The final case (Case 6) is one in which both the dispatch and subsystem sizing is optimized by RAVEN. All the tests performed are reported in section 7.

The conclusions section of the report highlights the high flexibility of the analysis framework that has been realized, possible shortcomings of the approach, and areas for further enhancement.

2. STATUS OF THE PHYSICAL REPRESENTATION OF N-R HES VIA MODELICA

This section summarizes the current status of the Modelica physical representation of an N-R HES. In particular, it focuses on the modular structure of the current approach, libraries used and the input/output structure, which is used for communication with RAVEN.

2.1 Modelica Hybrid System Library Infrastructure

The physical modeling of the N-R HES in Modelica employs a templated infrastructure approach. Generally, a templated approach enables the users to be able to generate personal "sandboxes" for subsystem development with little knowledge of other subsystems. In addition, users can employ component models (e.g., pumps, pipes, etc.) from a variety of sources, not just those located in the Nuclear HES (NHES) library. For example, alternate sources of components could include the Modelica Standard Library or any of the many open-source libraries available on sites like GitHub. To be more specific, the primary goals of the templated approach are to:

I. Develop templates for specific implementations of subsystems that handle all potential interrelated subsystem signals (e.g., control systems and fluid connections), thereby allowing the user to focus on modeling his or her respective subsystem

II. Implement templates for categories of subsystems that allow users to easily interchange a subsystem at the overall architecture level according to the case study of interest while minimizing required user knowledge of other subsystems

III. Provide the necessary handles and structure within the templates to implement and verify nominal and initial conditions at subsystem interfaces

IV. Ensure that all templates are as generic as possible to accommodate a large variety of subsystems, with minimal or no modification to the underlying infrastructure.

Nine potential subsystem categories have been identified for an HES in the current configuration. Figure 1 presents a top-level view or architecture of the templated Modelica-based model. Table 1 provides a textual summary of each of the nine subsystem categories as presented in Figure 1. The templated models and supporting file structure upon which Figure 1 is based are shown in Figure 2. Additional discussion on the development of this approach can be found in [1]. Instructions on the use of the infrastructure can be found within the NHES Library documentation.

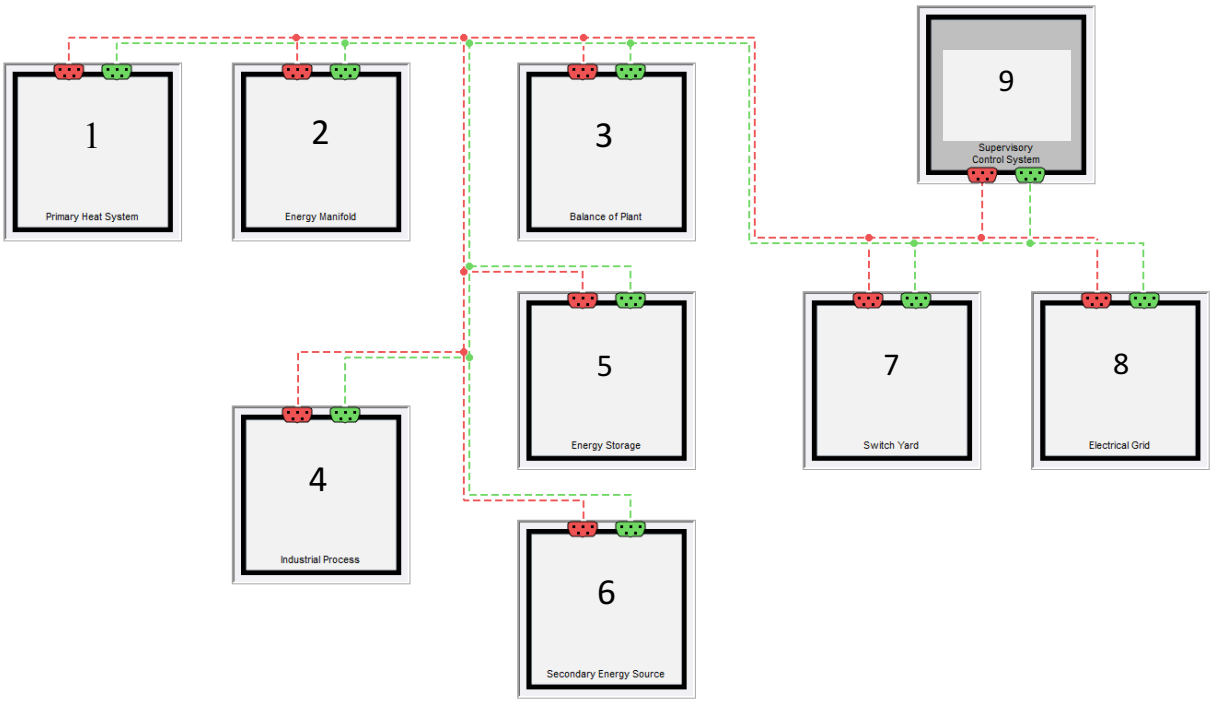

Figure 1. Top-level implementation of an HES using a templated approach: identifier numbers are defined in Table 1, and red and green dotted lines are connection lines for the signal busses described.

Table 1. Potential subsystems of an HES.

Identifier	Category	Description	Specific Example
1	Primary heat system (PHS)	Provides base load heat and power	Nuclear reactor
2	Energy manifold (EM)	Distributes thermal energy between subsystems	Steam distribution
3	Balance of plant (BOP)	Serves as primary electricity supply from energy not used in other subsystems	Turbine and condenser
4	Industrial process (IP)	Generates high-value product(s) using heat from energy manifold/secondary energy supply and electricity from switch yard	Steam electrolysis, gas to liquids, or reverse osmosis desalination
5	Energy storage (ES)	Serves as energy buffer to increase overall system robustness	Batteries and firebrick
6	Secondary energy source (SES)	Delivers small amounts of topping heat required by industrial processes or rapid dynamics in grid demand that cannot be met by the remainder of the system	Gas turbine makeup
7	Switch yard (SY)	Distributes electricity between subsystems, including the grid	Electricity distribution
8	Electrical grid (EG)	Sets the behavior of the grid connected to the NHES	Large grid behavior (not influenced by NHES)
9	Control system center (CS)	Provides proper system control, test scenarios, etc.	Control/supervisory systems and event drivers

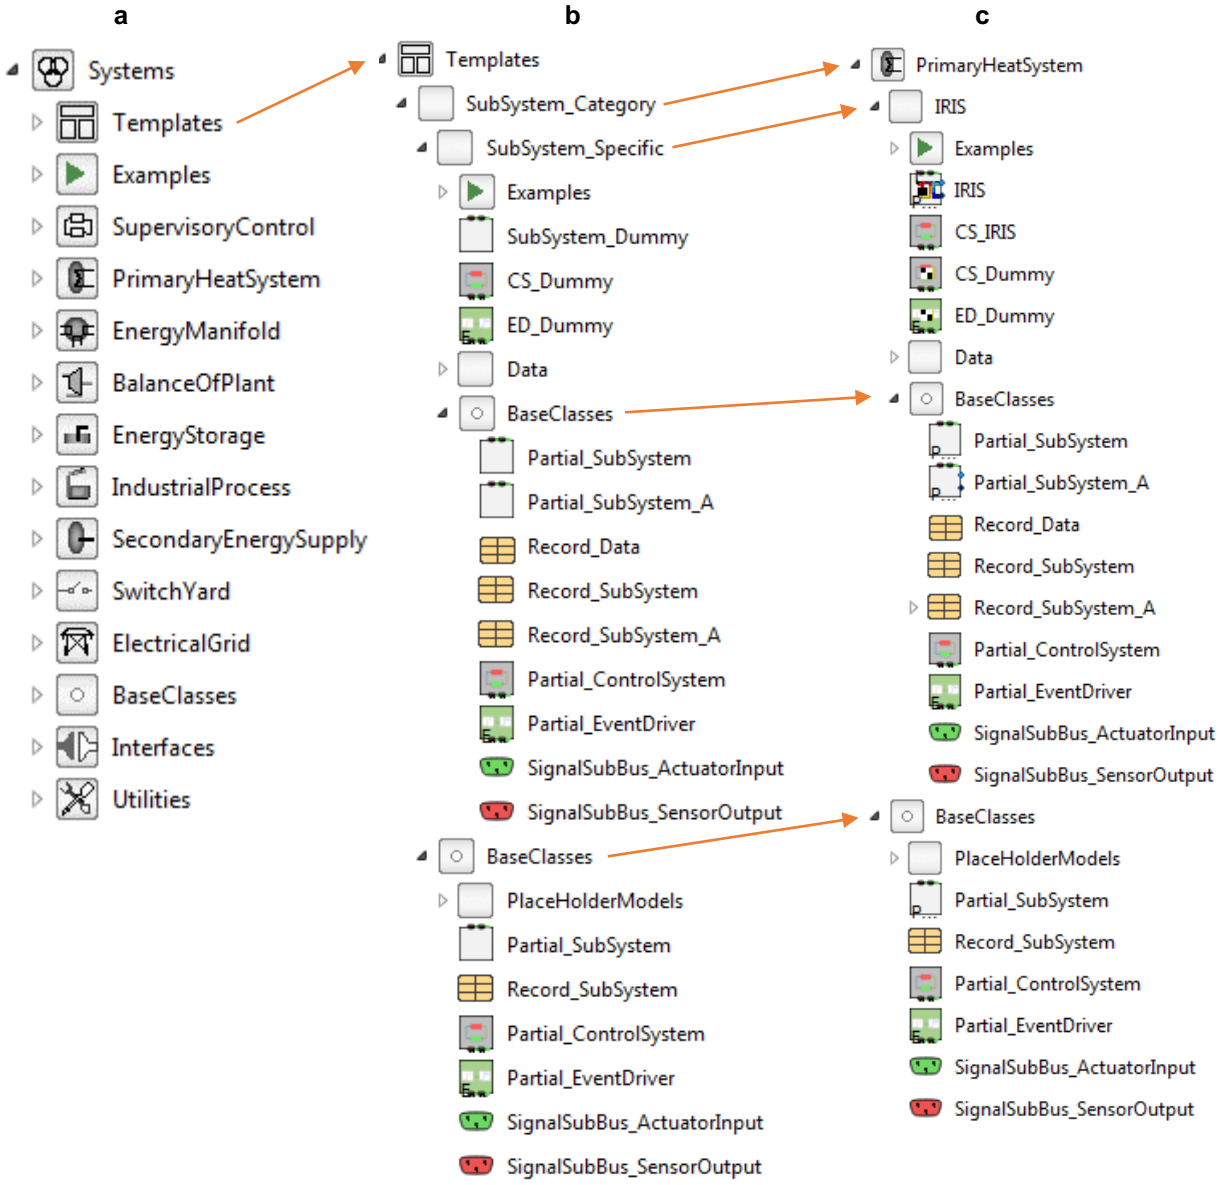

Figure 2. (a) Overall Modelica package for investigation of N-R HESs, (b) template structure for creating new subsystem categories and specific subsystem models within a category, and (c) example of a specific implementation of a primary heat system using the template approach.

2.2 Subsystem Models

The nominal subsystems that have been developed for this initial study are presented in Table 2. Subsystems may be added/modified as necessary using the templated approach previously discussed. The input and output variables identified for integration with RAVEN for this study are summarized in Table 3 for each subsystem. The variables indicated were selected as they represent the high-level performance of the system that RAVEN seeks to optimize. Additional variables can readily be passed to/from RAVEN as they are identified if the full variable name is known as all signals are, by default, available to the modeler in the Modelica input and output files.

Table 2. Identification of the nominal systems developed for this study.

Subsystem Category	Implemented Model	Description
PHS	International Reactor Innovative and Secure (IRIS) [6]	Light water small modular reactor
EM	Steam Manifold	Simple pipes and valves for redirection of steam to BOP/IP and two volumes for additional thermal inertia
BOP	Simple Steam Rankine	Steam Rankine cycle with thermodynamic equation based turbine, ideal condenser, and return pump
IP	High-Temperature Steam Electrolysis (HTSE)	Industrial process unit that utilizes a combination of thermal energy and electricity to split water into hydrogen and oxygen in planar solid oxide electrolysis cells
ES	Logical Battery	Battery with a simple logic based implementation of battery performance
SES	Natural-Gas Fired Turbine	Natural-gas fired turbine with open-ended Brayton cycle
SY	Simple Yard	Ideal electrical distribution to subsystems (i.e., no electrical component models)
EG	Infinite Grid	Infinitely large grid that sets the electricity frequency
CS	Input Setpoint Data	Reads electricity demand signals from RAVEN and sends to subsystems. No additional control logic

Table 3. List of input and output variables for the nominal subsystems of Table 2 that are accessed by RAVEN.

	Variable	Subsystem	Variable Name	Description	Range Limits	Units
Inputs	Power Demand Setpoints	BOP	timeSeriesData.txt file - BOP	Electrical power production setpoint	x>0	W
		IP	timeSeriesData.txt file - IP	Electrical power consumption setpoint	x>0	W
		ES	timeSeriesData.txt file - ES	Electrical power production (+) or consumption (-) setpoint	-	W
		SES	timeSeriesData.txt file - SES	Electrical power production setpoint	x>0	W
		EG	timeSeriesData.txt file - NetDemand	Overall electrical setpoint of the system. Currently only for comparisons and not for altering simulation performance	-	W
	Capacity Sizing	IP	dataCapacity.IP_capacity	Scales the maximum output of the IP	0.9*51e6 < x < 2*51e6	W
		ES	dataCapacity.ES_capacity	Scales the maximum output of the ES	x > 0	W-hr
		SES	dataCapacity.SES_capacity	Scales the maximum output of the SES	0.5*35e6 < x < 2*35e6	W
		EM	EM.V_splitter	Volume of component that sends streams to BOP and IP	10 < x < ~100	m^3
			EM.V_mixer	Volume of component that mixes streams from BOP and IP	10 < x < ~100	m^3
			EM.length_To_BOP	Pipe length to/from BOP	1 < x < 100	m
			EM.length_To_IP	Pipe length to/from IP	1 < x < 100	m

Outputs	Waste Production or Fuel Consumption	PHS	PHS.sensorBus.subBus_PHS.m_flow_fuelConsumption	UO2 fuel consumption rate	-	kg/s
		IP	IP.sensorBus.subBus_IP.HTSE.m_flow_H2_prod	Hydrogen production rate	-	kg/s
			IP.sensorBus.subBus_IP.HTSE.m_flow_O2_prod	Oxygen production rate	-	kg/s
		SES	SES.sensorBus.subBus_SES.GTPP.m_flow_fuel	Natural gas consumption rate	-	kg/s
			SES.sensorBus.subBus_SES.GTPP.m_flow_CO2	CO2 production rate	-	kg/s
	Electricity Production or Consumption	SY	SY.sensorBus.subBus_SY.W_BOP	Balance of Plant	-	W
			SY.sensorBus.subBus_SY.W_ES	Energy Storage	-	W
			SY.sensorBus.subBus_SY.W_SES	Secondary Energy Source	-	W
			SY.sensorBus.subBus_SY.W_IP	Industrial Process	-	W
			SY.sensorBus.subBus_SY.W_EG	Electrical Grid	-	W

2.3 Individual Components and Reference System

The current version of the NHES library employs both third party components from the Modelica Standard Library [7], ThermoPower [8], and TRANSFORM [9] and components developed internal to the project for specific subsystems. For example, the NHES library contains a large variety of models for the development of a high-temperature steam electrolysis plant, a gas turbine, a basic Rankine cycle balance of plant, and a light water nuclear reactor. Components included in the library that support the development of these systems include 1-D pipes, pressurizers, condensers, turbine (steam and gas), heat exchangers, a simple logic based battery, a nuclear fuel subchannel, etc. Third party models include numerous additional models including source/sink components (e.g., fluid boundary conditions), additional heat exchanger models, logical components for control system development, multi-body components, additional supporting functions (e.g., LAPACK, interpolation, smoothing), etc. The reader is directed to specific libraries for additional information.

The reference system used in all the tests performed, if not specified otherwise, is constituted by all the unique subsystems in Table 1 with a hydrogen production plant as the industrial process. This configuration has been chosen for testing the enhanced toolset due to its completeness and relevance to possible realistic applications.

2.4 Ongoing Development

Once the overall template infrastructure and nominal subsystems were developed and tested to adequately perform for the purposes of this initial study, development was frozen on all components and models. While this set of component models was frozen to ensure stability for the subsequent analysis, the enhancement of the system components has proceeded further in a different branch of the repository. For example, the distributed 1-D pipe component model was revisited to improve initialization and flexibility for modeling of different geometries, heat transfer models, and pressure loss models. These upgrades were accomplished by improving the code structure to simplify the development of new models and to enlarge the application range of the components. Additional components such as condenser, pump, mixing volumes, etc. were added to prepare for future needs studies that will require more realistic subsystem models (e.g., balance of plant). Components will continue to be added as needs are identified (e.g., transport delay pipes for long distances that may be encountered in a hybrid system, additional industrial processes, etc.). Also, as models are revisited, nomenclature is being checked to improve standard naming techniques for all components.

2.5 Reliability Testing

The stochastic optimization search requires a large number of runs changing both the capacity and the dispatch of the different resources. This being said, given that the optimization is an intrinsic sequential problem (the next set of values describing the physical systems is derived from the previous set), it is highly important that the software representing the physical system is highly reliable.

2.5.1 Test Description

A RAVEN-based parametric and robustness study of the Modelica model was performed to assess software reliability. Given the large number of simulations required, a High-Performance Computing (HPC) system is necessary. The ORNL Reactor and Nuclear Systems Division (RNSD) has HPC resources available that use the Linux Red Hat operating system.[a]

[a] 64-bit Linux Red Hat 2.6.32 (x86_64)

The analysis presented herein was performed on the ORNL cluster Romulus, which includes a head node and 40 compute nodes.[b] The calculations were accomplished using the parallel management infrastructure provided by RAVEN and using the license-free binary file dymosim generated by the Dymola user interface to execute the simulations. For successful execution, the shared libraries that come with the Dymola distribution were made available.

The Modelica NHES model was tested for a five-hour period and a week-long period of simulated time with the synthetic time histories generated by RAVEN/Auto Regressive Moving Average (ARMA). Test steps include:

I. ARMA model is trained with the available wind and demand data[c] at an hourly resolution. The ARMAs for wind and demand are then trained with: HYBRun_trainARMA_1week.xml

II. RAVEN uses the trained ARMA to create a random realization of the system dispatch

III. Modelica is run for the trained ARMA

IV. Point II and III are repeated as many times as deemed necessary to test the reliability of the model.

The N-R HES model tested was a complete configuration, comprised of a nuclear reactor, hydrogen production plant, steam manifold, gas turbine and battery.

A dispatch model is created as part of a RAVEN ensemble model such that, given the time histories of demand and wind, an hourly dispatch sequence is sent to the N-R HES Dymola model. The runs are driven by a Monte Carlo sampler named "DYMOLA_MC" and the number of samples is controlled by the RAVEN input node <limit></limit>, e.g. 100 for 100 runs:

```
<Samplers>
        <MonteCarlo name="DYMOLA_MC">
         <samplerInit>
           <limit>100</limit>
```

Given the availability of cores all the runs were executed in parallel by setting the batch size to 100, e.g.:

```
<RunInfo>
    <WorkingDir>.</WorkingDir>
    <Sequence> read_ARMAdemand,read_ARMAwind,MCh,dumpCSV</Sequence>
    <batchSize>100</batchSize>
```

2.5.2 Reliability Testing Results

Successful execution is defined in this section as the ability of the model to complete a simulation session successfully and to provide consistent solutions and simulation statistics (central processing unit time, number of time steps, number of Jacobian calculations, etc.) over a variety of forcing functions or parameters.

---

[b] The head node has 16 2800-MHz CPUs with a total of 32-GB memory. Each compute node has 32 2400-MHz CPUs with a total of 132-GB memory.

[c] Available demand data at http://www.ercot.com/gridinfo/load/load_hist/

Available wind data at http://www.nrel.gov/electricity/transmission/eastern_wind_dataset.html

N-R HES Modelica simulation integration with RAVEN/ARMA is able to perform optimization studies for a number of selected model parameters and simulation control options, demand profiles, and subsystem capacities. However, this capability is not used in the implementation discussed here since the scope was to test stability of the model under random excitation (dispatch time series).

Two sets of simulations were performed to investigate the limits of successful simulation of the N-R HES model. The first simulation set focused on a broader range of parameter setting values to sample the simulation space and to understand simulation problems that may arise using a five-hour simulation period. The second set of simulations reduced the number of parameters varied, refined the parameter values, and extended the simulation period to a full week. Results from each study and the information gleaned are further discussed below.

The results of simulations run are examined using high-level results such as success rate, failure rate, simulation time, and memory used, as summarized in Table 4.

Table 4. Modelica model reliability test results (five hours real time).

Model "dsin.txt" Runs	**CPU Time [hh:mm:ss]**	**Run #**	**CPU time/ run [s]**	**Memory [kb]**	**Success Rate [%]**	**Failure Reason**
Cluster Parallel Runs	00:38:59	10	233.9	842028	100	~/dymosim terminated since signal 15 raised (= SIGTERM; kill <process-ID>)*
	01:25:01	20	255.1	779932	100	N/A
	07:25:39	100	267.4	7061304	100	N/A
Local Run[d]	16:34:28	1000	59.7		100	N/A

* ORNL cluster using TORQUE Resource Manager to manage batch job submission and distributed computing. The runs terminated due to queuing system errors are not considered as failed runs but noted as an issue to be resolved for future runs.

[d] Local runs performed on a Mac OS X Yosemite 10.10.5 with 4-GHz Intel Core i7 processor with 32-GB DDR3 RAM at 1600 MHz.

Table 5. Modelica model reliability test results (one week real time).

Model "dsin.txt" Runs	CPU Time [hh:mm:ss]	Run #	CPU time/ run [s]	Memory (resources_used.mem) [kb]	Success Rate [%]	Failure Reason
Cluster Parallel Runs	03:17:29	10	1184.9	744856	98	Model Error*
	06:02:14	20	1086.7	871172	95	Model Error
	33:26:26	100	1203.9	2381456	98	Model Error

* In addition to the queuing system errors, a few tests began simulating and then failed during the simulation. The Modelica log files of the failed tests shared the same error, which indicated a model numerical issue. Specifically, the argument with a natural log function call within the industrial process electrolysis stack (solid oxide electrolysis cells, SOEC) caused the simulation to fail (i.e., natural log of a negative number is undefined) due to a very large transient event defined as model error. Limiting the argument to positive values corrected the error and the simulations that had failed are now working. (See Merge 68 and 69 in the HYBRID software repository history).

2.6 Scale-up of the N-R HES Optimization Framework for HPC Platforms

Activities have been initiated to port the N-R HES library to an HPC platform. Titan, which is hosted at ORNL, was selected as the target demonstration vehicle. With a theoretical peak performance of more than 20 petaflops, Titan, a Cray XK7 supercomputer located at the Oak Ridge Leadership Computing Facility (OLCF), gives computational scientists unprecedented resolution for studying a whole range of natural phenomena, from climate change to energy assurance, to nanotechnology and nuclear energy.

The N-R HES project was awarded an allocation on the following computational resources:

- Titan: 3,000,000 core hours
- Rhea: 2,500 node hours

Titan is a world class supercomputer while Rhea is a smaller cluster used for large data processing. The computational environment is described in appendix A.

2.7 Deployment Status on Titan

Currently, the necessary N-R HES optimization framework packages, i.e., the Modelica dynamic simulation libraries and RAVEN, have been installed on Titan. Dymola, which is the simulation environment that executes models written in the Modelica language, has been installed at the project level, so project participants with account access to Titan can generate new binaries and run their cases. However, Dymola cannot be run in full mode in the current configuration as the license server that hosts the licenses cannot be accessed from within the Titan subnet. This issue is being addressed; one of the licenses will be temporarily transferred over to Titan to enable generation of the binary that can be executed on compute nodes. Unfortunately, binaries generated on other Linux platforms are not compatible with the kernel configuration on Titan.

The ORNL team was able to test the execution of Dymola using the example models provided within the Modelica Standard Library. A binary was successfully generated and the simulation executed. However, because the Dymola license could not be acquired, it is not possible to execute the binaries from the console directly. Once the license issue is resolved, it is expected that it will be possible to distribute it to compute nodes and run cases in parallel. In upcoming work, ORNL researchers will also test the compilers provided on Titan with different optimization options to attain the shortest execution time.

3. ECONOMIC DATA FOR N-R HES

This section of the report contains information relevant to the economic analysis of a hybrid system. The economic numbers for the nuclear plant and for the H_2 electrolysis plant have already been collected and discussed in detail in [1]. A summary of this data is provided at the beginning of this section. In addition to the nuclear plant and the hydrogen electrolysis plant, already available in the previous report [1], data for the following components is provided here:

- Gas turbine and associated plant
- Natural gas markets and recommended historical data for the modeling of natural gas costs
- Steam turbines and BOP
- Manifold
- Grid scale battery (80-120 MWh)
- Switchyard/electrical connection.

3.1 Nuclear Economics Data Summary

Detailed information on the nuclear-related costs are provided in [39].

The current study assumes an overnight capital cost for a reference 1100 MW_e Pressurized Water Reactor (PWR) to be 4100 $/$kW_e$ as reported in [11, 16]. This estimate corresponds to a well-executed construction project without regulatory interventions or design changes after the construction has begun. The weighted average scaling exponent for nuclear plants is 0.64. The commonly used formula [14, 15] to scale CAPital EXpenditure (CAPEX) is:

$$CAPEX = CAPEX_{ref}\left(\frac{Capacity_{new}}{Capacity_{ref}}\right)^{exp=0.64} \quad (1)$$

Light water reactor (LWR) average variable costs based on the U.S. experience in the late 1980s has been reported in [1]. The recommended fixed operations and maintenance (O&M) cost is 85 $/kWe/year. This is similar to the recent cost average for the Exelon fleet [12]. Total variable cost is evaluated as a fraction of total fixed cost. Total variable costs as a fraction of total fixed costs is between 4.6% and 5.9% [16], which results in 0.5 $/MWh as an approximate intermediate value. Additionally, fuel cost is estimated between 8.0 $/MWh and 8.8 $/MWh, averaging 8.4 $/MWh.

3.2 Hydrogen Economics Data Summary

Reference [1] reports that the total (i.e. worldwide) hydrogen market in 2016 was $118 billion, expected to grow about 5% year-over-year over the next few years. Although various technologies to produce hydrogen exist, most of the generation is from steam methane reforming of natural gas: reacting CH_4 with H_2O in the presence of a catalytic converter. Steam methane reforming accounted for more than one third of total generation, while other technologies are partial oxidation, gasification and electrolysis.

The U.S. merchant hydrogen market rose from $3.36 billion in 2010 to $3.44 billion in 2011. This corresponds to about 3% of the total (worldwide) H_2 market. By 2016, the merchant market was expected to have reached $3.87 billion (more updated information was not available in the literature). In terms of volume, the market totaled 1.22 million tons in 2010. In 2011, the U.S. merchant hydrogen market totaled 1.24 billion kilos, and by 2016, the net amount was expected expected to reach 1.23 billion kilos (it was not possible to find more updated information). By taking a ratio of these two numbers, it is possible to calculate the price of H_2/kg (i.e. for 2010: $3.36E9/1.22E9kg = 2.75 $/kg in 2010; for 2016 $3.87E9/1.23E9kg = 3.1 $/kg in 2016). Other references mentioned in [1] place the sale price of hydrogen between $1.26/kg and $2.51/kg.

The economic values applied for a hydrogen production facility in the N-R HES analysis are based on high temperature steam electrolysis (HTSE). The cost of H_2 production applied in the present study is assumed to be 2.56 \$/kg (in 2012 dollars) at a production rate of 1.85 kgH_2/sec in the optimized case.

The total plant cost (including indirect costs) is assumed to be \$153 million as detailed in [1], with all the needed systems to feed H_2 into a pipeline, including the cost of the SOEC module that will need relatively frequent replacements. This value is for a plant capacity of 231 MW_e.

References [1, 17] report that the SOEC cells will likely have to be replaced several times during the life of the plant (assumed to be 40 years). Since the degradation of the SOEC cells depends on the capacity factor of the plants, the fixed O&M cost is a function of the plant capacity factor. Table 6 shows a summary of the fixed and variable O&M costs of the hydrogen plant as a function of plant capacity factor.

Table 6. Summary of the Fixed and Variable O&M costs as a function of the unit capacity factor.

Capacity factor	Full time production (hours/year)	H2 produced (kgH2/year)	Fixed O&M ($/kgH2)	Variable O&M ($/kgH2)
10%	876.6	5.84E+06	0.61	0.048
20%	1753.2	1.17E+07	0.30	0.048
30%	2629.8	1.75E+07	0.20	0.048
40%	3506.4	2.34E+07	0.15	0.048
50%	4383	2.92E+07	0.12	0.048
60%	5259.6	3.50E+07	0.10	0.048
70%	6136.2	4.09E+07	0.09	0.048
80%	7012.8	4.67E+07	0.08	0.048
90%	7889.4	5.25E+07	0.07	0.048
100%	8766	5.84E+07	0.06	0.048

3.3 Gas Turbine

Reference [18] contains information on the cost of natural gas plants (both Combustion Turbines [CT] and Combined Cycles Gas Turbines [CCGT]) generated by PJM[e] to inform new entrants into their markets. This information is based on a large base of actual construction projects and should therefore be defensible. Figure 3 shows an example of a gas turbine.

PJM reports in [18] installed costs of between 947 \$/kW to 1061 \$/kW (average of 1004 \$/kW, approximated as 1000 \$/kWe) for CT of 380-390 MWe (approximated as 400 MWe), depending on the area, using the GE-7FA as reference turbine. Installed costs for CCGT are reported to be 1168 \$/kW to 1326 \$/kW for a 576-595 MWe system (approximated as 600 MWe), also using a GE-7FA turbine, where the most expensive area for installation is near New Jersey and Delaware.

[e] PJM is a regional transmission organization; see www.pjm.com for more information.

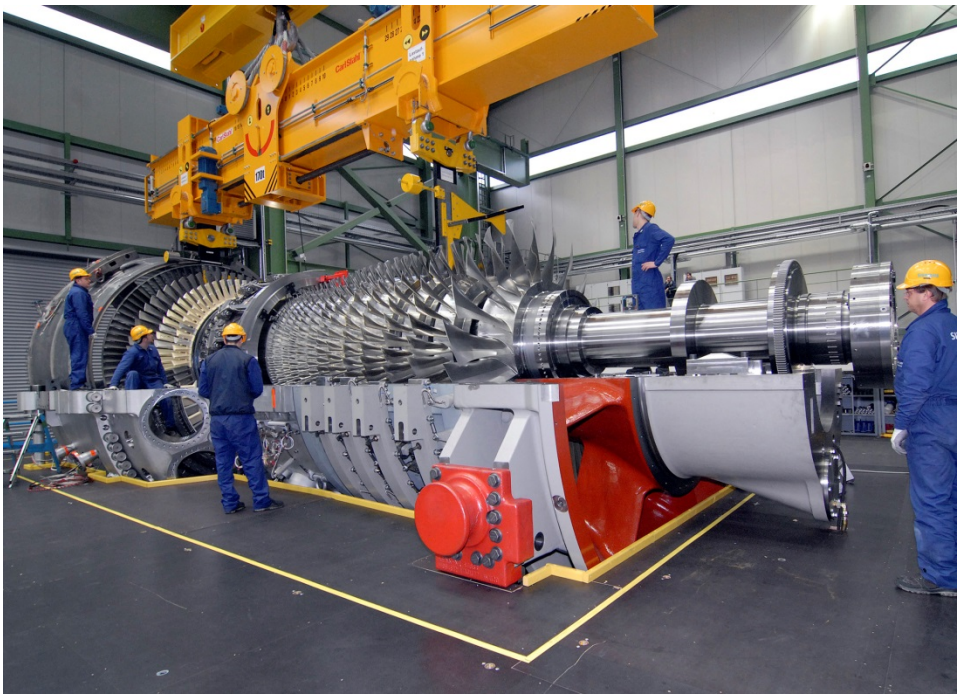

Figure 3. Siemens SGT5-800H rated at 375MW (similar to what used in the hybrid concept), located at the Siemens Berlin facility ready for shipment.

Reference [18] includes a breakdown of the installation cost, which will be used to develop a better understanding of capital costs. For a 400 MW_e CT in the EMAAC PJM zone (NJ and Delaware) as an example, the direct equipment purchase (called "Owner's furnished Equipment", terminology that can lead to confusion with "owner's costs"), are $125.9M, dominated by the turbine at approximately $100M. The Engineering, Orocurement and Construction (EPC) contract includes equipment purchased indirectly through the architect engineer, as well as labor (direct construction labor is $50M-$70M). Other labor includes $20M of "indirect costs" in nuclear construction. The direct costs could therefore be approximated by the sum of "Owner's furnished Equipment" and "EPC costs" without the "non-construction labor" and the "EPC fees and contingency", which can be interpreted as indirect cost. Of course, some of the "material" will also be indirect costs, such as rented construction equipment, temporary construction facilities, etc., but the amount of "materials" is a small number that can be neglected in the first approximation. Hence, the overall direct costs are estimated at 126+191-21-26-29=$249M. This equates to 62% of the total costs, which is substantially higher than for a nuclear plant. This result likely reflects the greater simplicity of gas plant designs, and the lower cost burden associated with regulatory compliance.

A summary table of the capital cost breakdown for CT is provided in Table 7.

Table 7. Summary table of capital costs for CT gas turbine, from [18].

	CONE Area				
	1	2	3	4	5
	EMAAC	SWMAAC	Rest of RTO	WMAAC	Dominion
Capital Costs (in $millions)	*396 MW*	*393 MW*	*385 MW*	*383 MW*	*391 MW*
Owner Furnished Equipment					
Gas Turbines	$98.8	$98.4	$94.0	$98.7	$98.6
SCR	$18.9	$18.7	$17.9	$18.8	$18.8
Sales Tax	$8.2	$7.0	$6.7	$7.1	$7.3
Total Owner Furnished Equipment	**$125.9**	**$124.1**	**$118.6**	**$124.6**	**$124.8**
EPC Costs					
Equipment	$30.9	$30.5	$25.5	$30.8	$30.7
Construction Labor	$71.7	$55.4	$55.3	$54.5	$48.2
Other Labor	$21.2	$19.6	$18.6	$19.6	$19.0
Materials	$9.7	$9.0	$8.6	$9.6	$9.4
Sales Tax	$2.8	$2.4	$2.0	$2.4	$2.5
EPC Contractor Fee	$26.2	$24.1	$22.9	$24.1	$23.5
EPC Contingency	$28.8	$26.5	$25.2	$26.6	$25.8
Total EPC Costs	**$191.4**	**$167.4**	**$158.1**	**$167.6**	**$159.2**
Non-EPC Costs					
Project Development	$15.9	$14.6	$13.8	$14.6	$14.2
Mobilization and Start-Up	$3.2	$2.9	$2.8	$2.9	$2.8
Net Start-Up Fuel Costs	$4.0	$4.7	$3.2	$4.6	$4.7
Electrical Interconnection	$13.0	$12.9	$12.7	$12.6	$12.9
Gas Interconnection	$22.6	$22.6	$22.6	$22.6	$22.6
Land	$2.0	$2.2	$1.1	$1.2	$1.6
Fuel Inventories	$5.3	$5.3	$0.0	$5.1	$5.2
Non-Fuel Inventories	$1.6	$1.5	$1.4	$1.5	$1.4
Owner's Contingency	$6.1	$6.0	$5.2	$5.9	$5.9
Financing Fees	$9.4	$8.7	$8.1	$8.7	$8.5
Total Non-EPC Costs	**$82.9**	**$81.4**	**$70.9**	**$79.6**	**$79.8**
Total Capital Costs	**$400.2**	**$372.9**	**$347.6**	**$371.8**	**$363.8**
Overnight Capital Costs ($million)	**$400**	**$373**	**$348**	**$372**	**$364**
Overnight Capital Costs ($/kW)	**$1,012**	**$948**	**$903**	**$971**	**$931**
Installed Cost ($/kW)	**$1,061**	**$994**	**$947**	**$1,018**	**$977**

Owner's costs are a significant percentage for a CT plant, at $83M over a total of $400M, or about 20%, as compared to a typical value of 8%-10% for a nuclear plant. These costs are dominated by gas and electrical interconnections, and project development.

3.3.1 Scaling Factors

Figure 4, which is taken from reference [19], illustrates the impact of size on the cost for simple cycle gas-fired combustion turbine plants. The interpolated specific cost of gas turbines is shown as a function of the gas-fired plant power level, yielding a cost exponent of 0.78 for "heavy" gas turbines and of 0.77 for "aero-derivative" ones.

The calculated specific costs for the turbine-only portion of the plant in [18] is then compared in Table 8 to the cost of the turbine as would be calculated for each power level using the scaling law of Figure 4, from [19], and then updating the cost from 2008 to 2018 dollars (as shown in [18]) using the Consumer Price Inflation (CPI) from 2008 to 2016, and afterwards an assumed inflation of 2%/year for the years 2017 and 2018. It is noted that the calculated values from the two sources are consistent within an accuracy of about 10% for the power level of CT plants reported in [18].

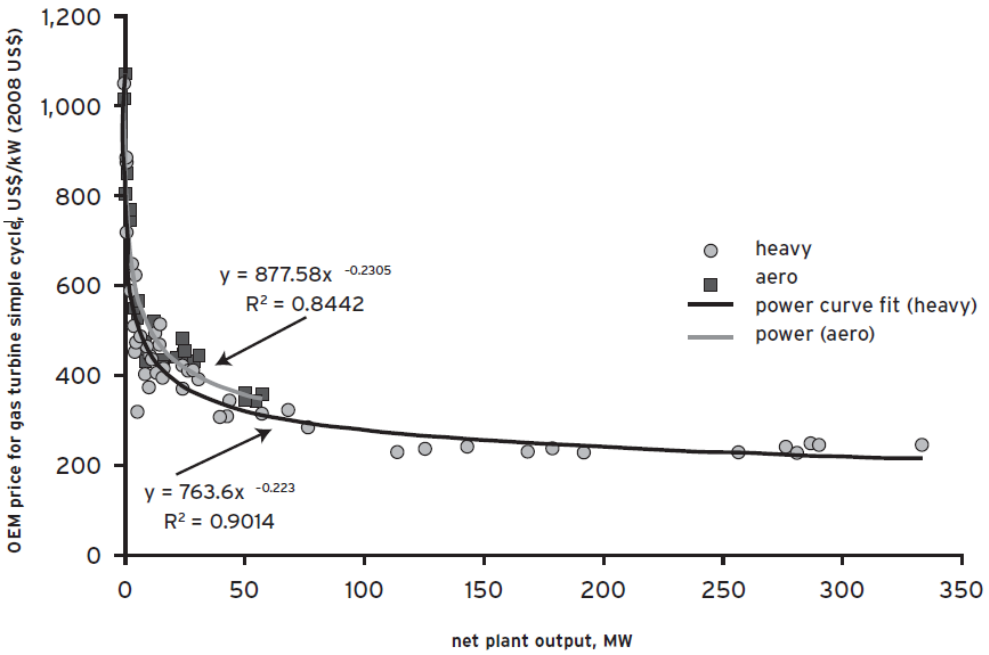

Figure 4. Impact of Size on original equipment manufacturer cost for simple cycle units from [19].

Table 8. Gas turbine equipment costs, comparison of specific cost ($/kW) calculated from data in [18] to specific cost ($/kW) calculated from the scaling law shown in Figure 4 [19].

Turbine power (MW), from [18]	Turbine OEM cost (2017 M$), from [18]	Specific cost ($/kW), calculated from data in [18]	Specific cost ($/kW), calculated from the scaling law of Figure 4 from [19]
396	102.57	259.01	237.39
393	102.15	259.93	237.79
385	97.59	253.47	238.89
383	102.47	267.53	239.16
391	102.36	261.79	238.06

The data in Figure 4 and in Table 8 are for the gas turbine portion of the plant (including directly associated mechanical equipment, such as exhaust ducts, silencers, fuel systems, filters, standard controls and starting systems and NOx emission control systems). However, reference [18] also contains information on total plant costs of different sizes. When interpolating the data using a power law, a scaling factor of 0.72 (with an R^2 of 0.986) is obtained, confirming that likely scaling factors for CT plants are in the range of 0.72 to 0.78. The exponent for entire gas plants is expected to be lower than that of the gas turbine only, since it also includes land, connections, buildings etc. that will scale with a lower exponent (it could approach zero for the land, for example). For this reason, it is recommended to use an exponent of 0.72 to scale the entire plant costs.

It is also noted that other references report substantially lower costs for installed gas plants. For example, [20] reports an overnight unit cost of $973 $/kWe for "conventional" CT plants of 85 MW capacity and $676 $/kWe for "advanced" CT plants of 210 MW capacity. The scaling factor between those two numbers, assuming comparability, would be approximately 0.6, substantially lower than what was found in [19]. However, it is also noted that the heat rate of the PJM [18] CT plant is approximately 10,300 BTU/kWh, while the "conventional" CT plant in [20] has a heat rate of approximately 10,850 BTU/kWh. The "advanced" CT plant in [20] has a heat rate of approximately 9,750 BTU/kWh, placing

the heat rate of the PJM plant of reference [18] in between the two cited in [20]. It is also noted that the unit cost of the 210 MW "advanced" CT plant is $676/kW in [21], substantially below the unit cost of the larger 400 MW plant reported by PJM based on their experience. It is therefore recommended to use the values of [18], from PJM, which may include a degree of conservativeness. A summary table of the referenced plants and their primary characteristics is shown in Table 9.

Table 9. Summary table of the gas turbine characteristics from data in [18] and [20].

Power	Overnight cost ($/kWe)	CT plant type	Reference	Heat Rate (BTU/kWh)
396 MW	1012	Unspecified	[18]	10300
85 MW	973	Conventional	[20]	10850
210 MW	676	Advanced	[20]	9750

3.3.2 O&M Costs

O&M costs for a 400 MW CT plant are reported in [18] at $5.2 M/y to $10.1 M/y fixed, and 4.25 $/MWh to 4.29 $/MWh variable. A detailed discussion of the O&M costs of CT plants was performed in the context of hydrogen cost estimations in [17].

3.3.3 Summary of Gas Turbine Costs

In conclusion, it is recommended that users adopt the unit specific cost of the reference 400 MW CT plants from [18] and exclude the electrical connection costs of $13 million discussed separately in this report (see section 3.8). This assumption results in a unit plant cost "as installed," in 2018 dollars, of $967/kW, scaled with an exponent of 0.72 for CT gas plants of different sizes.

Considering the large variability of fixed O&M costs across regions of the U.S. just based on geography, it is recommended that the fixed O&M costs not be scaled with the size of the plant but to instead perform sensitivities at 5 M$/y and 10 M$/y. Similarly, it is recommended that the variable O&M costs not be scaled for plants of different sizes, instead adopting a constant amount of 4.27 $/MWh, which is the average of the range reported in [18].

3.4 Gas Markets

This section discusses reasonable ranges for natural gas costs, measured in $/MMBTU (US dollars per million BTUs). As shown in Table 10, the Henry Hub gas prices have ranged from approximately 2 $/MMBTU to 4 $/MMBTU from November 2016 to March 2017.

Table 10. Screen capture of the Henry Hub natural gas cost prices for Monday-Friday of each week over a 3-month period, as seen on 3/7/2017 at https://www.eia.gov/dnav/ng/hist/rngwhhdd.htm.

2016 Nov-14 to Nov-18	2.26	2.49	2.49	2.37	2.60
2016 Nov-21 to Nov-25	2.81	2.73	2.76	2.76	2.76
2016 Nov-28 to Dec- 2	2.96	3.02	3.32	3.32	3.44
2016 Dec- 5 to Dec- 9	3.44	3.75	3.80	3.68	3.75
2016 Dec-12 to Dec-16	3.60	3.65	3.55	3.56	3.51
2016 Dec-19 to Dec-23	3.55	3.39	3.50	3.50	3.60
2016 Dec-26 to Dec-30	3.60	3.70	3.70	3.71	3.71
2017 Jan- 2 to Jan- 6	3.71	3.41	3.42	3.42	3.38
2017 Jan- 9 to Jan-13	3.14	3.21	3.27	3.36	3.36
2017 Jan-16 to Jan-20	3.36	3.37	3.26	3.25	3.23
2017 Jan-23 to Jan-27	3.16	3.16	3.26	3.44	3.31
2017 Jan-30 to Feb- 3	3.22	3.00	3.15	3.13	3.13
2017 Feb- 6 to Feb-10	2.92	3.03	3.05	3.11	3.11
2017 Feb-13 to Feb-17	2.93	2.85	2.95	2.88	2.75
2017 Feb-20 to Feb-24	2.75	2.56	2.52	2.63	2.63
2017 Feb-27 to Mar- 3	2.44				

The cost of natural gas can vary considerably at different delivery points. For example, Figure 5 shows the gas prices at various delivery points within the PJM territory in the years 2015 and 2016 [21]. Large spikes are visible in the January-March 2015 period. Additional detailed historical nodal prices for gas were found at the following website: https://www.eia.gov/electricity/wholesale/#history. The map in Figure 7 shows the eight major trading hubs for which historical gas *daily prices* were found for 2014, 2015, 2016 and 2017 up-to-date. The names of those eight hubs are shown in Table 11, while Figure 6 shows a map of the U.S. pipelines. A higher pipeline density should lead to a lower price of gas since access to reliable delivery is supposedly greater.

Additional data on nodal prices of gas can be found on the ICE (intercontinental exchange) web site, for the more than 120 hubs at which natural gas prices can currently be traded. However, it does not appear that this data is sufficiently organized to allow an efficient collection of the data.

The daily weighted average gas prices for 2016 are plotted as a function of time in Figure 8, as an example, to show the daily fluctuation in prices. Simple averages were performed for each year and for each hub, yielding the yearly average gas prices shown in Table 12.

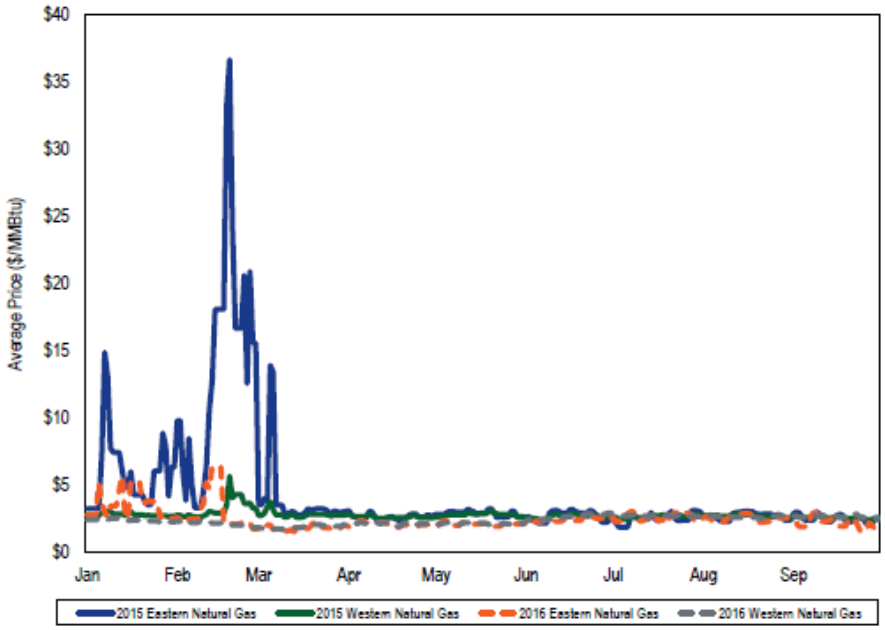

Figure 5. Daily delivered natural gas prices for January through September, 2015 and 2016, where price spikes are due to scarcity driven by cold weather events; this plot appears as Figure 3-50 in reference [18].

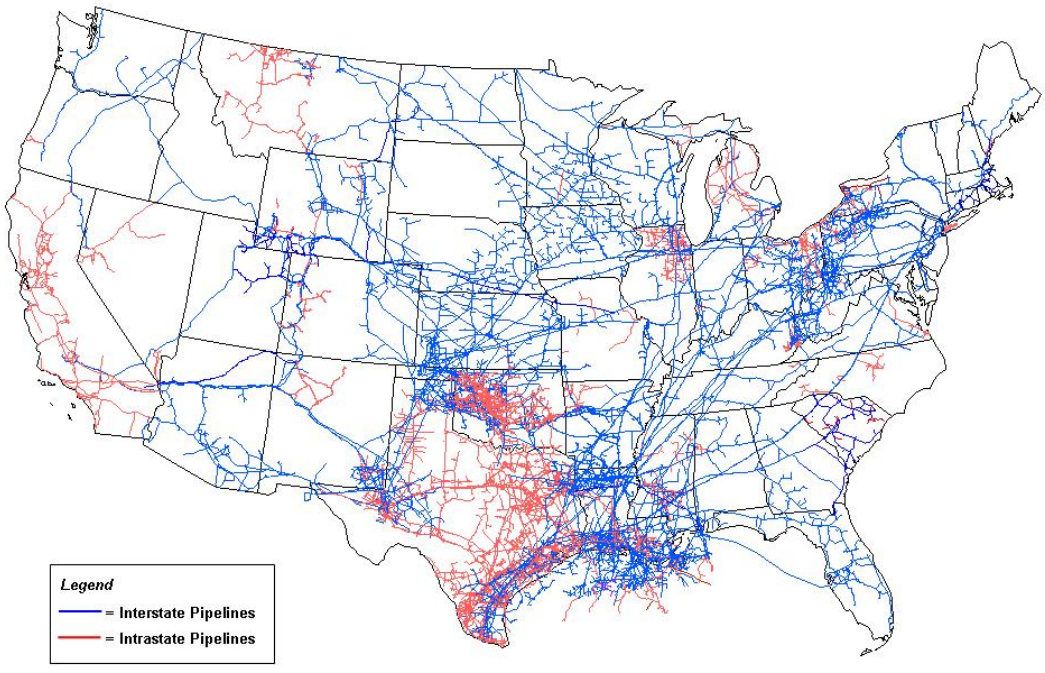

Figure 6. Map of U.S. gas pipelines [downloaded from https://www.eia.gov/energyexplained/index.cfm?page=natural_gas_pipelines].

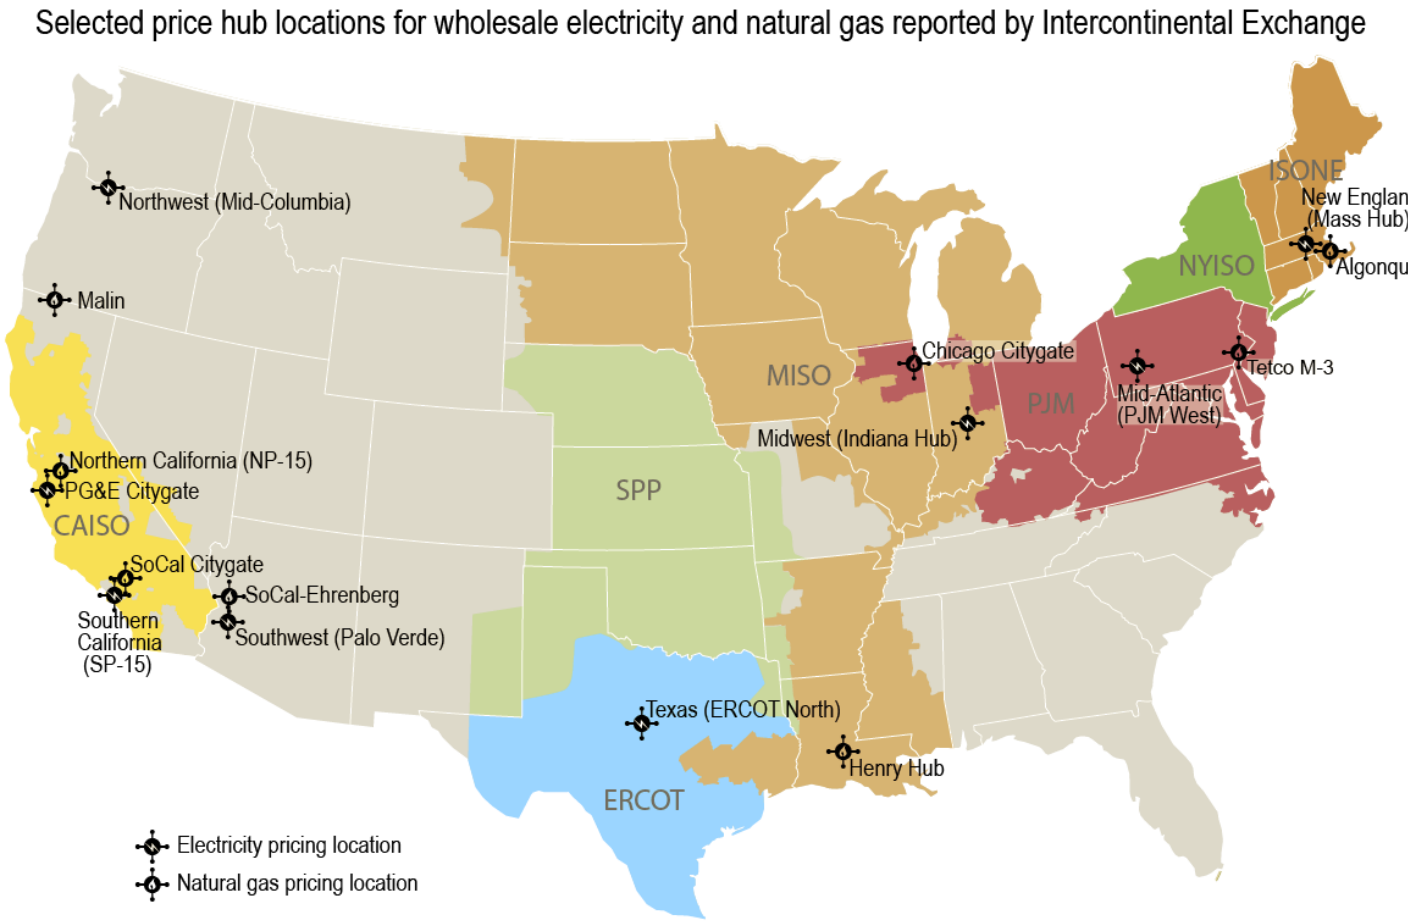

Figure 7. Nodal information on natural gas cost prices [downloaded from https://www.eia.gov/electricity/wholesale/].

Table 11. Natural gas hubs and wholesale daily spot price names [from https://www.eia.gov/electricity/wholesale/].

Natural Gas Hub Name	ICE Natural Gas Product Name
Algonquin	Algonquin Citygates
TETCO-M3	TETCO-M3
Chicago Citygates	Chicago Citygates
Henry Hub	Henry
Malin	Malin
PG&E - Citygate	PG&E - Citygate
Socal-Ehrenberg	Socal-Ehrenberg
Socal-Citygate	Socal-Citygate

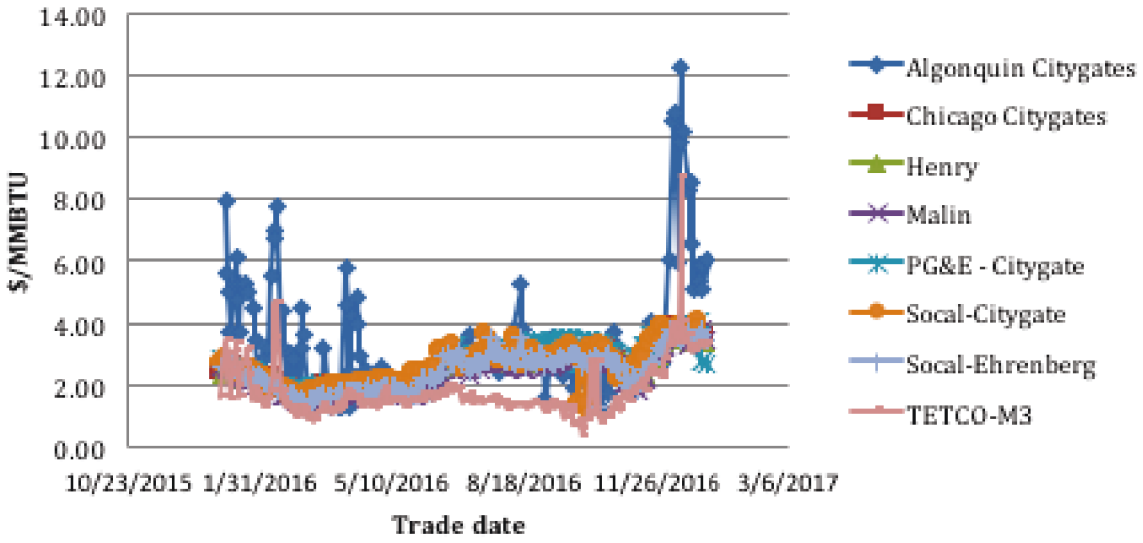

Figure 8. Weighted average daily natural gas prices at the eight trading hubs shown in Figure 7, plotted for the year 2016.

Table 12. Yearly average gas costs in $/MMBTUs and in $/kWh for the 8 trading hubs of Figure 7.

	$/MMBTU 2014	$/kWh 2014	$/MMBTU 2015	$/kWh 2015	$/MMBTU 2016	$/kWh 2016	$/MMBTU 2017-YTD	$/kWh 2017-YTD
Algonquin Citygates	4.40	0.0150	4.81	0.0164	3.56	0.0121	4.37	0.0149
Chicago Citygates	4.23	0.0144	2.75	0.0094	2.49	0.0085	3.23	0.0110
Henry	4.11	0.0140	2.61	0.0089	2.50	0.0085	3.25	0.0111
Malin	4.10	0.0140	2.53	0.0086	2.36	0.0081	3.22	0.0110
PG&E - Citygate	4.64	0.0158	2.99	0.0102	2.72	0.0093	3.57	0.0122
Socal-Citygate	4.50	0.0153	2.8	0.0095	2.57	0.0088	3.53	0.0120
Socal-Ehrenberg	4.34	0.0148	2.64	0.0090	2.44	0.0083	3.27	0.0111
TETCO-M3	3.07	0.0105	2.54	0.0087	1.76	0.0060	3.32	0.0113

Summary of Gas Markets

It is recommended that economic analyses of N-R HES that incorporate a gas turbine or compare costs to natural gas driven systems apply daily gas nodal prices for 2014 to 2016 from https://www.eia.gov/electricity/wholesale/#history. In addition, "Henry Hub" prices that are closest (geographically) to the values used for the daily electricity costs reported at a 5-minute resolution provided in reference [1] for Southern Louisiana, the location of the longest U.S. H_2 pipeline as discussed in [1].

3.5 Steam Turbines and Balance of Plant

The section reports the cost of the balance of plant, broken down by the cost of the steam turbine and other components. They are grouped together since in literature BOP is difficult to find separately and may need to be computed as the difference of the turbine cost and the total balance of plant (including turbine).

3.5.1 Steam Turbine Generator Costs and Scaling Laws

This section describes the cost basis and escalation factors for the nuclear steam turbine generator and BOP. Two quotes were found for nuclear steam turbine/generator pricing: cost for the 1100 MW PWR12-BE in [23] and for the 1,000 MWe Molten Salt Breeder Reactor (MSBR) in [24]. The costs (escalated to January 2017 USD, using the methodology described in [11]), together with the power level and operating conditions of the two machines, are shown in Table 13. The MSBR, with a lower unit cost, relies on a supercritical Rankine cycle, which has a higher thermal cycle temperature and pressure than a standard Rankine cycle. As the steam pressure and temperature are higher, the size of the turbine is substantially smaller than that of an LWR cycle of the same power, although the casing will have thicker walls. The smaller size accounts for the smaller specific costs of the MSBR turbine. An example steam turbine is shown in Figure 9.

Table 13. Turbine equipment cost (from [23] and [24], escalated to 2017 using the methodology described in [11]).

	Power (MW)	Equipment cost (2017 M$)	Specific equipment cost ($/kW)	Inlet turbine pressure (bar)	Inlet turbine temperature (°C)	Type
MSBR	1,035	211.75	259.17	248	538	Supercritical
PWR12-BE	1,192	362.62	304.21	67	283	Subcritical

Figure 9. Steam turbine Siemens SST-800 (https://www.energy.siemens.com/co/en/fossil-power-generation/steam-turbines/sst-800.htm#content=Description).

It was found that the cost of Rankine turbo-generators is a function of both the power level and of the operating conditions, and the scaling laws have different exponents for machines designed to operate at different pressures. For this reason, the large amount of cost data for steam turbines for coal-fired power plant and combined cycle systems cannot be used directly for nuclear turbo-generators. However, information on these systems is briefly summarized here to provide information on the scaling of nuclear turbines. Subcritical steam turbines employed in light water reactors operate at lower pressures than those used in typical fossil power plants, such as coal and CCGT plants (please see Table 14 for a comparison of the typical operating conditions of the two different systems).

Table 14. PWR12-LWR and fossil plants Rankine cycle pressures and temperatures [28].

	LWR	**Fossil**
Pressure (bar)	50-75	150-300
Temperature (°C)	265-290	540-600

A lower Rankine cycle operating pressure results in a lower enthalpy difference across the turbine. To obtain the same amount of power from the cycle, a higher steam mass flow rate is used, which results in the need for higher flow areas. Also, as the density increases with pressure, to obtain the same mass flow rates at lower pressures, high cross-sectional areas are needed. Moreover, the efficiency of each stage of the turbine increases with the length of the blades, which further justifies the adoption of large flow areas. Steam turbine generator equipment costs from references [18, 19, 27, 28] were combined and classified by ranges of operating pressure and temperature. For each range of operating conditions, equipment costs as a function of power were interpolated using a power function in order to estimate the scaling exponents.

For consistency, all costs were converted to January 2017 USD using the methodology described in [11]. Figure 10, Figure 11, and Figure 12 show the costs as a function of power and the interpolated scaling exponents, respectively for the following pressure ranges: 115 bar and 125 bar, 166-167 bar and for supercritical 242.3 bar. The scaling-law exponents found for those pressure ranges are, respectively, 0.87, 0.70 and 0.48, although the supercritical data are based on only two data points.

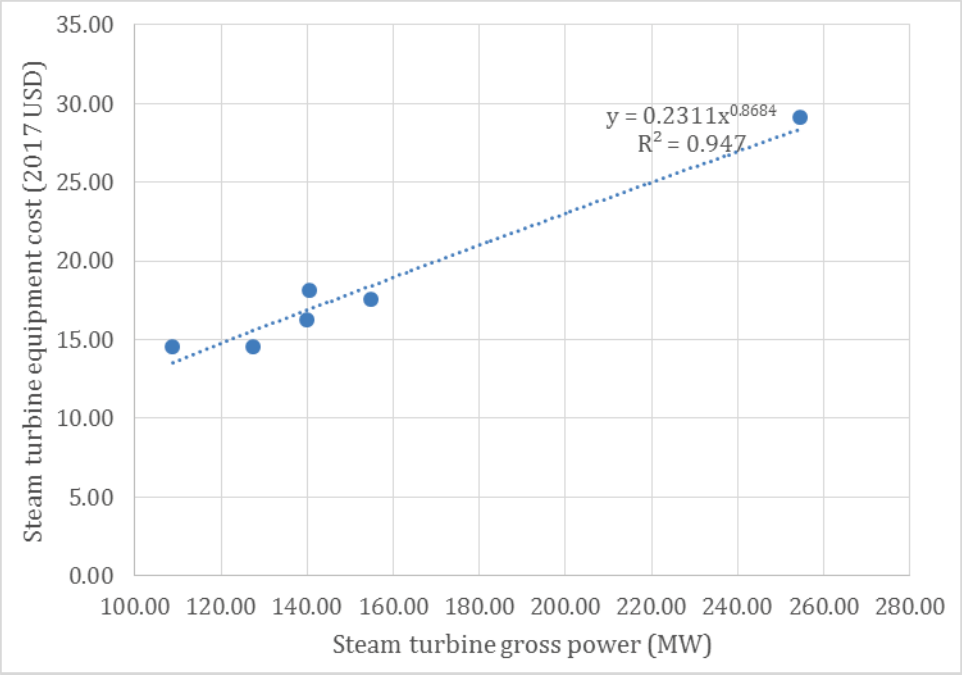

Figure 10. Steam turbine-generator equipment costs (115-125 bar); costs escalated to 2017 USD.

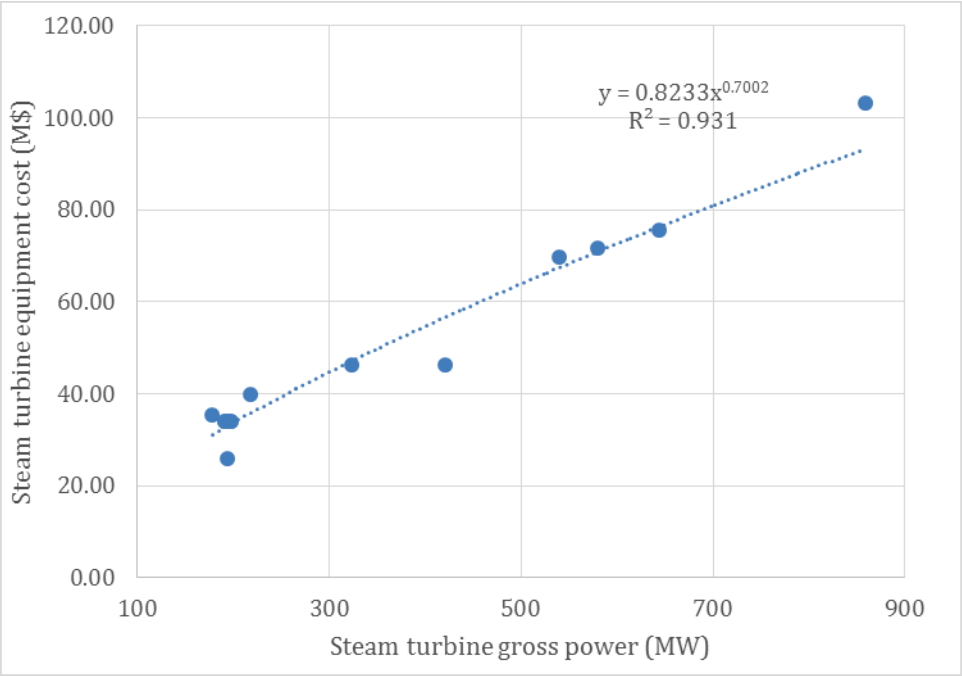

Figure 11. Steam turbine-generator equipment costs (166-167 bar); costs escalated to 2017 USD.

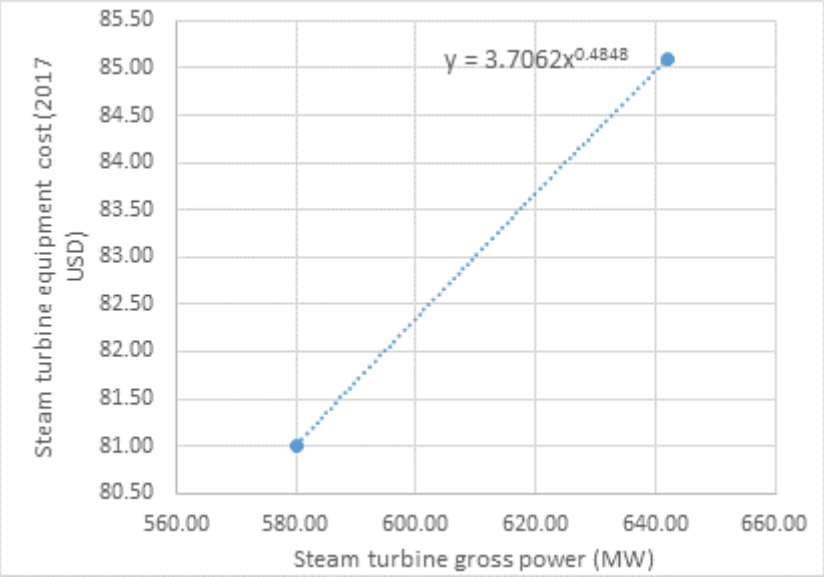

Figure 12. Steam turbine-generator equipment costs (supercritical, 593 °C); costs escalated to 2017 USD.

Since steam turbine pricing for nuclear power plants is limited, a regression analysis similar to those performed for fossil plant turbines is not viable. Instead, the scaling law for NPP steam turbines was estimated from the scaling laws derived for turbines of different pressures discussed in section 3.3. The data points shown in Figure 13 were interpolated using a linear function, as shown in Table 15.

Table 15. Scaling law exponents as a function of the HP steam turbine inlet pressure.

Pressure (bar)	Exponent
242.3	0.4848
166.39	0.7002
123.33	0.8684

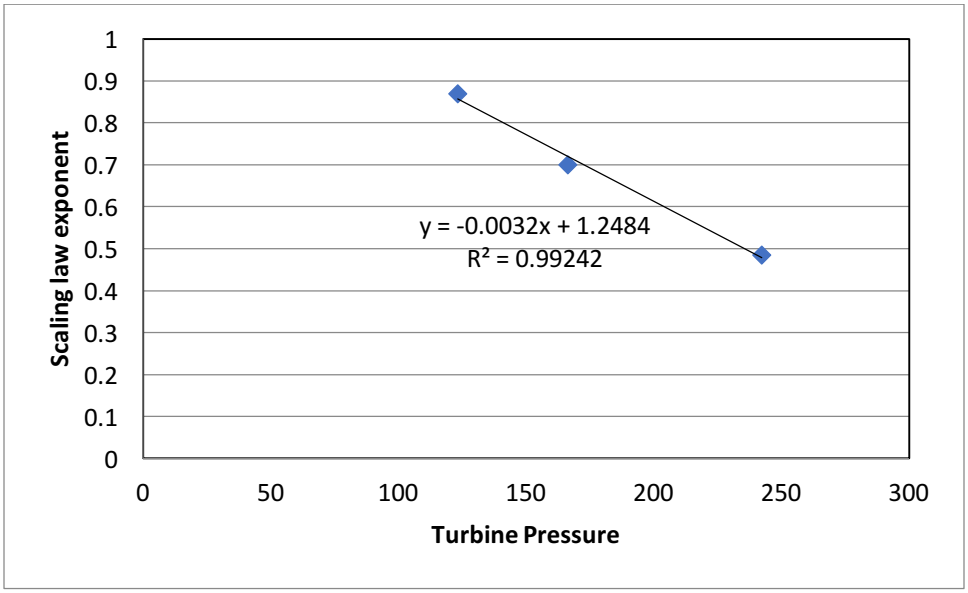

Figure 13. Steam turbine-generator scaling law exponent as a function of inlet pressure.

Table 16. PWR scaling law exponent.

	Pressure (bar)	Exponent
PWR12	67.2	1.035

It is clear that the power function exponent decreases with the pressure, showing that the economy of scale becomes more relevant at higher pressures. For a generic PWR, the calculated scaling law exponent is shown in Table 16. The value of the exponent shows a slight diseconomy of scale, suggesting that a single turbine is marginally more expensive than two turbines of half size working in parallel. However, the scaling law does not take into account the economy of scale of other turbine-related equipment, such as steam piping, auxiliaries and accessories.

3.5.2 Other Balance of Plant Components

The entire BOP, according to [23], has a total cost of $610 million (in 2017 dollars escalated according to the methodology described in [11]) including:

- Turbine generator
- Condensing system
- Feed heating system
- Other turbine-related plant equipment
- Instrumentation and control of the turbine and associated equipment
- Turbine plant miscellaneous items.

These BOP costs are dominated by the turbo-generator at $363 million, which is approximately 60% of the total cost. Additionally, the building cost is approximately $63 million. Overall, therefore, the total direct cost of the turbine and associated equipment and building is approximately $673 million (in 2017 dollars), or 26.3% of direct costs for the PWR12-BE.

3.5.3 Summary of Steam Turbine and BOP Costs

It is recommended that the fractional cost reported in [23] for the PWR12-BE for the turbine generator and associated equipment and building, 26.3% of direct costs, be used in the current N-R HES system analyses. Additionally, it is recommended that indirect costs, owner costs and contingencies be assumed to be directly proportional to the direct costs. Under these assumptions, the BOP will contribute approximately 26.3% of the overnight cost of the entire nuclear plant (i.e., 26.3% of 4100 $/kWe, according to the derivation described in [1]), or approximately 1080 $/kWe, while the "rest" of the nuclear plant without the balance of plant is then assumed to be 3020 $/kWe for the reference 1192 MWe PWR12-BE.

Regarding scaling, since it was found that nuclear turbines likely have no economies of scale, it is recommended that the cost of the entire BOP be scaled linearly with the power level. These approximations may be revised as more information becomes available on other parts of the BOP and their scaling laws, with further research in the future.

Information on the O&M costs of the BOP portion of nuclear plants was not available in the literature as a separate cost. However, assuming direct proportionality between capital and O&M costs, O&M costs can be estimated as 26.3% of the O&M of the reactor (as described in [1]).

3.6 Manifold

A schematic of the manifold equipment is shown in Figure 14, including the dimensions of the components, where Φ is the inside diameter of the pipes, L is the length of the pipes, t is the thickness of pipes and tanks, V is the volume of the tanks, p is the operating pressure and $\dot{m}$ is the mass flow rate. An example of a manifold is illustrated in Figure 15.

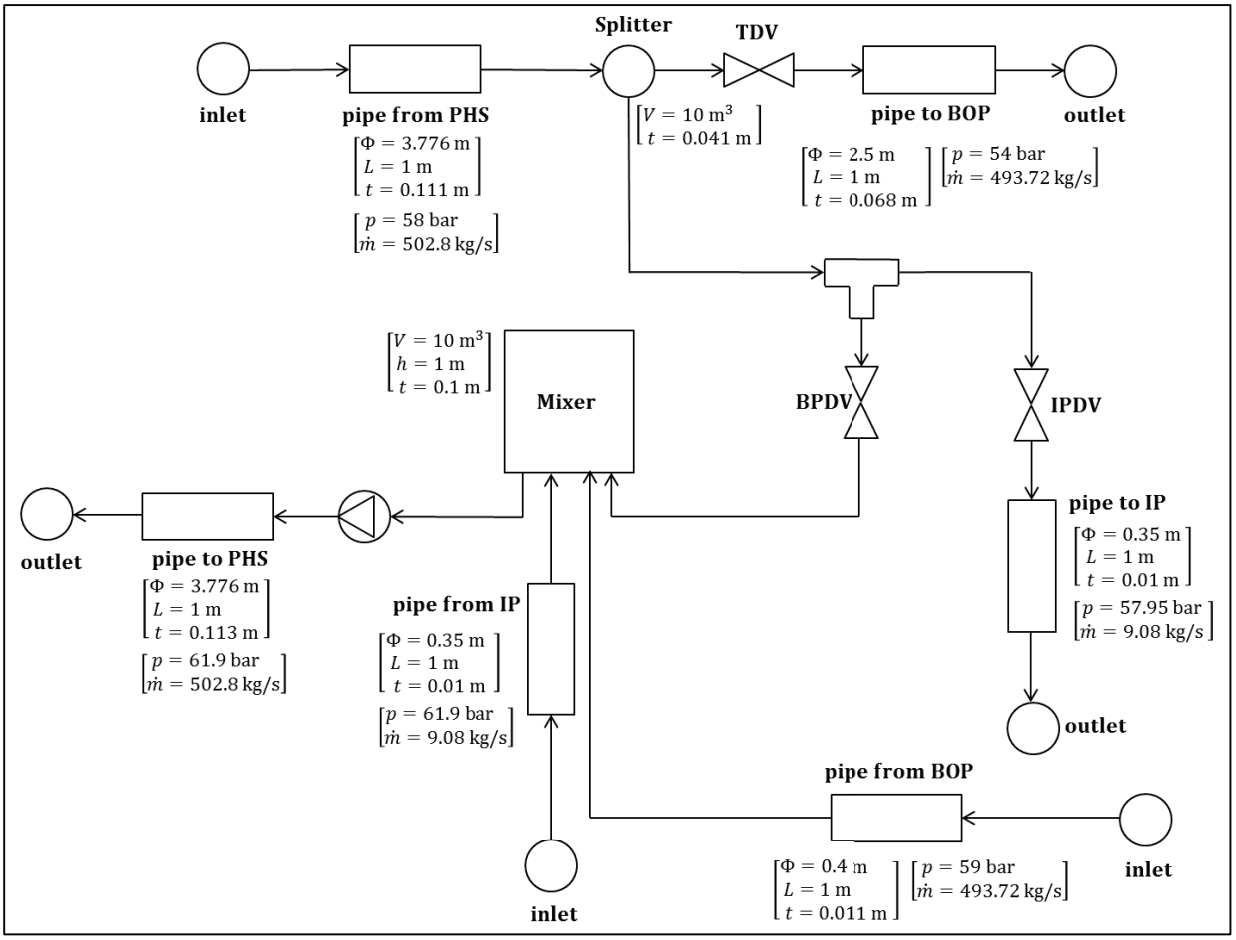

Figure 14. Schematic of the manifold structure and main components.

The total installed cost of a component is defined as the sum of the equipment cost (C_{eq}, or purchased cost) and installation cost (C_{inst}):

$$C_{tot} = C_{eq} + C_{inst} \quad (2)$$

Additionally, the cost of the building housing the manifold equipment needs to be considered.

Figure 15. Example of a manifold for steam distribution.

3.6.1 Purchased Costs

This subsection presents the cost of pipes of smaller and larger diameters, as well as the cost of tanks for mixers and splitters. Finally, a summary of all the purchased costs is provided.

3.6.1.1 Pipes of smaller diameter

From the schematic provided in Figure 14, the pipe parameters are shown in Table 17.

Table 17. Pipe parameters.

Component	Design pressure (MPa)	Diameter (m)	Thickness	Length (m)
pipe from PHS	6.38	3.776	0.111	1
pipe from BOP	6.49	0.400	0.011	1
pipe from IP	6.81	0.350	0.010	1
pipe to PHS	6.81	3.776	0.113	1
pipe to BOP	5.94	2.500	0.068	1
pipe to IP	6.37	0.350	0.010	1

The purchased costs of pipe of smaller diameters, i.e.:

- Pipe from IP
- Pipe from BOP
- Pipe to IP

were estimated using Figure 12-4 on page 503 of reference [14]. The figure, reported here in Figure 16, shows purchased costs of welded pipe per unit length as a function of diameter for both stainless and carbon steel, welded and screwed pipes, and for "*schedules*" 10, 40 and 80. *Schedule* 40 and *schedule* 80 pipes have the same outside diameters, but *schedule* 80 pipes have a larger wall thickness. Table 18 shows the thicknesses for *schedule* 40 pipes of different diameters; it is observed that all the pipes of interest here fall within the thickness of *schedule* 40. Consequently, it was immediately possible to obtain the purchase cost of the pipes of interest from [14], shown in Table 19.

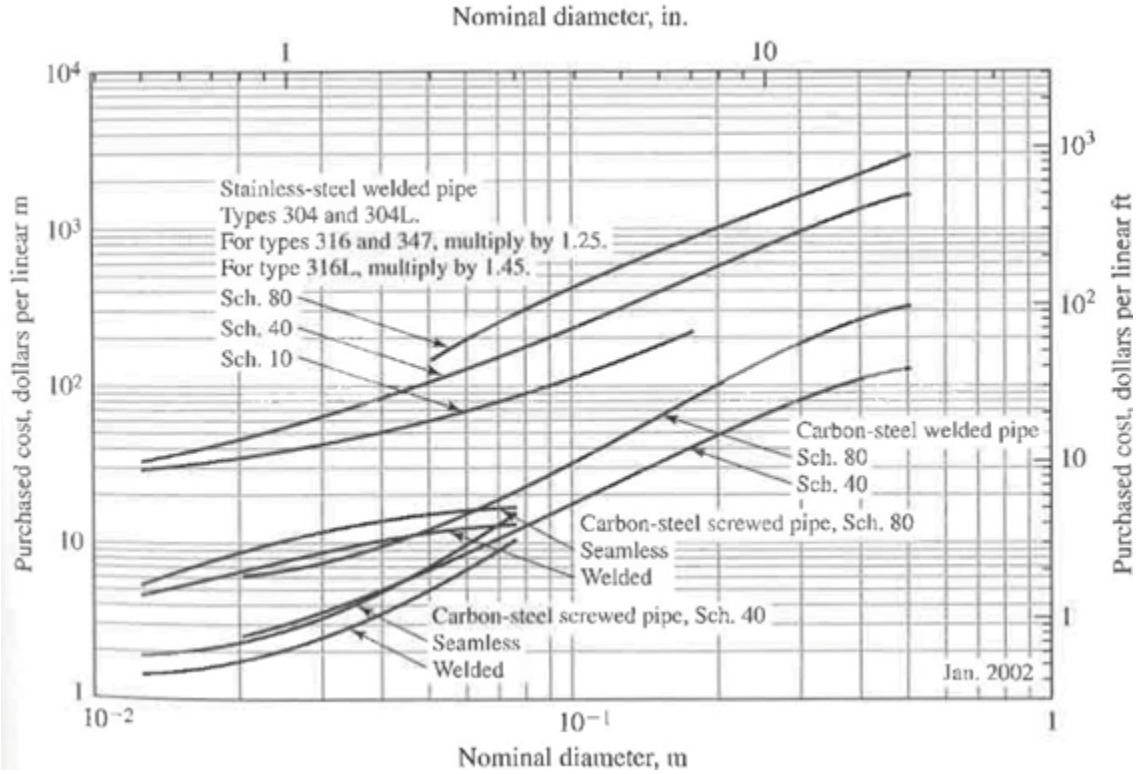

Figure 16. Purchased cost (in 2002 USD) of welded and screwed pipe per unit length, extracted from [14].

Table 18. Chart of schedule 40 Pipes (see http://www.pipeflowcalculations.com/tables/schedule-40.php).

Nominal size (inches)	Outside diameter (inches)	Outside diameter (mm)	Wall thickness (inches)	Wall thickness (mm)	Weight (lb/ft)	Weight (kg/m)
12	12.750	323.8	0.406	10.31	53.52	79.73
14	14	355.6	0.375	11.13	54.57	94.55
16	16	406.4	0.500	12.70	82.77	123.30
18	18	457.0	0.562	14.27	104.67	155.80
20	20	508.0	0.594	15.09	123.11	183.42
24	24	610.0	0.688	17.48	171.29	255.41
32	32	813.0	0.688	17.48	230.08	342.91

Table 19. Purchased costs of small diameter pipes per unit length (meters).

Component	Cost (2002 $) ($/m)	Escalated Cost (2017 $) ($/m)
Pipe from BOP	150.00	205.68
Pipe from IP	100.00	137.12
Pipe to IP	150.00	205.68

3.6.1.2 Pipes of larger diameter

Pipes to/from PHS and the *pipe to BOP* are characterized by large dimeters, substantially larger than the maximum diameter for pipes (typically 1 meter or less) available in the cost literature for chemical

plants. For this reason, the costs of these pipes were approximated using the cost of tanks. Figure 12-54 on page 558 of [14] (reported here as Figure 17) shows purchased cost (in 2002 USD) for stainless steel and carbon steel horizontal tanks of different diameters, as a function of length.

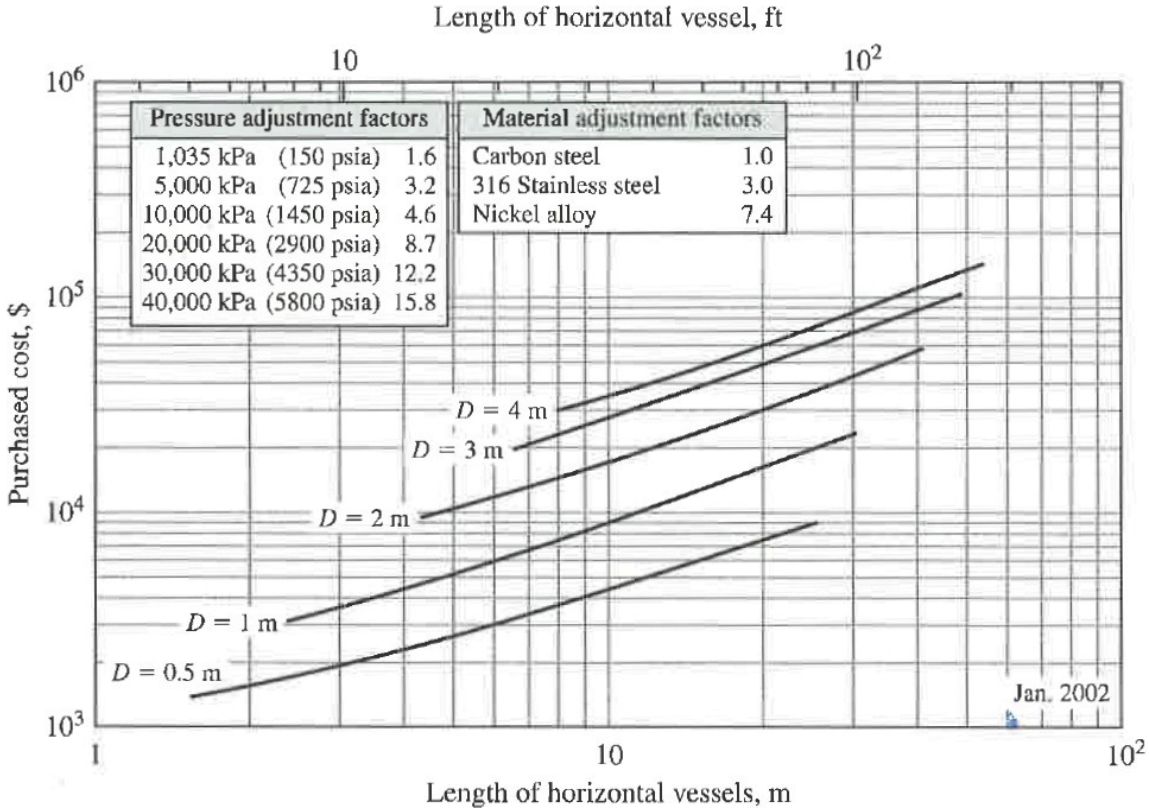

Figure 17. Purchase cost for horizontal storage vessels from [14].

The curves were linearly extrapolated to the origin of the x-axis to reflect the length of one meter (which is shown in Figure 14 as the reference length for these pipes). The extrapolated purchased cost of one-meter long tanks, as a function of the diameter, is shown in Figure 18. As the reference shows costs for tank diameters of 2, 3 and 4 meters, the data was interpolated using linear functions. Two different linear functions were used for the ranges 2-3 meters and 3-4 meters.

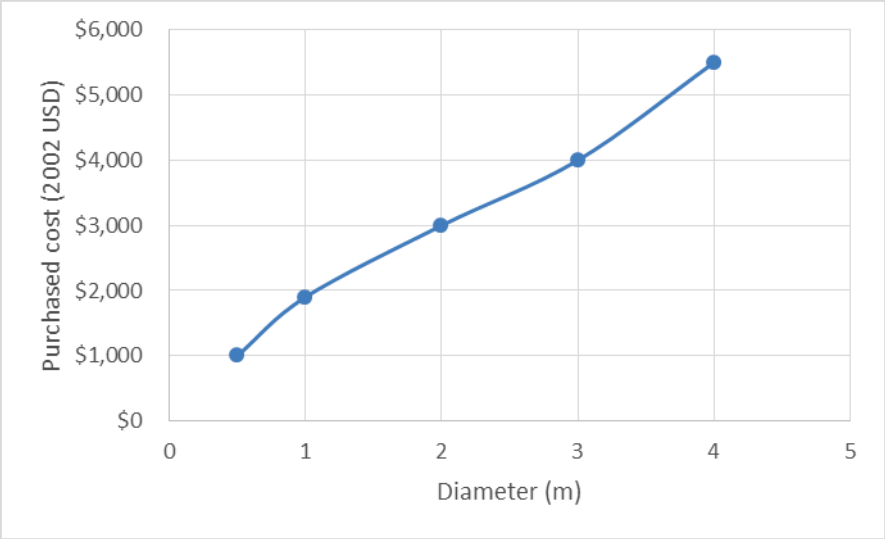

Figure 18. Purchased cost for horizontal storage vessels as a function of diameter.

Additionally, the cost of the tanks needs to be adjusted to account for different design pressures: this is accomplished by multiplying the purchased cost of a tank designed for a pressure p by a factor M_p, according to the following equation:

$$C_{eq,p} = M_P \cdot C_{eq} \quad (3)$$

Reference [14] gives the value of the factor M_p as a function of pressure (shown in Table 20, extracted from the table within Figure 17). As the pressure values of the equipment are not tabulated, pressure factors were interpolated using a linear function. The data points, along with the interpolating line, are shown in Figure 19.

Table 20. Pressure adjustment factors from [14], extracted from the table within Figure 17.

P (kPa)	Mp
1,035	1.6
5,000	3.2
10,000	4.6
20,000	8.7
30,000	12.2
40,000	15.8

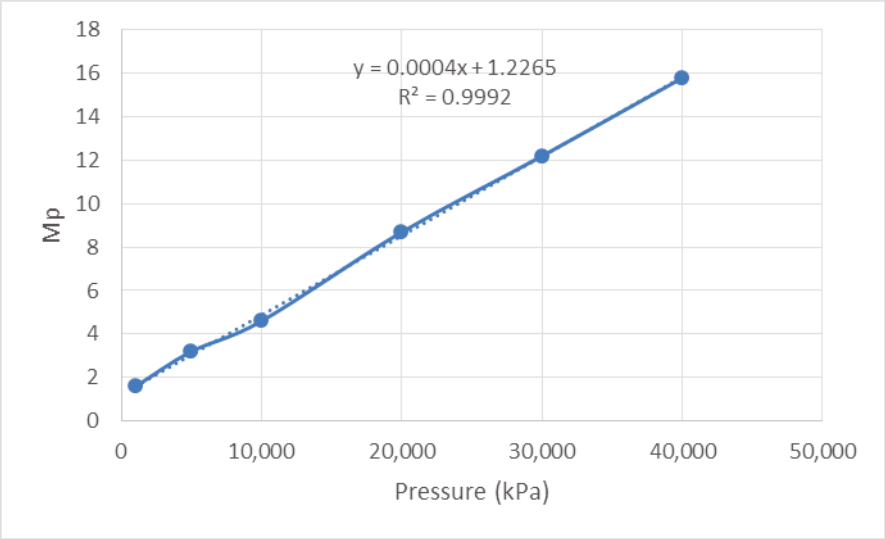

Figure 19. Pressure adjustment factors as a function of pressure, interpolated linearly to obtain the values for the exact pipe diameters of Figure 14.

The calculated costs for each component, along with the calculated and interpolated relevant parameters, are shown in Table 21. Costs were converted from the value of 2002 USD to 2017 USD using the ratio between consumer price indexes (CPI) in the period from January 2002 to January 2017.

Costs for longer lengths can easily be obtained from Figure 19, or from the values reported in Table 21, and by noticing that the cost does not increase linearly with the length of the pipe, but rather with an exponent of approximately 0.85.

Table 21. Interpolated pressure adjustment factors and costs.

Component	Diameter (m)	Cost (2002 USD)	Design Pressure (MPa)	Pressure adjustment factor Mp	Adjusted Cost (2002 $)	Escalated Cost (2017 $)
Pipe from PHS	3.776	5,164.00	6.38	3.56	18,372.23	25,192.00
Pipe to PHS	3.776	5,164.00	6.81	3.71	19,181.72	26,301.98
Pipe to BOP	2.5	3,500.00	5.94	3.40	11,889.42	16,302.77

3.6.1.3 Tanks: mixers and splitters

Purchased cost of the mixer and the splitter components were estimated using the data reported in Figure 17. Pressure adjustment factors (*Mp*), discussed in the previous section, are tabulated in Table 20 and plotted along with the interpolating line in Figure 19. The pressure factors at the working conditions of the mixer and splitters tanks (which operate at 6.38 and 6.81 MPa, respectively) were found to be 3.56 and 3.71, respectively.

The volume of the mixer and splitter were specified in Figure 14 as 10 cubic meters. Both the mixer and splitters were assumed to have cylindrical shapes, since spherical tank costs were found to be substantially greater (about a factor of 3) as compared to horizontal cylindrical tanks of the same volume and design pressure.

In order to extract the cost from the curves of Figure 18, it was assumed that each tank would have a diameter of 2 m; consequently, the height of the tank would be approximately 3.2 m. Extrapolating from

the D=2m curve of Figure 19, it is seen that the cost of the tanks, in 2002 USD unadjusted for pressure, would be approximately $8000 each. Adjusting for the pressure adjustment factors and for the CPI between 2002 and 2008, it was found that the cost of the mixer and splitter are approximately $39000 and $40700 in 2017 dollars. The assumed purchased cost for the equipment and the splitter are shown in Table 22.

Table 22. Mixer and splitter purchased costs.

Component	Vol (m^3)	D (m)	H (m)	Cost (2002 $)	Press. adjust. factor Mp	Adjusted Cost (2002 $)	Escalated Cost (2017 $)
Mixer	10	2	3.2	8000	3.56	28480	39052
Splitter	10	2	3.2	8000	3.71	29680	40697

3.6.1.4 Summary of Manifold Purchase Costs

Table 23 shows a summary of the manifold equipment purchasing costs. The costs for pipes are provided in cost per unit length ($/m), while the costs of the mixer and splitter tanks are provided for a volume of 10 m^3 and for a reference diameter of 2 meters. The cost for pipes of different lengths can be calculated from the unit cost and an exponent of 0.85. A similar exponent can be used to calculate the cost of the mixer and splitter tanks of larger volumes than 10 m^3, by preserving the diameter of 2 meters. Should a different diameter be chosen, the costs will have to be re-estimated using a similar procedure to the one described in this report, starting from data in Table 23.

Table 23.Purchase cost summary for manifold equipment.

Component	Escalated Cost (2017 $)
Pipe from PHS ($/m)	**25,192.00**
Pipe from BOP ($/m)	**205.68**
Pipe from IP ($/m)	**137.12**
Pipe to PHS ($/m)	**26,301.98**
Pipe to BOP ($/m)	**16,302.77**
Pipe to IP ($14/m)	**205.68**
Mixer (10 m^3, 2m D)	**39,052.00**
Splitter (10 m^3, 2m D)	**40,697.00**

3.6.2 Installation and Total Installed Costs

Reference [14] reports the installation costs for different types of equipment. Labor for pipe installation is estimated to be approximately 40 to 50% of the total installed cost of piping. Rearranging Eq. 2, we obtain:

$$\frac{C_{inst}}{C_{eq}} = \frac{C_{tot}}{C_{eq}} - 1 \tag{3}$$

which gives:

$$\frac{C_{inst}}{C_{eq}} = \frac{1}{60\%} - 1 = 67\% \tag{4}$$

$$\frac{C_{inst}}{C_{eq}} = \frac{1}{50\%} - 1 = 100\% \quad (5)$$

where a ratio of 50-60% was used for the equipment cost relative to total cost. Consequently, the piping installation cost results in the range of 67% to 100% (with an average of 83.5%) of the purchased equipment cost. For metal tanks, the installation cost was estimated at 30% to 60% (with an average of 45%) of the equipment cost.

3.6.3 Building Cost

The turbine building of the reference PWR12-BE [29] is utilized to estimate the cost per unit volume of a building that could be representative of a building that could host the manifold equipment. Cost of the turbine building and heater bay was extracted from [23] and is summarized in Table 24. The building cost then can be estimated using the turbine building and heater bay specific cost of 265 $/m^3, as:

$$C_{building} = 265 \cdot V \quad (6)$$

where V is the building volume in cubic meters.

Table 24. Volume, total cost and specific costs of the turbine building and heater bay, from references [23] and [29].

Turbine building and heater bay	
Volume (m^3)	207,573
Cost (escalated to 2017 USD)	54,989,900
Specific cost ($/$m^3$)	265

3.6.4 Summary of Manifold Costs

It is recommended that the unit cost data for purchased equipment shown in Table 23 be applied in initial N-R HES analyses. The estimated pipe costs are provided in cost per unit length ($/m), while the costs of the mixer and splitter tanks are provided for a volume of 10 m^3 and for a reference diameter of 2 meters. The cost for pipes of different lengths can be calculated from the unit cost and an exponent of 0.85. A similar exponent can be used to calculate the cost of the mixer and splitter tanks of larger volumes than 10 m^3 by preserving the diameter of 2 meters. Should a different diameter be chosen, the costs must be re-estimated using a procedure similar to the one described in this report, starting from data in Table 23. The installation cost can then be calculated as 83.5% of the purchased equipment costs for pipes and as 45% of the purchased equipment costs for the mixer and splitter tanks. Finally, the building cost can be calculated once its volume has been assessed, at a unit cost of 265 $/$m^3$. O&M costs for the manifold can be approximated as 1% of CAPEX, analogous to the O&M costs of the battery (see section 3.7).

3.7 Battery

The unit costs for pumped-hydro or compressed air in underground caverns typically are measured in $/MW, i.e. by cost per maximum charge/discharge speed. Battery unit cost is instead typically measured by the total electrical storage capacity in $/kWh. Reference [16] contains a summary of a literature review on the unit cost of batteries. Historically, battery storage has been on the order of 400-500 $/kWh [16]. However, there are battery makers that forecast redox batteries for 150 $/kWh, and the electric automotive manufacturer Tesla has promised a battery at 350 $/kWh [16], with recent contracts for car

manufacturers suggesting prices as low as 300 $/kWh. Battery costs are forecasted to drop further towards 200 $/kWh [31].

Separately, reference [32] contains an analysis of the levelized cost of storage as performed by Lazard Co. The investment cost of lithium batteries is expected to be between 500 $/kWh and 1000 $/kWh, similar in cost to that of sodium and of lead batteries, and of flywheels. The cost of flow batteries, based on vanadium or on zinc-bromine, is expected to be higher at 700 $/kWh to 1100 $/kWh. The useful project life for lithium batteries is reported in [32] at 5-10 years, depending on the number and depth of typical daily cycles.

Figure 20. San Diego Gas and Electric Escondido storage facility.

Two recently-built large scale grid storage projects in California [33] provide some approximate guidance on the relationship between storage and power. The San Diego Gas and Electric's Escondido storage facility (Figure 20) has a capacity of 120 MWh and maximum power of 30 MW, while Tesla's Ontario, California storage facility has a capacity of 80 MWh and maximum power of 20 MW. In both cases, the minimum time to go from full charge to full discharge (and vice-versa) is 4 hours. The official costs for the two projects were not found, but it is known [33] that a 2 MW Tesla "PowerPack" (i.e. a battery pack) costs approximately $2.9 million. Consequently, the 20 MW Ontario battery system should have an approximate cost of $29 million, or 362 $/kWh. However, acquisition contracts larger than 2.5 MW can be negotiated directly with the company, according to Tesla's website, indicating that the procurement cost of this contract could have been lower than 362 $/kWh. The time between procurement and completion for the Ontario project was on the order of months, so the interest during construction is approximately zero and the overnight and total costs are approximately the same.

Summary of Battery Costs

In summary, it is recommended that a reference unit cost for the battery of 350 $/kWh with a useful life of 10 years be applied, with possible sensitivity studies at 300 $/kWh and 5 years of useful life. Annual O&M costs for are reported in [32] at 1% of CAPEX. The maximum power should be such that the battery can be fully charged or discharged in 4 hours.

3.8 Electrical Connection

This section summarizes the cost of the electrical connection. Since the nuclear power plant is expected to have a power of about 300 MW_e, and the gas turbine a power level less than 100 MW_e, it is expected that the most informative electrical connection costs will be for a power of about 400 MW_e. Scaling laws will be provided to adjust the costs of this component for higher and lower electrical connection capacity. Figure 21 shows an example of an electrical connection.

Figure 21. Example electrical grid connection

Reference [18] includes the cost of the electrical connection as a separate account in the breakdown of the cost of new CT and CCGT plants, shown in Table 25. It is noted in [18] that the electrical connection costs also include the cost of the typical network upgrades, where necessary. It is observed that the cost of the electrical connection for a 400 MW_e plant is about $13 million, and the scaling law (derived from the average values of CT and CCGT plants) is more than linear, with an exponent of 1.26. This likely reflects the larger necessary network upgrades to accommodate larger power plants, or simply the fact that the necessary equipment for larger plants may not be entirely available "off the shelf".

While the cost of the network upgrades is not broken down separately, and therefore cannot be excluded from the analysis, it is also noted that the overall costs of the electrical connections are small as compared to the cost of other major plant equipment. Consequently, retaining the network upgrade costs will result in a negligible error in the overall calculations even when those upgrades may not be necessary.

Table 25. Cost of electrical connections for CT and CCGT plants, from [18].

Plant type	Power level (MW_e)	Cost of electrical connection ($M in 2017 USD)
CT	400	13
CCGT	600	22

Summary of Electrical Connection Costs

In summary, it is recommended that an electrical connection cost of $13 million for a 400 MWe plant be used (this cost needs to be excluded from the overall plant cost of for the gas turbine, discussed in section 3.3 of this report). This value should be scaled with an exponent of 1.26 for different required power levels. O&M costs for the electrical connection can be approximated as 1% of CAPEX, analogous to the O&M costs of the battery (see section 3.7).

4. IMPROVEMENTS TO THE REGRESSION TESTING FRAMEWORK

Regression testing is a development methodology that is used to verify the correct performance of software after modifications are made. As software is created, developers create additional small programs (*tests*) to exercise its various features. The tests and their expected results are stored as part of the repository. Before new software changes are accepted into the repository, all of the tests are run against the modified code. The outputs of these tests are then compared against expected results stored in the repository. If a result is different than expected for any test, or if any test does not run successfully, then the test is considered to have failed. When the tests do not pass, a developer must determine why the failure occurred before a proposed change may be accepted. Regressions tests are one way to assure the quality of software produced.

As described in the HYBRID project report for FY16 [34], the RAVEN code is used to optimize system economics. A key RAVEN capability used in performing this analysis is the ability to incorporate multiple user-developed *external models*. RAVEN is able to combine multiple such models in a single run to solve a larger combined problem [35]. For this project two types of external models are used:

- Those written in the Modelica programming language, used in HYBRID to simulate the behavior of physical systems that make up the N-R HES under varying conditions.
- Additional code written in the Python scripting language that describes economic aspects of the problem including electricity dispatch and costing.

Automated regression testing of Modelica was successfully implemented in FY16 and is now an integrated part of the HYBRID physical model development process. Added in FY17 is the ability to perform test runs of RAVEN from outside its normal file system location. This allows external models stored in the HYBRID repository (such as the economic and dispatch modules) to be tested using RAVEN without having to be made a part of RAVEN itself.

Significant effort has been invested in developing the economic and dispatch models. The RAVEN code used to execute them is under active development. Having the ability to test these modules in an automated manner provides a mechanism to detect any regression problems inadvertently introduced by ongoing RAVEN development. Each time the HYBRID repository is to be updated with a new RAVEN version, these modules will be tested with it before the update is accepted. Any failures in those tests provide a warning about a regression issue before the update to RAVEN becomes a part of the HYBRID repository.

4.1 Testing System Prerequisites

The module test system consists of scripts written in the bash shell language and Python. It may be used on any platform that is supported by RAVEN (Linux, Mac, and Windows). The following conditions must be satisfied for the module test system to function properly:

- The RAVEN submodule of the HYBRID repository used to run the module tests must be initialized and fully updated. It is this installation of RAVEN that will be used to run the module tests.
- The MOOSE [36] submodule of the RAVEN submodule must also be initialized and updated. RAVEN (and by extension Hybrid) makes use of the testing system contained in MOOSE.
- The system running the tests must be configured with the software prerequisites necessary to build and run RAVEN [4]. These include a Python interpreter, Python libraries (*h5py*, *matplotlib*, *numpy*, *scipy*, and *scikit-learn*), and development tools (C++ compiler, *Miniconda* package manager for Python, and *git* source code control).

- The RAVEN submodule must be built with the appropriate compiler before it can be used to run the tests.

4.2 Test Location and Definition

The RAVEN submodule that is already part of the HYBRID repository contains a complete testing system used to provide regression testing for itself. HYBRID module testing is achieved by extending the RAVEN system so that it can run tests stored in the parent repository (which it would not normally be able to do). Files associated with RAVEN module testing reside in the *raven_module_testing* directory located immediately below the root of the HYBRID project repository. The main script file that configures the RAVEN test system for Hybrid-specific functionality is placed in this directory (see Figure 22) Inside the root folder "hybrid" it is defined a "raven_module_testing" folder which contains the scripting for the testing and the folder "tests". The folder "tests" contains the input of the tests and the corresponding output (in the "gold" folders) which are used to verify that the behavior of the code is not changed at each regression test. For organizational purposes the tests are usually collected in subfolders based on their characteristics.

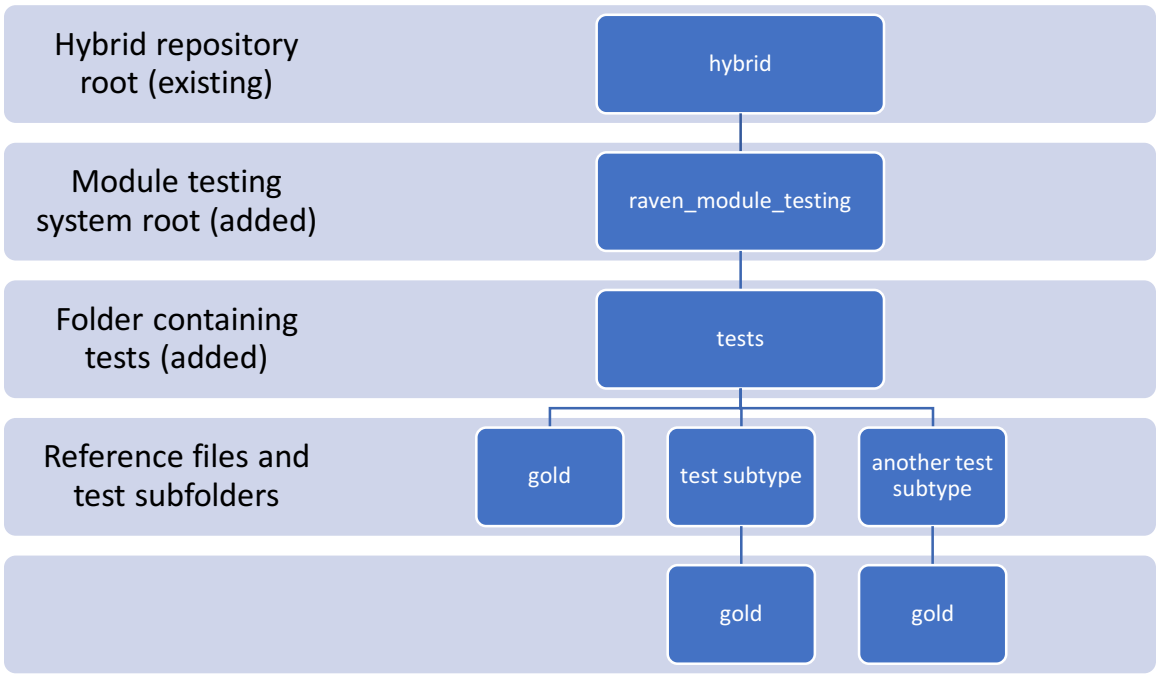

Figure 22. Folder tree used by the regression system.

Hybrid-specific module tests are defined in the same manner as they are for RAVEN. A single test consists of a RAVEN input file along with associated data needed to perform that run. That can include input data, external models, and Python files. These may be placed in the *tests* directory, or they may also be placed in subdirectories of *tests*. Every directory that contains tests to be run by the framework must contain a *test specification file* named "tests". The syntax of these files is defined by the MOOSE test framework [36], which controls how each test is run and sets the criteria used to determine whether it passed or failed. An example of a test specification file is presented here:

```
[Tests]
 [./pcaResponseSurface]
  type = 'RavenFramework'
  input = 'test_pca_responseSurface.xml'
  csv = 'rsPCA/BBRDesign_dump.csv rsPCA/CCRDesign_dump.csv'
  max_time = 500
  rel_err = 0.001
 [../]

 [./2steps_same_db]
  type = 'RavenFramework'
  input = 'test_2steps_same_db.xml'
  csv = 'test_2steps/fullprint_HistorySet_3.csv'
  UnorderedXml = 'test_2steps/fullprint_HistorySet.xml'
 [../]
       .      .
   (Other tests may be defined here)
       .      .
[]
```

In the above example, two tests are defined, named "pcaResponseSurface" and "2steps_same_db". They are both defined as tests that use the RAVEN-specific test module. In most cases, this RAVEN test module (defined in the test file as "*type = 'RavenFramework'* ") will also be used for Hybrid module testing.

Comparison criteria are also defined in the "tests" file. In most cases, one or more output files generated by running the specified input file with RAVEN are compared against a *gold standard* provided by the developer and stored in the repository. Typically, comparisons are performed on numeric values contained in *Comma-Separated Values (CSV)* files to a defined tolerance. It is also possible to compare output files written in *eXtensible Markup Language (XML)*. When these file comparisons are specified by the test developer, reference files must have the same name and be placed in the *gold* subdirectory below that containing the "tests" file.

4.3 Running the Tests

The test system is accessed by running an operating system shell script located in the top directory of the HYBRID project repository. This file *run_raven_module_tests* has the same parameters controlling its function as the RAVEN test script. The format of the command with the most used parameters is provided below:

```
run_raven_module_tests [-h | --help]
                       [-j <number of processes> |
                        --jobs <number of processes>]
                       [--re <regular expression>]
```

Where:

-h or –help	Causes the complete list of parameters the script will accept to be printed and then exit without running any tests. Note that there are many other parameters provided by the RAVEN/Moose test framework, many of which do not apply to Hybrid module testing.
-j or –jobs <number of processes>	Runs multiple tests in parallel, allowing at most <number of processes> of them to be active concurrently. Using this option when the computer running tests has several logical processors will reduce the total amount of time needed to run multiple tests.

-re <regular expression>	When specified, this parameter will only execute tests whose name matches the *regular expression* provided. This is typically used to run single tests.

4.4 Repository Test System (Civet)

As described in the FY 16 report [34], the purpose of having a single top-level test script is to provide a way for an automated system to run all of the module tests together. The HYBRID and RAVEN projects make use of a tool called *Civet* [38], developed at INL, that provides regression testing services. Each time a developer proposes modification of the contents of the HYBRID repository, *Civet* will cause the automated tests to be run on the modified version. These tests (which are performed separately on both the Modelica models and the HYBRID add-on modules) must all pass before a proposed change may become part of the official repository. In this way the Hybrid project is protected from the accidental introduction of flaws into the software that required significant investment of resources to develop.

5. OPTMIZATION IMPROVEMENTS

As already mentioned one of the cornerstone of the framework developed is the capability to optimize the N-R HES economic performance. The parameter space to be optimized is the capacity (size) of each sub-system of the N-R HES. This could not be done without considering the optimal dispatch strategy of the system resources to meet the electricity demand. Consequently, optimal sizing of the system is connected to the dispatching problem and the dispatching should be optimized given the stochastic nature of net electricity demand (or electricity prices). The whole problem could be therefore classified as a stochastic optimization problem which parameter space is both the system sub-compoenent size and the dispatch.

5.1 Pre-existing Features

The first part of section 5 provides a summary of the already available capabilities of the RAVEN code to perform optimization in stochastic systems. This should also serve as a short recall of the development already performed and reported in [1, 39].

5.1.1 SPSA

The algorithm implemented in RAVEN used for optimization problems in this work is based on the Simultaneous Perturbation Stochastic Approximation (SPSA) technique. This gradient descent method approximates a N-dimensional local gradient by comparing the change of value between only two points. Because a traditional gradient calculation requires N+1 evaluations, this can save a large amount of calculations at the cost of introducing some error into the gradient. If enough successive gradient evaluations are performed, it is expected that the gradient descent algorithm will move towards the optimal point on average. Features of the RAVEN implementation of SPSA include the choice between minimization and maximization problems, fine control over step sizing parameters, and limits for each variable being optimized, limiting value searches to the hypercube formed by the input variable limits.

5.1.2 Normalized Input

It is common for inputs to a model to vary widely in magnitude. For instance, the inputs to the models in this work range over eight orders of magnitude. As a result, the input domain for the optimizer search can be very irregular; as a result, the correct step size to take in a gradient descent can depend strongly on what direction the step is going. To mitigate this complication, all input data is normalized, mapping original domains to the 0 to 1 domain. The result is a hypercube with length of 1 on any side, making it more practical to traverse the space in a consistent manner. We do note, however, that the distances from point to point in the hypercube tend to increase as the number of dimensions increase. Regardless, a step size of 0.5 always means a step size equal to half the length of one dimension.

5.1.3 Stochastic De-noising

One way to combat the complication of stochastic noise in an optimization environment is to find expected values of points and gradients. The SPSA algorithm in RAVEN allows the user to request multiple evaluations of each sample taken, which will then be used to calculate expected values and reduces the stochastic noise. This helps prevent the optimizer becoming stuck on a point that was optimal once, but on average is less optimal than the points around it.

5.1.4 Multi-trajectory

A significant pitfall of gradient descent algorithms is local minima. Because gradient descent follows gradients, it is easily trapped by plateaus or valleys that are ideal locally but may not be a global optimum. To combat this phenomenon, RAVEN allows multiple initial starting points to be identified by

the user, each of which will follow an independent optimization walk, or "trajectory". The global optimum is assumed to be the most ideal of the endpoints of each trajectory. Alternatively, if one trajectory begins following another trajectory, the first is eliminated and only the second continues. The threshold for termination is a parameter available to the user.

5.2 New Developments

This section covers the additional developments introduced to improve the optimization convergence along with some academic examples that illustrate the effect of such improvements.

5.2.1 Adaptive Stepping

One of the challenges with the original SPSA implementation is a tendency to slow convergence across a domain. The search step size in this algorithm can be adjusted somewhat by a few parameters, then is determined largely by the magnitude of the gradient as well as the number of steps taken. Large gradient magnitudes result in large steps taken, and the number of search steps taken tempers the size of step. As a result, if gradients are fairly steep near optimal points, it takes a great number of search iterations to reduce the step size enough to achieve convergence.

To mitigate this issue, we implemented an adaptive step sizing determination algorithm in RAVEN's SPSA algorithm. This adaptive step feature uses the estimated gradient for direction information only, and determines the step size based on previous performance. To determine the size of the next step, the algorithm considers the scalar product of the two previous step directions, ignoring the step sizes themselves. If the scalar product is positive, the algorithm infers a larger step size is needed in order to search for a change in gradient direction, while if the scalar product is negative, the step size needs to be cut in order to converge on the optimal point. If the scalar product is zero, the step size should remain unchanged.

The mechanical change in step size is based on raising a gain factor to the scalar product of the previous gradient directions. This assures that at most, the step is multiplied by the gain factor, and at smallest, the step is divided by the gain factor. For example, given a gain factor of 2, if the SPSA algorithm determines the gradient moves in the same direction twice, then the next step will double in size, while if the gradient exactly reverses direction, the next step will be halved. This allows for quickly moving across the input space as well as converging rapidly in valleys. We found for many test-cases a value of 1.5 for the gain factor performed well by exploring the domain but still converging quickly, without experiencing false convergence from noise in the zeroth-order gradient approximations that are key to SPSA.

However, we observed that using the same gain value for both growing and shrinking the step size is not always ideal. While a reasonably large value for the gain factor assures rapid movement across the domain, the large value causes the step size to be cut too quickly in some cases and result in false convergence. To this end, we divided the gain factors into a gain growth factor and gain shrink factor. When the scalar product of the gradients is positive, the step size is multiplied by the gain growth factor raised to the gradient scalar product. When the product is negative, the step size is divided by the grain shrink factor raised to the product. In general, we found success with a gain growth factor of 1.5 and gain shrink factor of 1.25, although these values are problem dependent.

In Figure 23 and Figure 25 the input variable is on the x-axis and the target on the y-axis, and the optimizer is seeking a minimum point. The initial point for the optimizer is the maximal value for x. In Figure 23, the old method, many small steps are required to find the minimum point. In Figure 25, the step size adaptively grows until the turning point is found, then the step size shrinks to find the minimum point.

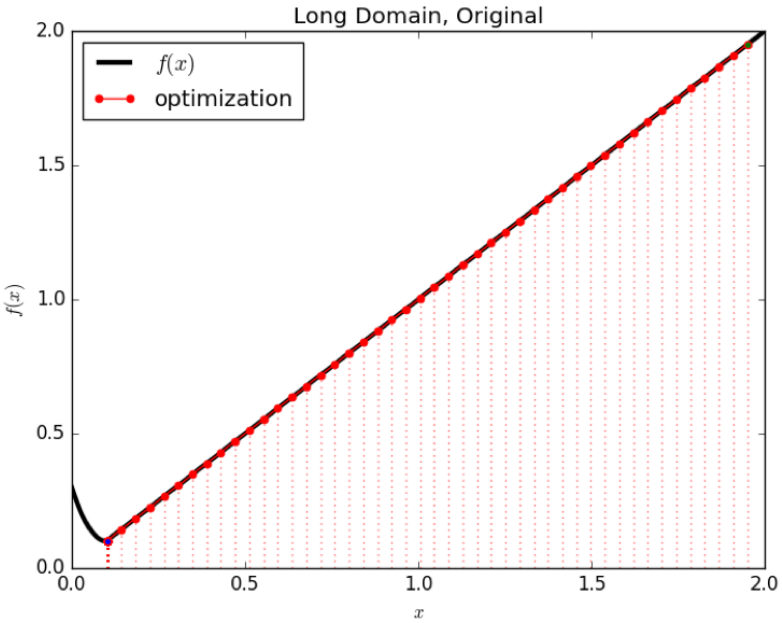

Figure 23. Comparison of adaptive versus not adaptive stepping.

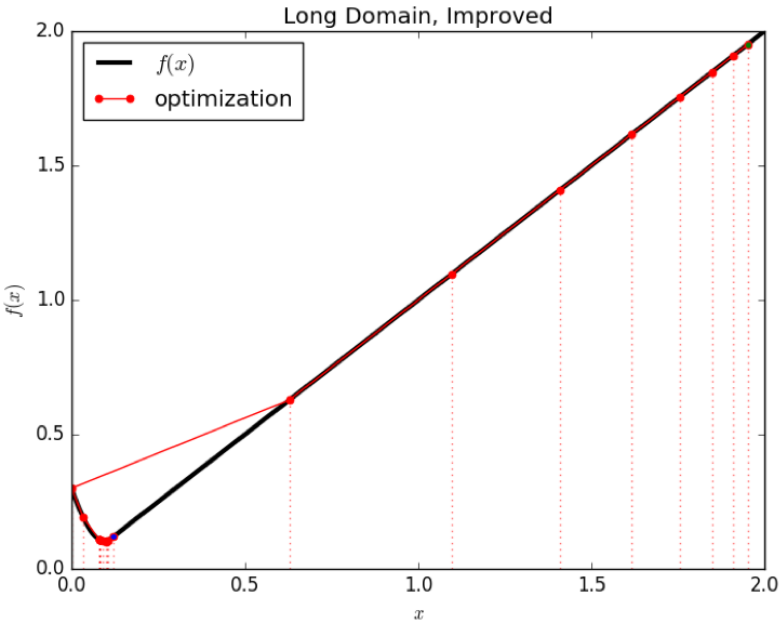

Figure 24. Comparison of adaptive versus not adaptive stepping.

5.2.2 Restriction to Improvements

Another reason SPSA can be slow to converge is that once it estimates a gradient and determines a direction to move in, there is never a check to see if the new proposed optimal point is actually more optimal than the point recently left by the search algorithm. In order to speed up convergence, we implemented an "only improvements" adjustment that will reject proposed optimal points that on inspection are less favorable than the current optimal point. In the event a proposed point is rejected, the step size is cut as per the adaptive stepping algorithm, and the gradient is re-evaluated before a step is taken. One weakness of this restriction is if the optimization algorithm gets "lucky" with a particular evaluation of the stochastic space, it may never find a more optimal input point, even if other points are better on average. Thus, it is necessary to carefully consider de-noising when using a stochastic model and the SPSA optimizer in RAVEN.

In Figure 24 two inputs are on the x-axis and y-axis, and the color represents the target value, with blue as the minimum and red as the maximum. The target value is calculated using Beale's function. On the left is the old method, where we see the path of the optimization algorithm move in an erratic fashion around the domain, closing in on the minimum without ever really settling on it. On the right, we see the new restricted movements, which in the same number of iterations not only reaches a better minimum, but with far fewer movements.

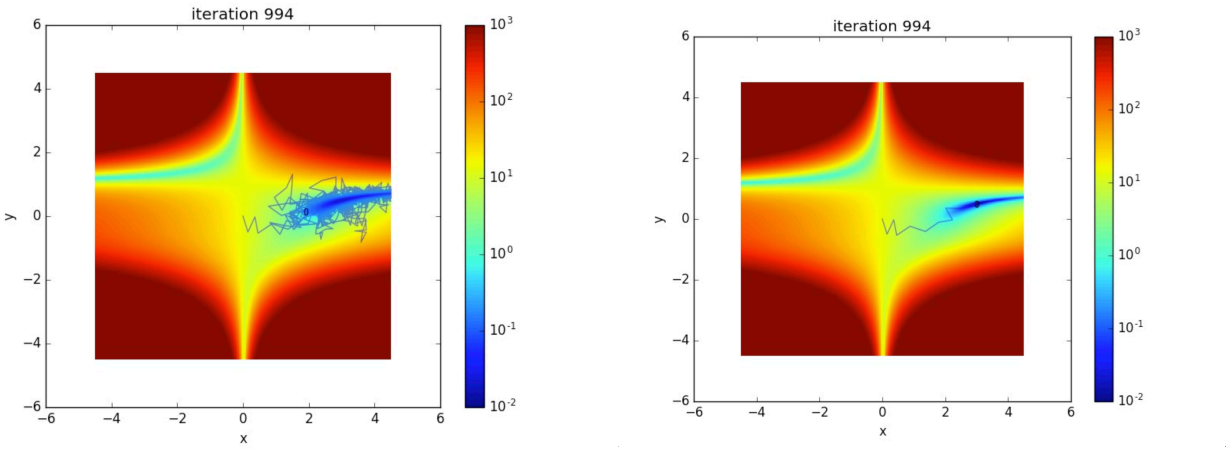

Figure 25. Effect of the "restriction to improvement" on convergence, where the old method is implemented in the left plot while the option is activated in the right plot.

5.2.3 Intra-domain Constraints

Under the original SPSA implementation, we attempted to implement a constraints system that allows the user to specify portions of the input space that are off limits to the optimization search. Because of the step sizing algorithm, there were complications that could not reasonably be resolved until the adaptive stepping was implemented, at which point user input constraints were made a normal part of the optimization algorithm. The user constraints are treated similarly to the boundary constraints; if the optimization algorithm takes it to an area that violates the constraints, a projection is taken angling away from the algorithm's preferred direction of travel, seeking an acceptable point in the input space. As with boundary constraints, whenever the optimization algorithm attempts to violate intra-domain constraints, the step size is reset to that suggested by the original SPSA algorithm before restarting the adaptive stepping.

Figure 25 shows the Beale function described above, with the addition of an arbitrary constrained section marked by the black circle. We can see how the optimization path moves around the constrained circle to find the minimum.

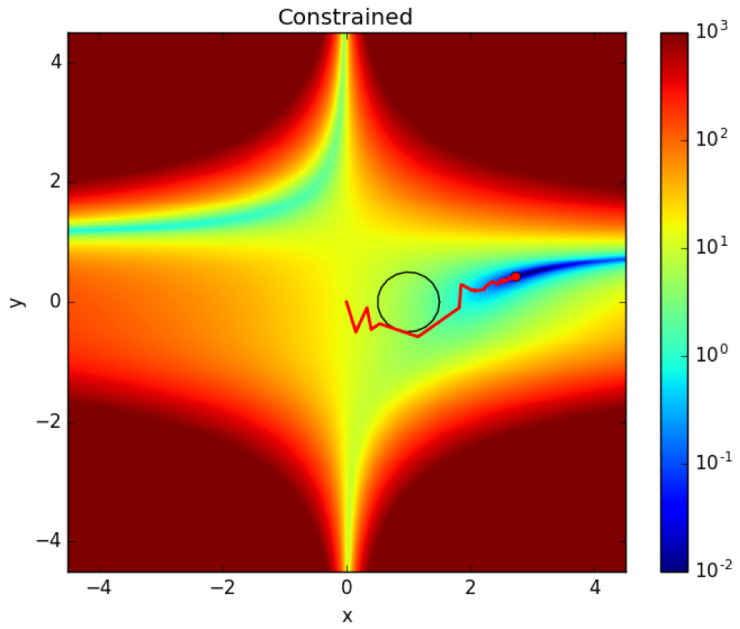

Figure 26. Effect of an intra-domain constraint on the optimization path.

5.2.4 Multi-Level Optimization

One of the major improvements to the stochastic optimization capability, or more generally to any optimization scheme in RAVEN, is the capability to separate the optimization of different variables. To clarify the idea, we can focus on the problem of optimizing both the dispatch and the capacity of the subsystems. The problem is intrinsically stochastic due to the nature of the wind speed and the electricity demand and, therefore, is very difficult to solve. There are a few considerations that indicate that a better strategy than a brute force approach could be identified.

- The convergence speed of any optimization problem usually tends to decrease more than linearly with the number of variables.
- The dispatch problem needs to converge in average (stochastic optimization converges to the average optimum), but it could be characterized by a high noise to signal ratio. A different time history of demand generates a completely different corresponding optimal dispatch.
- The economic performance of the systems needs to be good on average (the size of the system should not be changed depending on a specific realization of the demand).
- The economic performance of the system is rather stable when the demand and wind time histories change, provided that the dispatch scheme is optimal.

Given the above considerations, it is normal to seek two different approaches for the optimization of the dispatch and subsystem capacity. This has been achieved by allowing RAVEN to split the optimization in two levels, where the outermost level (master) optimizes the subsystem components,

while the innermost level optimizes the system dispatch for a given time history of wind speed and electricity demand and for given subsystem sizes.

The computational flow could be summarized as it follows:

1. The subsystem capacities are fixed
2. A perturbation of the capacity is generated
3. Wind and electricity demand profile are generated
4. The SPSA algorithm is used to determine the optimal dispatch. It is relevant to notice that:
    - at this level, there is no more stochastic behavior of the system,
    - consequently, no de-noising is needed and SPSA becomes a standard steepest descent algorithm with a zero-order approximation of the gradient
5. Steps 3 and 4 are repeated, as many de-noising iterations are requested for the capacity optimization
6. The new assessment of the economic performance of the system is used to move the subsystem sizes toward the optimal point.

5.2.5 Acceleration of the Dispatch Optimization

Any physical system possesses characteristic time scales, which are representative of the time a forcing function (input) takes to get reflected in the response (output). The longer this time is, the higher the "inertia" of the system or, equivalently, the memory of the system. For times much longer than the characteristic time of the system the initial conditions are completely forgotten and the response of the system is only related to the changes in the input. At those time scales the system is, therefore, memoryless.

When evaluating the optimal dispatch (the one corresponding to the best economic performance) for a N-R HES, if the time scale is such that the system can be considered memoryless, then the optimal strategy is given by the minimization of the marginal cost at each time step. This is due to the fact that, to meet a given net demand, the best option is the one corresponding to the least cost to satisfy that request (i.e. least marginal cost) unless this would unfavorably impact the future performance of the system. This is clearly impossible under the assumption of no inertia of the system.

To give an example, nuclear power plants are willing to pay someone to take electricity (i.e. accept negative electricity prices) for short periods of time, incurring very high marginal costs (variable costs plus price paid to sell). The reason is that nuclear plants have a very long characteristic time, and, if they respond to negative prices in the short term, they would not be able to sell electricity in the future when the electricity price is again positive. On the other hand, plants that provide power to meet peak demand ("peaker" plants) stop production when electricity prices are low. This response is due to their very short characteristic time that allows them to respond quickly in the event of future higher electricity prices.

Not considering the reactor, which is nominally operated at steady state, all the components of the N-R HES have a very short memory. Therefore, it is natural to seed the dispatch optimization with a guess generated using a marginal cost analytical formulation. Hence, the optimization scheme is changed as follows:

1. The subsystem capacities are fixed
2. A perturbation of the capacity is generated
3. Wind and electricity demand profile are generated
4. A marginal cost based dispatch is generated

5. The SPSA algorithm is used to determine the optimal dispatch, using the dispatch determined by the marginal cost based dispatch as the initial point. It is relevant to notice that:
    - at this level, there is no more stochastic behavior of the system,
    - consequently, no de-noising is needed and SPSA becomes a standard steepest descent algorithm with a zero order approximation of the gradient
6. Steps 3 and 4 are repeated, as many de-noising iterations are requested for the capacity optimization
7. The new assessment of the economic performance of the system is used to move the subsystem sizes toward the optimal point.

5.2.6 Summary of New Developments for Optimization

Several improvements have been implemented to the original optimization algorithms, which have decreased the number of necessary iterations to achieve the solution. Unfortunately, a large number of parameters are needed to control the optimal performance of all the algorithms. Consequently, several assessment tests will be required to decide the best configurations. Nonetheless, the degree of flexibility achieved is remarkable and provides confidence that the right tool has been developed to solve a very difficult problem.

6. MODEL AND COMPUTATIONAL FRAMEWORK

This section describes the N-R HES configuration that was modeled in the example problem presented in this report, as well as the computational framework to analyze the system. The N-R HES is comprised of several components, namely:

- Nuclear reactor (light water cooled)
- Gas turbine
- Battery storage
- Steam Manifold
- Industrial process (HTSE H_2 plant).

The different components of the N-R HES are described in more detail in section 2. The computational framework is set up to analyze the economics of the N-R HES, as previously described. The framework allows analysis and optimization of two aspects of the NHES system, per the detailed discussion in section 5:

- Capacity planning (subsystem sizing)
- Dispatch.

The capacity-planning problem tries to determine the energy mix leading to the least cost of electricity, i.e. the different plant sizes (except the nuclear reactor, which is assumed to be always 300 MW_e) necessary to satisfy a net demand profile. For all cases considered here, the demand profile the N-R HES system tries to satisfy is the net demand profile, i.e. the demand after subtracting all renewable contributions. As a consequence, the levelized cost of electricity (LCOE) computed here, is the LCOE of the N-R HES to cover the net demand only and not a global LCOE to cover the gross demand, which would have to include the renewable sources and the part of demand covered by them. The dispatch problem, on the other hand, tries to determine how the different plants need to be operated (for given capacities) to minimize the (LCOE). The dispatch determines for every time step how productive a plant is, e.g. at 50% capacity at time 0, at 85% capacity at time step 1, etc.

A schema of the computational framework is shown in Figure 27. As one can see, the RAVEN 'outer level' has the possibility to sample the capacities of all components of the N-R HES. The "sampler" can be any sampler available in RAVEN, e.g. a grid (i.e. parametric sweep) or Monte Carlo sampler, or an optimizer. The outer level generates synthetic time histories (ARMA) for gross demand and wind speeds that are fed into the ARMA post processor (ARMA PosP). This post processor converts the wind speeds into electricity (scaled by the desired renewable penetration) and generates the net demand profile. Capacities and net demand are then passed to the 'Initial dispatch' module. This module computes a time dependent dispatch for all the components in the N-R HES. This dispatch is based on marginal cost. It is passed along with the sampled component capacities and the net demand profile to the 'inner level' of RAVEN. The inner level samples the dispatch for all components in the N-R HES. As mentioned above, this means that the utilization factor of each time step and each component of the N-R HES are variables to be sampled. For example, running the simulation for one day having a time step size of one hour leads to 24 variables for each component of the N-R HES. This sampler can be, as for the outer loop, any sampler available in RAVEN, e.g. a grid or Monte Carlo sampler, but it will most likely be an optimizer.

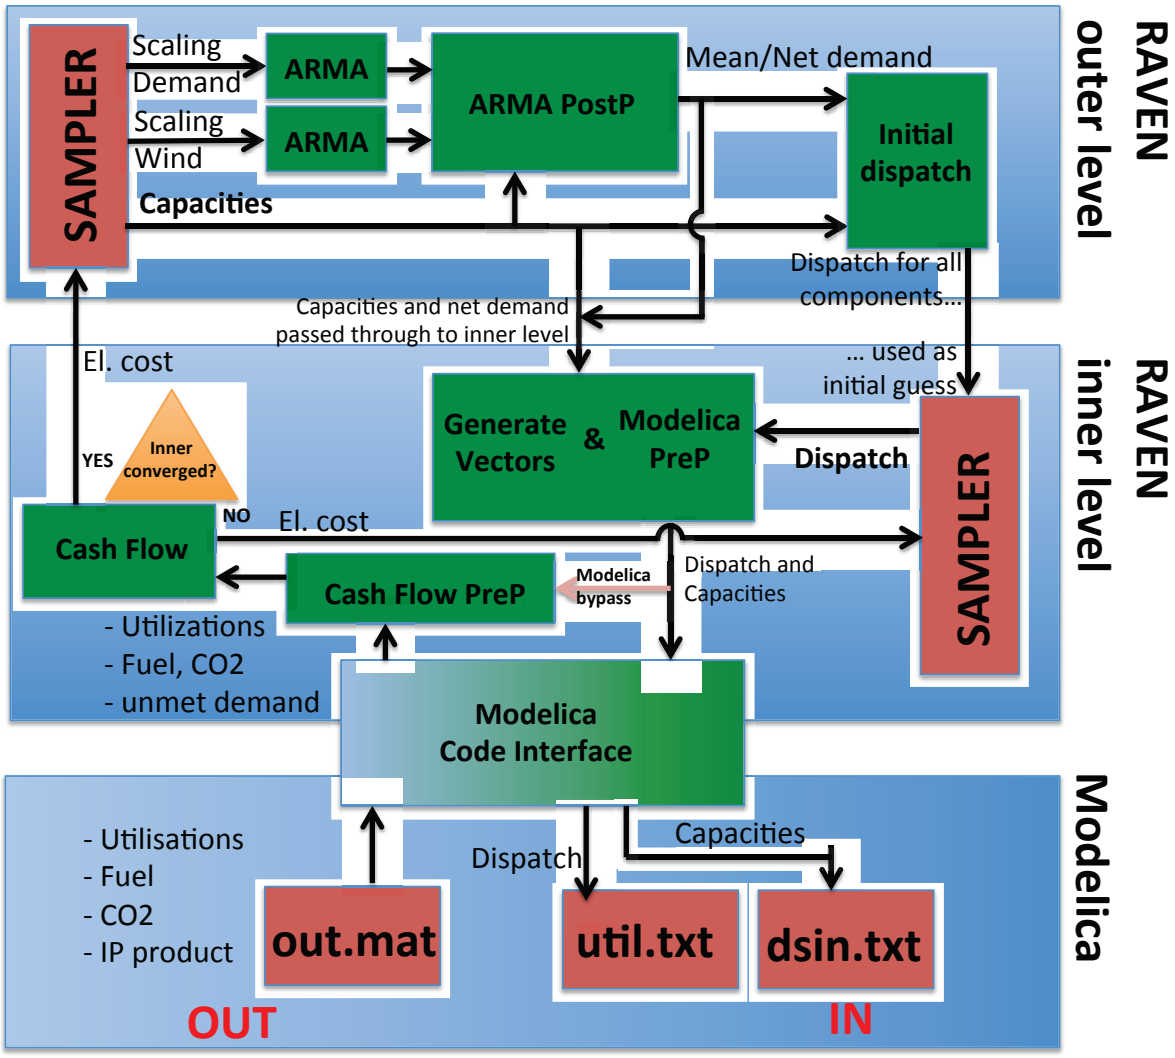

Figure 27. Computational framework for economics analysis of N-RHES: Capacity and dispatch optimization.

This separation of inner and outer loop allows separation of the stochastic aspects of the model (ARMA) into the outer loop, i.e. only the outer loop needs to deal with the effect for the same set of component capacities and mean demand, but for a different demand and wind profile, the LCOE will change. The outer loop needs to average the LCOE for multiple evaluations of the same capacities with different system demand and wind histories, while the inner loop optimization is done for a fixed demand history.

The dispatch passed inside from the outer loop ('Initial dispatch') is used as initial guess for the 'inner' sampler. The sampled dispatch, the capacities sampled in the outer loop, and some constants are passed to the 'Generate Vectors' and 'Modelica PreP' modules that put all the data in the form that Modelica needs. The rearranged data is then passed to the 'Modelica Code interface' that runs the Modelica code and collects the outputs from the Modelica run. These outputs are the actual possible dispatch, fuel consumption and CO_2 production as well as the commodity production from the industrial process. The Modelica model tries to follow the dispatch provided by RAVEN. The actual possible dispatch output by Modelica might be different due to physical constraints in the dynamics of the N-R HES components that make it impossible to reach the requested dispatch point. In that case, Modelica outputs the closest physically possible dispatch. Once all output is collected, it is passed to the 'Cash

Flow PreP' module of RAVEN that prepares all data needed for the economics analysis, i.e. the 'Cash Flow' module in RAVEN. This last module computes the cost of electricity and returns it to RAVEN. Depending on user defined criteria RAVEN runs another 'inner iteration' with another dispatch using the same component capacities as in the previous iteration or goes back to the outer loop and samples a new set of capacities. All the different modules in the framework are explained in more detail in the following sections. For the N-R HES considered here, Table 26 shows all the variables managed by the framework. Green indicates that the module provides the variable and red indicates that the module requires the variable.

It is worth noting that this framework allows for different workflows. For example, if it is assumed that the 'Initial dispatch' is already the optimal dispatch (which might be true depending on the time step size considered following discussion in 5.2.4), and no 'inner' optimization is needed, the number of iterations can be set to 1 in the inner loop. In this manner, the cost of electricity is only evaluated for the dispatch coming from the 'Initial dispatch' module and no further inner iterations are performed. Similarly, if it is assumed that the sampled dispatch is the actual possible dispatch (which might be true depending on the time step size considered), Modelica does not need to be run. In that case, Modelica can be bypassed and the data assembled by the 'Modelica PreP' module (the dispatch and component capacities) can be passed directly to the 'Cash Flow PreP' module. The workflows used in the calculations presented in this report are described in Section 0.

Table 26. RAVEN variables managed by the framework. Green indicates that the module provides the variable and red indicates that the module requires the variable.

	Models											
(RAVEN names)	Sampler Outer	ARMA	ARMA PosP	Dispatch Init	extOpt Outer Level	extOpt Inner Level	Sampler Inner	Generate Vectors	Dymola PreP	NHES	CashFlow PreP	Cash Flow
Constants												
BOP_eff	Const					pass	through					
IP_EL	Const					pass	through					
IP_TH	Const					pass	through					
IP_KG	Const					pass	through					
H2_price	Const					pass	through					
Penalty_mult	Const					pass	through					
Start_time	Const					pass	through					
DYMOLA_tot_time	Const					pass	through					
The_constant_one	Const			for NPV								
Pow2Capa_battery	Const											
Component needs (NPV)												
NPV_mult (LCOE)												
Time												
Time1												
DYMOTime												
Demand												
Demand												
scaling_demand						pass	through					
Demand_time												
Demand_time_net												
Demand_time_net_XXXX						pass	through					
Grid_DYMO_productionEL												
DEMAND_TOT_productionEL												
Imbalance_Price												
Renewable												
Speed												
scaling_wind												
Renewable_capacity												
BOP												
BOP_001 to BOP_XXX				first guess	first guess							
BOP_SAMP_productionEL												
BOP_DYMO_productionEL												
BOP_capacity						pass	through					
BOP_TOT_productionEL												
SES												
SES_001 to SES_XXX				first guess	first guess							
SES_SAMP_productionEL												
SES_DYMO_productionEL												
SES_capacity						pass	through					
SES_TOT_productionEL												
ES												
ES_001 to ES_XXX				first guess	first guess		constants					
ES_SAMP_productionEL												
ES_DYMO_productionEL												
ES_capacity						pass	through					
ES_TOT_productionEL												
IP												
IP_001 to IP_XXX				first guess	first guess							
IP_SAMP_productionEL												
IP_DYMO_productionEL												
IP_capacity						pass	through					
IP_TOT_productionEL												
IP_DYMO_productionBY												
IP_TOT_productionBY												
IP_TOT_revenueBY												

6.1 Outer RAVEN sampler

First, the sampler in the outer RAVEN level samples the component capacities, i.e. the capacity of the gas turbine [W_e], the industrial process [W_e] and the battery [Wh]. Note that the capacities of the gas turbine and the industrial process are in Watts (electric) while the battery capacity is energy in Watt-hours. The sampler can be any sampler available in RAVEN, e.g. a grid or Monte Carlo sampler, but also an optimizer. The reactor capacity is fixed at 300 MWe. The model computes the cost of electricity (LCOE) that satisfies a given net demand profile (generated by the ARMA module, as discussed in section 6.2). Therefore, given the fixed reactor size, the mean demand is also a parameter that is sampled and can be optimized. In addition, the sampler provides additional parameters for sensitivity studies that are not part of the optimization space, such as:

- Wind penetration factor
- Price of hydrogen
- Other constants needed by the different sub modules in RAVEN, such as the thermal efficiency of the reactor, the nominal hydrogen production rate of the industrial process or the maximum battery charge and discharge capacities.

6.2 Synthetic Time History Generation (ARMA)

After RAVEN samples the component capacities, the first modules run are the two ARMA models for wind speed and demand followed by the ARMA postprocessor ARMA PostP, which generates the net demand profile (see [1]). This procedure is briefly explained in this section.

The synthetic time histories generation algorithm (ARMA) [39, 40] available in RAVEN is used to generate synthetic demand and wind speed histories. This includes multiple steps. First, historical measurement data must be collected. The data source and the frequency of the data are reported in Table 27. In the simulations reported in section 7, a total simulation time of one day, unless differently specified, is used with a time discretization of one hour. The 10-minute wind speed data noted in Table 27 was therefore averaged to obtain an hourly resolution.

Table 27. Data base information.

Data Type	Time Span	Resolution	Region	Source
Wind Speed	2004-2006	10 min (collapsed to 1 hour)	Site 3247	NREL [41]
Load	2011-2015	Hourly	West	ERCOT [42]

Second, from all the historical measurements data available, a ‘typical’ history must be created and used to train the ARMA in RAVEN. As mentioned, the simulations reported in section 7 use a total simulation time of one day. The ‘TypicalHistory’ postprocessor in RAVEN has been used to create the prototypical day. There are different approaches that can be taken to create this prototypical day. One is that the prototypical day is constructed from typical hours during the year. In this approach, the prototypical day will consist of a combination of average hours where outliers have been averaged out. The approach that is taken in the simulations presented here is that a typical day is chosen out of the year. To do so, the multiple years of training data are first cut into one-day slices. Out of these slices, the one that is closest to the average of all of them is then taken to be the prototypical day. The corresponding RAVEN input is shown in Figure 28.

```
<PostProcessor name = 'TypicalHistory' subType = 'InterfacedPostProcessor'>
    <method>TypicalHistoryFromHistorySet</method>
    <subseqLen>86400</subseqLen> <!-- 1 day-->
    <outputLen>86400</outputLen>  <!-- 1 day -->
    <pivotParameter>Time</pivotParameter>
</PostProcessor>
```

Figure 28. RAVEN TypicalHistory input used to create typical histories from measurement data to be used to train the ARMA rom.

In the next step, the ARMA is trained so that the later generated time series statistically conform to the actual measurement but possess different temporal profiles. In particular, a combined model with Fourier series and ARMA is utilized to de-trend the measurement and to characterize the autocorrelation of the residues. The synthetic data generation consists of generating independent white noise for each time step, utilizing the ARMA model to compute residues for each time step, and then adding the Fourier series representing seasonal trends. For both the wind speed and demand ARMA, the minimum Fourier frequency considered is 3 hours. Different ARMA expansion parameters have been chosen for the two ARMAs as shown in Figure 29. These parameters were determined in earlier studies [39, 40].

In order to construct the net demand (done in 'ARMA PostP'), a transfer function from wind speed to electricity generated by wind is needed. The corresponding function is shown in Figure 30. The exact equation and parameters used were reported in [40]. The value of 3.6 MW_e is chosen as the nominal value (corresponding to the Siemens SWT-3.6-120 onshore turbine) for computing the installed wind capacity. The user can input a desired wind penetration (in installed capacity) in the RAVEN input deck. This wind penetration is then used to scale the electricity generated from the wind. Finally, the electricity produced by the wind is subtracted from the demand to generate the net demand for that given wind penetration.

```
    <ROM name = 'demand_ARMA' subType = 'ARMA'>
        <Target>Demand,Time</Target>
        <Features>scaling_demand</Features>
        <pivotParameter>Time</pivotParameter>
        <Pmax>6</Pmax>
        <Pmin>6</Pmin>
        <Qmax>1</Qmax>
        <Qmin>1</Qmin>
        <outTruncation>positive</outTruncation>
        <Fourier>31536000,604800,86400,10800</Fourier>
        <FourierOrder>12, 2, 6, 3</FourierOrder>
    </ROM>

    <ROM name = 'wind_ARMA' subType = 'ARMA'>
        <Target>Speed,Time</Target>
        <Features>scaling_wind</Features>
        <pivotParameter>Time</pivotParameter>
        <Pmax>1</Pmax>
        <Pmin>1</Pmin>
        <Qmax>2</Qmax>
        <Qmin>2</Qmin>
        <outTruncation>positive</outTruncation>
        <Fourier>31536000,604800,86400,10800</Fourier>
        <FourierOrder>12, 2, 6, 3</FourierOrder>
    </ROM>
</Models>
```

Figure 29. RAVEN ARMA rom input used to train the wind speed and demand ARMAs.

For the data considered in this study (see Table 27), a detailed statistical analysis of synthetic versus measured data, i.e. a quality assessment of the synthetic data as well as a detailed study on the impact of wind penetration on the statistical properties of the demand have already been reported in [40] and are not repeated here.

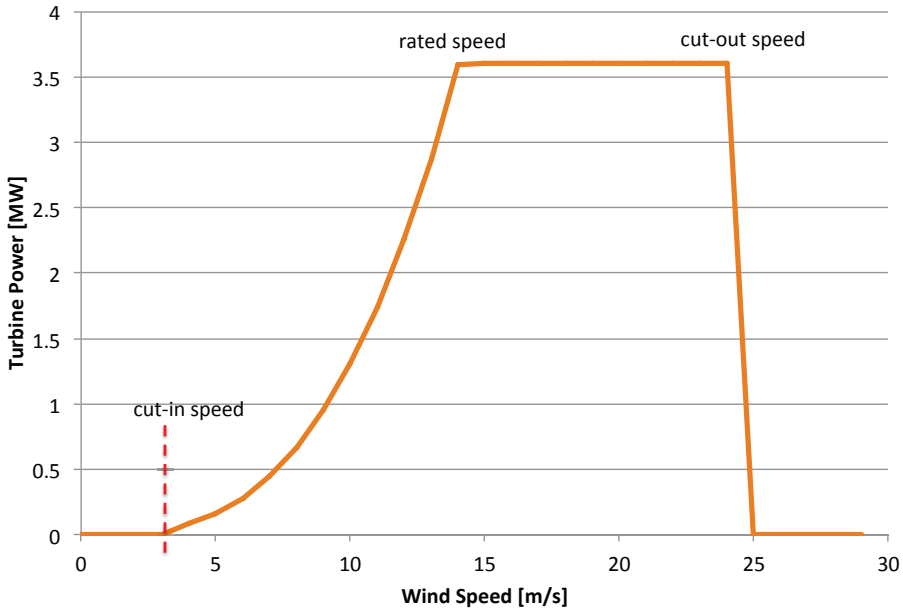

Figure 30. Wind speed to electricity conversion.

6.3 Initial Dispatch

As can be seen in Figure 27, once the net demand is generated, it is passed to the 'Initial Dispatch' module together with the sampled capacities. The 'Initial Dispatch' module creates from this information a dispatch schedule for each component of the N-R HES. The dispatch assumes no memory or predictive capability of the future. Therefore, for every time step, the dispatch is generated according to the different marginal costs of the different N-R HES components only. Since the marginal cost of the battery cannot be assessed given the assumption of no memory (the marginal cost of the battery would be the average price at which the battery purchased electricity to charge), it is used as a last resort in the dispatch.

The nuclear reactor is assumed to work at nominal capacity all the time. It can dispatch steam to the steam turbine to produce electricity or divert steam to the industrial process. Assuming that the industrial process has to get all the steam and electricity it needs from the reactor, i.e. it cannot buy electricity from the grid, there is a minimum of electricity that the reactor has to provide to the grid. This minimum electricity is the difference between the reactor capacity and the total (steam and electricity) energy needed by the industrial process when it runs at its maximum capacity (see Figure 31).

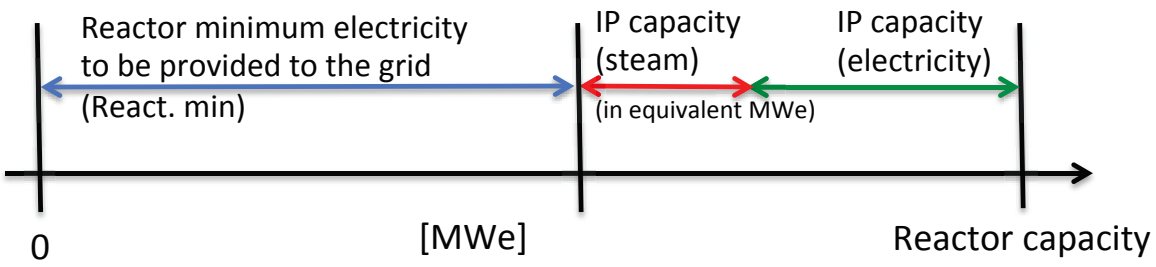

Figure 31. Minimum electricity the reactor has to provide to the grid as a function of the industrial process capacity.

The flow diagram of the 'Marginal cost & battery as last resort' dispatch is shown in Figure 32. As one can see, two main cases exist: either the net demand is less than the minimum electricity the reactor has to provide to the grid (React. min) or not.

The case in which the net demand is less than React. min is fairly simple: the reactor provides the minimum electricity it can while the industrial process runs at its capacity and the gas turbine does not produce anything. The battery tries to absorb the overproduction from the reactor as much as possible according to its maximum charge rate and current charge level. In this case, the amount of electricity provided to the grid might be larger than the net demand. In that case, the economic analysis of the system adds a penalty to be paid for the overproduction.

The second case, in which the net demand is greater than React. min, is more complex. In this case, first the reactor dispatches the fraction it has to (React. min) since the marginal cost for that portion of then reactor capacity is zero. For the rest of the demand (Net demand - React. min), the component with the least marginal cost (but bigger than zero) is dispatched first. The marginal cost of the gas turbine is assessed as the variable O&M cost including the fuel cost. The marginal cost of the portion of the reactor capacity that could be used to produce hydrogen in the industrial process (IP capacity steam and IP capacity electricity in Figure 31) is assessed as the opportunity cost of not producing hydrogen. If whoever dispatches first cannot cover the rest of the demand completely, the other will attempt to cover it. If both the reactor and the gas turbine at full capacity (and therefore the industrial process shut down) cannot cover the demand completely, the battery will try to compensate for the underproduction. The battery discharges depending on its maximum discharge rate and current charge level. As for the overproduction, if the battery cannot cover the whole demand, the economic analysis of the system adds a penalty to be paid for the missed demand.

It is worth mentioning that in all cases studied in this report, the battery is initially (at time zero) considered 50% full. This is considered to be the optimal state of the battery, since it can help absorb or provide electricity if needed. Furthermore, the maximum change and discharge rates are assumed to be 25% of the battery capacity per hour, as suggested by the data presented in Section 3.7.

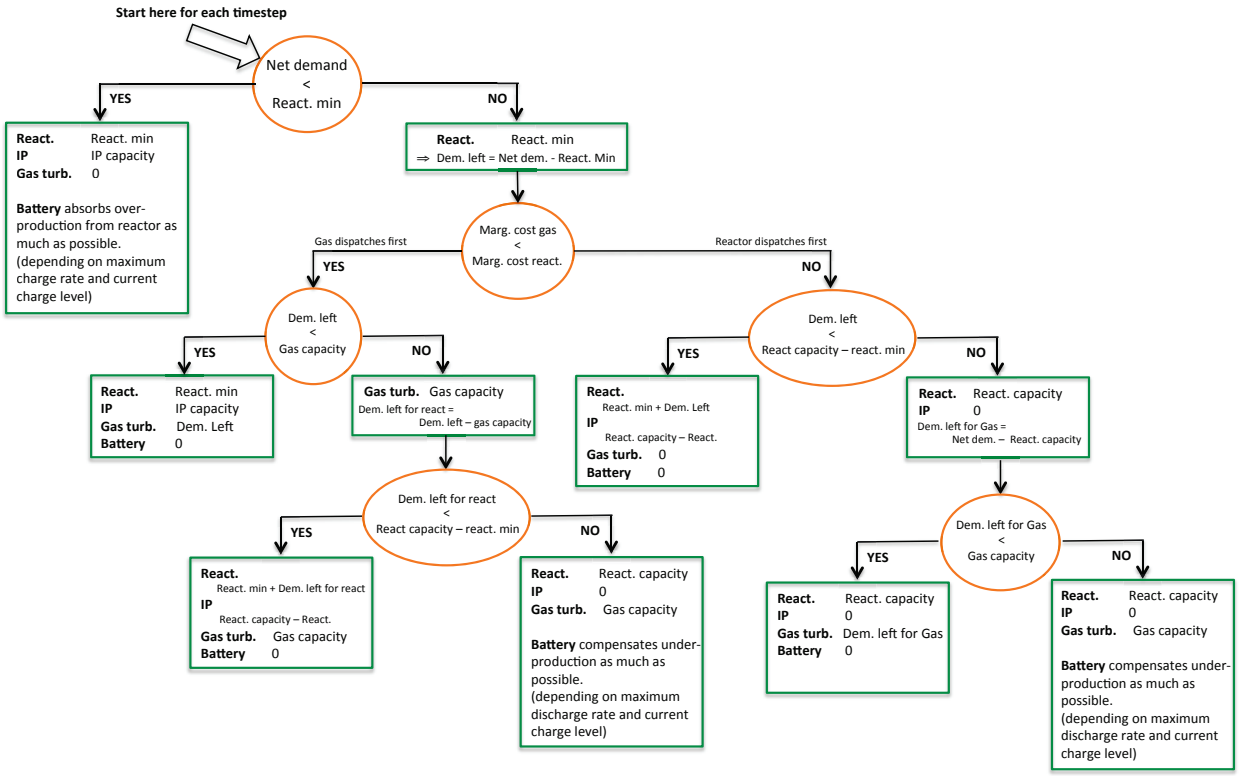

Figure 32. Dispatch according to marginal cost for the reactor and the gas turbine; the battery dispatches as last resort.

6.4 Inner RAVEN sampler

The dispatch schedules for all the components, i.e. the reactor, the gas turbine, the industrial process and the battery, are passed from the outer loop RAVEN into the inner loop sampler. The inner loop sampler can be used in two ways:

- **Pass through:** If it can be assumed that the 'initial dispatch' provided from the outer RAVEN level is already the optimal dispatch, the inner sampler can be set to 'pass through'. This assumption is true if there is no inertia in the system for the considered time step size, i.e. the system can reach any arbitrary state at the end of any time step. This is true if the system dynamics of Modelica are bypassed and the assumption is made that the system can reach any state proposed by the 'initial dispatch'. In the 'pass through' case, the initial dispatch is just passed to the subsequent Modelica and Cash Flow modules; once the LCOE is computed, the inner level is left and the LCOE passed back to the outer level.
- **Optimizer:** If the 'initial dispatch' is not the optimal one, e.g. the system might not be able reach the states described by the 'initial dispatch' due to dynamic restrictions (ramp rates), the inner sampler can further optimize the dispatch schedule for the components to lower the LCOE. In particular, the inner RAVEN optimizer can change the production level for each time step of the simulation of the gas turbine and industrial process. The Reactor levels depend on the industrial process level (reactor level = reactor capacity – industrial process level) and are computed in the 'ModelicaPreP' module. The battery charge and discharge levels for each

hour, i.e. the dispatch of the battery, are fixed to the values from the 'initial dispatch' and the optimizer cannot change them. Allowing the optimizer to vary the battery dispatch schedule will lead to the optimizer try to play the market by charging and discharging the battery to maximize profit. That is not the assumed purpose of the battery in this exercise; the battery is instead assumed to help the system satisfy the net demand. The inner optimizer will iterate attempts to minimize the LCOE by changing the dispatch until it reaches the maximum number of iterations or the convergence criteria are satisfied. In that case, the inner loop is left and the LCOE passed back to the outer loop.

6.5 Modelica Preprocessors

The modules 'Generate Vectors' and 'Modelica PreP' in Figure 27 collect all the needed data passed through from the upper level and inner level samplers and combine this into the format needed by the Modelica code interface. In particular, the component capacities and the net demand profile are collected. Furthermore, the 'Modelica PreP' also computes the reactor electricity production levels (dispatch) from the sampled industrial process dispatch.

If one decides to bypass Modelica by assuming that the system has no inertia or to run some quick scoping calculations without the expensive overhead of the full dynamics of Modelica, the 'Modelica PerP' module can pass its output directly to the 'Cash Flow PerP' and 'Cash Flow' modules to compute the LCOE assuming the dispatch from the inner sampler can be met by the system.

The Modelica Code Interface is not an external model, but is an integrated part of RAVEN. It collects all the needed data to run the Modelica simulation, runs it, and collects the required output values from the output files of Modelica.

6.6 Modelica

The Modelica model was described in section 2, and it is more specifically visualized in Figure 1 with reference to Table 1 and Table 2, noting that the IP selected for the current set of analyses is a hydrogen production plant. The reader will recall that the Modelica model includes detailed representations (sub-models) for each component and subsystem in the N-R HES in order to properly characterize the dynamic interactions of the integrated subsystems.

6.7 Cash Flow

The 'Cash Flow PreP' and 'Cash Flow' modules compute the LCOE from the component capacities and dispatch schedules. For all cases studied in this report, the same economic numbers and associated cash flows were used. This section describes them briefly. In the 'Cash Flow' module a set of cash flows can be defined, some of which depend on the LCOE and some not. The LCOE is then computed so that the resulting Net Present Value (NPV) of all cash flows is zero. How the Cash Flow module computes the LCOE is described in detail in [1]. The following sections describe the cash flows used and their associated economic values.

The tax and inflation rates are assumed to be 39.2% and 3%, respectively. For the computation of the NPV, a WACC of 5% (real) was assumed. The whole cash flow is computed in real terms (compared to nominal).

6.7.1 Nuclear Reactor

The cash flows and associated economic numbers used for the nuclear reactor are described in detail in [1]. This section provides a brief summary for the reader's convenience.

- The nuclear reactor has an assumed lifetime of 60 years.
- The capital expenditure (CAPEX) cash flow for the nuclear reactor is called BOP_CA (for balance of plant). An overnight capital cost is considered. The driver for this cash flow is the nuclear reactor capacity (electric). The reference for this cash flow is a 1100 MWe plant that has an overnight construction cost of $4.51 billion. The weighted average scaling exponent (X), i.e. the economy of scale factor for nuclear plants, is 0.64.
- A cost for the connection to the grid has been considered. Although this grid connection is for the whole NHES system, i.e. all components together, it is considered as part of the nuclear reactor, since its lifetime is also 60 years. The driver for this cash flow is the nuclear reactor capacity (electric). As suggested in the economics section above (see Section 3), the reference for this cash flow is a 400 MWe connection that has a construction cost of $13 million. The weighted average scaling exponent (*exp* in Eq. 1), i.e. the economy of scale factor for nuclear plants, is 1.26.
- O&M for the nuclear reactor is split into three cash flows, i.e. fixed O&M, variable O&M and fuel cost. Since the reactor is always operates at nominal capacity, the driver for all three cash flows is the nuclear reactor capacity (electric). The reference for the fixed O&M cash flow is a 1100 MWe plant that has yearly fixed O&M cost of $93.5 million. The weighted average scaling exponent (*exp*), i.e. the economy of scale factor for nuclear plants is 0.64. Taxes are applied to this cash flow. The variable O&M for the reactor is 0.5 $/MWh and the fuel cost 8.4$/MWh.
- A 15 year Modified Accelerated Cost Recovery System (MACRS) is applied to the nuclear reactor. (for details, see [1])
- A cash flow for a salvage value/decommissioning has been defined. This is for possible parametric analysis to study the influence of the salvage value on the profit. For the reference case, the salvage value has been set to zero.

6.7.2 Industrial process (hydrogen production plant)

The cash flows and associated economic numbers used for the industrial process are described in detail in [1]. This section provides a brief summary for the reader's convenience.

- The industrial process has an assumed lifetime of 40 years.
- The CAPEX cash flow for the industrial process is called IP_CA. An overnight capital cost is considered. The driver for this cash flow is the hydrogen plant capacity (in energy input, i.e. electric and thermal). The reference for this cash flow is a 231 MW hydrogen plant that has an overnight construction cost of $153 million. The weighted average scaling exponent, i.e. the economy of scale factor for hydrogen plants, is assumed to be 1.0.
- The revenue from hydrogen sales from the industrial plant for one year is calculated in the CashFlow preprocessor.
- O&M for the industrial H_2 plant is split into fixed O&M and variable O&M. The driver for the fixed O&M cash flow is the H_2 plant capacity. The reference for this cash flow is a 231 MW plant that has yearly fixed O&M cost of $3.5 million. The variable O&M cash flow driver is the amount of H_2 produced in a year. This value is calculated in the CashFlow preprocessor. The variable O&M for the H_2 plant is 0.048 $/kg of H_2.

- A 15-year MACRS is applied to the industrial plant. The MACRS applicable percentages for the 15-year property class is the same as for the nuclear reactor.
- A cash flow for a salvage value has been defined. This is for possible parametric analysis to study the influence of the salvage value on the profit. For the reference case, the salvage value has been set to zero.

6.7.3 Non-component Cash Flows

Two cash flows are defined that are not part of any component: revenue from electricity and over/under production penalty, as described in the following sections.

6.7.3.1 Revenue from electricity

The goal of the studies presented in this report is to find the minimum LCOE that satisfies an average net demand profile for a given mean demand and renewable penetration. If the net demand is not met by the components of the N-R HES, a penalty must be paid. This penalty can be seen as a price to be paid to somebody else to absorb the electricity or to compensate for the missed demand. In this manner, the net demand can always be considered met and the penalty becomes part of the cost to meet that demand. Therefore, the revenue from electricity does not depend on the individual production rates from the different components in the N-R HES, but only on the net demand.

The driver for this cash flow is the total energy in one year requested by the net demand profile, i.e. the integral of the net demand over one year. Since this driver is generated in the preprocessor for the one year of simulation time and assumed to be the same for every year of the lifetime of the plant, the **<alpha>** vector for this cash flow is 1.0 for all years except the year zero. **<reference>** and **<X>** are 1.0.

Multiplying this cash flow with the LCOE would generate the electricity revenue for the whole N-R HES. However, the LCOE that results in an NPV of zero for the whole system is what is desired. The mult_target attribute of the cash flow is therefore set to 'true'. This is the only cash flow that multiplies the LCOE.

6.7.3.2 Over/under production penalty

If the system produces more or less electricity than required by the net demand, a penalty must be paid. An exponential penalty with a cap is considered, as shown in Figure 33 [43]. The function shows the price for an over- or under-produced MWh for a given missed demand in MW. The penalty to be paid is then computed by multiplying the penalty function with the missed electricity for every hour, i.e.

$$penalty\ [\$] = \sum_{t=1}^{hours\ in\ simu} 1[h] * AbsMisDem_t[MW] * f_{penalty}(AbsMisDem_t\ [MW])[\$/MWh] \quad (7)$$

where $AbsMisDem_t = abs(NetDem_t - \sum_{components} production(component)_t)$

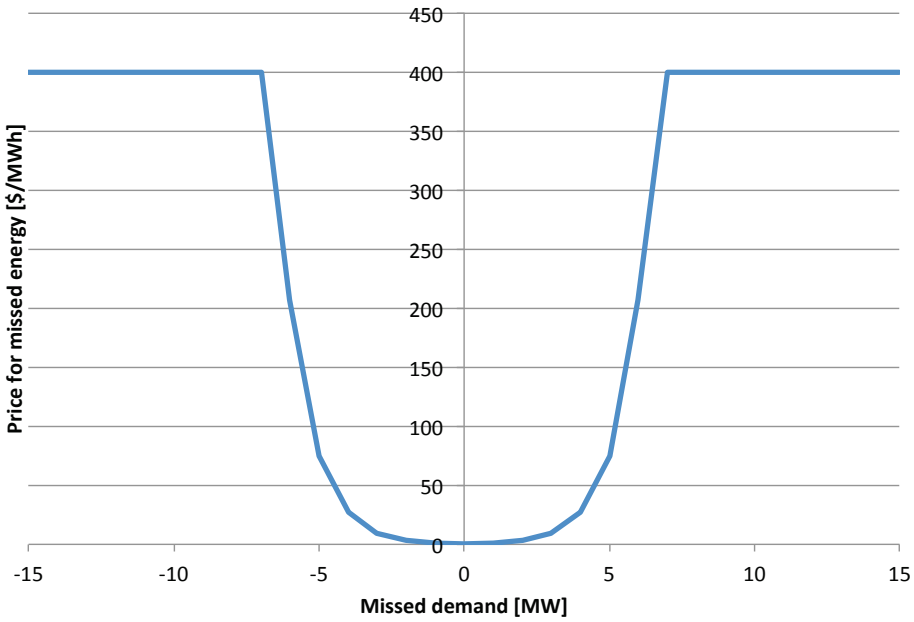

Figure 33. Penalty function for missed (net) demand.

The penalty is generated in the preprocessor for the simulation time and extrapolated for one year. It is assumed to be the same for every year of the lifetime of the project; the **<alpha>** vector for this cash flow is 1.0 for all years except the year zero. **<reference>** and **<X>** are 1.0.

6.7.4 Manifold

The manifold component is a stem manifold between the reactor and the gas turbine that can divert steam to the industrial process. It has an assumed lifetime of 60 years.

Capital expenditure (CAPEX)

The CAPEX cash flow for the manifold is called MANI_CA. An overnight capital cost is considered. It is assumed that the manifold size does not scale with the industrial process, but only with the reactor. Since the reactor has a constant size of 300 MWe in this exercise, the corresponding manifold cost is estimated at $55 million (see Section 3.6).

O&M

O&M for the manifold is assumed to be 1% of the CAPEX, as suggested in Section 3.6. The manifold therefore has a yearly O&M cost of $5.5 million. The O&M cash flow for the manifold is called MANI_OMperCap.

Depreciation

A 15-year MACRS is applied to the manifold. The MACRS applicable percentages for the 15-year property class are given in Table 28. The depreciation for year y is $DA_y(1+i)^{-y}ta$. This is constructed inside the CashFlow module as $DA_y(1+i)^{-y} - DA_y(1+i)^{-y}(1-ta$). For this reason, two cash flows have been defined. The first (MANI_DA1) computes the depreciation for year y. The second cash flow (MANI_DA2) uses the first one as a driver, has all **<alpha>** equal -1.0 (**<reference>** and **<X>** are 1.0) and includes the tax. In this way, MANI_DA1 + MANI_DA2 leads to the desired depreciation equation as shown above.

Salvage value

A cash flow for a salvage value has been defined. This is for possible parametric analysis to study the influence of the salvage value on the profit. For the reference case, the salvage value is set to zero.

Table 28. MACRS applicable percentage for 15-year property class.

Recovery Year	Percentage	Recovery Year	Percentage
1	5.00	9	5.91
2	9.50	10	5.90
3	8.55	11	5.91
4	7.70	12	5.90
5	6.93	13	5.91
6	6.23	14	5.90
7	5.90	15	5.91
8	5.90	16	2.95

6.7.5 Gas Turbine

The gas turbine has an assumed lifetime of 40 years.

Capital expenditure (CAPEX)

The CAPEX cash flow for the gas turbine is called SES_CA (SES for Secondary Energy System). An overnight capital cost is considered. The driver for this cash flow is the gas turbine capacity. The reference for this cash flow is a 400 MWe turbine that has an overnight construction cost of $386.8 million. The weighted average scaling exponent, i.e. the economy of scale factor for gas turbine, is 0.72.

O&M

O&M for the gas turbine is split into fixed O&M, variable O&M and fuel cost:

- **Fixed O&M**: The fixed O&M cash flow for the gas turbine is called SES_OmperCap. It is suggested not to scale this cash flow with the plant size. The fixed O&M is therefore a constant value of 7.5 million per year.
- **Variable O&M**: The variable O&M cash flow for the nuclear reactor is called SES_OmperProduction. The driver for this cash flow is the total produced electricity from the gas turbine. Same as the fixed O&M, it is suggested not to scale the variable O&M with the reactor size. The suggested constant value of 4.27 $/MWh is used.
- **Fuel cost**: The fuel cost cash flow for the gas turbine is called SES_OmperFuel. The driver for this cash flow is the total produced electricity from the gas turbine. The fuel cost for the gas turbine is 33.636$/MWh. This number has been computed by taking the average gas price [3.1 $/MMBTU] of the 'Henry' hub (see Section 3.4) and then multiplying by the gas turbine heat rate [10850 MMBTU/kWe].

Depreciation

A 15-year MACRS is applied to the nuclear reactor. The MACRS applicable percentages for the 15-year property class are the same than for the manifold given in Table 28. The depreciation for year y is $DA_y(1+i)^{-y}ta$. This is constructed inside the CashFlow module as $DA_y(1+i)^{-y} - DA_y(1+$

$i)^{-y}(1 - ta\ )$. For this reason, two cash flows have been defined. The first (SES_DA1) computes the depreciation for year y. The second cash flow (SES_DA2) uses the first one as a driver, has all **<alpha>** equal -1.0 (**<reference>** and **<X>** are 1.0) and includes the tax. In this case, SES_DA1 + SES_DA2 leads to the desired depreciation equation as shown above.

Salvage value

A cash flow for a salvage value has been defined. This is for possible parametric analysis to study the influence of the salvage value on the profit. For the reference case, the salvage value is set to zero.

6.7.6 Battery

The battery has an assumed lifetime of 10 years.

Capital expenditure (CAPEX)

The CAPEX cash flow for the manifold is called ES_CA (ES for Energy Storage). An overnight capital cost is considered. The driver for this cash flow is the battery capacity. As suggested in Section 3.7, a value of 350 $/kWh has been taken. The **<reference>** and **<X>** values are 1.0.

O&M

As for the manifold, O&M for the battery is assumed to be 1% of the CAPEX as suggested in Section 3.7. The battery has therefore a O&M cost of $3.5 $/kWh. The driver for this cash flow is the battery capacity and the O&M cash flow for the manifold is called ES_OMperCap.

Depreciation

A 7-year MACRS is applied to the battery. The MACRS applicable percentages for the 7-year property class are given in Table 29. The depreciation for year y is $DA_y(1 + i)^{-y}ta\ $. This is constructed inside the CashFlow module as $DA_y(1 + i)^{-y} - DA_y(1 + i)^{-y}(1 - ta\ )$. For this reason, two cash flows have been defined. The first (ES_DA1) computes the depreciation for year y. The second cash flow (ES_DA2) uses the first one as a driver, has all **<alpha>** equal -1.0 (**<reference>** and **<X>** are 1.0) and includes the tax. In this case, ES_DA1 + ES_DA2 leads to the desired depreciation equation as shown above.

Salvage value

A cash flow for a salvage value has been defined. This is for possible parametric analysis to study the influence of the salvage value on the profit. For the reference case, the salvage value is set to zero.

Table 29. MACRS applicable percentage for 7 year property class.

Recovery Year	Percentage
1	14.29
2	24.49
3	17.49
4	12.49
5	8.93
6	8.92
7	8.93
8	4.46

7. SIMULATION CASES

As mentioned at the beginning of Section 1, the goal of the simulation is to optimize the size and dispatch for a N-R HES. The size should be such that the LCOE is minimized under the constraint to satisfy a given net demand profile. As mentioned earlier, the demand profile the N-R HES system tries to satisfy is the net demand profile, i.e. the demand after subtracting all renewable contributions. As a consequence, the LCOE computed is the LCOE of the N-R HES to cover the net demand only and not a global LCOE to cover the gross demand including the renewable sources. The demand is treated stochastically, i.e. for a given size of the N-R HES, the LCOE is not a constant, because the net demand profile can change. By evaluating the same N-R HES size multiple times and averaging the obtained LCOEs during the optimization, one can find the optimum size for an average net demand profile.

In addition, to find the optimum size of the N-R HES system, one has to assure that for each size of the system, the different components are utilized in an optimal way, i.e. the dispatch of the different components is optimal with respect to minimizing the LCOE. For a copper plate model, it can be proven that the most efficient, economical way to supply the electricity demand is based on the least marginal cost (as occurs in current deregulated markets). In the presence of system inertia, this is no longer true, since prediction of future demand has to be taken in account to justify, for example, selling at a loss (e.g., nuclear power plant coping with a negative price of electricity without reducing plant power, since the decision takes into account the profit that will be made a few hours later when the electricity price will spike again). By introducing Modelica, inertia is introduced in the system and the marginal cost dispatch is no longer optimal. Nevertheless, to accelerate the optimization for the dispatch, the marginal cost can be used as an initial guess.

In order to understand the model and progress in a stepwise fashion to the final result of having the N-R HES size optimized with a full dispatch optimization using Modelica underneath, several simulations have been performed with increasing complexity. Hourly time resolution was selected for all simulation cases. All cases are summarized in Table 30.

Case 1: Exploration of the capacity space using a grid sampler (parametric sweep). No dispatch optimization, no Modelica, i.e. the LCOE is computed from the marginal dispatch that is 'passed through' the inner loop. This allows one to develop an understanding of the model. The simulation time is one day.

Case 2: Same as case 1, but with a simulation time of one year. This will allow conclusions on how accurate the LCOE is when extrapolating from one day simulation time to one year.

Case 3: This case is based on a linearization of the original Modelica model of the system. In addition, instead of looking for the lowest cost of electricity to meet net demand, the simulation aims to compute the dispatch, which maximizes the NPV of the N-R HES. This represents a more classical approach (maximization of the profit), which is commonly adopted in the industry due to a focus on the plant owner's point of view rather than societal perspectives. For the framework developed it is important to have the flexibility to identify a solution that achieves maximum profit.

Case 4: Same as case 1, but the dispatch is optimized. Cases without and including Modelica in the dispatch optimization are shown. This will allow quantification of the cost of the inertia by comparison with case 1.

Case 5: Only in the capacity (i.e. subsystem size) of the component is optimized to minimize LCOE. No dispatch optimization, no Modelica. This will demonstrate capacity optimization for an N-R HES configuration. The optimizer should find the same minimum identified in the parametric sweep (grid) found in case 1 and 4.

Case 6: Final case. Running the analysis to optimize capacity and dispatch with Modelica based representation of the physical system.

Table 30. Simulation cases.

Case	Component capacity seeking algorithm	Dispatch	Modelica	Simulation time
1	Grid	Marginal cost	No	1 day
2	Grid	Marginal cost	No	1 year
3	Fixed	Optimized	Surrogate	1 year
4	Grid	Optimized	Yes	1 day
5	Optimized	Marginal cost	no	1 day
6	Optimized	Optimized	Yes	1 day
6	Optimized	Optimized	Yes	1 week

7.1 Case 1

This case is characterized by use of a marginal cost dispatch (no Modelica model), and performing a parametric sweep over the following variables:

- Hydrogen price. Values are defined according to [1]: 0 $/kg, 1.75 $/kg, 3.5 $/kg
- Wind penetration[f] (nominal). Values: 0%, 50%, 100%, 200%
- Industrial process capacity (Hydrogen production). Values: 0 MW_e, 120 MW_e, 240 MW_e, 360 MW_e, 480 MW_e, 600 MW_e
- Demand. Values: 100 MW_e, 200 MW_e, 300 MW_e, 400 MW_e, 500 MW_e, 600 MW_e
- Battery (SE in figures). Values: 100 MWh, 200 MWh, 300 MWh, 400 MWh, 500 MWh, 600 MWh
- Gas turbine (SES in figures). Values: 100 MW_e, 200 MW_e, 300 MW_e, 400 MW_e, 500 MW_e, 600 MW_e
- Reactor power is kept constant at 300 MW_e

For each of the 15,552 points, 19 different random time series of wind and demand have been generated, for a total of 294,488 simulations, in attempt to represent an average behavior for each point of the grid. Results (Figure 34 to Figure 45) are reported in a 5-dimensional axis plot were the major axes are IP capacity and Demand, while the minor axes are gas turbine capacity and battery size. The color map is determined by the effective LCOE (i.e. the cost of meeting net demand). Low electricity cost (low LCOE) is highlighted by a deep blue color while deep red/brown indicate high LCOE. The objective is to identify the region of low electricity cost.

There are global trends that can be identified in the results. First, there is an almost constant valley of low prices along the lower left to the upper right diagonal. This diagonal expresses an optimal relationship between IP capacity and demand. The effect of the presence of the IP is even felt at 0 value for the hydrogen price. This is due to the volatility absorption of the IP. Even if there is no revenue from selling hydrogen there is a gain in not paying the penalty for over or under production which, in certain ranges of the ratio IP capacity to demand, exceeds the capital cost of constructing the plant. This behavior

[f] Wind penetration is defined as a percent of the mean electricity demand. Hence "100% penetration" equates to an installed wind capacity equal to mean demand (which ranges from 100 to 600 MW_e). Note, however, that the mean production from wind is about 27% of the installed capacity given the high variability of wind.

is, of course, something that should not be sought but is symptomatic of the importance of volatility absorption. It is worth noticing that the above-mentioned behavior is smoothed by the presence of the gas turbine as a secondary energy system (SES).

In general, increasing the capacity of the gas turbine always has a positive effect, most likely due to the ranges of sampled capacity. It could be possible that we never achieve a range in which the gas turbine capacity is in excess. It is also worth considering that capital costs for the turbine are relatively low and therefore this type of idle capacity is not too impactful on the total cost of electricity.

Global trends additionally indicate that the wind penetration has a strong effect in raising the LCOE due to the added volatility and consequently higher penalty costs (Figure 37, Figure 41, Figure 45) which is globally mitigated only by when we operate with high margins in the hydrogen production (Figure 45).It is interesting to note that the battery effect, while generally small, tends to be dependent on the other parameters. In fact, one may notice that in Figure 41 (second line from the top) in the center of the plot (with low gas turbine size) the increasing battery size correlates with an increasing LCOE while toward the right of the figure this behavior is inversed.

In conclusion, we have identified a very complex trend in the cost of meeting net demand (indicated as LCOE). The overall picture confirms that the hydrogen plant offers two values: for high hydrogen prices it is a strong source of cash; for low hydrogen prices mitigation of the volatility effect becomes more evident, but clearly it is always present.

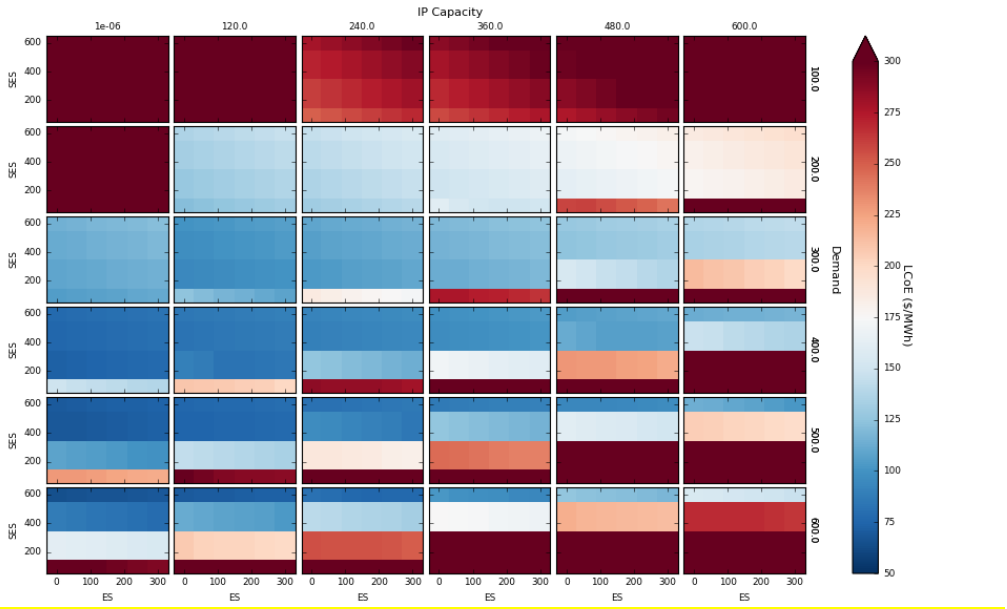

Figure 34. Effective LCOE color map for Case 1. H_2 price 0 $/kg, wind penetration 0%.

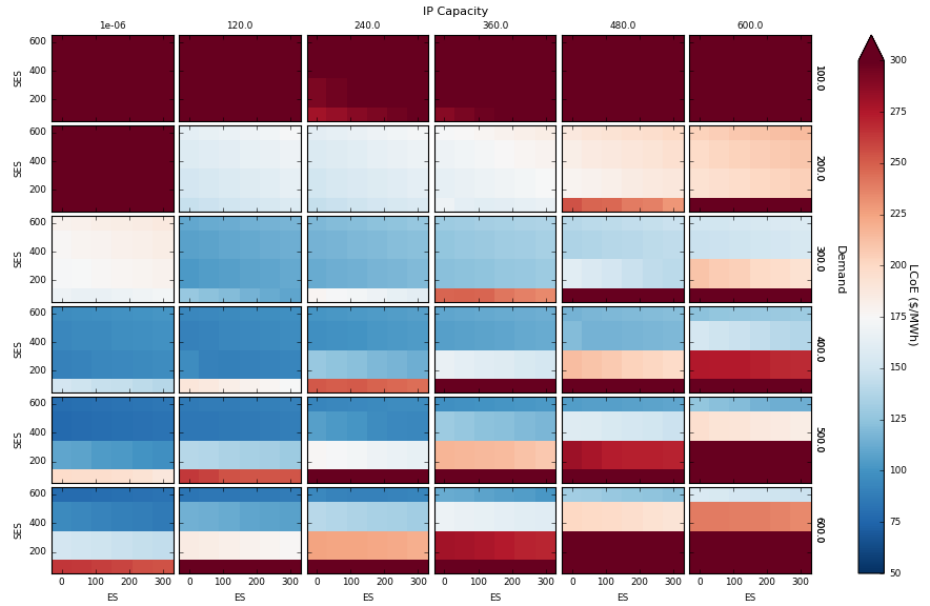

Figure 35. Effective LCOE color map for Case 1. H_2 price 0 $/kg, wind penetration 50%.

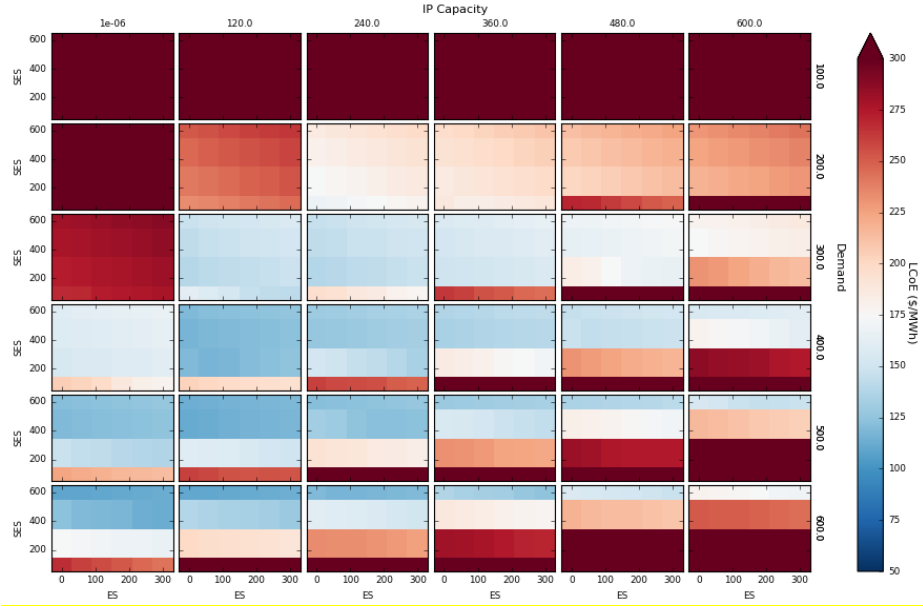

Figure 36. Effective LCOE color map for Case 1. H_2 price 0 $/kg, wind penetration 100%.

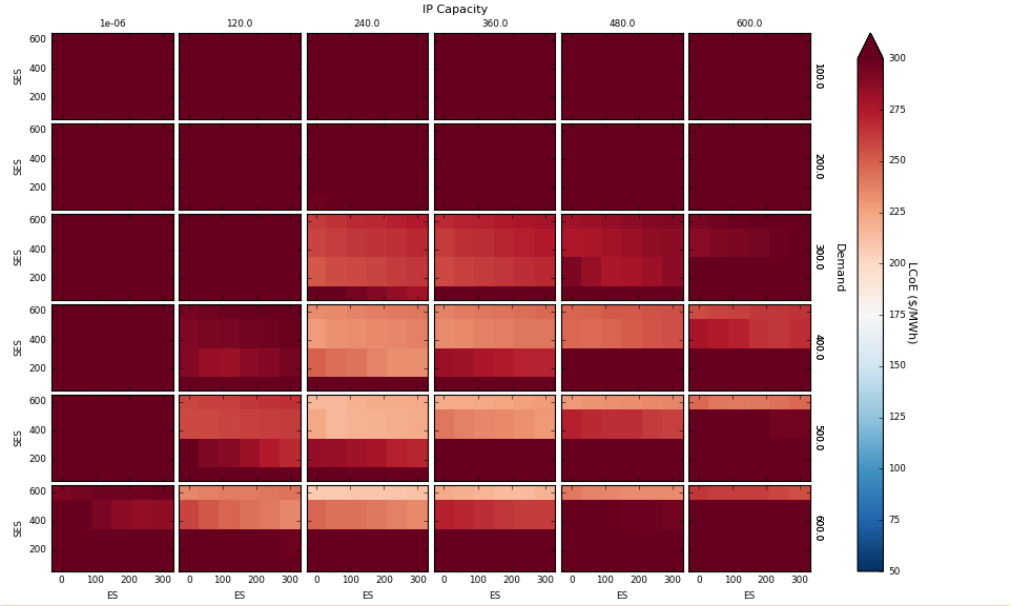

Figure 37. Effective LCOE color map for Case 1. H_2 price 0 $/kg, wind penetration 200%.

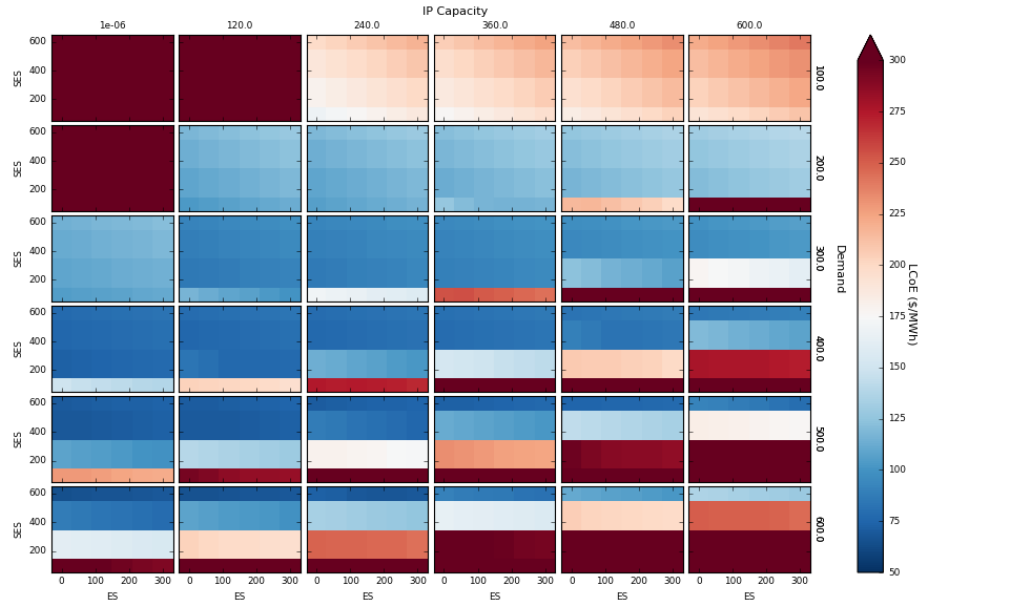

Figure 38. Effective LCOE color map for Case 1. H_2 price 1.75 $/kg, wind penetration 0%.

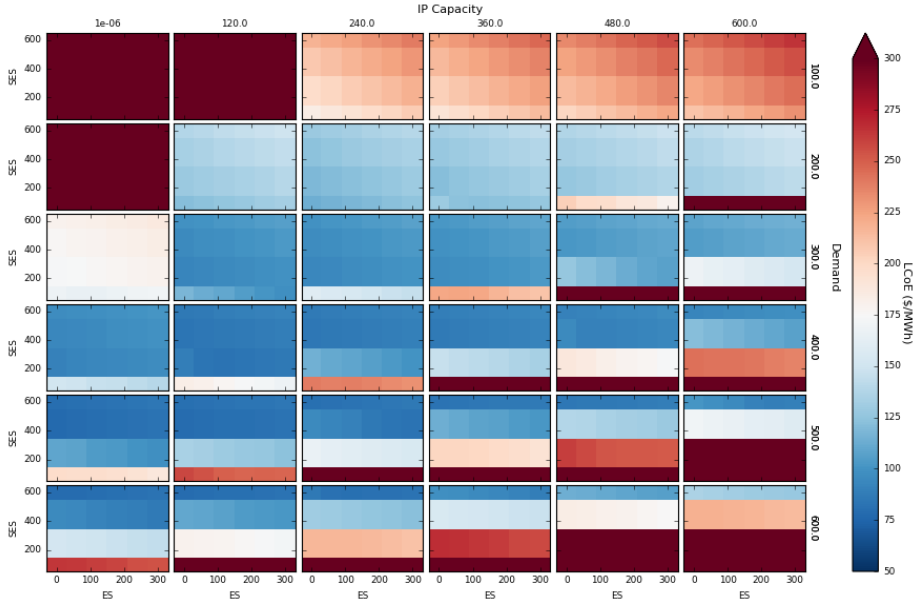

Figure 39. Effective LCOE color map for Case 1. H_2 price 1.75 $/kg, wind penetration 50%.

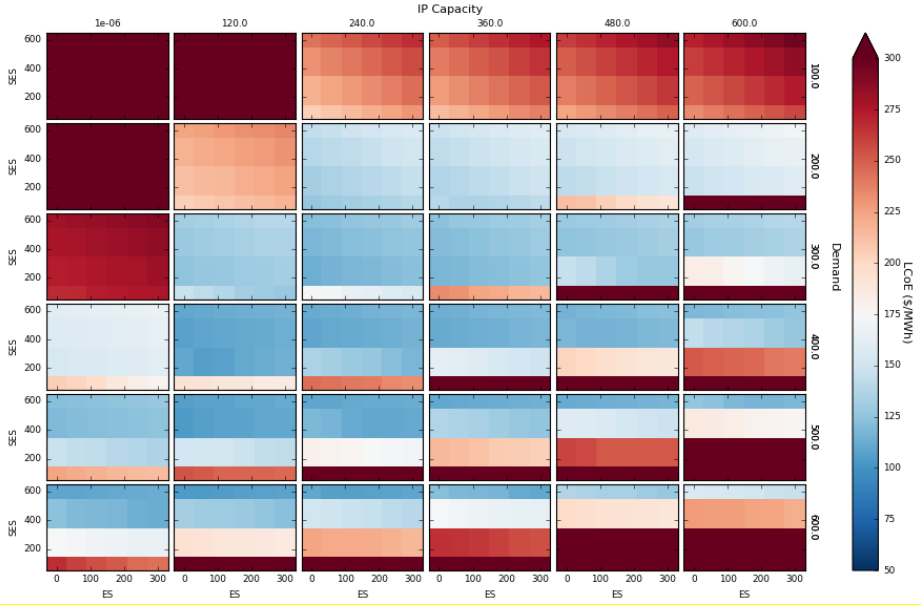

Figure 40. Effective LCOE color map for Case 1. H_2 price 1.75 $/Kg, wind penetration 100%.

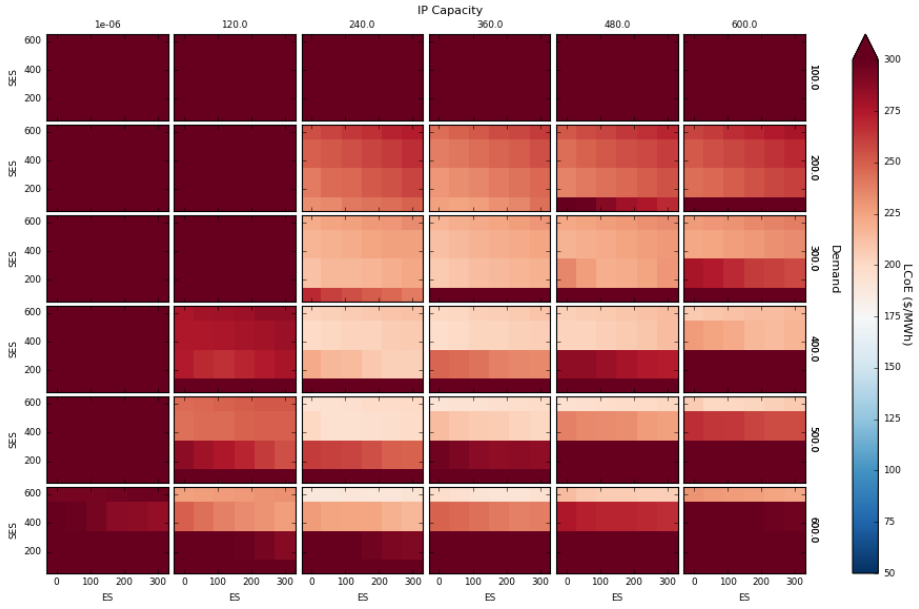

Figure 41. Effective LCOE color map for Case 1. H_2 price 1.75 $/kg, wind penetration 200%.

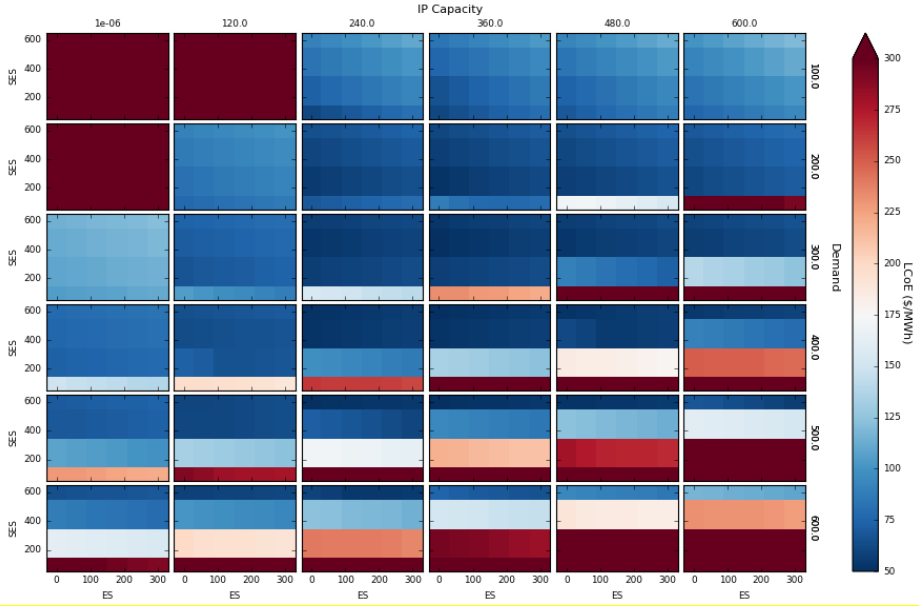

Figure 42. Effective LCOE color map for Case 1. H_2 price 3.5 $/kg, wind penetration 0%.

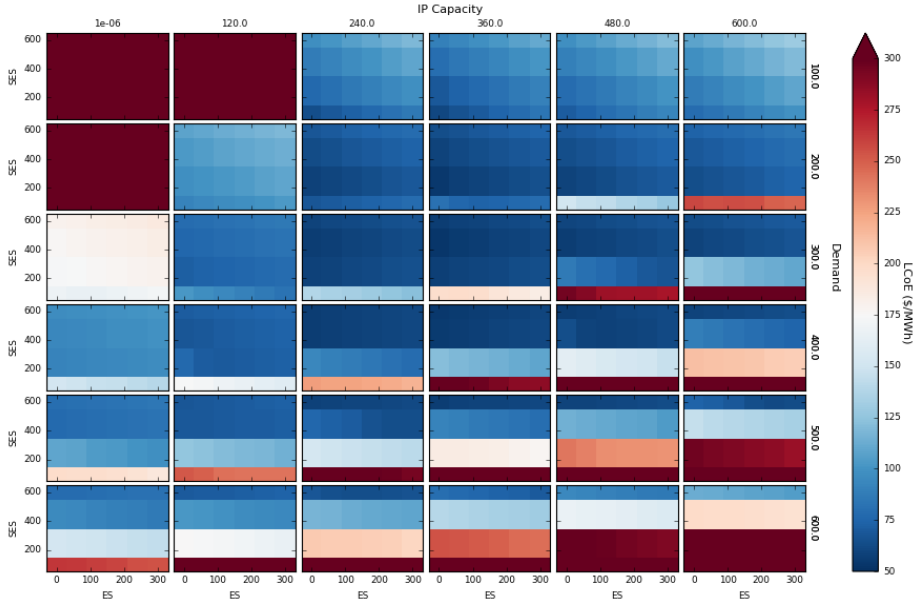

Figure 43. Effective LCOE color map for Case 1. H_2 price 3.5 $/kg, wind penetration 50%.

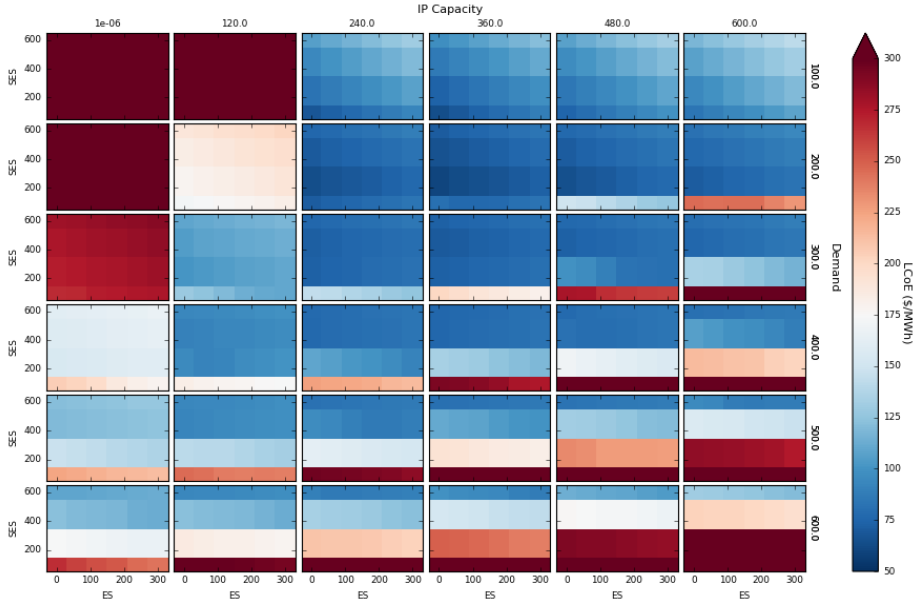

Figure 44. Effective LCOE color map for Case 1. H_2 price 3.5 $/kg, wind penetration 100%.

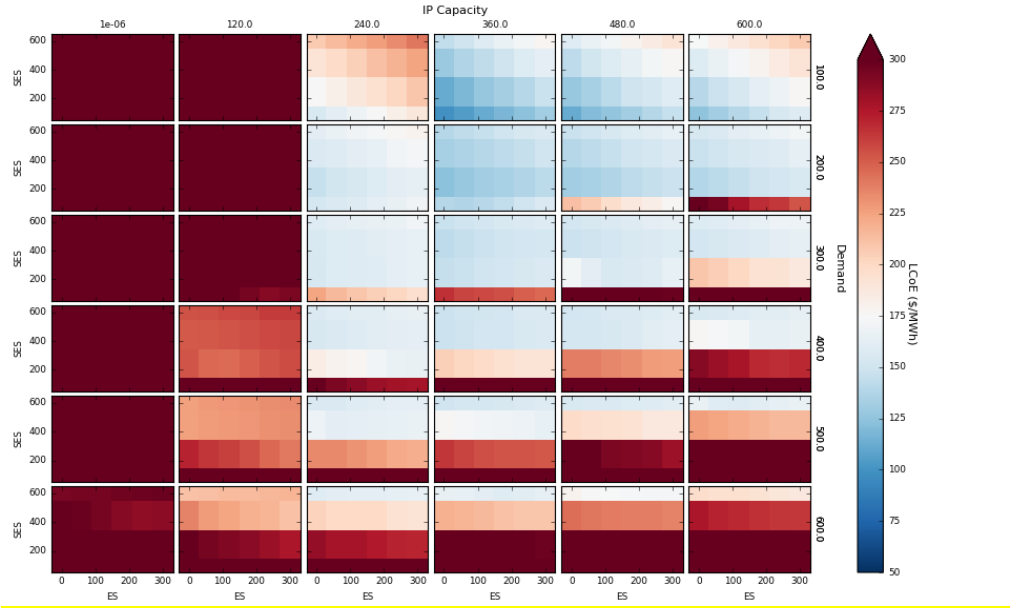

Figure 45. Effective LCOE color map for Case 1. H_2 price 3.5 $/kg, wind penetration 200%.

7.2 Case 2

This case is equivalent to Case 1 except for the duration of the time for which the system is monitored before extrapolating the cash flow to the lifetime of the project. The goal of this test is to assess the impact of a complete year history accounting for seasonal effects. In the same sense, it should be equivalent to a one day simulation using 365 de-noising repetitions, but, since the previous calculation was done with a surrogate trained on the average day (ARMA), differences should be contained. The simulation was run using 29 de-noising repetitions (multiple runs on the same point of the grid using different random time series of demand and wind speed, to converge toward a mean value).

For convenience, we also report the grid points sampled:

- Hydrogen price. Values are defined according to [1]: 0 $/kg, 1.75 $/kg, 3.5 $/kg
- Wind penetration (nominal). Values: 0%, 50%, 100%, 200%
- Industrial process capacity (Hydrogen production). Values: 0 MW_e, 120 MW_e, 240 MW_e, 360 MW_e, 480 MW_e, 600 MW_e
- Demand. Values: 100 MW_e, 200 MW_e, 300 MW_e, 400 MW_e, 500 MW_e, 600 MW_e
- Battery (SE in figures). Values: 100 MWh, 200 MWh, 300 MWh, 400 MWh, 500 MWh, 600 MWh
- Gas turbine (SES in figures). Values: 100 MW_e, 200 MW_e, 300 MW_e, 400 MW_e, 500 MW_e, 600 MW_e
- Reactor power is kept constant at 300 MW_e

The resulting plots are very similar to those presented in Case 1 for low wind penetration. The high wind penetration cases show much higher values of the cost of electricity than when one day is chosen as a representative time period. The reason for this behavior can be found in the much larger statistical sampling that the one year represents with respect the one day case. The representative day cannot be used to construct the representative cost of electricity when high volatility is present. Mathematically we could think of this as a non-negative, highly non-linear (exponential) penalty function that causes the cost

of outliers (days with very high or very low wind speed) that more significantly impacts the cost of electricity on outlier days than than on average days. For convenience, we report in the following results only the case with the highest price for hydrogen; the full series of results is reported in appendix B.

Overall the trends are the same as shown in Case 1. This is naturally due to the magnifying effect of the penalty function, but given that it is a monotonically growing function of the unmet demand its effect does not change the shape of the LCOE.

In conclusion, this simulation case highlights the importance of not relying on average values for net demand with large volatility. It is important to consider larger time intervals (e.g. one year) than one day to construct the characteristic response and evaluate the residual sigma after de-trending [40] to assess how much exposure the system might have to non-linear feedback. The larger the residual volatility after the de-trending of the signal (removal of the Fourier component), the more there will be a noisy component that might excite non-linear feedback of the system (in this case, this is embodied by the penalty function). These results illustrate the value of having the capability to perform this type of analysis as part of an integrated framework that allow analysis on large number of simulation from one day to one year. This in particular will allow study of various strategies to properly capture these effects and account for them even when shorter time intervals are considered due to the computational limitations associated with analysis of complex representations of the physical N-R HES.

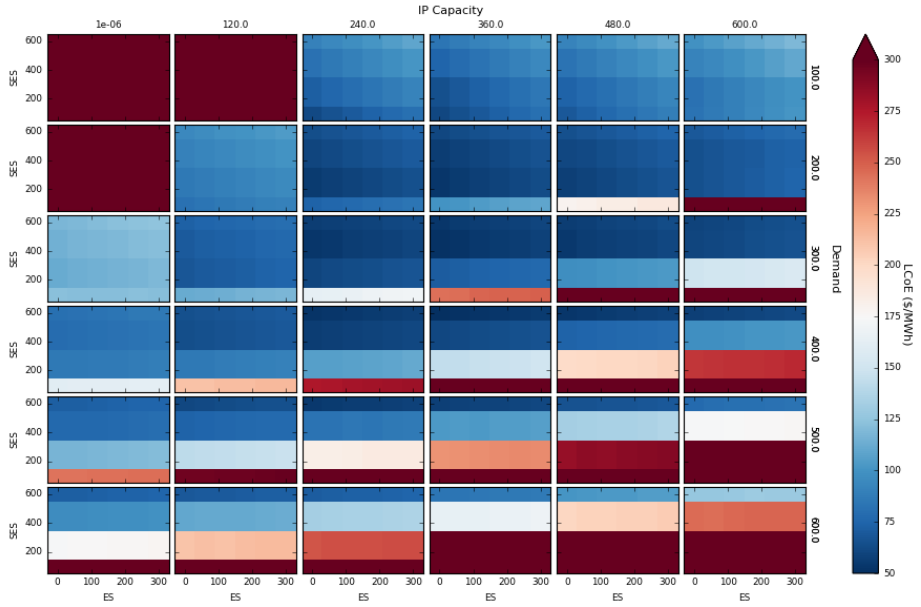

Figure 46: Effective LCOE color map for Case 2. H_2 price 3.5 $/kg, wind penetration 0%.

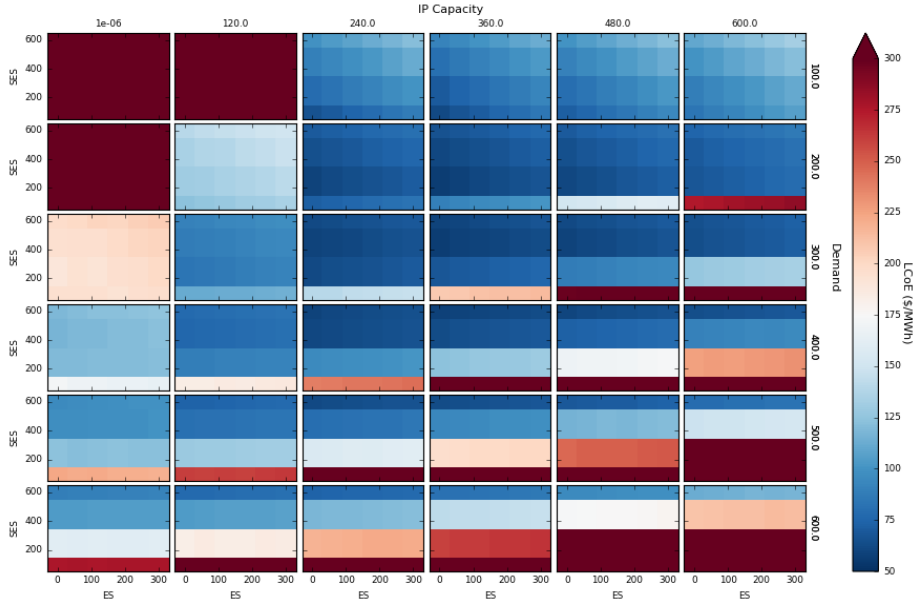

Figure 47: Effective LCOE color map for Case 2. H_2 price 3.5 $/kg, wind penetration 50%.

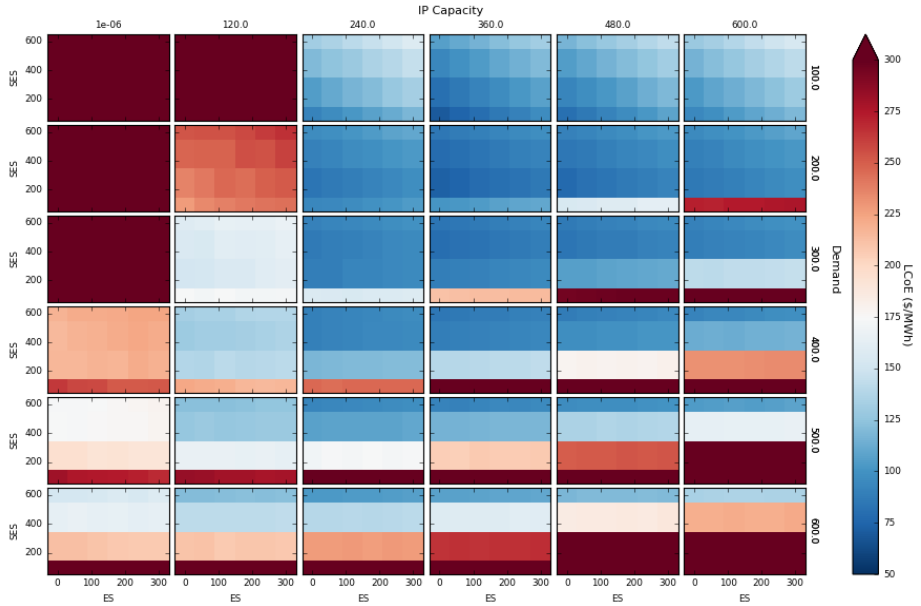

Figure 48: Effective LCOE color map for Case 2. H_2 price 3.5 $/kg, wind penetration 100%.

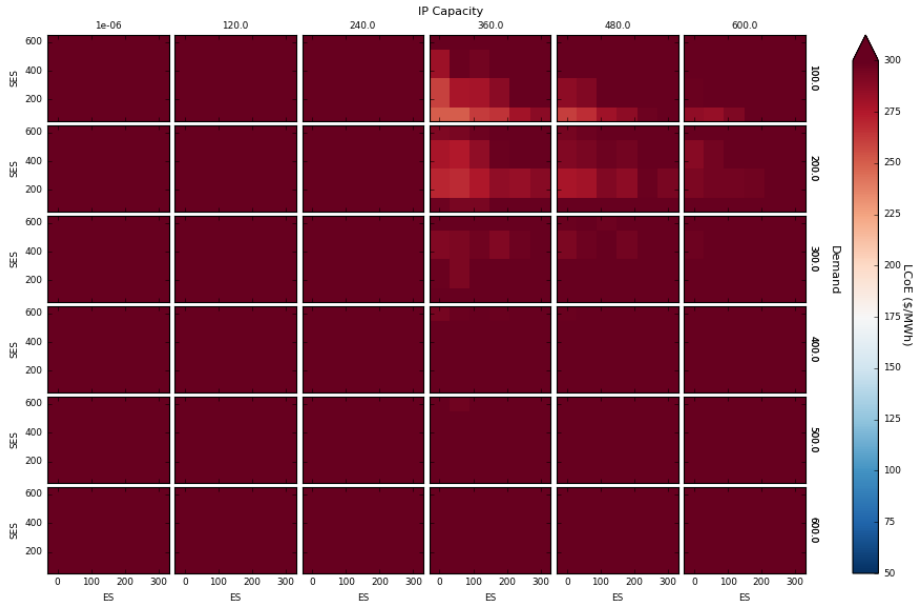

Figure 49: Effective LCOE color map for Case 2. H_2 price 3.5 $/kg, wind penetration 200%.

7.3 Case 3

This section is dedicated to an economic analysis of an N-R HES configuration to be located in the Midwest region in the United States. This region has abundant resources (wind, natural gas, coal, water); effective infrastructures (estuaries, rail, oil and gas pipelines, and electric transmission); a population who needs products and energy services; and both an agricultural and industrial history [44]. In particular, this region has a large demand for hydrogen (H_2), which serves as a precursor for ammonia/fertilizer production. Production of ammonia and fertilizer both require large amounts of electricity to separate nitrogen from air and large amounts of hydrogen to convert that nitrogen to ammonia. In addition, large iron deposits exist in Minnesota, Wisconsin, and Michigan and direct iron reduction, like ammonia production, has a large demand for hydrogen [44]. Thus, an HTSE system, which employs planar SOECs, was proposed as the industrial process to be integrated with a Light Water Reactor (LWR) for the Midwest region

This region has an excellent wind resource; hence, land-based wind power is proposed as the renewable resource. However, renewable energy generation was not considered in the case studies performed in this section. This case study is focused on profit analysis, such that accounting for the wind contribution is not necessary. In fact, if the system is considered a price taker (i.e. the presence of the system does not cause a feedback in the electricity prices), the coupling between the wind and the hybrid system is not active and, consequently, the economic performance of the system is simply the sum of the two performances considered separately. Given the interest of the project in analyzing the performance of the coupled nuclear and industrial system it is outside of the scope to analyze just the economic performance of the wind electricity supplier.

This case is also intended for analysis of the feasibility of replacing the full Modelica representation of each subsystem and the overall integrated system with linear surrogates.

As already mentioned, instead of looking for the lowest cost of electricity to meet net demand in Case 3, the simulation aims to compute the dispatch that maximizes the NPV of the N-R HES. This represents a more classical approach (maximization of the profit), which is commonly adopted in the industry due to

a focus on the plant owner's point of view rather than societal perspectives. For the framework developed it is important to maintain flexibility in analysis approaches.

7.3.1 Surrogate Model Development

A high-fidelity model (i.e., Modelica model) may provide an accurate reflection of reality but requires significant computational power; thus, approximation model (i.e., surrogate model) are constructed to mimic the behavior of the high-fidelity models as closely as possible while being maintaining computational efficiency.

First, the *linear regressor*[g] is proposed to characterize the relationship shown in Eq. (7) between the power consumption in the IP P_{HTSE} (i.e., a decision variable in the profit optimization problem) and the hydrogen production w_{H2} at the steady state:

$$w_{H2} = k_0 + k_1 P_{HTSE} \tag{7}$$

where k_1 and k_0 are the model-fitting parameters. Likewise, Eq. (8) is proposed to characterize the relationship between P_{IP} and the electricity generation to the grid P_{EG}:

$$P_{EG} = k_0 + k_1 P_{HTSE} \tag{8}$$

Several simulations were conducted to estimate the model estimates by linear regression. Regression results for Eq. (7) and Eq. (8) are plotted in Figure 50 and Figure 51, respectively. The estimated model-fitting parameters and goodness of model fits (R^2 values) are listed in Table 31. The quality of the surrogate model fits compared to the Modelica model outputs indicate excellent model fits.

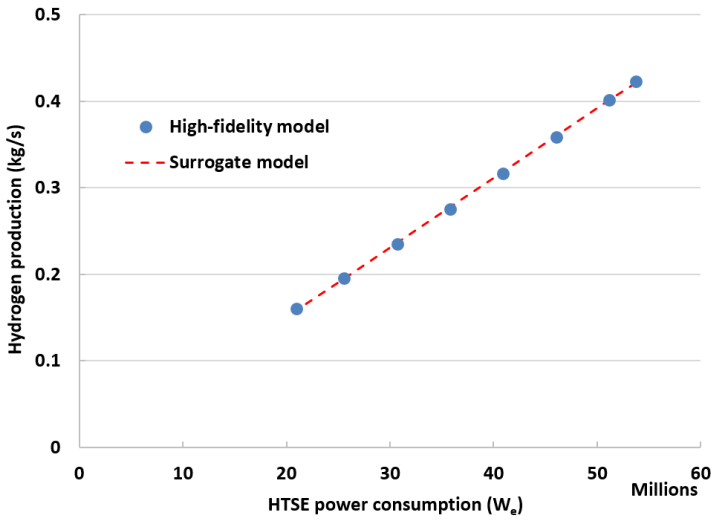

Figure 50. Hydrogen production versus power consumption in the HTSE plant – Eq. (7).

[g] The *linear regressor* is a least-squares fitting of the response of the system for a linear representation (linear regression).

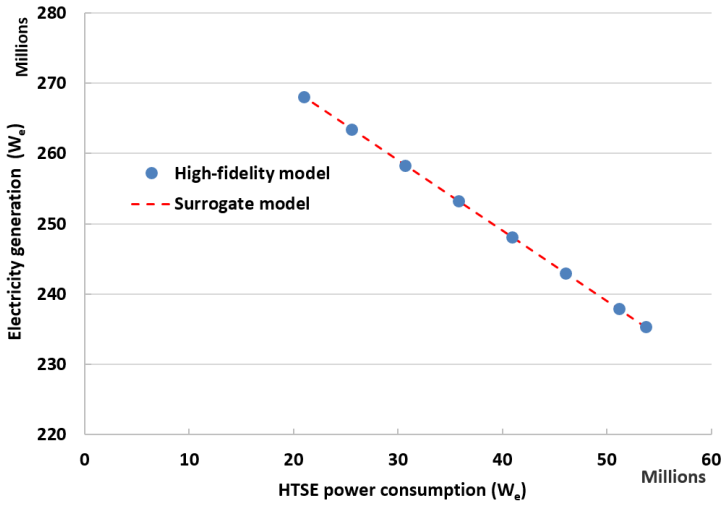

Figure 51. Electricity generation to the grid versus power consumption in the HTSE plant – Eq. (8).

Table 31. Model parameter estimates for Eq. 7 and Eq. 8.

Symbol	Description	Value	
		Eq. 7	Eq. 8
k_0	Model parameter	-9.81885×10^{-3} (kg/s)	2.89020×10^{8} (W_e)
k_1	Model parameter	8.02320×10^{-9} (kg/s·W_e)	-9.99929×10^{-1} (–)
R^2	Goodness of fit	0.9997 (–)	1.000 (–)

Next, the dynamic relationship of two process variables (i.e., an input variable P_{HTSE} and an output variable w_{H2} shown in Eq. (7)) is approximated by a model based on Laplace transforms (also referred to as a *transfer function model*):

$$G(s)=\frac{\tau_3 s+1}{(\tau_1 s+1)(\tau_2 s+1)} \tag{9}$$

where $G(s)$ is a transfer function in the Laplace domain, s is a complex independent variable, and τ_1, τ_2 and τ_3 are time constants (i.e., model parameters to be estimated by regression). Note that the transfer function models are only directly applicable to processes that exhibit linear dynamic behavior, such as a process that as can be modeled by a linear ordinary differential equation. However, the relationship between the input and output variables described in Eq. (7) is linear and thus does not depend on the operating regime. For such a condition, Eq. (9) is sufficiently accurate. The response of this system in the time domain $y(t)$ to a step change in input is

$$y(t)=KM\left(1+\frac{\tau_3-\tau_1}{\tau_1-\tau_2}e^{-t/\tau_1}+\frac{\tau_3-\tau_2}{\tau_2-\tau_1}e^{-t/\tau_2}\right) \tag{10}$$

where t is a time variable, K is a process gain and M is an input step change.

Figure 52 shows a plot of the output response (w_{H2}) with the high-fidelity model (Modelica model) to a step change in input (P_{HTSE}), and the corresponding surrogate model fitted by nonlinear regression. The transient was initiated at 100 s following a 10% step decrease in P_{HTSE} from an initial load level of 51.1 MW_e. The estimated model-fitting parameters and R^2 value are listed in Table 32. Results indicate that the surrogate model is exact in terms of the goodness of fit. In fact, the surrogate model is indistinguishable from the response with the Modelica model. Generally, this is not always true, but is instead a function of the system to be simulated and the amplitude and steepness of the transient. Applicability of surrogate models generally cannot be determined *a priori* but only *a posterior* by comparison with the high-fidelity model (Modelica). Nonetheless, this analysis confirms the opportunity to construct surrogate models which may be sufficiently accurate for economic analysis.

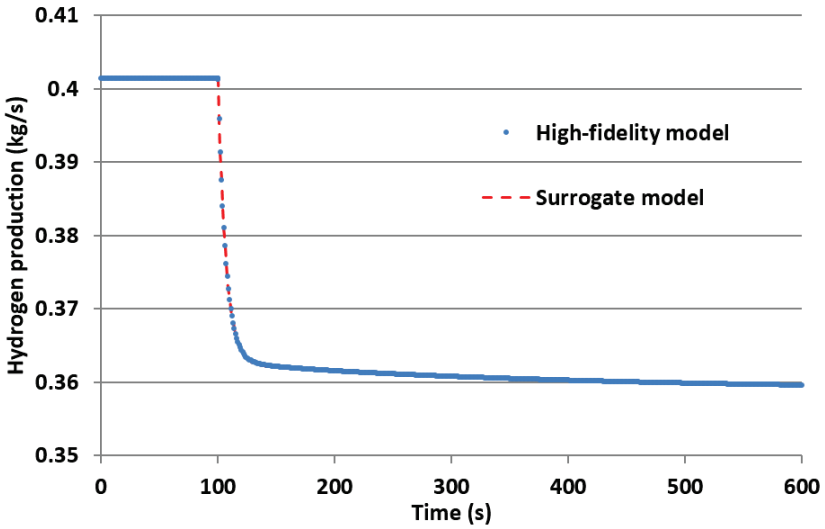

Figure 52. Step response of the process described in Eq. 10.

Table 32. Model parameter estimates for Eq. (10)

Symbol	Description	Unit	Value
K	Process gain	$kg/s \cdot W_e$	8.353×10^{-9}
M	Step input	W_e	-5.11×10^{6}
τ_1	Time constant	s	353.8
τ_2	Time constant	s	6.806
τ_3	Time constant	s	323.6
R^2	Goodness of fit	–	0.9994

7.3.2 Case Studies

Four different sub-case (case 3-1, case 3-2, case 3-3 case 3-4) studies were conducted to analyze the economic performance of the proposed N-R HES configuration to be located in the Midwest region. In all cases, the following assumptions were made for the economic analysis:

- The size (capacity) of the subsystem(s) is fixed.

- The nuclear plant (LWR) operates at its full capacity (300 MW_e) at all times.
- The initial load level for the industrial plant (HTSE plant) corresponds to its nominal operation condition (i.e., 51.1 MW_e).
- The simulation time is one day. All subsequent days in the lifetime of the nuclear plant are assumed to have the same revenue from electricity sales. Similarly, all subsequent days in the lifetime of the industrial plant are assumed to have the same revenue from hydrogen sales.
- The 10-minute time resolution is selected, leading to 144 decision variables for running a full-day simulation.
- The nuclear plant has an assumed lifetime of 60 years.
- The industrial plant and the corresponding EM unit have a lifetime of 40 years.
- The NPV is calculated at the fixed hydrogen price of $3/kg.
- In each case, the obtained NPVs from ten thousand simulation runs (each of which uses a different temporal profile for the electricity price) are averaged to find a statistically meaningful measure of the NPV.
- The system is assumed to be a price taker; hence, it does not influence the market (neither electrical nor hydrogen).
- The participation in the ancillary service market is not considered.
- Neither an LWR nor an HTSE plant produces carbon dioxide.

Case 3-1 is the LWR/HTSE integration case, which seeks to find the optimal dispatch schedule for the system depending on market conditions (i.e., electricity prices). Case 3-2 considers only an LWR, which is operated at full production mode and sells all of its electricity generated to the electric grid; no optimization is involved in calculating the NPV. Both cases assume the construction of a "new" nuclear plant. The analysis period for the economic evaluation for these cases is 120 years – i.e., the least common multiple of all subsystem lifetimes involved, guaranteeing that the NPV is computed for a time span that all subsystems reach their end of life in the same year.

Case 3-3 is the LWR/HTSE integration case as in Case 1 but considers the integration of the IP with an "existing" nuclear plant; this case assumes that the capital cost for a nuclear reactor has been paid off. Case 3-4 is the same as Case 3 but considers an "existing" LWR. In these cases, the nuclear plant has been operating for 20 years and has a remaining lifetime of 40 years, which coincides with a lifetime of a new HTSE plant. Therefore, the analysis period for the economic evaluation for Cases 3 and 4 is 40 years.

Table 33 summarizes all cases considered.

Table 33. Variations of Case 3 simulation cases.

Case No.	Nuclear plant (LWR)		Industrial plant (HTSE)		Simulation output interval (s)	Simulation time (s)
	New	Lifetime (years)	Integration	Lifetime (years)		
3-1	Yes	60	Yes	40	600	86400
3-2	Yes	60	No	N/A	600	86400
3-3	No	40	No	40	600	86400
3-4	No	40	Yes	N/A	600	86400

7.3.3 New Nuclear Plant – Case 3-1 (LWR/HTSE) vs. Case 3-2 (LWR)

The resulting optimal electricity generation to the grid and electricity consumption in the IP for Case 1 are shown in Figure 53(a) and Figure 53(b), respectively, corresponding to the maximization of NPV and assuming the fixed hydrogen price of $3/kg. Figure 53(c) shows the corresponding electricity price data (i.e., one ARMA realization out of ten thousand ARMA realizations).

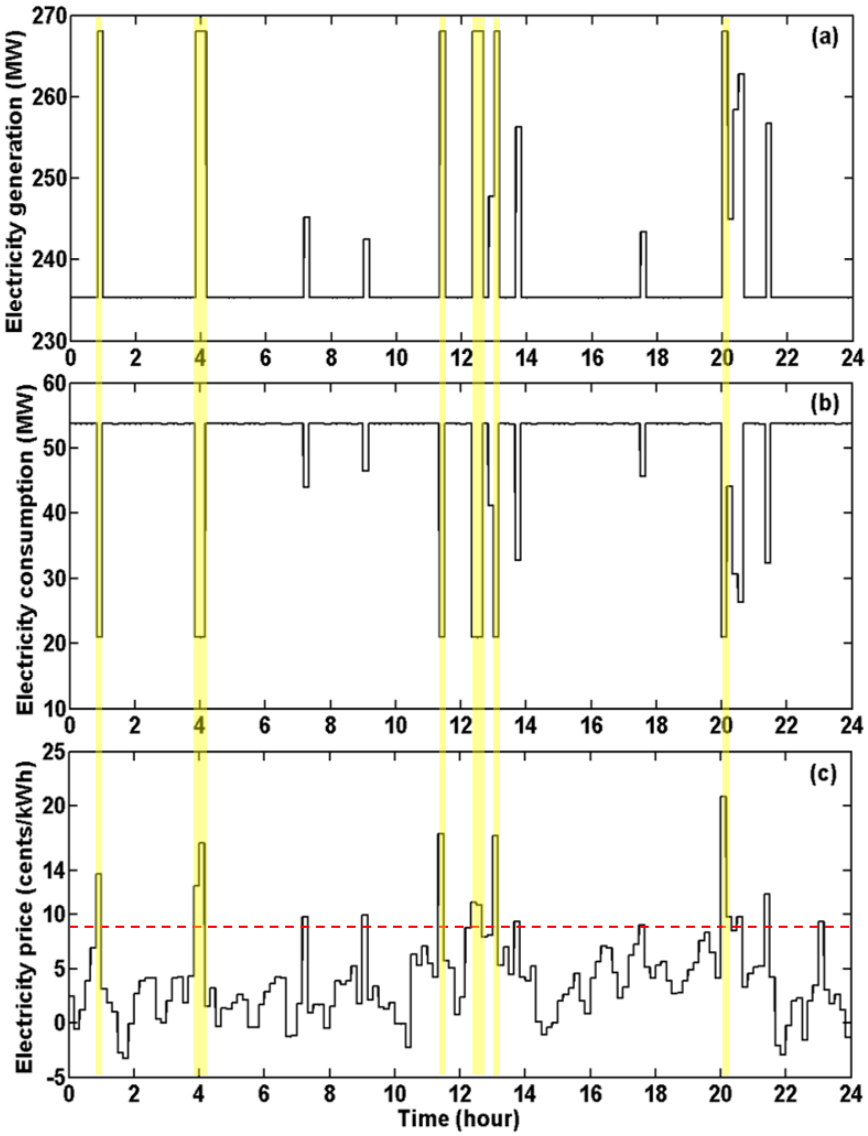

Figure 53. Case 3-1 results: (a) optimal electricity generation to the grid, (b) optimal electricity consumption in the IP, and (c) electricity price. The red dashed line indicates an electricity price threshold below which it was preferred to sell electricity to the grid rather than producing hydrogen. The yellow rectangles indicate the hours during which electricity generation was maximized, taking advantage of the expensive electricity prices.

As can be seen in Figure 52, when the electricity prices are below about 9 cents/kWh (i.e., price range below the red dashed line shown in Figure 53(c)), the operations optimizer (supporting economic optimization for operations) diverted electricity to maximize the hydrogen production instead of selling it

to the electric grid. For the electricity prices above 9 cents/kWh, the economics favored electricity production over hydrogen production. In particular in hours 1, 4, 11.5, 13, and 20 (highlighted with yellow rectangles in Figure 53), the plant maximized its power generation and sold as much power as possible to the grid due to the high electricity prices (<11 cents/kWh). In other words, under the current formulation, the electrical contribution delivered to the electric grid by the nuclear plant increased only as a response to high electricity prices.

Table 34 summarizes the mean NPV (with a 95% confidence interval) and STD for 120 years of operation under the NPV-optimization mode (Case 3-1) and those at constant electricity generation mode (Case 3-2), assuming new nuclear plants are built.

Table 34. Economic values for Cases 3-1 and 3-2.

Description	Unit	Values		
		Case 3-1 (LWR/HTSE)	Case 3-2 (LWR)	Gain
Mean NPV	$ million [%]	-1079 ± 1.45[a]	-1220 ± 1.72[a]	141 [11.6]
STD	$ million [%]	73.0 [6.76]	87.9 [7.20]	N/A

[a] 95% confidence interval, using a normal distribution.

The results indicate that neither case is economically viable as negative NPVs are expected (i.e., $-1.079 billion and $-1.22 billion for Cases 3-1 and 3-2, respectively); it is not economical to build a new nuclear plant (rated at 300 MW_e) in the Midwest region at the current market conditions, regardless of whether or not the IP is integrated with a nuclear plant. However, an expected NPV gain by integrating an LWR with the HTSE process is about $141 million (a 11.6% gain). Thus, these results suggest that the LWR/HTSE integration case has the flexibility to be controlled for economic optimization, while supporting hydrogen production. This attractive performance is further magnified when considering that oxygen could possibly be recovered from the HTSE plant by condensation and sold in the market if steam is used as a sweep gas. Furthermore, participation in the ancillary service market, which is not considered in this report, could increase the NPV of the system. For example, considering that the maximum and minimum power consumption of the proposed IP is 53.7 MW_e and 21 MW_e, respectively, for most of the time the LWR/HTSE integration case has a capacity of 32.7 MW_e to participate in operating reserve services (such as regulating, ramping, and load following), bringing revenue from providing operating reserve service on top of the sale of electricity.

7.3.4 Existing Nuclear Plant – Case 3-3 (LWR/HTSE) vs. Case 3-4 (LWR)

The optimal dispatch schedule for Case 3-3 is essentially the same as in Case 3-1 (see Figure 53), therefore it is not presented here. This is expected since the optimal operational strategy only depends on variable costs (such as commodity price and operational costs), provided that the fixed cost (capital cost) is the same for both cases.

Table 35 presents the resulting economics for 40 years of operation under the NPV-optimization mode (Case 3-3) and those at constant electricity generation mode (Case 3-4), assuming that the IP is to be integrated with an existing nuclear plant.

The results for Case 3-4 show that the expected NPV is $340 million when considering only the nuclear plant, which sold all of its electricity generated to the MISO grid regardless of the electricity price. When considering the integration of the IP with an existing nuclear plant, an expected NPV was 464 million (a $124 million [14.2%] gain); hence, such an option appears to be economically viable at the current market conditions.

Table 35. Economic values for Cases 3-3 and 3-4.

Description	Unit	Values		
		Case 3-3 (LWR/HTSE)	Case 3-4 (LWR)	Gain
Mean NPV	$ million [%]	464 ± 1.45[a]	340 ± 1.48[a]	124 [14.2]
STD	$ million [%]	66.0 [6.76]	75.6 [22.2]	N/A

[a] 95% confidence interval, using a normal distribution.

7.4 Case 4

The goal of Case 4 is to investigate optimization of the dispatch for fixed capacities. A parametric grid sweep has been performed for the capacity combinations of the components and the dispatch has been optimized of each of them. As mentioned, we assume that without inertia, i.e. without Modelica, the marginal dispatch is the optimum. Before introducing Modelica into the calculation, this assumption was checked by optimizing the dispatch of selected capacity points without Modelica. Table 36 shows three different dispatch optimization runs without Modelica. All runs shown are for a gas turbine capacity of 100 MWe, a battery capacity of 60 MWh, an IP capacity of 300 MWe, a mean demand of 400 MWe and a wind penetration of 100%. As one can see from the first run, if the penalty is more expensive than the marginal cost of the gas turbine (SES) and balance of plant (BOP), the LCOE from the marginal cost cannot be optimized further and the optimizer finds the same price. In the second run, the maximum penalty cost for a MWh missed demand is reduced to $50, which is lower that the marginal cost of the BOP (which is assumed to be the opportunity cost of not producing hydrogen). As one can see, the optimizer can find a slightly better solution than the marginal cost dispatch (77.51$/MWh vs 76.21%/MWh). In this case the optimizer decides not to produce some electricity, preferring to pay the penalty. For example, 1 MWh produced from the BOP cost $88.94 but the penalty cost is only $50. One can see in the last run, if the price of hydrogen is zero and therefore the marginal cost of the BOP is zero as well, the optimizer cannot find a better solution than the LCOE from the marginal cost dispatch. If the penalty cost is below the marginal cost of the gas turbine, the same behavior that was observed for the BOP is also seen, i.e. the optimizer chooses to pay the penalty instead of producing electricity with the gas turbine. In all simulation cases shown in this document, the penalty cost is always higher than the highest marginal cost of all components, i.e. the LCOE from the marginal cost dispatch is the optimum, assuming no inertia in the system.

Table 36. Marginal cost dispatch vs optimized dispatch for different maximum penalty costs.

		Marginal Cost	Marginal Cost	Marginal cost dispatch		Optimizer	
Max penalty	H_2 price	SES	BOB	penalty paied	LCOE	penalty paid	LCOE
[$/MWh missed]	[$/kg]	[$/MWh]	[$/MWh]	[M$/year]	[$/MWh]	[$/year]	[$/MWh]
400	3.5	37.9	88.94	390.4	261.7	390.4	261.7
50	3.5	37.9	88.94	49.9	77.51	64.6	76.21
50	0	37.9	0	49.9	117.02	49.9	117.02

With respect to Case 1 the number of points in the parametric grid sweep was reduced given the much larger computational effort that results from the high fidelity Modelica representation of the N-R HES used to determine the optimal dispatch over one day and one week with one hour time resolution. Moreover, the span of the parameter range was reduced to be close to the nominal value for which the Modelica model of the N-R HES has been designed for.

The parametric grid is reported below:

- Hydrogen prices: 1.75 $/kg, 3.5 $/kg
- Wind penetration (nominal). Values: 0%, 50%, 100%
- Industrial process capacity (Hydrogen production). Values: 48.59 MWe, 51.15 MWe, 53.7 MWe
- Mean demand. Values: 200 MW, 300 MW, 400 MW
- Battery (SE in figures). Values: 15 MWh, 20 MWh, 25 MWh
- Gas turbine (SES in figures). Value: 28 MWe, 35 MWe, 42 MWe
- Reactor power is kept constant at 300 MWe.

This simulation highlights one of the major problems identified thus far. Even if a very small number of Modelica simulations failed to successfully complete the one day (~1%), the overall simulation also fails due to the intrinsic sequential nature of the optimization algorithm used. This is a problem that is currently being investigated to determine a mitigation path. Failure appears to be associated with a high penetration of wind (no failure detected at 0% wind penetration). This result appears to indicate that the failures are associated with large oscillation within the dispatch time history. The mitigation strategies rely on the fact that the stochastic optimization is based on multiple random samplings at the same optimization point, so it is possible to discharge a limited number of simulations. Second, the failed runs could be resubmitted using a newly generated random electricity demand and wind history. Work is also ongoing to make the Modelica model more robust, which will come at the expense of being more computationally demanding. This issue will be addressed more in detail in the report conclusions.

In the following selected representative plots of the LCOE are reported (Figure 54, Figure 55) for the one-day case and on a smaller grid (Figure 56, Figure 57) for the one-week dispatch optimization with the Modelica case. It is interesting to notice that, even if in the narrow ranges considered, the wind penetration changes, the minimum LCOE occurs at much higher levels of the demand. This is due to the fact that the system attempts to overcome the effect of the volatility introduced by the wind. The system attempts to maintain the reactor as baseload, i.e. tries to minimize the time when the net demand is lower than the reactor capacity. With increasing wind penetration, i.e. increasing volatility, this can be achieved by increasing the mean demand. This also explains, in those ranges, the beneficial effect of the presence of the gas turbine for which increased turbine size always decreases the cost of electricity. Table 37 reports for selected cases (1 day) the number of Modelica model failure, which is clearly correlated with the wind presence.

Table 38 illustrates, for a given combination of wind penetration and hydrogen price and for the point in the component size grid that leads to the lower LCOE, the differences between the LCOE obtained using the dispatch optimization with the Modelica representation of the N-R HES and the dispatch based on marginal cost (for the 1 day cases). The differences are quite small, which justifies the usage of the marginal cost based dispatch (referred as initial dispatch in section 6.3) as a means of accelerating the optimization. Moreover, as expected, the differences increase for increasing wind penetration. This is due to the fact that the wind introduces larger changes in the dispatch making the natural inertia of the system more relevant. This effect is properly captured by Modelica, where the results usually lead to higher LCOE, likely due to higher penalties paid for mismatching the demand (further investigation is needed to fully assess causality).

Table 37. Modelica number of failures per given wind penetration and hydrogen price.

H2 price ($/Kg)	Wind (%)	Failed runs (%)
1.75	0	0.00
1.75	50	1.89
1.75	100	1.71
3.5	0	0.00
3.5	50	1.93
3.5	100	1.20

Table 38. Differences between dispatch optimization using Modelica N-R HES model and marginal cost dispatch.

H2 price ($/kg)	Wind (%)	Demand (kWe)	SES (kWe)	ES (kWe)	IP (kWe)	LCOE ($/kWh)	LCOE, no Modelica ($/kWh)	Difference (%)
1.75	0	300	42	20	48.59	137.13	134.05	-2.30
1.75	50	300	42	15	53.70	149.2	149.98	0.52
1.75	100	300	42	20	53.70	229.45	213.21	-7.62
3.5	0	300	42	25	48.59	132.51	129.03	-2.70
3.5	50	300	42	20	53.70	142.79	143.74	0.66
3.5	100	300	42	20	53.70	221.44	205.08	-7.80

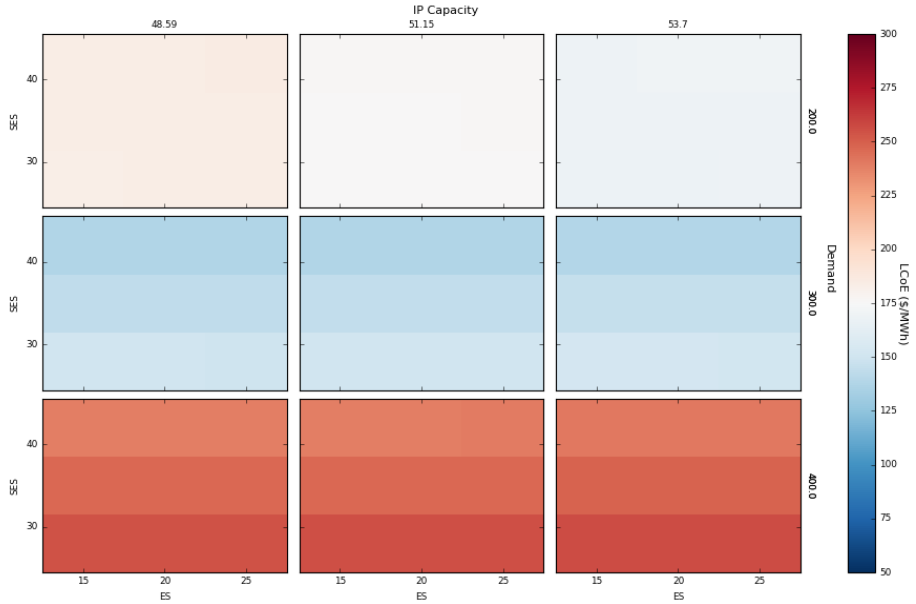

Figure 54. Effective LCOE color map for Case 4. H_2 price 1.75 $/kg, wind penetration 0%, dispatch optimization with 1 day Modelica simulation.

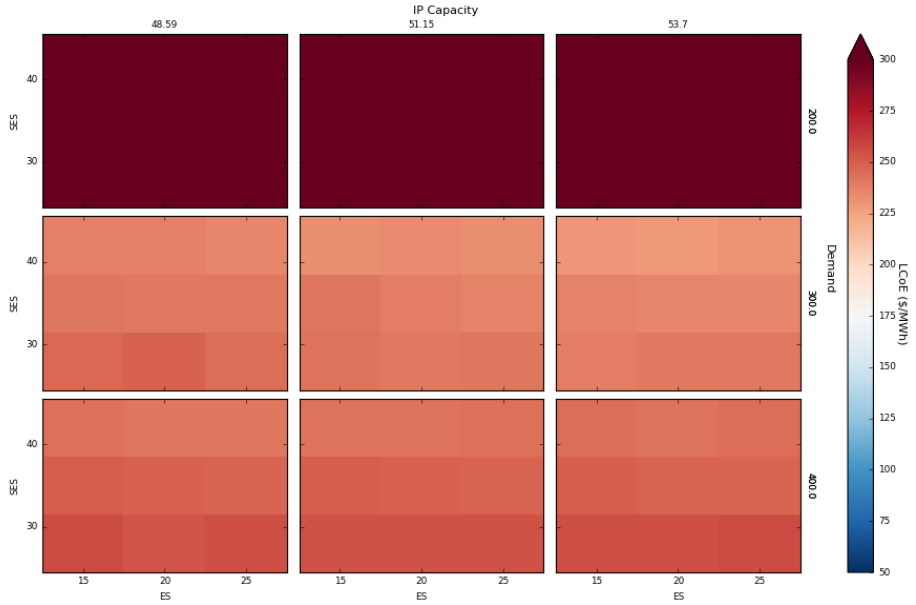

Figure 55. Effective LCOE color map for Case 4. H_2 price 1.75 $/kg, wind penetration 100%, dispatch optimization with 1 day Modelica simulation.

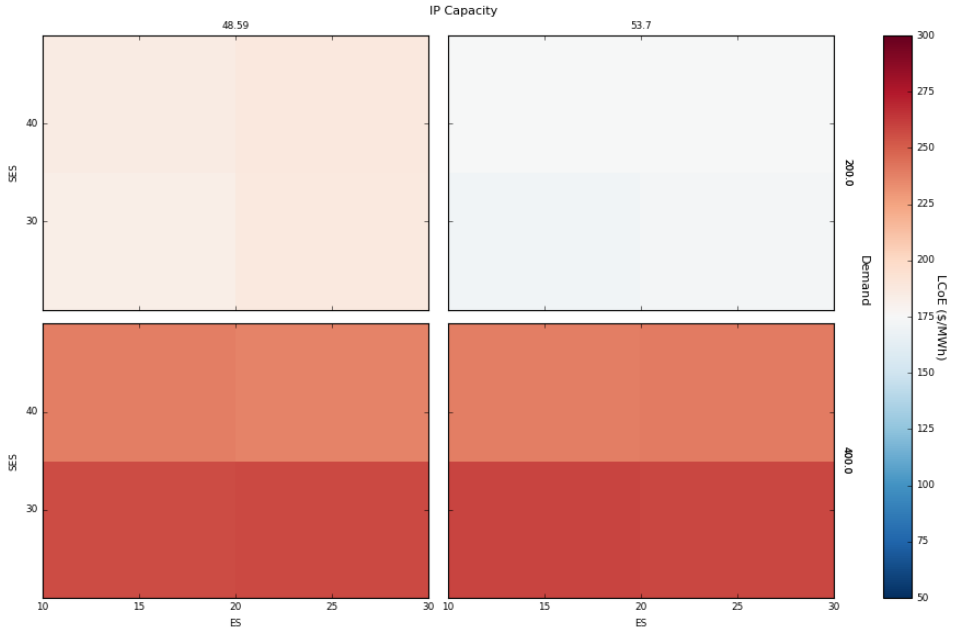

Figure 56. Effective LCOE color map for Case 4. H_2 price 1.75 $/kg, wind penetration 0%, dispatch optimization with 1 week Modelica simulation.

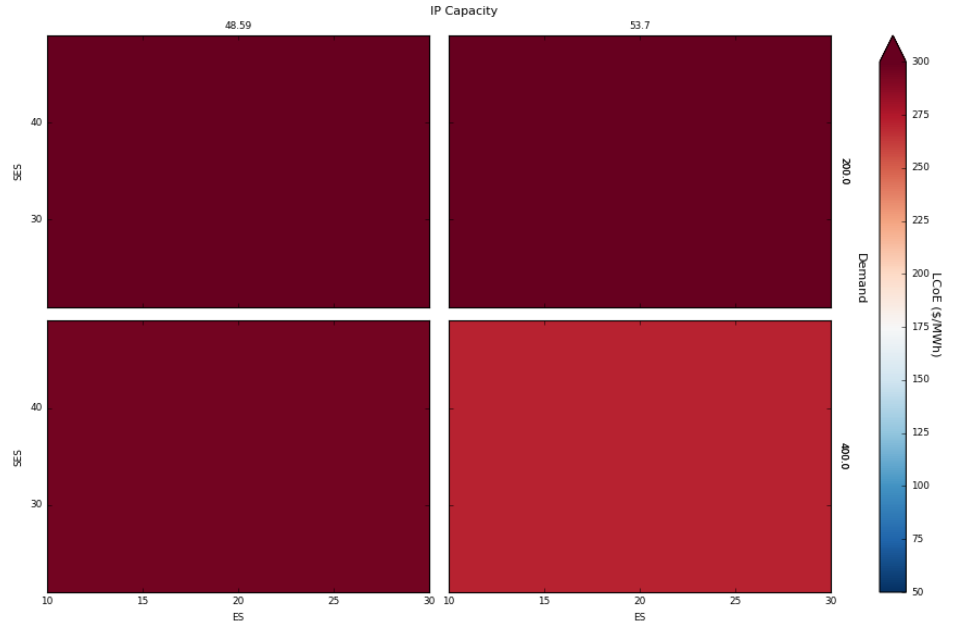

Figure 57. Effective LCOE color map for Case 4. H_2 price 1.75 $/kg, wind penetration 100%, dispatch optimization with 1 week Modelica simulation.

7.5 Case 5

This case tests the capability of the optimizer to properly size the components of the N-R HES for six different combinations of wind penetration and hydrogen price:

- Hydrogen prices: 1.75 $/kg, 3.5 $/kg
- Wind penetration (nominal). Values: 0%, 50%, 100%

The convergence history for one of these cases is reported in Figure 58 (all the convergence histories are reported in appendix C). It is important to notice that the optimizer converges fairly quickly with respect to the demand size and the LCOE while the other component capacities remain more "noisy". This is due to the fact that the slowly converging capacities are less influential in the determination of the LCOE. As mentioned previously, the large number of iterations required for the overall search could be very expensive if the lower level of the optimization run the full Modelica model for the N-R HES. Moreover, for a tighter convergence the slowly converging component is even more computationally expensive. A solution to this problem has been implemented but not yet tested. The upper level of the optimization, which takes care of converging the capacities, would be based on a full first order evaluation of the gradient rather than a zero order as in the SPSA approach, which would be more effective. The full gradient approach is feasible given the low number of variables describing the capacities of the system.

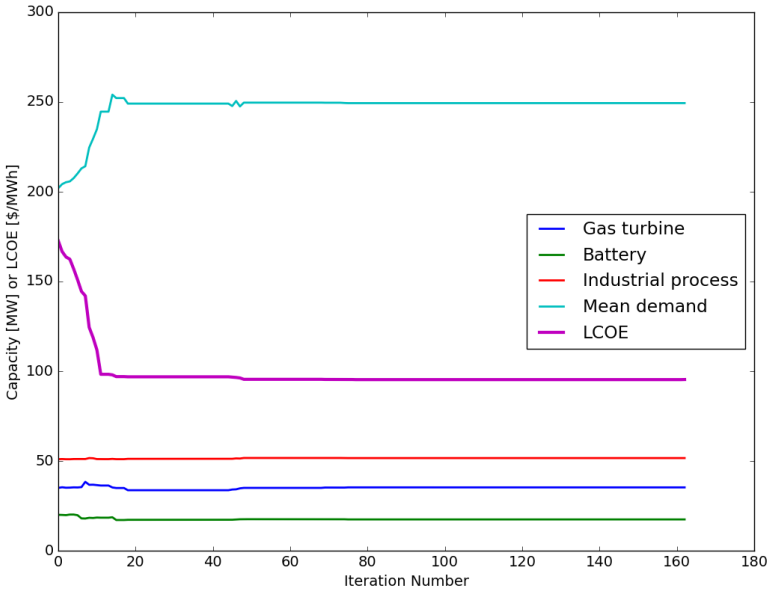

Figure 58. Convergence history for Case 5 (hydrogen price 3.5 $/Kg, wind penetration 0%).

7.6 Case 6

Case 6 combines the optimization of the component capacities and the optimization of the dispatch for 1-day simulations with Modelica. For each new combination of capacities, a full dispatch optimization including Modelica is run. This dispatch optimization is always initiated with the marginal cost dispatch. As we already know from earlier cases and can be seen in Table 39 (Inner iterations), this marginal cost dispatch leads to the minimum LCOE and only three iterations are required for the inner optimization.

The optimal capacities are known from Case 5. These capacities are used to initialize the capacity optimization. Table 39 shows the convergence history for one selected case: 1.75 $/kg of hydrogen and 50% wind penetration. One can see that the initial capacity combination (iteration 0) leads to an LCOE of 156.37$. The optimizer then computes the gradient and takes a step in that direction. The optimizer computes the LCOE for that point and finds it to be higher than for the original point (iteration 1). Since the gradient points in that direction, the only possibility is that the step taken was too big and the optimizer 'overshoots' the minimum. Therefore, the optimizer goes back to the initial point, re-evaluates the gradient and tries a smaller step in that direction. In the case shown, it repeats that behavior 5 times before it finds a better point, which is actually very close to the original point (iteration 7). The optimizer evaluates the gradient again and finds a better point with a very small step size (iteration 8). The step size is so small that the solution is considered converged (iteration 8). Since the problem is stochastic, this point may not be the optimum considering another wind and demand profile. To check this, the solution has to be persistent, i.e. a point with the same step size is evaluated and checked to see if the LCOE is lower. In this case it in not, and the solutions is considered converged.

At this point, the very low number of iterations allows us to cope well with the expensive Modelica simulations. This is due to the very good agreement between the copper plate marginal dispatch model and the Modelica simulation which allows selection of very good initial guesses for both the initial dispatch and capacities. This advantage will be reduced as simulation moves to finer time resolution (e.g. 5-min versus 1-hour).

Table 39. Convergence history for the simultaneous optimization of capacity and dispatch with Modelica simulations for 1 day.

	current optimum point					convergence			
Iteration	Battery capacity [MWh]	Demand [MW]	IP capacity [MW]	Gas turbine capacity [MW] [MW]	LCOE [$/MWh]	Rel conv.	Step size	Inner iterations	comment
0	18.75	261.409	51.02	34.54	156.37423	-	-1.0000E+00	3	Initial point
1	"	"	"	"	"	-	3.3333E-02	3	Initial step size =>Reject point for worse value
2	"	"	"	"	"	-	1.1111E-02	3	Reduce step size =>Reject point for worse value
3	"	"	"	"	"	-	3.7037E-03	3	Reduce step size =>Reject point for worse value
4	"	"	"	"	"	-	1.2346E-03	3	Reduce step size =>Reject point for worse value
5	"	"	"	"	"	-	1.3717E-04	3	Reduce step size =>Reject point for worse value
6	"	"	"	"	"	-	4.5725E-05	3	Reduce step size =>Reject point for worse value
7	18.749986	261.4074	51.01999612	34.53996	156.35589	0.000117332	1.5242E-05	3	Reduce step size =>Found better point
8	18.749988	261.4075	51.01999962	34.53995	156.24609	0.000702223	2.0728E-06	3	Move => Found better point => converged on step size
9	18.749989	261.4075	51.02000022	34.53991	157.43107	-	3.0330E-06	3	Pesistence: Move by same step size => don't find better => converged twice

8. CONCLUSIONS

The software framework for the economic analysis of N-R HES has been tested in its completeness for the first time. The four main components of this simulation system are:

- The capability to generate random time histories that statically equivalent to a given training set but always different from each other (ARMA).
- A global driver (RAVEN) capable of performing statistical sampling, parametric and sensitivity studies, and stochastic optimization. These capabilities are deployed taking advantages of parallel computing.
- A general economics framework allowing calculation of different economic indicators such as NPV, IRR, etc. starting from a generic set of cost drivers, which in our case are the sizes of the components (capital cost), and the flow of commodities (e.g. electricity and industrial product).
- A multi-level fidelity representation of the physical response of a N-R HES:
    - Low level: the system is considered without any inertia and any component of the system could be dispatched at will.
    - Middle level: linear surrogates are used to simulate the response of the system components.
    - High level: Modelica based model of the integrated N-R HES.

A sizable amount of work has been dedicated to collect and organize the economica data to perform a cash flow based analysis of the system. This has been done not only for the reference system, but cost scaling laws have been derived for all the components that could be resized by seeking the overall optimal financial performance of the system.

A large series of tests has been performed to confirm the proper functioning of the complete infrastructure. Those tests not only assessed a good degree of maturity of the framework, but the large amount of data they generated is already providing useful information on the financial performance of N-R HESs and their dynamic performance aspects.

A few shortcomings of the framework have been identified through this process. Each Modelica run that corresponds to a week of real time requires about two hours to complete. Multiplied by the number of needed iterations (in the hundreds without acceleration) requires too long of a computational time. In addition, some failures of the Modelica runs (<1%) result in failure of the whole optimization search. These issues have been identified and are currently being mitigated (Modelica will be run in parallel on an HPC platform, initial dispatch based on marginal costs, failure resilient optimization, and further development of surrogate models).

Overall the framework has been proven to be very flexible and capable of handling many different approaches and analysis needs. It is worth mentioning that all the grid based analyses have been dispatched in parallel by RAVEN, as has de-noising (multiple runs on the same points to remove statistical noises).

At this point of the development, it is clear that the next step should be the deployment of those capabilities to evaluate a case that is highly relevant to the industry, while tuning algorithms and improving overall performance.

9. REFERENCES

1. A. Epiney, C. Rabiti, A. Alfonsi, P. Talbot, F. Ganda, "Report on the Economic Optimization of a Demonstration Case for a Static N-R HES Configuration using RAVEN", Idaho National Laboratory, April 2017, INL/EXT-17-41915.
2. R. Ponciroli and R. B. Vilim "Testing of Strategies for the Acceleration of the Cost Optimization," Argonne National Laboratory August 2017, ANL/NE-17/21.
3. M. S. Greenwood M. S. Cetiner D. L. Fugate "Nuclear Hybrid Energy System Model Stability Testing" Oak Ridge National Laboratory, March 2017, ORNL/TM-2017/153.
4. C. Rabiti, et al., "RAVEN User Manual", INL report INL/EXT-15-34123 Revision 5, February 2016.
5. M.S. Greenwood, S.M. Cetiner, T.J. Harrison, D. L. Fugate, "A Templated Approach for Multi-Physics Modeling of Hybrid Energy Systems in Modelica." (In Press)
6. Gary D. Storrick, Bojan Petrovic & Luca Oriani. Computer Models for IRIS Control System Transient Analysis. (Westinghouse, 2007). doi:10.2172/933157
7. Modelica Association, "Modelica Standard Library." https://github.com/modelica/Modelica
8. Casella, F., "ThermoPower: A Modelica library for thermal power generation system modelling." https://github.com/casella/ThermoPower
9. Greenwood, M.S., "TRANSFORM-Library: A Modelica based library for modeling thermal hydraulic energy systems and other multi-physics systems." Oak Ridge National Laboratory, https://github.com/ORNL-TRANSFORM/TRANSFORM-Library
10. F. Ganda, J. Hansen, T. K. Kim, T. A. Taiwo, R. Wigeland, "Reactor Capital Costs Breakdown And Statistical Analysis Of Historical U.S. Construction Costs", Proceedings of ICAPP 2016, April 19th, 2016, San Francisco, CA, Paper 16829.
11. F. Ganda, T. K. Kim, T. A. Taiwo and R. Wigeland, "Analysis of reactor capital costs and correlated sampling of economic input variables", Proceedings of ICAPP 2015, May 03-06, 2015, Nice (France), Paper 15342.
12. F. Ganda, T. K. Kim, T. A. Taiwo, J. Hansen, R. Wigeland, "Economic Evaluation of Promising Options", FCRD-FCO-2015-000013, September 30th, 2015.
13. F. Ganda, E. Schneider, K. Williams, T. K. Kim, T. A. Taiwo, "Identification and Analyses of Fuel Cycle Economics Issues", FCRD-FCO-2014-000402, September 30th, 2014.
14. Peters, M.S., K.D. Timmerhaus, and R.E. West, Plant Design and Economics for Chemical Engineers, ed. M. Hill. 2003, New York, NY, U.S.
15. F. Ganda, R. Wigeland, T. K. Kim, T. A. Taiwo, B. Dixon, "Advances in Developing Improved Fuel Cycle Cost Estimates", FCRD-FCO-2016-000498, June 1st, 2016.
16. Nuclear Energy Cost Database (EEDB): DOE/NE-0095.
17. P. Krull, J. Roll, R. D. Varrin, Jr., "HTSE Plant Cost Model for the INL HTSE Optimization Study," R-6828-00-01, Revision 1, March 2013.
18. S. A. Newell, J. M. Hagerty, K. Spees, J. P. Pfeifenberger, Q. Liao, C. D. Ungate, J. Wroble, *"Cost of New Entry Estimates for Combustion Turbine and Combined Cycle Plants in PJM With June 1, 2018 Online Date",* Prepared For PJM Interconnection, L.L.C., by The Brattle Group and by Sargent & Lundy, May 15, 2014.
19. D. Pauschert, *"Study of Equipment Prices in the Power Sector, Energy Sector Management Assistance Program",* ESMAP Technical Paper 122/09, The International Bank for Reconstruction and Development/THE WORLD BANK GROUP, Washington, D.C. 20433, U.S.A., 2009.
20. *"Updated Capital Cost Estimates for Utility Scale Electricity Generating Plant"*, U.S. Energy Information Administration, April 2013.

21. *"State of the Market Report for PJM",* Monitoring Analytics, LLC, Independent Market Monitor, for PJM Q3, 11.10.2016, January through September.
22. T. Randall, "Tesla Wins Massive Contract to Help Power the California Grid," September, 15, 2016, Bloomberg.
23. United Engineers and Constructors, *"Complete concise printout for model 184 - 1144 MWe Pressurized Water Reactor - PWR BE (Best Experience Basis)"*, 1988.
24. *"Conceptual design study of a single-fluid molten-salt breeder reactor"*, Oak Ridge National Laboratory, 1971.
25. F. Ganda, T. K. Kim, T. A. Taiwo and R. Wigeland, "Analysis of reactor capital costs and correlated sampling of economic input variables", Proceedings of ICAPP 2015, May 03-06, 2015, Nice (France), Paper 15342.
26. *Cost* and *Performance Baseline for Fossil Energy Plants Volume 1a: Bituminous Coal (PC) and Natural Gas to Electricity, Revision 3* National Energy Technology Laboratory (NETL).
27. *Market-Based* Advanced *Coal Power Systems, Final Report*. U.S. Department of Energy, Office of Fossil Energy.
28. P. Anglaret, *Nuclear Power Plants: The Turbine Island*. 2013, Alstom Nuclear.
29. D.E. Holcomb, F.J. Peretz;, and A.L. Qualls, Advanced High Temperature Reactor Systems and Economic Analysis, ORNL/TM-2011/364. 2011, Oak Ridge National Laboratory.
30. H. Geoffrey, "What Would It Take To Reduce Us Greenhouse Gas Emissions 80% By 2050?", Working Paper 22525, National Bureau Of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA 02138, August 2016, http://www.nber.org/papers/w22525.
31. B. Nykvist, and M. Nilsson, "Rapidly falling costs of battery packs for electric vehicles, Nature Climate Change" 5, 329–332 (2015), available at https://www.seiinternational.org/mediamanager/documents/Publications/SEI-Nature-pre-pub-2015-falling-costs-battery-packs-BEVs.pdf
32. Lazard's "levelized cost Of Storage—Version2.0," December 2016, available at https://www.lazard.com/media/438042/lazard-levelized-cost-of-storage-v20.pdf
33. M. Geuss, "Largest grid-tied lithium ion battery system deployed today in San Diego," Ars Technica Addendum, 2/24/2017.
34. A. Epiney, et al., "Software development infrastructure for the HYBRID modeling and simulation project", INL report INL/EXT-16-40004, September 2016.
35. A. Alfonsi, C. Rabiti, D. Mandelli, "Raven Facing the Problem of assembling Multiple Models to Speed up the Uncertainty Quantification and Probabilistic Risk Assessment Analyses", Proceedings of *13th Iternational Conference on Probabilistic Safety Assessment and Management (PSAM 13)*, 2~7 October, 2016, Seoul, South Korea.
36. "MOOSE Framework: Testing", (August 29, 2017). Retrieved from http://mooseframework.org/wiki/MooseTraining/testing/
37. RAVEN Team, "RAVEN Test Plan-DRAFT", (August 29, 2017). Retrieved from raven/doc/qa_docs/test_plan.pdf
38. Slaughter, A E et al "Continuous Integration for Concurrent MOOSE Framework and Application Development on GitHub." Journal of Open Research Software, 3: e14, DOI: http://dx.doi.org/10.5334/jors.bx, 2015
39. A. Epiney, J. Chen, C. Rabiti, "Status on the Development of a Modeling and Simulation Framework for the Economic Assessment of Nuclear Hybrid Energy Systems (FY 16)", Idaho National Laboratory, September 2016, INL/EXT-16-39832.
40. J. Chen, C. Rabiti, "Synthetic wind speed scenarios generation for probabilistic analysis of hybrid

energy systems", Energy 120, 2017, p507-517.
41. http://www.nrel.gov/electricity/transmission/eastern_wind_dataset.html
42. http://www.ercot.com/gridinfo/load/load_hist/
43. T. McJunkin, A. Epiney, C. Rabiti, "Report on Integration of Existing Grid odels for N-R HES Interaction Focused on Balancing Authorities for Sub-hour Penalties and Opportunities", Idaho National Laboratory, June 2017, INL/EXT-17-42534
44. Bragg-Sitton S, Boardman R, Ruth M, Zinaman O, Forsberg C, Collins J. Integrated Nuclear-Renewable Energy Systems: Foundational Workshop Report. INL Report INL/EXT-14-32857, Rev.1, NREL Report NREL/TP-6A20-62778, August 2014

APPENDIX A: TITAN AND RHEA COMPUTATIONAL ENVIRONMENT SPECIFICATIONS

Titan contains 18,688 physical compute nodes, each with a processor, physical memory, and a connection to the Cray custom high-speed interconnect. Each compute node contains a 16-core 2.2 GHz AMD Opteron™ 6274 (Interlagos) processor and 32 GB of RAM. Two nodes share a Gemini™ high-speed interconnect router. The resulting partition contains 299,008 traditional processor cores, and 598 TB of memory. In addition to the Opteron CPU, all of Titan's 18,688 physical compute nodes contain an NVIDIA Kepler™ accelerator (GPU) with 6 GB of DDR5 memory. Upon login, users are placed onto login nodes by default. Each Titan login node houses an 8-core AMD Opteron™ 6140-series CPU and 256 GB of RAM. Nodes within the compute partition are connected in a three-dimensional torus. This provides a very scalable network with low latency and high bandwidth.

The OLCF's center-wide Lustre® file system, named Spider, is available on Titan for computational work. Spider contains over 26,000 clients and 32 PB of disk space. A separate, NFS (Network File System)-based file system provides `$HOME` storage areas, and an HPSS-based file system provides Titan users with archival spaces.

Titan employs the Cray Linux Environment as its operating system. This consists of a full-featured version of Linux on the login nodes, and a Compute Node Linux microkernel on compute nodes. The microkernel is designed to minimize partition overhead allowing scalable, low-latency global communications.

Rhea is a 521-node commodity-type Linux cluster. The primary purpose of Rhea is to provide a conduit for large-scale scientific discovery via pre/post processing and analysis of simulation data generated on Titan. Users with accounts on Titan will automatically be given an account on Rhea.

Rhea contains 521 Dell PowerEdge compute nodes. The compute nodes are separated into two partitions as shown in Table 1. Both compute partitions are accessible through the same batch queue from Rhea's login nodes.

Table 40. Specifications of Rhea compute node partitions.

Partition	Node Count	Memory	GPU	CPU
Rhea (default)	512	128GB	–	Dual Intel® Xeon® E5-2650 @ 2.0 GHz 16 cores, (32) HT
GPU	9	1TB	2 NVIDIA® K80	Dual Intel® Xeon® E5-2695 @ 2.3 GHz 28 cores, (56) HT

Each CPU in the Rhea partition features 8 physical cores, for a total of 16 physical cores per node. With Intel® Hyper-Threading Technology enabled the node has 32 logical cores capable of executing 32 hardware threads for increased parallelism. On the GPU partition, there are 14 physical cores, for a total of 28 physical cores per node. With Hyper-Threading enabled, these nodes have 56 logical cores that can execute 56 hardware threads for increased parallelism. This GPU partition also has 1TB of memory and 2 K80 GPUs per node. Rhea also features a 4X FDR Infiniband interconnect, with a maximum theoretical transfer rate of 56 Gb/s.

Rhea features 4 login nodes which are identical to the compute nodes, but with 32 GB of RAM. The login nodes provide an environment for editing, compiling, and launching codes onto the compute nodes.

The OLCF's center-wide Lustre® file system, named Spider, is available on Rhea for computational work. With over 26,000 clients and 32 PB of disk space, it is one of the largest-scale Lustre file system in

the world. A separate, NFS-based file system provides `$HOME` storage areas, and an HPSS-based file system provides Rhea users with archival spaces.

A.1 Compiling on Titan

Compiling code on Titan (and other Cray machines) is different than compiling code for commodity or Beowulf-style HPC linux clusters. Among the most prominent differences are the following:

1. Cray provides a sophisticated set of compiler wrappers to ensure that the compile environment is set up correctly.
2. In general, linking/using shared object libraries on compute partitions is not supported.
3. Cray systems include many different types of nodes, so some compiles are, in fact, cross-compiles.

The following compilers are available on Titan:

1. PGI, the Portland Group Compiler Suite (default)
2. GCC, the GNU Compiler Collection
3. CCE, the Cray Compiling Environment
4. Intel, Intel Composer XE

Cray provides a number of compiler wrappers that substitute for the traditional compiler invocation commands. The wrappers call the appropriate compiler, add the appropriate header files, and link against the appropriate libraries based on the currently loaded programming environment module. To build codes for the compute nodes, the user invokes the Cray wrappers using the following tools:

- cc To use the C compiler
- CC To use the C++ compiler
- ftn To use the FORTRAN 90 compiler

The `-craype-verbose` option can be used to view the compile line built by the compiler wrapper:

```
titan-ext$ cc -craype-verbose ./a.out
pgcc -tp=bulldozer -Bstatic ...
```

Titan is comprised of different types of nodes:

1. Login nodes running traditional Linux
2. Service nodes running traditional Linux
3. Compute nodes running the Cray Node Linux (CNL) microkernel

The type of work performed dictates the type of node for which the code is built.

A.2 Compiling for Compute Nodes (Cross Compilation)

Titan compute nodes are the nodes that carry out the vast majority of computation on the system. Compute nodes are running the CNL microkernel, which is markedly different than the OS running on the login and service nodes. Most code that runs on Titan will be built targeting the compute nodes. All parallel codes should run on the compute nodes. Compute nodes are accessible only by invoking `aprun` within a batch job. To build codes for the compute nodes, the Cray compiler wrappers are invoked using the following commands:

```
titan-ext$ cc code.c
```

```
titan-ext$ CC code.cc
titan-ext$ ftn code.f90
```

A.3 Controlling the Programming Environment

Upon login, the default versions of the PGI compiler and associated Message Passing Interface (MPI) libraries are added to each user's environment through a programming environment module. Users do not need to make any environment changes to use the default version of PGI and MPI. If a different compiler is required, it is important to use the correct environment for each compiler. To aid users in pairing the correct compiler and environment, programming environment modules are provided. The programming environment modules will load the correct pairing of compiler version, message passing libraries, and other items required to build and run. It is highly recommended that the programming environment modules be used when changing compiler vendors. The following programming environment modules are available on Titan:

- PrgEnv-pgi
- PrgEnv-gnu
- PrgEnv-cray
- PrgEnv-intel

To change the default loaded PGI environment to the default GCC environment use:

```
$ module unload PrgEnv-pgi
$ module load PrgEnv-gnu
```

Or alternatively:

```
$ module swap PrgEnv-pgi PrgEnv-gnu
```

To use a specific compiler version, the compiler's PrgEnv module must be loaded, and the use must switch to the correct compiler version. For example, the following will configure the environment to use the GCC compilers, then load a non-default GCC compiler version:

```
$ module swap PrgEnv-pgi PrgEnv-gnu
$ module swap gcc gcc/4.6.1
```

A.4 Running Jobs on Titan

Computational work is performed using HPC by jobs. Individual jobs produce data that lend relevant insight into grand challenges in science and engineering. As such, the timely, efficient execution of jobs is the primary concern in the operation of any HPC system. A job on Titan typically comprises a few different components:

- a batch submission script,
- a statically-linked binary executable,

- a set of input files for the executable, and
- a set of output files created by the executable.

The process for running a job is typically accomplished in the following order:

1. prepare executables and input files,
2. write a batch script,
3. submit the batch script to the batch scheduler, and
4. monitor the job before and during execution.

On Cray machines, when the `aprun` command is issued within a job script (or on the command line within an interactive batch job), the binary passed to `aprun` is copied to and executed in parallel on a set of compute nodes. Compute nodes run a Linux microkernel for reduced overhead and improved performance. It must be noted that the only way to access the compute nodes on Cray machines is via the `aprun` command.

A.5 Writing Batch Scripts

Batch scripts, or job submission scripts, are the mechanism by which a user submits and configures a job for eventual execution. A batch script is simply a shell script which contains:

- Commands that can be interpreted by batch scheduling software (e.g. PBS)
- Commands that can be interpreted by a shell

The batch script is submitted to the batch scheduler where it is parsed. Based on the parsed data, the batch scheduler places the script in the scheduler queue as a batch job. Once the batch job makes its way through the queue, the script will be executed on a service node within the set of allocated computational resources. Batch scripts are parsed into the following three sections:

A.6 Interpreter Line

The first line of a script can be used to specify the script's interpreter. This line is optional. If not used, the submitter's default shell will be used. The line uses the "hash-bang-shell" syntax:

```
#!/path/to/shell
```

A.7 Scheduler Options

The batch scheduler options are preceded by #PBS, making them appear as comments to a shell. PBS will look for #PBS options in a batch script from the script's first line through the first non-comment line. A comment line begins with #. #PBS options entered after the first non-comment line will not be read by PBS. All batch scheduler options must appear at the beginning of the batch script.

A.8 Executable Commands

The shell commands follow the last #PBS option and represent the main content of the batch job. If any #PBS lines follow executable statements, they will be ignored as comments. The execution section of a script will be interpreted by a shell and can contain multiple lines of executable invocations, shell commands, and comments. When the job's queue wait time is finished, commands within this section will be executed on a service node (sometimes called a "head node") from the set of the job's allocated

resources. Under normal circumstances, the batch job will exit the queue after the last line of the script is executed. An example batch script is shown below:

```
 1: #!/bin/bash
 2: #    Begin PBS directives
 3: #PBS -A pjt000
 4: #PBS -N test
 5: #PBS -j oe
 6: #PBS -l walltime=1:00:00,nodes=1500
 7: #PBS -l gres=atlas1%atlas2
 8: #    End PBS directives and begin shell commands
 9: cd $MEMBERWORK/pjt000
10: date
11: aprun -n 24000 ./a.out
```

This introductory information was provided to familiarize potential users with the specifics of the Titan HPC platform. Further information can be found on the OLCF system user guide documentation.[h]

[h] https://www.olcf.ornl.gov/support/system-user-guides/titan-user-guide

APPENDIX B: COMPLETE SERIES OF PLOTS FROM TEST CASE 2

The grid points sampled for Case 2 are summarized here:

- Hydrogen price. Values are defined according to [1]: 0 $/kg, 1.75 $/Kg, 3.5 $/Kg
- Wind penetration (nominal). Values: 0%, 50%, 100%, 200%
- Industrial process capacity (Hydrogen production). Values: 0 MWe, 120 MWe, 240 MWe, 360 MWe, 480 MWe, 600 MWe
- Demand. Values: 100 MW, 200 MW, 300 MW, 400 MW, 500 MW, 600 MW
- Battery (SE in figures). Values: 100 MWh, 200 MWh, 300 MWh, 400 MWh, 500 MWh, 600 MWh
- Gas turbine (SES in figures). Values: 100 MWe, 200 MWe, 300 MWe, 400 MWe, 500 MWe, 600 MWe
- Reactor power is kept constant at 300 MWe

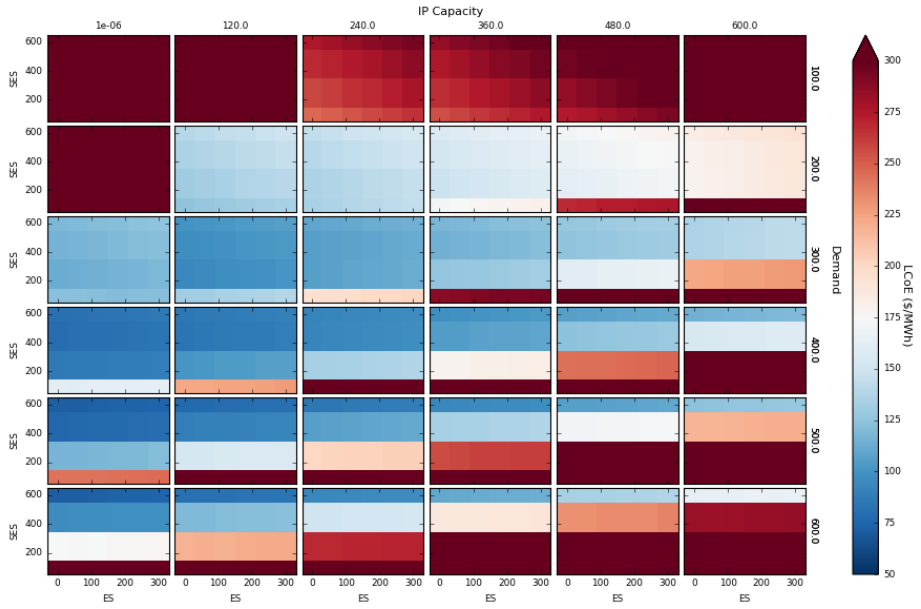

Figure 59. Effective LCOE color map for Case 2. H_2 price 0 $/kg, wind penetration 0%.

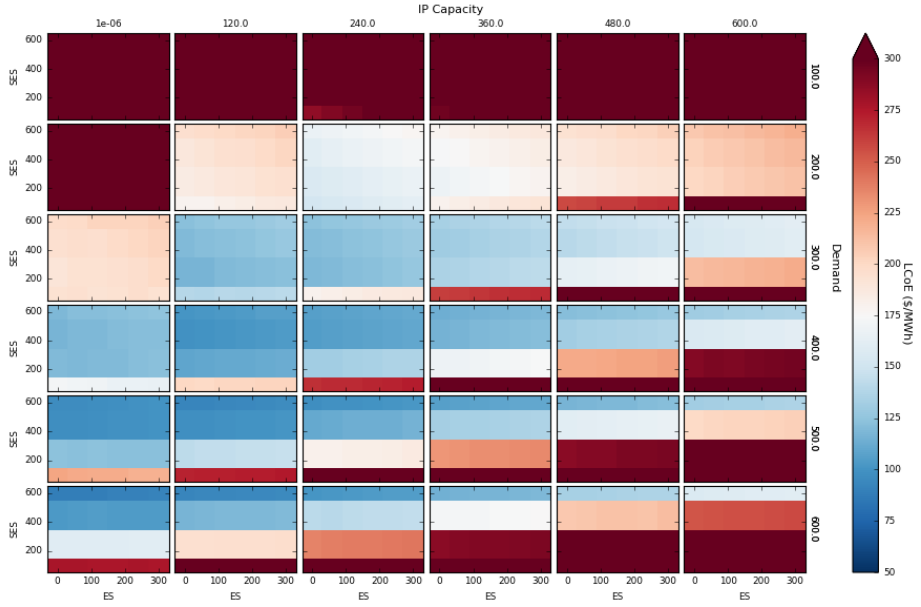

Figure 60. Effective LCOE color map for Case 2. H_2 price 0 $/kg, wind penetration 50%.

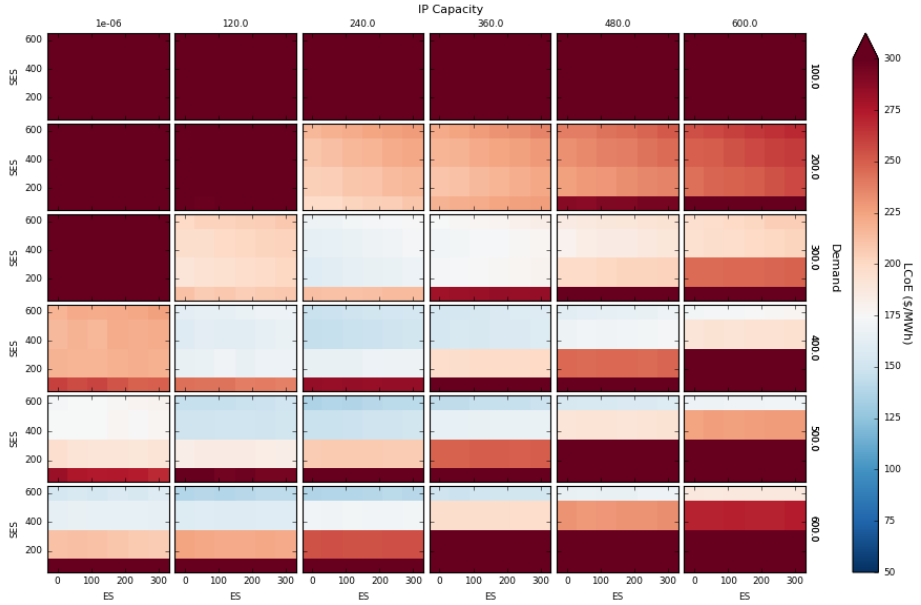

Figure 61. Effective LCOE color map for Case 2. H_2 price 0 $/kg, wind penetration 100%.

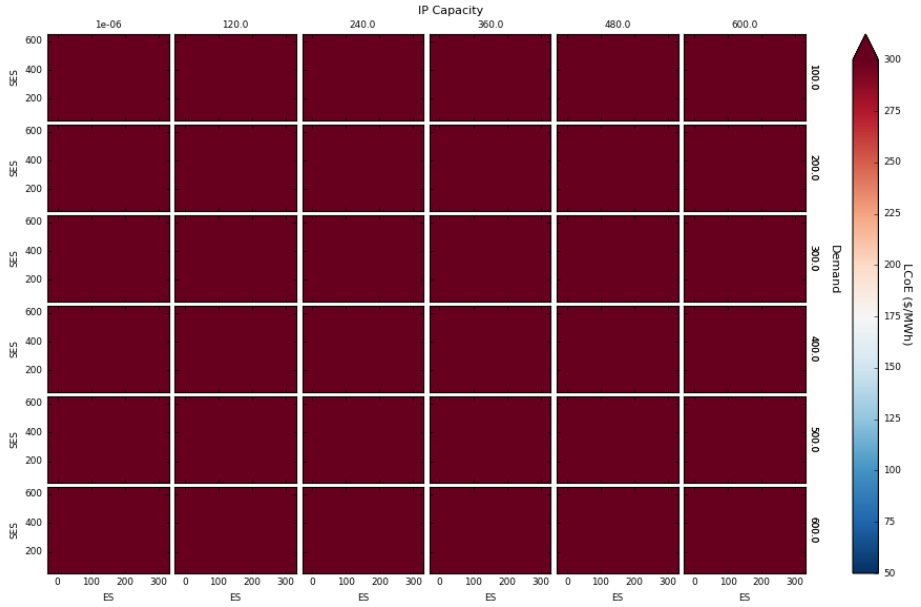

Figure 62. Effective LCOE color map for Case 2. H_2 price 3.5 $/kg, wind penetration 200%.

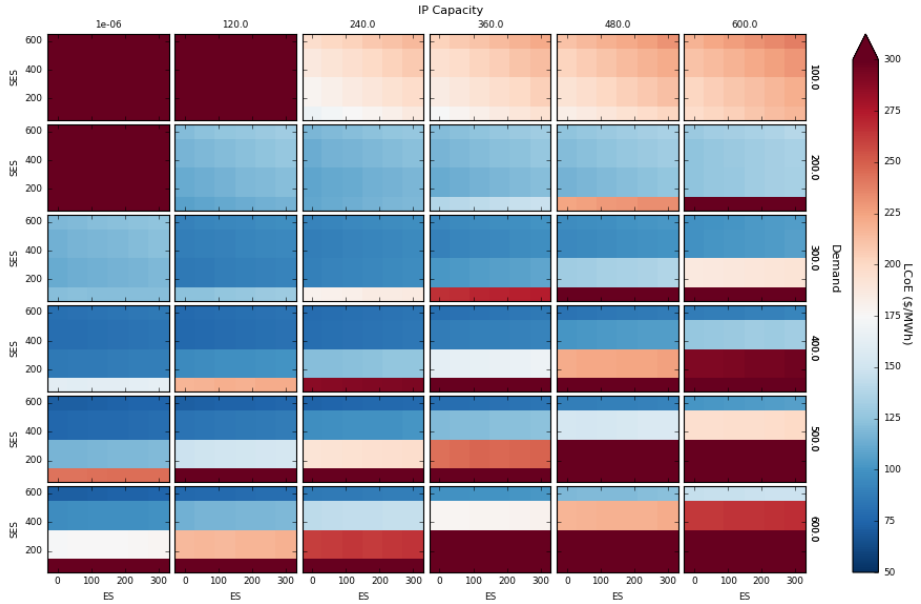

Figure 63. Effective LCOE color map for Case 2. H_2 price 1.75 $/kg, wind penetration 0%.

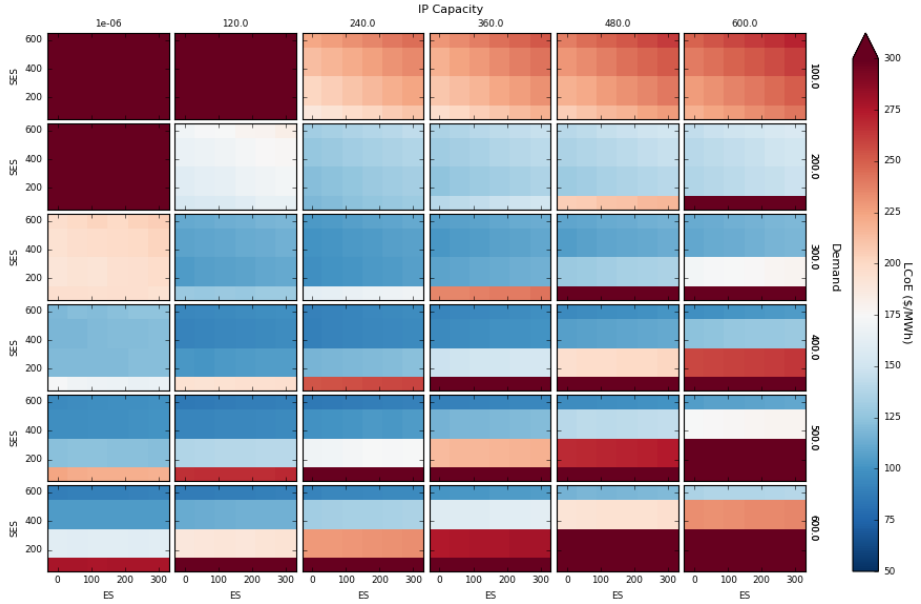

Figure 64. Effective LCOE color map for Case 2. H_2 price 1.75 $/kg, wind penetration 50%.

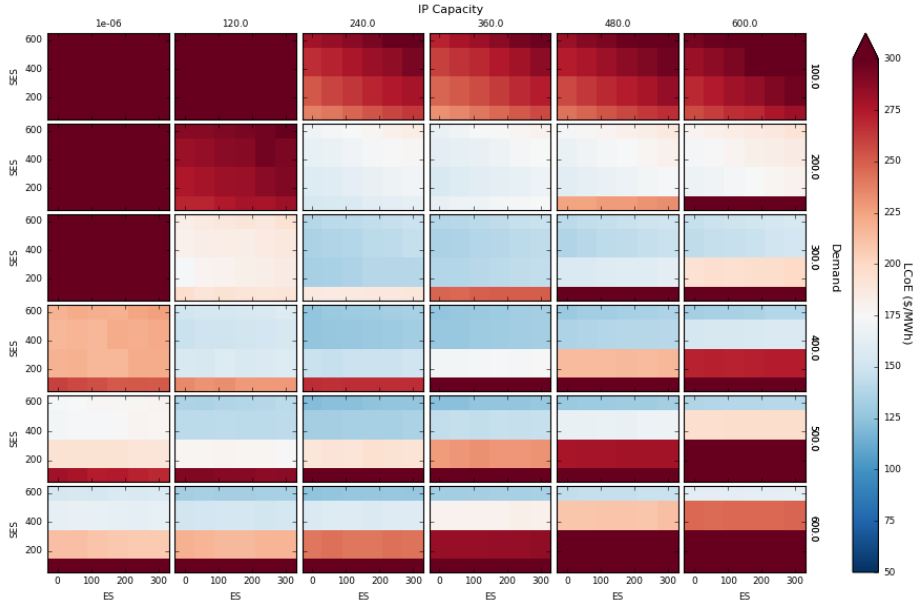

Figure 65. Effective LCOE color map for Case 2. H_2 price 1.75 $/kg, wind penetration 100%.

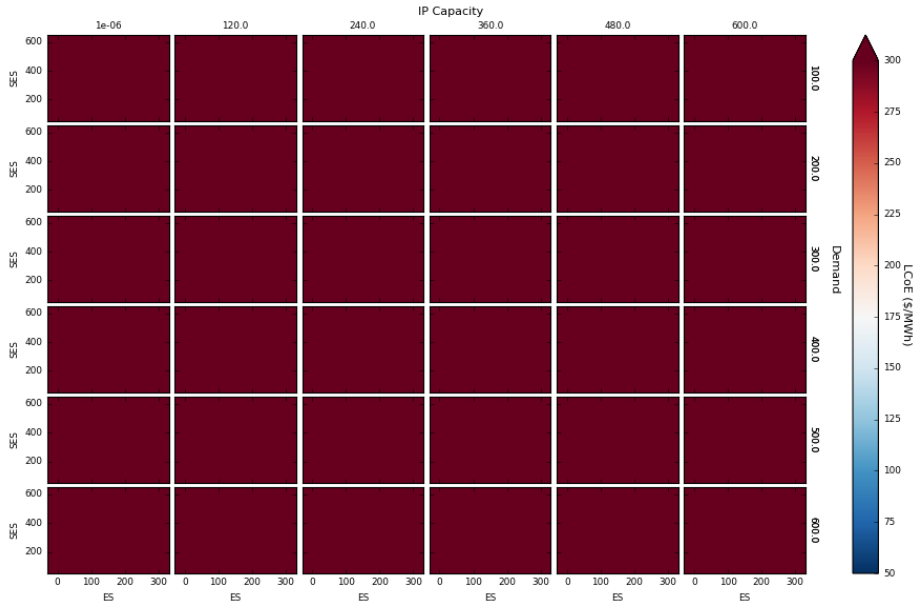

Figure 66. Effective LCOE color map for Case 2. H_2 price 1.75 $/kg, wind penetration 200%.

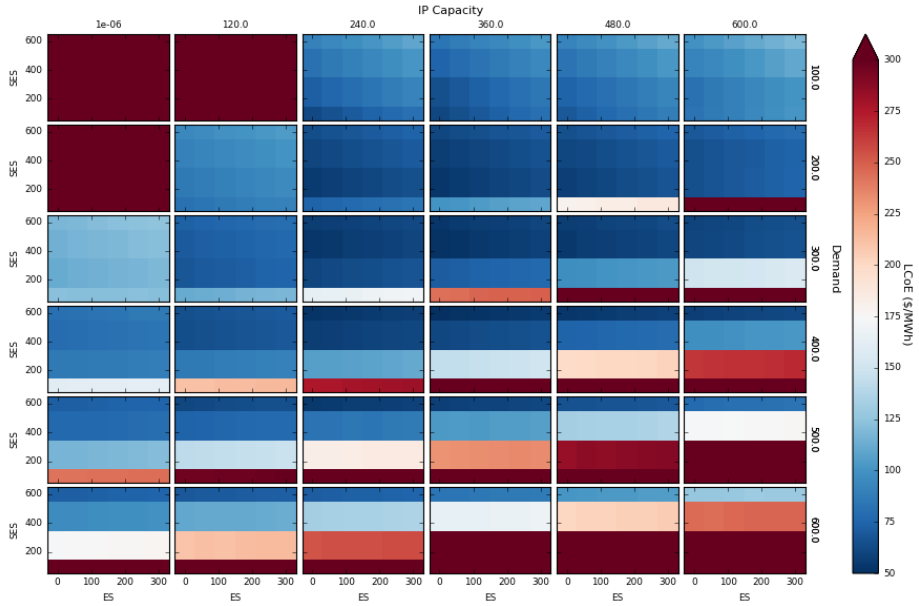

Figure 67. Effective LCOE color map for Case 2. H_2 price 3.5 $/kg, wind penetration 0%.

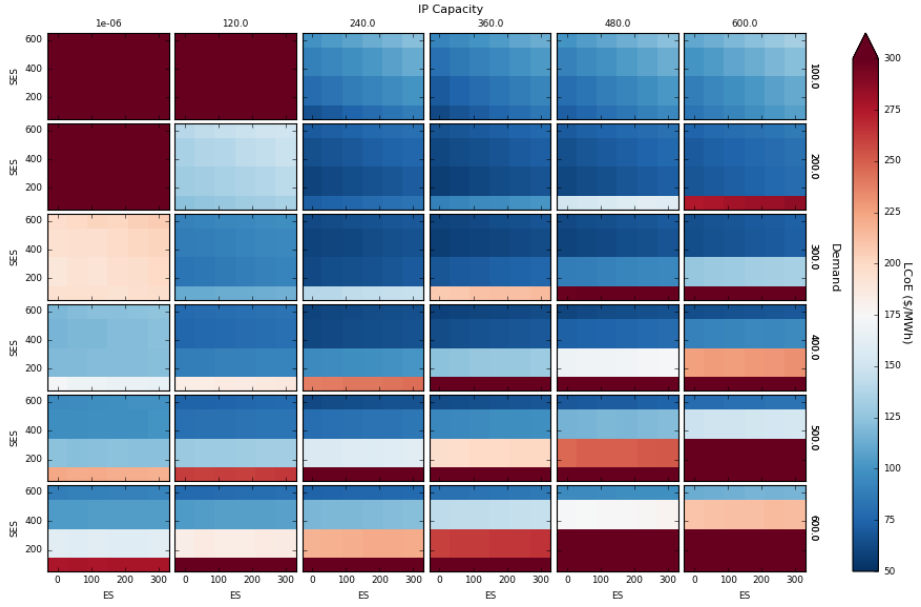

Figure 68. Effective LCOE color map for Case 2. H_2 price 3.5 $/kg, wind penetration 50%.

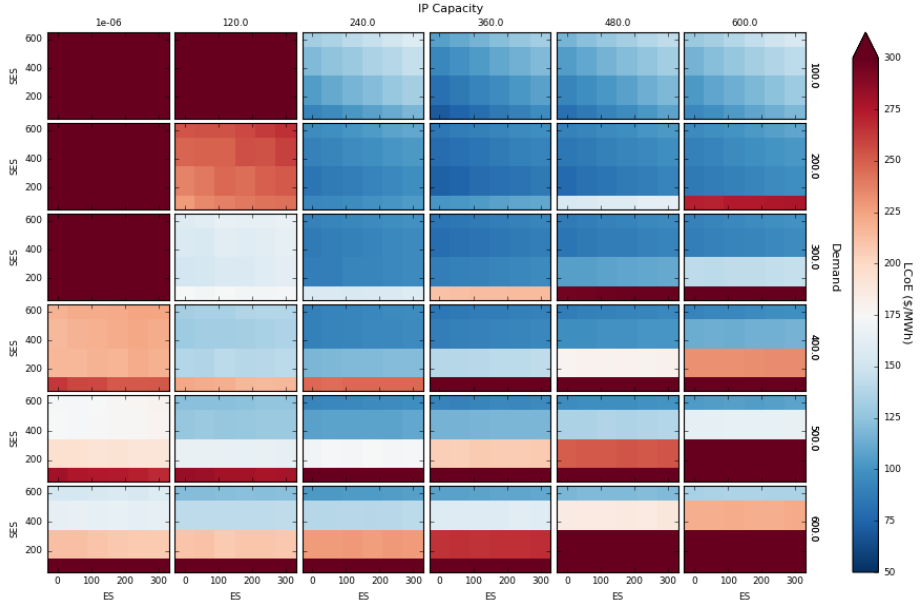

Figure 69. Effective LCOE color map for Case 2. H_2 price 3.5 $/kg, wind penetration 100%.

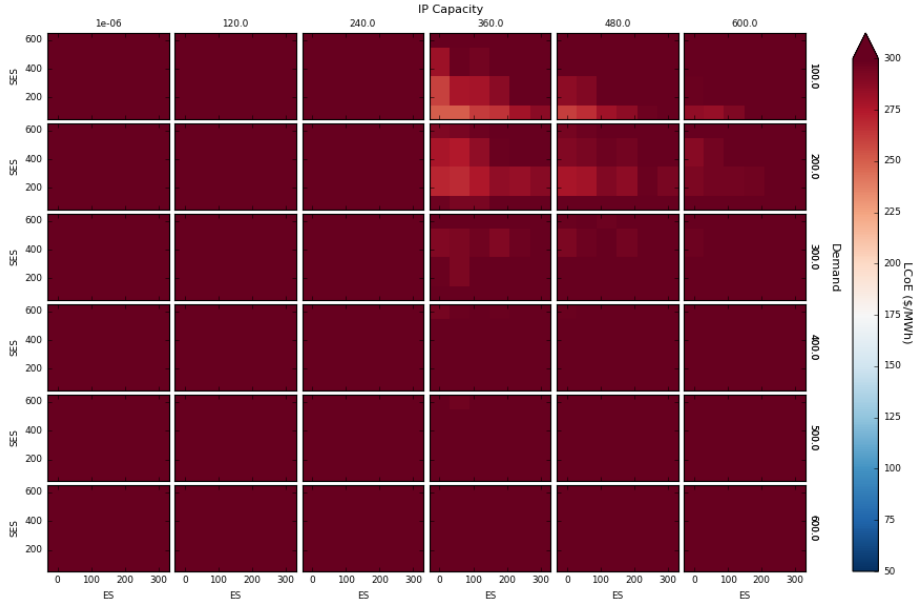

Figure 70. Effective LCOE color map for Case 2. H_2 price 3.5 $/kg, wind penetration 200%.

APPENDIX C: COMPLETE SERIES OF CONVERGENCE HISTORIES FOR CASE 5

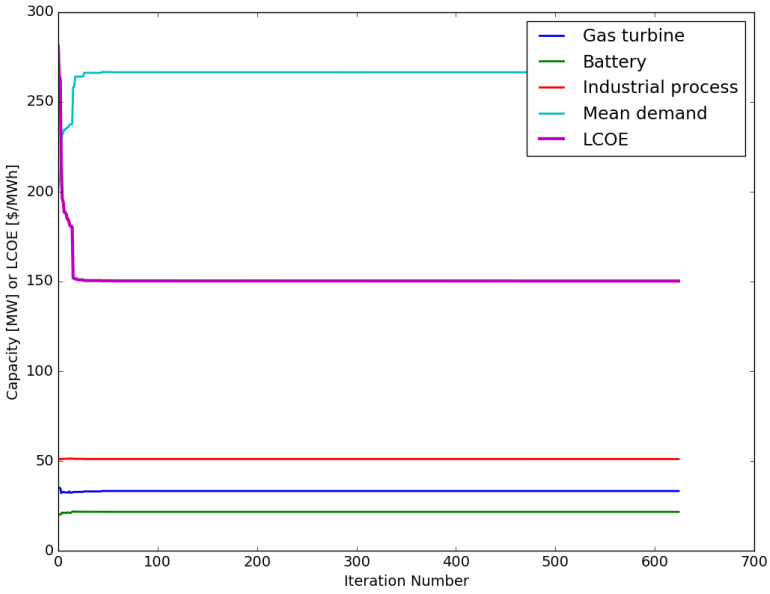

Figure 71. Convergence history for Case 5 (hydrogen price 1.75 $/Kg, wind penetration 0%).

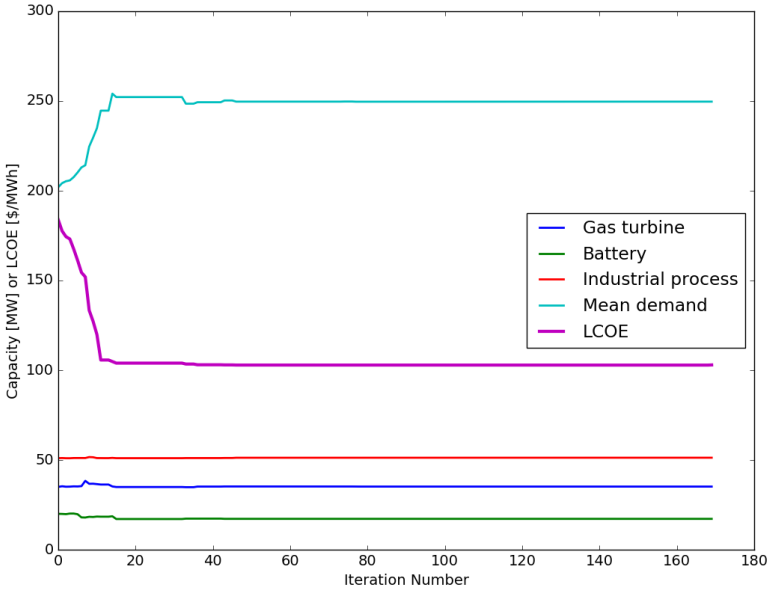

Figure 72. Convergence history for Case 5 (hydrogen price 1.75 $/Kg, wind penetration 50%)

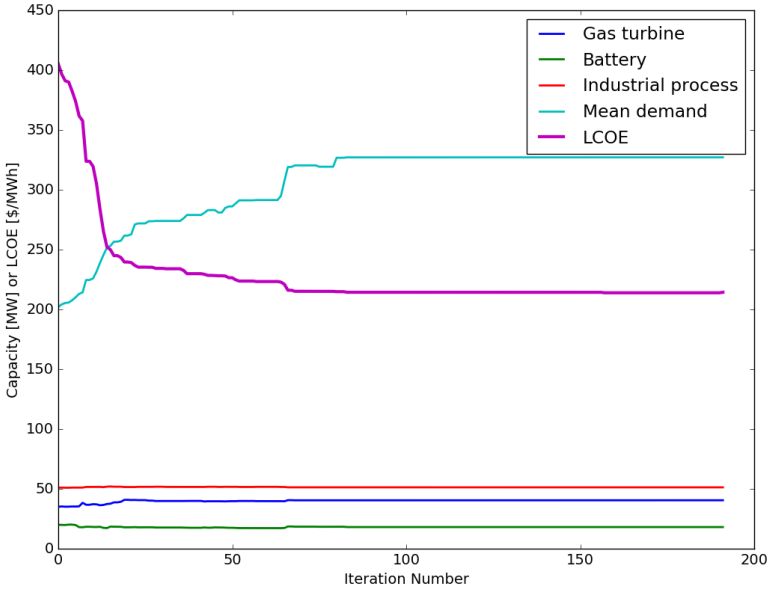

Figure 73. Convergence history for Case 5 (hydrogen price 1.75 $/kg, wind penetration 100%).

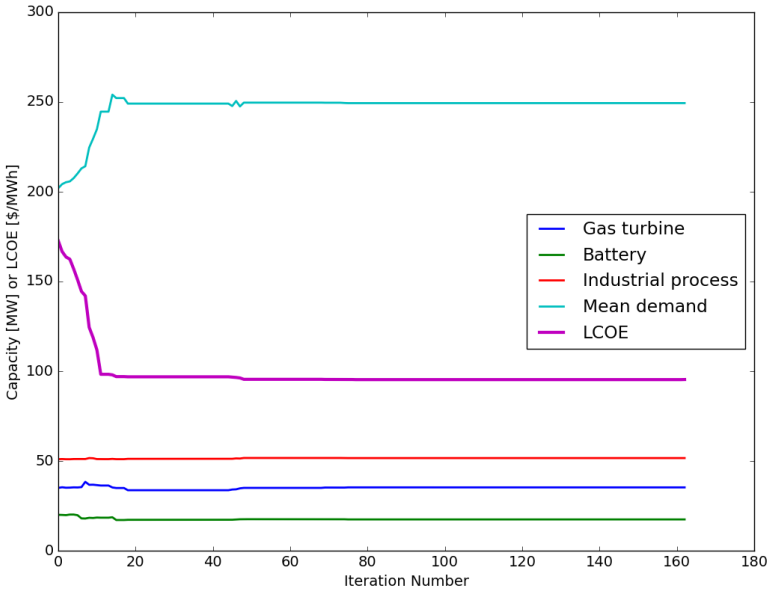

Figure 74. Convergence history for Case 5 (hydrogen price 3.5 $/kg, wind penetration 0%).

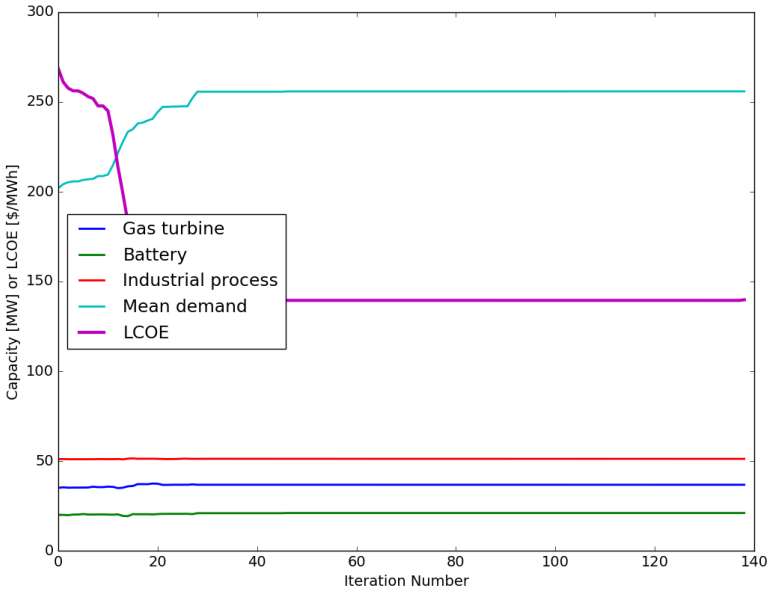

Figure 75. Convergence history for Case 2 (hydrogen price 3.5 $/Kg, wind penetration 50%).

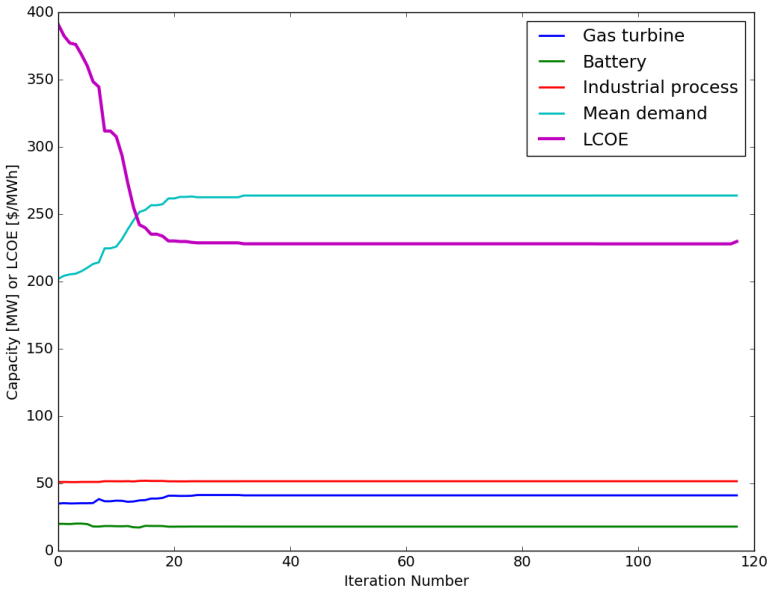

Figure 76. Convergence history for Case 2 (hydrogen price 3.5 $/Kg, wind penetration 100%).