Co-Simulating Integrated Energy Systems with Heterogeneous Digital Twins

By Peter Palensky*, Pierluigi Mancarella, Trevor Hardy, and Milos Cvetkovic

palensky@ieee.org

Energy system integration promises increased resiliency and the unlocking of synergies, while also contributing to our goal of decarbonization. It is enabled by both old and new technologies, glued together with data and digital services. Hydrolyzers, heat-pumps, distributed renewable generation, smart buildings, the digital grid edge: all currently the subject of integration with the power system and the energy sector at large. To plan and operate such a multi-disciplinary and multi-sectoral system properly, insight, tools, and expertise are all needed. This is exactly where the state of the art fails to deliver: Tools for integrated energy systems (IES) are still in their infancy and many times, even academia treats these sectors separately, producing experts in each of them but not across.

Heterogeneous digital twins (DTs), based on co-simulation, are currently a pragmatic and useful approach to working with such complex and interdisciplinary systems. They can host models and data, coming from entirely different schools of thinking, and bring together what is already connected in the real world.

Section 1 – Introduction

Energy systems can be integrated in several directions (Table 1). One dimension is within one sector in a vertical or horizontal way; transmission system and distribution system integration in electricity systems is one example of that. Most of the time, however, we think of cross-sectoral integration where, for instance, heat and electricity are somehow “integrated”. This integration can, and in parts should, happen in several phases. Integrated planning ensures that infrastructure supports each other, while integrated operations can unlock synergies when it comes to flexibility. There is integration in terms of automation, by using shared communication channels or information technology, operational technology (IT/OT) infrastructure. Even markets can be integrated so that bids may be placed for combined products, such as heat and electricity for combined heat and power (CHP) plants. The effort of producing this increased integration is intended to improve resiliency, flexibility, and/or efficiency.

Network type	*Carrier*	*Sector*	*Policy*	*Active Assets*
Transmission network Distribution network Microgrid	Electricity Heat Gas Steam	Industry Housing Transport Water	Infrastructure Markets Resiliency People	Storage Demand Response (Distributed) Generation

Table1: A non-exhaustive list of examples of dimensions of energy system integration

The operation of the system, however, does not become simpler with higher levels of integration. Previously separated processes (whether in planning or operations, physical assets or digital workflows) need to be unified or at least made interoperable. Decisions in one domain will have an impact in another that will only surface if the impacted domain is fully analyzed with its own specialized tools and domain experts; performing multi-domain analysis one domain at a time and in a serial fashion is likely to be time consuming and inefficient. Interdependent properties (think of sizing infrastructure for a thermo-electric system with multiple crossover points such as heat pumps, CHP plants, etc.) lead to chicken and egg situations that are expensive to solve. Cases with complex, cross-domain dynamics or market interactions are even more extreme and yet common. Grid operators that run heat and electricity grids, municipalities that want to plan the next decades of energy investments, or utilities that run CHP and sell both products, are facing the need for more integrated processes.

The complexity of the problem and the fact that experts from different domains must cooperate on the same integrated problem calls for an integrated model: a digital twin. Numerical help, digital twins, are frequently used when analytical methods face their limits, when systems become too complex to be understood, or when borders of disciplines are crossed. Integrated energy systems lead to exactly this case. The lack of multi-domain tools and methods in the energy sector, combined with the need to integrate domains such as heat, electricity, urban planning, markets, automation, communication networks, and even cyber-security, are why multi-domain digital twins are on the rise. They are expected to be one central platform, one central “source of truth”, where diverse teams can work on multi-domain, multi-objective, and multi-time-scale solutions.

Digital twins for integrated energy systems need to be able to:

- Represent multi-physics systems (thermal, electric, gas, building physics, etc.)
- Execute different models with heterogeneous data structures and solvers (ordinary differential equations, 3D such as finite elements or computational fluid dynamics, discrete models, multi-agent models), in order to
- Represent multi-domain systems (markets, energy, automation, communication, policy, spatial planning)
- Handle different magnitudes of time (microseconds for power electronics, minutes for heat flow), and
- Work with multi-domain scenarios in one unified way (e.g., one language to describe all parts of the system)

Further, the model and computation engine need to provide the “usual” features of a digital twin, such as numerical performance and stability, good scalability, and standardized interfaces. The numerical model of the multi-domain digital twin can either be based on a monolithic solver using a flexible multi-domain specification language or combine several models and solvers to a co-simulation.

The digital twin in this multi-energy setting is a platform that hosts models and parameters of all processes involved, takes time series and other environmental/exogenous input data, and then delivers the dynamic behavior of these coupled models within a certain scenario. The workflow is important, as well as how experts interact with the twin: Which interfaces are available? How can the digital twin system be optimized? How can parameters be estimated or updated? How can uncertainty be represented and traced?

In the case of black box models, one scenario gives one snapshot, and many snapshots are required to see the bigger picture (Figure 1). Without derivatives or other analytical insight into the shape of the problem, though, the location of an optimum (e.g., stability or costs) can only be found via searching.

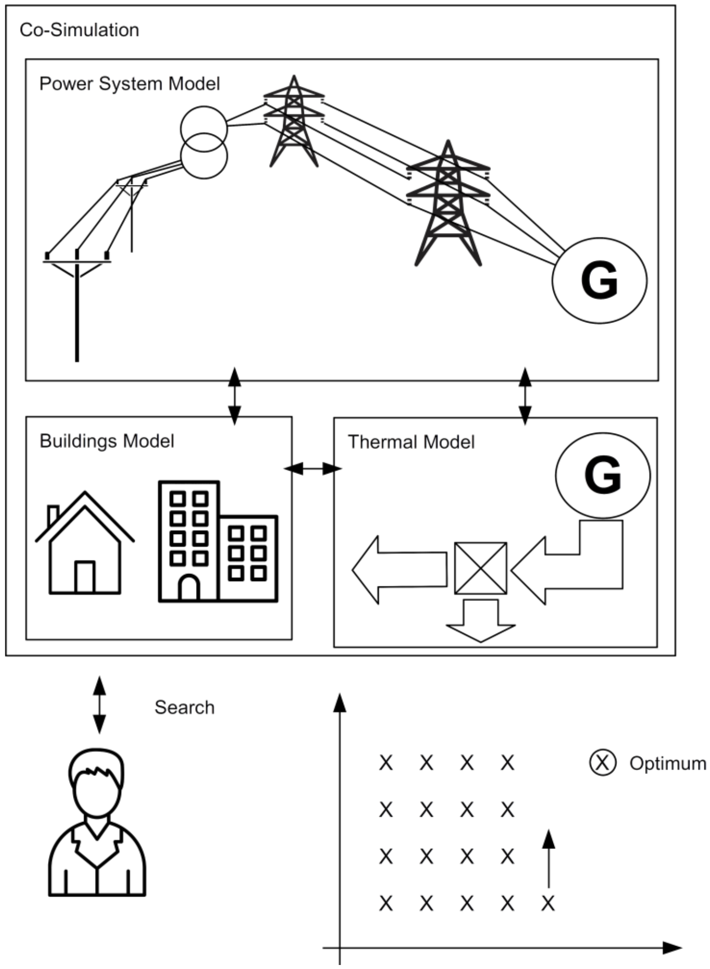

Figure 1: Co-Simulation (see also section 3) of an integrated energy system: a black box needs to be searched.

Sensitivity toward certain parameters or robustness towards others cannot be calculated and, instead, needs to be estimated over a broad range of input parameters. Smart parameter choices and smart sampling with Monte-Carlo or Latin Hypercube methods might reduce a potentially astronomical number of interesting cases to a manageable subset, but the fundamental problem stays: The models do not expose details such as derivatives that would allow for smart optimization and, therefore, the solution must be searched for.

If these processes and details were defined in one model (Figure 2), through the use of one language, and executed with one solver, optimization beyond heuristic searching methods would be possible.

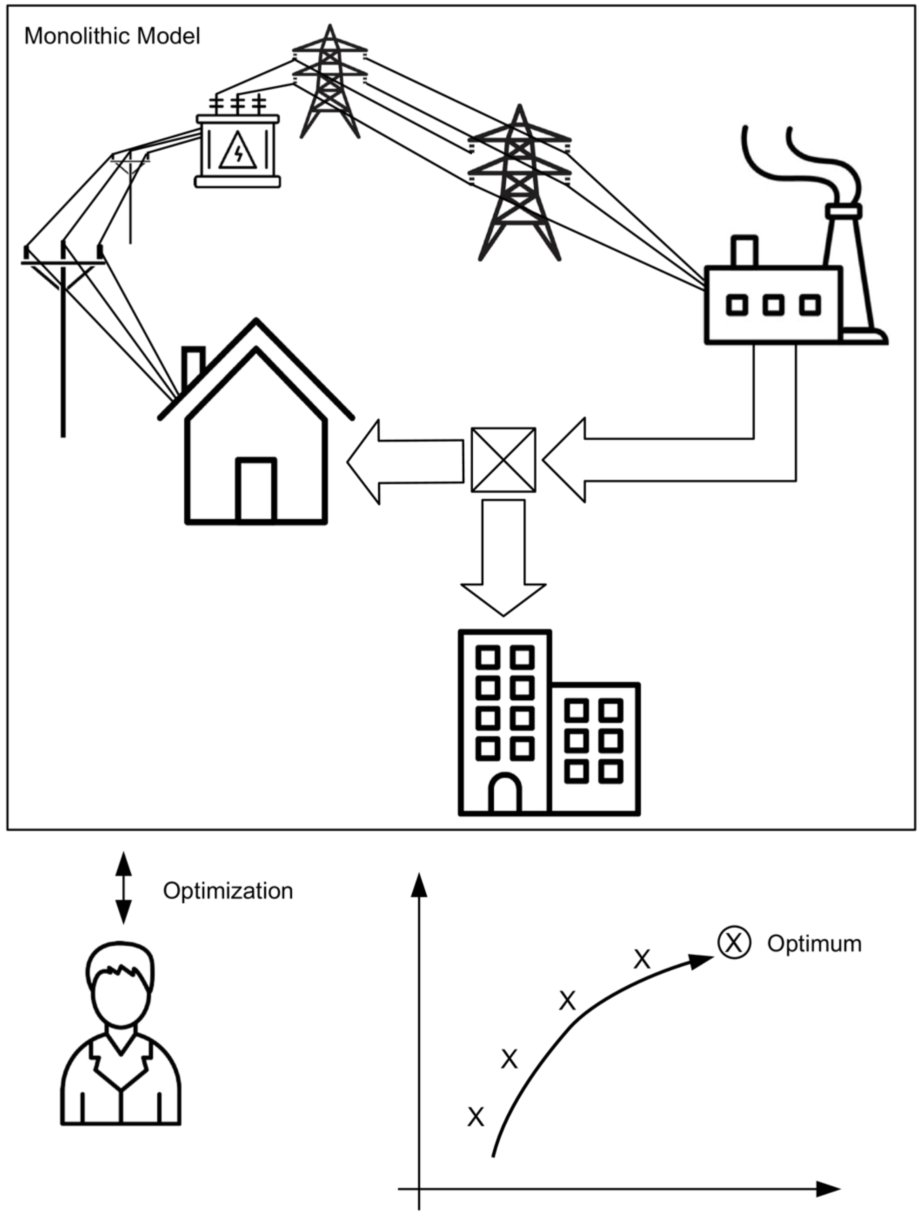

Figure 2: A monolithic model that allows for smart optimization.

Existing languages and tools, however, are not sufficient (in either model coverage or computational power) and legacy models/tools must be integrated, as well. Currently, the standard twin setup for integrated energy systems is, therefore, co-simulation-based, whose sub-models appear as black box to the simulation master, even if they are white-box internally.

Both options are computationally expensive. The co-simulation method requires assets to be shared among simulators (i.e., a power station is part of the thermal model and the electric model), so variables and states need to be shared and exchanged. The monolithic method, however, generates exceptionally large equation systems that must be solved. This does not scale well, either. Still, combined, cross-sectoral models are needed, and both methods are subject to improvement and innovation, as we speak.

Section 2 - Purpose and Modelling Requirements of Digital Twins

When considering the development, design or adoption of a digital twin of any system, and particularly of complex integrated energy systems (IES), it is essential to identify its main purpose and from that, the DT's key essential features and requirements. For IESs, there are different domains of interest that could be considered, which closely link to the specific application purposes and use cases of a DT (see also Section 4). Some examples are provided below.

Geographical scale and resolution

The geographical scale of an IES DT could vary from the national to regional level, city or town level, industrial hubs such as industry parks, neighborhood level and finally, even down to the building level. Accordingly, relevant use cases could range from operation and planning of national-level energy infrastructure (electrical and gas networks) to modelling the operation of an individual building. For studies of large-scale network infrastructure, the geographical resolution that is sought is normally at the level of grid supply points and transmission or distribution interface substations. However, hierarchical schemes could also be developed that co-simulate both transmission and distribution networks, again with appropriate levels of details, especially in the context of emerging TSO-DSO technical and market interfaces. This is particularly relevant for the electricity network, especially with more and more distributed energy resources (DER), while less resolution is usually required for gas network studies. At the building level, typical and important IES applications are quantifying and extracting the potential flexibility that could be provided, such as for demand response purposes by smart appliances, virtual storage in the building fabric while optimally controlling air conditioning plants, etc. Noteworthy DT use cases that are emerging at intermediate levels of the geographical scale and resolution are associated with building digital twins of low voltage (LV) networks for fast connection assessment and capacity allocation to DER.

Temporal scale and resolution

In terms of temporal scale and resolution, DT use cases may vary from studies concerning power/energy system and market operation, to long-term, multi-year planning and investment studies (typically between ten and thirty years ahead). For example, typical steady-state problems such as power flow, contingency analysis and (for market operation purposes) security-constrained economic dispatch and unit commitment, may typically be run with time resolution of between five minutes and one hour, with temporal windows of a few time intervals ahead, to day-ahead and week-ahead (e.g., for renewables forecast and pumped hydro scheduling). A much shorter time resolution, down to microseconds, may be needed for power system electro-magnetic transient (EMT) studies, which may be simulated over a time window of up to a few minutes. On the other hand, hourly or multi-hourly resolutions are often assumed for DT that have been built for electricity infrastructure planning studies. In the context of other energy infrastructures, for example the gas network, the temporal resolution may be associated with both the objective of the study and geographical scale, given the much higher time constants involved. For example, for planning purposes, a relatively small gas distribution network may be simulated with daily resolution. On the other hand, a large gas transmission network should be simulated with resolution of down to minutes, if the purpose of the study was to optimize linepack storage operation in managing gas network flexibility in its interaction with the electricity system.

Sector and infrastructure scope

"Sector-coupling" studies are emerging as a key area for IES, including the use of simulations of sectors such as electricity, heat, gas, hydrogen, water, and transport, as well as the relevant infrastructure. For instance, in the context of decarbonizing future fuels for heating, transport, and other applications, modelling and simulations of integrated electricity-gas-hydrogen systems and networks are being considered to inform strategic industry and policy plans around the world. For example, Figure 3 shows the DT of the integrated electricity and gas transmission network for the east coast of Australia. This has been developed to assess the real-time operational impact of the injection of green hydrogen that is produced from renewables into an increasingly decarbonized gas network. This analysis supports the relevant development of integrated electricity-gas-hydrogen markets, infrastructure planning, and policy development. Different types of IES DT modelling might also be required to study the decarbonization of the transport sector including. For example, models for simulation of the road transport behavior of different types of vehicles, or the interaction of these with relevant recharging (electricity) or refueling (hydrogen) stations, along with the upstream infrastructures of said stations. It is clear that such a multi-sector DT presents a great degree of complexity in terms of the interactions of the different models and modules, the adoption of appropriate geographical, temporal, and modelling resolution, availability and efficiency of data exchange across modules, as is discussed in the other sections of this paper.

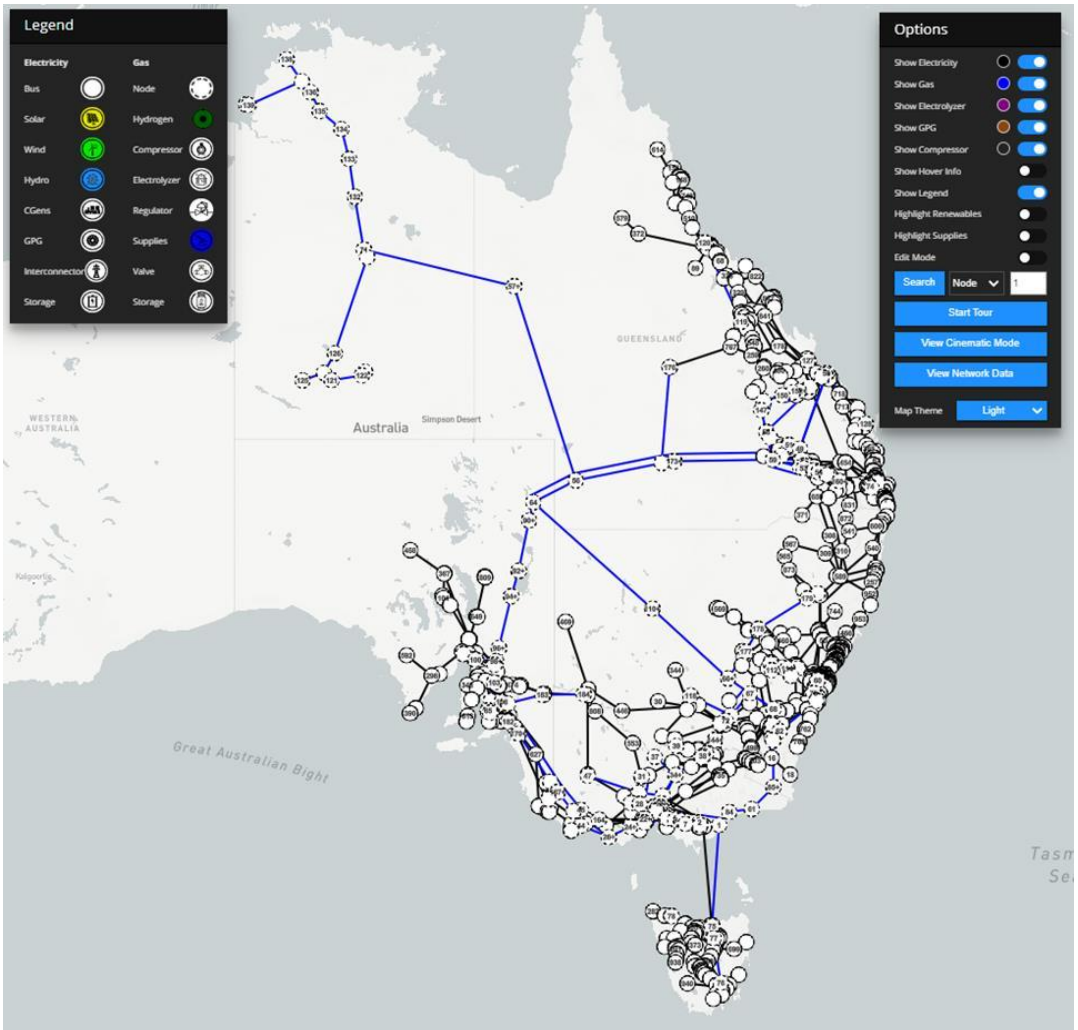

Figure 3: Integrated electricity-gas-hydrogen digital twin of the Australian east coast energy system, under development by the University of Melbourne.

Functional layers

Different functional layers may be considered in the development of a DT, depending on its purpose. For example, when focusing on the electricity network, a DT of only the physical infrastructure may be appropriate to study problems such as network capacity availability, security, and operational safety margins for different systems and components. A layer of the control infrastructure and architecture could then be added to better identify the potential response of the system to contingencies, such as introducing dynamic system simulations and including the impact of different control systems and layers, on top of steady-state analysis. This is also relevant to a more general incorporation of the information and communication (ICT) infrastructure and virtual layer on top of the physical energy infrastructure layer, as well as relevant studies pertaining to cyber-security and the study of cascaded impacts of disturbances across the cyber-physical system. Similarly, an energy market layer could be included to identify the starting point of steady-state and dynamic/control simulations. Again, as in the case of multi-sectorial integration of different DT modules, the inclusion of different functional layers requires a thorough and careful design of the modelling of each layer and their interactions.

General modelling functionalities and resolution

An important question that needs to be addressed across all scopes and domains of a DT is that of the functionalities and resolution of the modelling itself. For example, a DT of a power system might be developed with different levels of detail for simulation studies that include steady-state power flows, or root-mean-square (RMS) or EMT-type dynamic studies that are suitable to study new renewable generation connections. In this regard, the need for performing high-resolution dynamic studies with underlying EMT modelling has recently come to the forefront in several discussions across system operators around the globe, with the aim of quantifying the stability impact of deep penetration of variable renewables interfaced through inverter-based resources, particularly in low-inertia and weak grids.

Most notably, the Australian Energy Market Operator (AEMO) has developed an EMT simulator of the Australian east coast transmission interconnection. Given the computing complexity associated with it the question, therefore, arises as to when such detailed simulations are *actually* needed, and whether an “adaptable” digital twin of the system should be developed that allows (ideally automatic) for switching between models as needed, shifting from steady-state to RMS dynamic, and then EMT dynamic details, as required. While such modelling resolution is fundamentally driven by the presence of faster or slower dynamics and the relevant time constants involved in the different technologies and infrastructures, in reality the choice is also driven by what type of stability and general power system phenomena need to be specifically addressed and under which conditions. In other words, given the tradeoff with computational burden, it may not be desirable to have the highest possible resolution modelling running for any kind of study.

Linked to this is also the issue of data inputs: The more complex the DT of a system is, the more data it will require. It may be that the uncertainty or proprietary nature of much of this data (control schemes of renewable technologies) would produce inaccurate and/or misleading results, thus possibly defeating the purpose of the high-resolution model in the first place. Similar considerations, about both the time constants involved and the purpose of the study, also apply to other infrastructures such as heat and gas systems. For example, EMT modelling might always be required for connection studies of inverter-based resources in weak networks, while a DT should be able to identify the requirements to switch across level of complexity and possibly down to RMS dynamic simulations when making connections to strong networks. For IES, then, five-minute to hourly simulations might be suitable for steady-state power flow or optimal power flow studies, followed by hourly simulation for heat network studies, daily simulations for gas/hydrogen steady-state network studies, and so forth.

As digital twins may also be particularly useful in assessing the robustness of IES operation and planning against various degrees of uncertainties, it is important that the appropriate level of modelling complexity is adopted for such studies. In particular, modelling resolutions may be relaxed when, for example, studying sensitivities to different stochastic parameters in steady-state and dynamic system operations. This may include Monte Carlo-based time-ahead scheduling and stochastic simulations, considering different uncertain parameter inputs, or for long-term planning with scenario studies, in order to assess the robustness of investment solutions. In these cases, relatively faster modelling approaches for case screening and sensitivity assessment should be used, and it is desirable that a DT has such flexibility.

Section 3 – Co-Simulation as a Tool for Digital Twin Implementation

Digital twins often, but not always, model systems that cross traditional analysis and simulation and domain boundaries. Some digital twins can be assembled using a single simulation tool or a set of simulation tools that operate in series with the output of one affecting the input of the next, forming a linear simulation chain. Though this linear data exchange is possible for some DTs, there are others where it is too simplistic or not possible. In such cases, it is common for two or more simulation tools to share system boundaries with the outputs of one forming the inputs of the other and vice versa. To model these multiple domains, a co-simulation platform is often used to tie the simulators together, providing them the ability to exchange data during runtime and thus influence each other's operation. By providing data dynamically and enabling interaction between the simulation tools, larger and more complex DTs can more quickly be modeled and simulated using existing, best-in-class domain-specific tools without developing and validating a custom, integrated simulation tool.

For example, a digital twin could be designed to evaluate the impacts of distributed solar generation and an EV (electric vehicle) charge management scheme on total power system demand, as shown in Figure 4. Such a DT could, to decrease model computation and update time, make simplifying assumptions such as constant power residential loads and simple, voltage-independent inverter controls for the EV charger and distributed solar generation. In such a case, the distribution system simulator's power flow solution can be found very quickly because of this voltage independence.

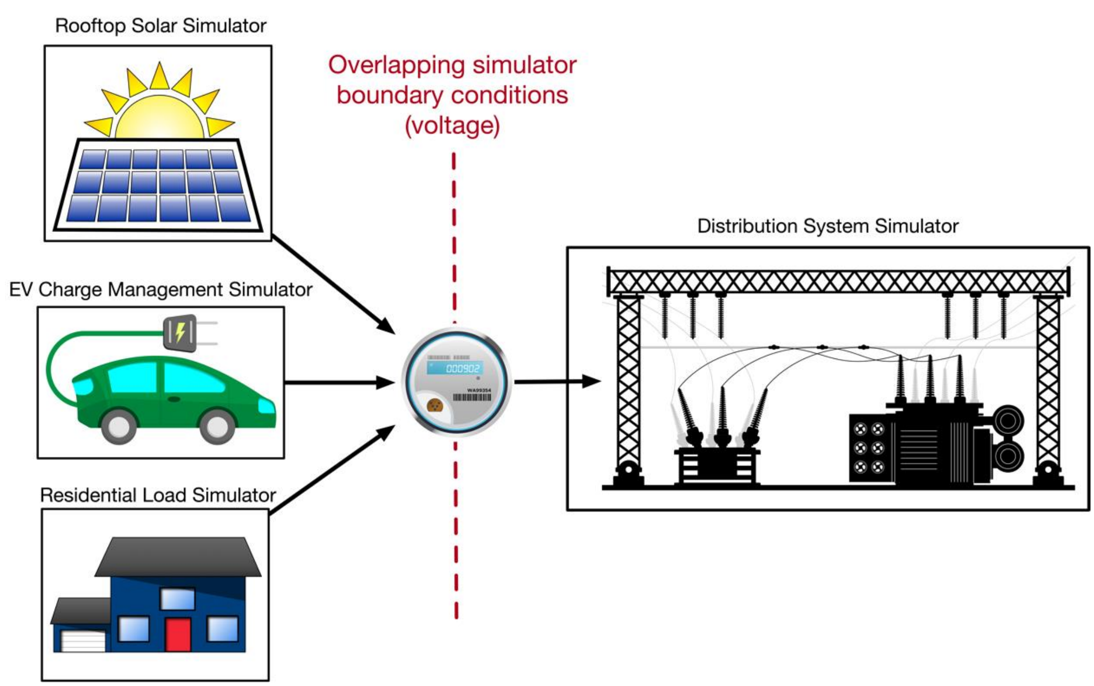

Figure 4: Linear digital twin simulation architecture when no voltage dependency is assumed

A more realistic digital twin, though, includes voltage-dependency which creates overlapping boundary conditions between the distribution system simulator and the behind-the-meter assets (Figure 5). The solution of the powerflow affects the voltage at the meter, affecting the operating state of the customer assets which, in turn, affects the solution of the powerflow. For these circular dependencies in a DT, a co-simulation platform is helpful in maintaining a consistent state across simulator tools.

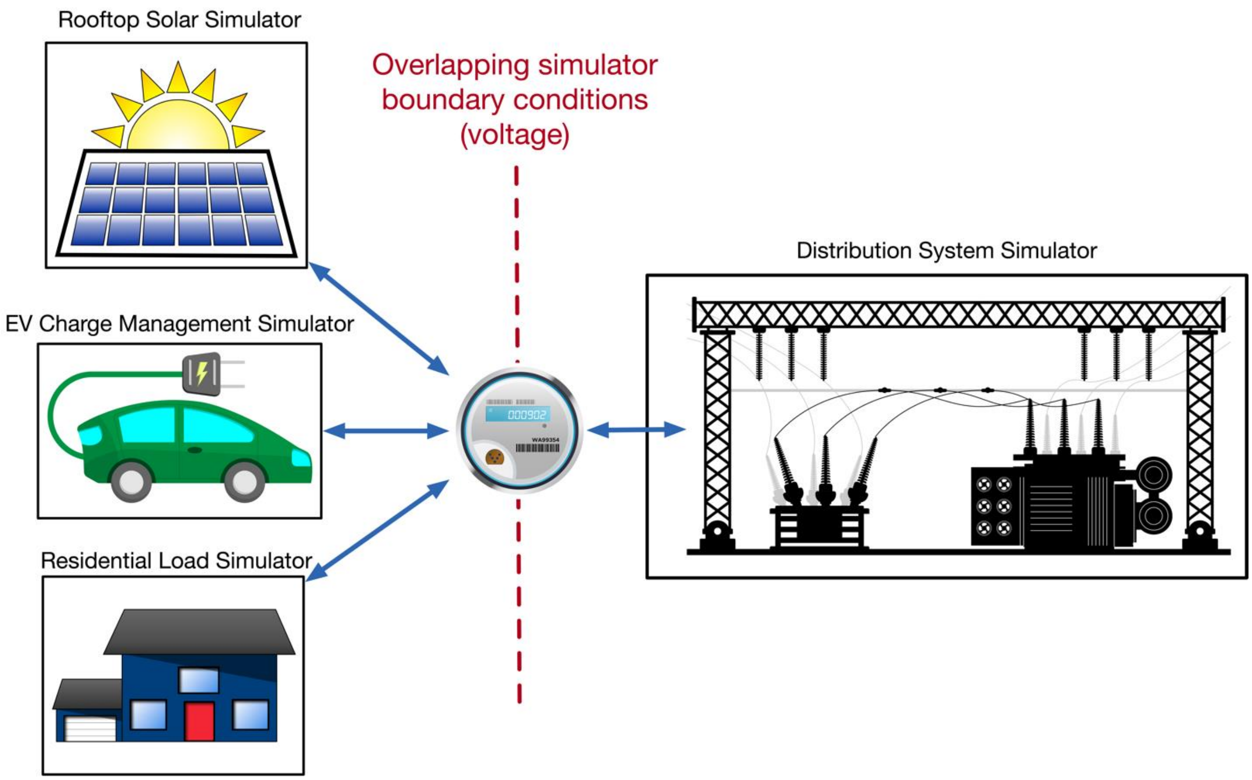

Figure 5: Circular digital twin simulation architecture with voltage dependent components where data flow is bi-directional

Co-simulation platforms generally have two primary, inter-related functions: synchronizing the individual simulation tools in simulated time and facilitating data exchanges at appropriate points of time. Time synchronization is necessary to ensure that the data being exchanged between simulation tools has the correct temporal context. Without regulation of the simulated time, individual tools could send data from the simulated future or past and the receiving tool would not necessarily handle it appropriately.

Though it is simply stated, managing these two functions can be challenging. Individual simulation tools may have different concepts of time (continuous vs discrete, for example), may be written in different languages and thus have different fundamental data types, and may not save previous model states to allow resolving and iterating at a particular time step. Furthermore, for a given DT there may be a need to sequence the data-exchange between particular software tools in a particular way, run in a wide variety of computing environments from laptops to high-performance computing clusters or even run in a disparate networking environment that includes multiple institutions and/or cloud computing resources. A robust co-simulation platform would be able to facilitate the creation of a digital twin despite these infrastructure challenges.

There are a number of candidate co-simulation platforms that could be used to create a digital twin. One of the earliest such platforms is the Hierarchical Language Architecture (HLA) pioneered by the United States Department of Defense and later codified as IEEE Standard 1516 in 2000. Over the past decade or two, the Functional Mock-up Interface has been developed, originally to allow a system integrator to evaluate components from a variety of vendors without exposing propriety details of the modeled components. It has since grown to facilitate more general co-simulation needs. More recently, other generic co-simulation platforms have been developed such as HELICS and Mosaik, both of which were designed with the energy sector in mind but can support a wide variety of simulation tools.

Creating integration of an existing simulation tool and a co-simulation platform typically happens in one of two ways. For open-source tools, co-simulation platforms today typically produce a library with the platform's APIs (application program interface) in a variety of popular programming languages (i.e., Python, C++, Java, MATLAB) and the source code of the simulation tool can be edited to include these API calls at the appropriate point in the tool's execution (Figure 6). For commercial tools where the codebase is not publicly available or when a tool is not written in a language with support from the co-simulation platform, integration relies on the tool developers producing an appropriately featured API. If such an API exists, a wrapper program written in a language supported by both the tool's API and the co-simulation platform's API can be written. This wrapper coordinates the execution of the tool and the co-simulation platform, utilizing the API calls of both.

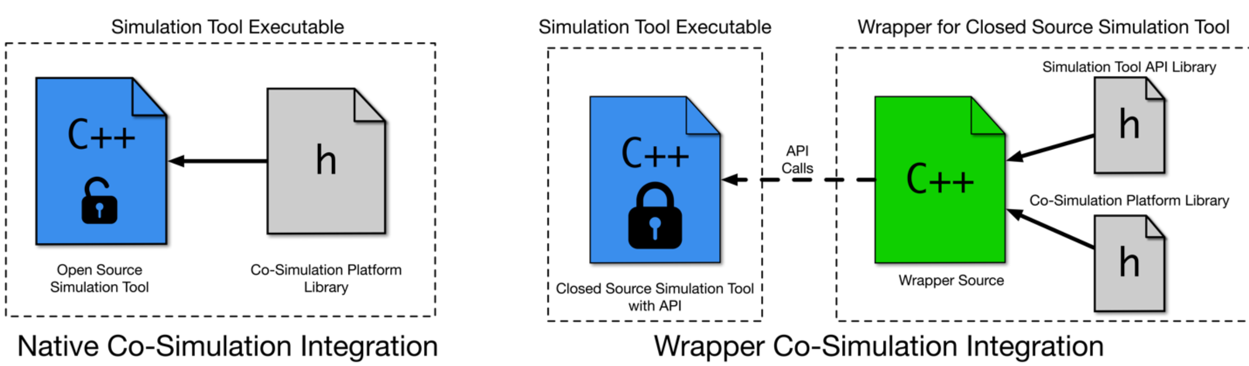

Figure 6: Native Co-Simulation Integration (left) and Wrapper Co-Simulation Integration (right)

Section 4 – Applications

Digital twins have been in use in various domains of science and technology. The concept first appeared in academia, as a proposal for the information and modeling technology to aid the product life cycle management. The first digital twin, not yet named as such, contained the full description of a physical product in virtual space, tracking it through its entire lifecycle from creation, through operation to the disposal. The twin was linked by the information flow from and to the physical product, keeping it updated and allowing its decisions to be implemented. The concept of linking a physical system to its exact digital replica was simple yet powerful. NASA (National Air and Space Administration) was the first to use it in their aircraft design and space exploration missions. The digital twins later evolved and deployed in various industries, while their specialization and differentiation became more prominent. Today we have the digital twin as a service concept or specialized digital twins for experimentation, among others.

Across different industries, DTs have found their place in supporting and even taking over many tasks (Figure 7). Planned and predictive maintenance of energy infrastructures and industrial processes have gained traction since digital twins provide necessary lifecycle tracking and advice to engineers and maintenance crews. The digital twins are used within these processes for prognostic health management and more widely, for the design and operations of industry 4.0 in which the focus of digital twins is on improving performance of industrial processes. Within the automotive industry, digital twins have been used in the vehicle design phase and for driving analytics and decision support. In the design phase, they are particularly suitable for crosschecking suitability, interoperability, and performance of subsystems within vehicle, speeding up the design and validating design choices. Similar is the situation in the aerospace industry. Within the healthcare sector digital twins have found their application in surgical interventions as well as predicting population behavior. The digital twins have also been used by governments to aid policy design and to train the new world leaders in making relevant policies.

In the field of power and energy systems, the digital twins can be used for both, design and operation of these systems (see further reading). Besides the improvement of the performance and validation of different design options, at the design phase, they could help reduce the search space of the feasible and optimal infrastructure design options. In addition to what has been said about their use for predictive maintenance, they also help visualize the flows of energy, and the potential of collective actions by all connected stakeholders. Alternatively, they are valuable for planning and operating of the energy consumption sector, helping to devise demand response policies and energy bill management practices necessary for energy cost reduction in households, neighborhoods, office buildings and business parks.

In the energy infrastructure sector, the digital twins have been proposed to augment decision making in the control room, providing necessary support to the grid operators by performing the online grid analysis and suggesting actions for grid reconfiguration (see further reading). The applications range from real-time power flow monitoring to real-time load shedding support. The advantage of digital twins in the control room also includes the possibility to embed trusted third-party and encrypted sub-models into decision support systems, leading to higher accuracy of the system, while preserving trade secrets of vendors. Such an approach increases opportunities for verification of control and management algorithms while improving interoperability of the technologies and leading further to mass deployment. The twins can be used for training and education of the grid and industrial system operators.

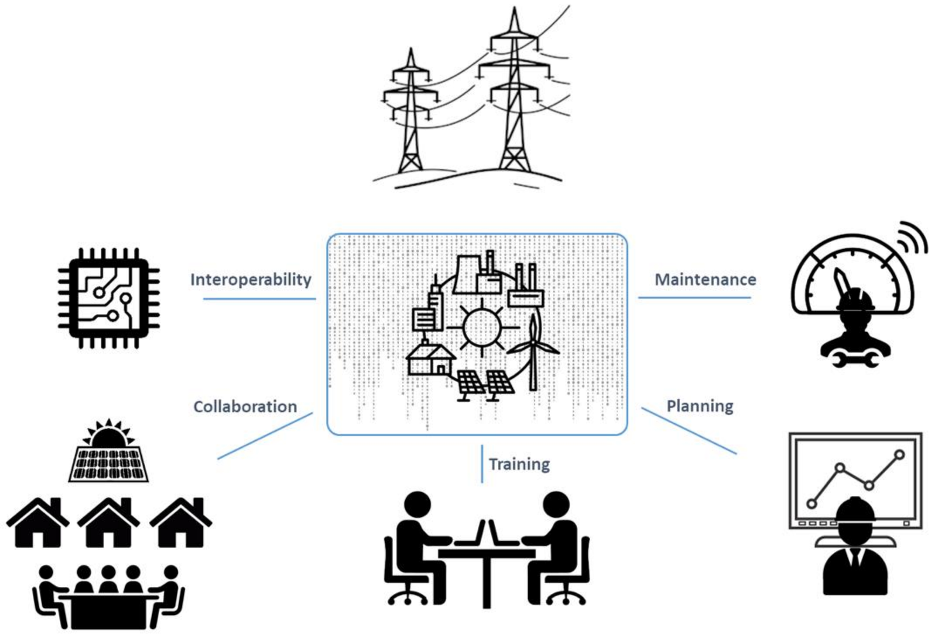

Figure 7: Applications of integrated energy system digital twins

Section 5 – Conclusion and Outlook

Integrated energy systems show all features that usually point towards using digital twins: complex behavior, trans-disciplinary nature, and too expensive for experimental mock-ups. The current implementations are usually based on co-simulation, which can suffer from mediocre numerical performance and complex handling of scenarios. There are promising developments with universal modeling languages such as Modelica, but scalability and simulation performance can still be an issue. Beside the technical limitations, it is mainly the workflow of creating, validating, operating, and updating such twins that requires further attention. Right now, most of the real-world examples in the energy domain are either academic or require intensive support by the platform creator to work with it.

The potential of digital twins in integrated energy systems is, however, great: synergies and risks that are invisible to the plain eye and can – due to the complexity of the system – also not be identified in an analytical way, can be tracked down with a digital twin. Be it in planning or in operations, such twins are a powerful tool for this complex business. They are a great platform to meet in projects: experts from different domains can join forces, basing their discussion on hard facts coming from a shared trans-disciplinary platform: integrated energy system digital twins. Their success will depend on standardized interfaces, standardized procedures for testing and validation, and – highly likely machine-learning-based – support for model and parameter identification and updating.

For Further Reading

- I. Saedi, S. Mhanna, P. Mancarella, “Integrated Electricity and Gas System Modelling with Hydrogen Injections and Gas Composition Tracking”, Applied Energy, 2021, doi: https://doi.org/10.1016/j.apenergy.2021.117598
- P. Palensky, et al, “Applied Cosimulation of intelligent power systems: Implementation, usage, and examples”, IEEE Industrial Electronics Magazine Vol. 11 (2), 2017, doi: http://dx.doi.org/10.1109/MIE.2017.2671198
- V. Stennikov, E. Barakhtenko, D. Sokolov, and G Mayorov, “Principles of Building Digital Twins to Design Integrated Energy Systems,” Computation 2022, 10, 222. https://doi.org/10.3390/computation10120222
- C. Brosinsky, D. Westermann, and R. Krebs, “Recent and prospective developments in power system control centers: Adapting the digital twin technology for application in power system control centers,” 2018 IEEE International Energy Conference (ENERGYCON), pp. 1–6.
- M. Zhou, J. Yan, and D. Feng, “Digital twin framework and its application to power grid online analysis,” CSEE Journal of Power and Energy Systems, vol. 5, no. 3, pp. 391–398, 2019
- J Naughton, H Wang, M Cantoni, and P. Mancarella, “Co-optimizing virtual power plant services under uncertainty: A robust scheduling and receding horizon dispatch approach,” IEEE Transactions on Power Systems 36 (5), 3960-3972

Biographies

Peter Palensky and Milos Cvetkovic are with TU Delft, the Netherlands.

Pierluigi Mancarella is with the University of Melbourne, Australia.

Trevor Hardy is with Pacific Northwest National Laboratory, United States.