MIS Quarterly Executive

Volume 24 | Issue 3 Article 6

September 2025

Transforming Energy Management with an AI-Enabled Digital Twin

Hadi Ghanbari

Petter Nissinen

MIS Quarterly Executive

Transforming Energy Management with an AI-Enabled Digital Twin

Digital twins (DTs) are increasingly adopted by organizations across various sectors. We report on how one of Europe's largest district heating providers implemented an AI-assisted DT in pursuit of energy efficiency and sustainability. The solution enabled the company to modernize its complex cyber-physical system (CPS) and tap into its rich data capabilities to gain a comprehensive real-time representation of the entire district heating network. Reflecting on the case study, we provide six recommendations for executives in other domains aiming to implement DTs.[1,2]

Hadi Ghanbari
Aalto University (Finland)

Petter Nissinen
Aalto University (Finland)

Digital Twins Are Transforming Cyber-Physical Systems

Organizations across sectors have shown increasing interest in digital twin (DT) technology. A DT is a software module that replicates physical objects and provides a comprehensive real-time understanding of organizational processes.[3,4] Initially used for small-scale objects, DTs are now applied to more complex entities, such as organizations, cities and large-scale cyber-physical systems (CPSs).[5] As the name implies, CPSs comprise cyber and physical components, where numerous sensors and embedded computers are used to monitor and control physical components and processes.[6] Smart energy grids, robotics systems, autonomous driving systems and district heating systems are examples of such systems. Due to their unique characteristics, DTs play a crucial role in transforming complex and large-scale CPSs across various domains.[7,8] (The text box provides a fuller description of digital twins and compares them with other digital models.)

While discussions around DTs are gaining momentum in information systems research, the DT literature is in its infancy and has not shed much light on DT implementation and

1 Mary Lacity is the senior accepting editor for this article.

2 We are grateful to Mary Lacity for her invaluable guidance and recommendations that helped shape and enhance this article. We also thank the anonymous reviewers and the research participants for their valuable insights.

3 Meske, C., Osmundsen, K. S. and Junglas, I. "Designing and Implementing Digital Twins in the Energy Grid Sector," *MIS Quarterly Executive* (20:3), September 2021, pp. 183-198.

4 Wurm, B., Becker, M., Pentland, B. T., Lyytinen, K., Weber, B., Grisold, T., Mendling, J. and Kremser, W. "Digital Twins of Organizations: A Socio-Technical View on Challenges and Opportunities for Future Research," *Communications of the Association for Information Systems* (52:1), June 2023, pp. 552-565.

5 Ibid.

6 Boyes, H. and Watson, T. "Digital Twins: An Analysis Framework and Open Issues," *Computers in Industry* (143), December 2022, pp. 1-19.

7 van der Aalst, W. M., Hinz, O. and Weinhardt, C. "Resilient Digital Twins: Organizations Need to Prepare for the Unexpected," *Business & Information Systems Engineering* (63:6), September 2021, pp. 615-619.

8 Grieves, M. W. "Virtually Intelligent Product Systems: Digital and Physical Twins," in Flumerfelt, S., Schwartz, K. G., Marvis, D. and Briceno, S. (eds.) *Complex Systems Engineering: Theory and Practice*, American Institute of Aeronautics and Astronautics, Online, 2019, pp. 175-200.

DOI: 10.17705/2msqe.00120

Background on Digital Twins and Other Digital Models

A digital twin is a software module that leverages a live flow of data and various technologies, including simulations, the Internet of Things and AI, to monitor, forecast and optimize the behavior of physical entities.[9,10] By offering a virtual environment for scenario-based experiments, DTs enable organizations to assess the status of their real-life counterparts and forecast the consequences of deliberate or accidental changes.[11,12]

Despite DTs' similarities with other related concepts, such as digital models and digital shadows, there are differences[13] (see Figure 1). Digital models have no automated data exchange between the physical object and its digital counterpart. A digital shadow involves an automated but one-directional flow from the physical object to its digital representation. In a DT, data flows automatically in two directions between the physical entity and its digital counterpart.

Figure 1: Digital Model, Digital Shadow and Digital Twin

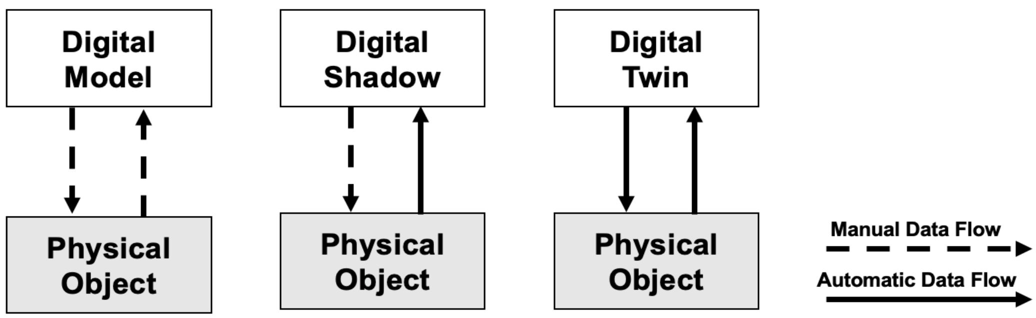

A DT has three defining characteristics: 1) twinning virtual and physical counterparts; 2) measuring the states of the counterparts; and 3) altering the states of the counterparts. As such, a DT mimics the behavior of a real-life entity in real time, allowing ongoing monitoring, performance analysis, scenario testing and preemptive problem detection in the physical object or system. Moreover, a DT's real-time nature, whereby the digital counterpart operates concurrently with the physical entity, ensures timely updates and impacts on the surrounding environment. Hence, bidirectionality is necessary in a DT, where interaction between the digital and physical counterparts occurs in both directions.

use in organizations. Further case studies are needed to answer the question: *How are organizations implementing and using DTs across sectors to address operational inefficiencies in complex cyber-physical systems?* To address this question, we examined the implementation of an AI-assisted DT solution for a leading Nordic energy company—referred to anonymously as "EnergyCo"—that operates one of the largest district heating networks in Europe. We selected this case because it illustrates a successful implementation of a DT in a highly complex CPS environment, where billions of data points are gathered from numerous internal systems and external sources.

9 Schmidt, D. "Digitalization of District Heating: Transforming Heat Networks for a Sustainable Future," *Proceedings of Energy Informatics Academy Conference*, Bali, Indonesia, October 23-25, 2024, pp. 351-356.
10 Wurm, B., Becker, M., Pentland, B. T., Lyytinen, K., Weber, B., Grisold, T., Mendling, J. and Kremser, W., op. cit., June 2023.
11 Djebali, S., Guerard, G. and Taleb, I. "Survey and Insights on Digital Twins Design and Smart Grid's Applications," *Future Generation Computer Systems* (153:C), November 2023, pp. 234-248.
12 Garske, B., Holz, W. and Ekardt, F. "Digital Twins in Sustainable Transition: Exploring the Role of EU Data Governance," *Frontiers in Research Metrics and Analytics* (9), March 2024, pp. 1-13.
13 For more details on these differences, see van der Aalst, W. M., Hinz, O. and Weinhardt, C. "Resilient Digital Twins: Organizations Need to Prepare for the Unexpected," *Business & Information Systems Engineering* (63:6), September 2021, pp. 615-619.

We collected data through interviews with key experts involved in the project, field observations made by the second author and publicly available information about EnergyCo and its DT vendor (the Appendix describes our research method). In this article, we show that the primary benefits EnergyCo gained were a comprehensive end-to-end and real-time representation of its entire district heating grid, as well as integrated optimization solutions and simulation capabilities. These benefits have enabled the company to improve operational efficiency and sustainable energy management, primarily through real-time optimization of the district heating network and secondarily through collaborative decision-making and enhanced network maintenance. The case yields six recommendations for other executives considering the adoption of DTs across industrial domains.

Implementing a Digital Twin in EnergyCo

EnergyCo is a major Nordic energy company operating one of the largest district heating networks in Europe. The network produces 7.2 terawatt hours (TWh) of energy annually, which is sufficient to power approximately 650,000 average U.S. homes for a year and is 20% of Finland's annual district heating consumption. The network mainly relies on burning a diverse mix of fossil fuels and emits a significant amount of carbon dioxide. This highly complex CPS, comprising numerous smaller systems and physical components, is used to distribute heat to thousands of buildings through over 2,800 kilometers of hot water pipes. EnergyCo gathers billions of data points from thousands of metering points, including supply and return temperatures, hydraulic information, asset states and consumer behavior, making it challenging for the company to comprehend and operate the network with traditional control systems. There are five main reasons for the challenges EnergyCo faces.

First, due to thermo-hydraulic delays in district heating, it takes considerable time for temperature changes initiated on the production side to reach consumers located kilometers away. Second, energy production previously relied on stable, easily predictable sources, but the current market is characterized by constant price fluctuations. Third, centralized district heating networks that rely on a few large fossil fuel-based plants are being replaced by more complex decentralized networks comprising numerous smaller, diverse heat sources, many of which are electrically powered. Fourth, managing the district heating network is complicated by changing weather conditions and fluctuating fuel prices. Fifth, because of its large size, EnergyCo's district heating unit is organizationally split into separate divisions, leading to inefficiencies: the customer team is responsible for ensuring adequate heat distribution and optimizing performance; the network team focuses on the physical infrastructure, performing thermo-hydraulic calculations to predict network behavior, heat distribution and potential bottlenecks and leaks, and the production team decides how much heat to produce, considering factors such as sustainability, fuel sources and insights from the trading team on power markets. However, managing the district heating network cannot be done in isolation and requires a comprehensive understanding of the entire network as well as coordination between the various divisions.

To deal with these challenges, EnergyCo's district heating managers decided to implement a cloud-based DT solution developed by "DigiTwinCo," (again referred to anonymously). The implementation timeline is summarized in Figure 2 and described in detail below.

Before 2021: Paving the Way for Implementing a Digital Twin

As a pioneer in green transformation and the transition to clean energy, EnergyCo has set three strategic priorities:

1. To become carbon-neutral by 2030 and phase out the burning of fossil fuels by 2040
2. To mitigate increasing electricity and fuel price fluctuations by maintaining exceptional operational flexibility
3. To achieve long-term profitability, ensuring that it can sustainably fund further investments in its green transition.

To achieve these objectives, EnergyCo has made large investments in the digital transformation of its district heating network

Figure 2: The Digital Twin Implementation Timeline

Before 2021	Phase 1 2021	Phase 2 2021-2023	Phase 3 2024	After 2024
EnergyCo invested heavily in advanced data analytics and AI. Several data teams assisted units in various areas, from process optimization to customer analysis.	EnergyCo tasked DigiTwinCo with developing a digital model of a limited area in the DH network. The model proved accurate when compared to EnergyCo's models.	EnergyCo tasked DigiTwinCo with developing a digital model of the entire DH network, combining a physical model with AI-enhanced learning.	The implementation of the AI-enabled DT solution was completed. The DT went live and enabled real-time data flow between the DH network and the DT solution.	The DT will be enhanced by developing real-time optimization features for different parts of the DH system and the data pipeline from the consumer side.

and in developing the necessary sociotechnical capabilities, including data analytics and AI capabilities. These capabilities have been used in multiple areas, from process optimization to customer analysis. The company's "We Love Data" slogan encapsulates its data strategy, emphasizing that the whole organization uses data. For instance, EnergyCo's data team uses "tribes," or representatives from the Business Intelligence and Analytics, Advanced Analytics, Data Engineering and Data Governance teams, to assist business units with data-related issues.

In addition, to foster its position and prominence as a leading European energy company, EnergyCo has established a corporate venture capital unit that invests in European startups in the energy, decarbonization and climate technology sectors. For instance, the unit has invested in dozens of European startups developing forward-looking solutions such as AI-enabled energy sharing and trading platforms, cloud-based software solutions for management of decentralized energy supply systems and even earth observation solutions. One of these companies is DigiTwinCo, a European startup with over 100 employees, which was founded in 2019 and has developed a cloud-based DT platform for district heating providers. So far, EnergyCo and a few other venture capitalists have invested approximately €40 million ($45.6 million)[14] in DigiTwinCo through three funding rounds: seed funding in 2020, Series A funding in 2022 and Series B funding in 2025.

These investments provided EnergyCo with the necessary sociotechnical capabilities to develop new data-driven digital solutions, paving the way for implementing a DT solution to enhance the company's operations.

Phase 1 (2021): Pilot Implementation and Demonstration

The DT project started in 2021 when EnergyCo's district heating team decided to test the performance of DigiTwinCo's DT platform. DigiTwinCo was tasked with modeling a limited grid area to demonstrate the proof of concept of its DT solution. The district heating team intentionally selected a specific neighborhood that was difficult to model because it was far from heat production sites. Comparing DigiTwinCo's model with EnergyCo's data from the area proved the accuracy of the model.

Even though EnergyCo's corporate venture capital unit is one of DigiTwinCo's investors, most EnergyCo employees we interviewed said they had initially approached the project with huge skepticism.

> *"At first, I was what you'd call a skeptic, but now I'm a firm believer and spreading the word. The thing is, many others who've offered similar solutions have only done some basic modeling without truly understanding or accounting for thermal and fluid dynamics."* Energy Systems Manager, EnergyCo

However, the skepticism slowly turned into optimism and enthusiasm after the success

14 Currency conversion as of August 2025.

of the initial proof-of-concept demonstration and subsequent phases. Most of our EnergyCo interviewees considered the cloud-based solution beneficial and noted that developing an internal DT solution would have been more expensive and less efficient.

> *"The amount of work required for us to develop such a system ourselves would have been of a completely different magnitude. Essentially, we would have had to establish our own unit here to do this and hire a group of people specialized in thermodynamics modeling. It would not be feasible; it would be a risky investment."* Development Manager, EnergyCo

EnergyCo's development manager for digital twins said that large companies often attempt to develop such software solutions in-house because of an overemphasis on how unique their needs are and a desire for control over the process. Additionally, they often underestimate the complexity of these projects, especially if the company has no prior experience with similar projects. As such, a big benefit of buying a DT as a service is that the service provider can focus all its energy on specializing in DT technology and developing a DT solution, gaining synergies from multiple similar projects and a growing team of highly specialized experts, resulting in an arguably higher quality and a more sophisticated product.

> *"In practice, the quality is always better if you have a company that creates the product and develops it for multiple clients. Intuitively, it's cheaper overall to do the work, so to speak, once and utilize it everywhere. Of course, it requires customization and there are certain frictions in getting it implemented ... but in the big picture, this drawback is very small compared to building it yourself."* Vice President of Corporate Venture Capital, EnergyCo

Understandably, technology providers like DigiTwinCo that develop niche innovations must protect their intellectual property and system security to ensure business continuity. Providing sufficient access to the client while protecting the core technology inside the DT solution is therefore a delicate balance for DigiTwinCo. According to interviewees from EnergyCo, this is not a new issue, as the company has previously used cloud-based solutions. However, the interviewees suggested that opening the inner workings of the DT solution as much as possible enables more customization, which is beneficial for both the client and the DT provider. In addition, the delivery of cloud-based DT solutions introduces challenges related to managing client expectations, validating deliverables, project leadership, scheduling and resource allocation. According to the interviewees, one especially challenging area was integrating the DT with other programs and computational models, which required alignment and prioritization. To overcome these challenges and maintain secure partnerships, the interviewees suggested that a long cooperation history, cases of overcoming obstacles together, regular testing, validation and certifications are important.

Phase 2 (2021-2023): Developing a Digital Model

Once EnergyCo's district heating team was satisfied with the Phase 1 proof of concept, it tasked DigiTwinCo with developing an offline digital model of the entire district heating network. DigiTwinCo set out to accurately model all components of EnergyCo's district heating network, including pipes, plants, substations, connection points and areas of use, combining a physical model with existing sensor data and AI-enhanced learning. Meanwhile, EnergyCo focused on addressing two challenges.

The first challenge was automating data integration from various sources within the district heating network into a format readable by the DT while ensuring data protection, privacy and resilience. The volume of data was huge and heterogeneous; it was gathered from approximately 13,000 metering points, with a variety of measurements taken hourly from each point over the years, so the number of individual observations was measured in billions. This data needed to flow automatically, securely and in the correct format to the cloud, and it was important that maintenance activities not disrupt data collection. The data also had to flow dynamically back from the DT to the

Figure 3: The Digital Twin Provides a Live Overview of the Entire District Heating Network

district heating grid. To address this challenge, EnergyCo's data team started developing a custom-made data integration module to allow efficient data streaming from various internal and external sources into a single secure data stream. Building this module was arduous but vital for the successful implementation of the DT.

The second challenge was to ensure employees' buy-in and address their legitimate concerns about the potential use and benefits of the DT solution. Although the DT project sparked interest and optimism among project leads and management, there was some initial resistance to AI-assisted decision-making and automation among operators, even before they understood the functions of the DT in detail. To address this issue, the interviewees emphasized the importance of effective communication and understanding conflicting views (e.g., on prioritizing optimization, short-term operational goals and long-term strategic goals), as well as listening to the most resistant or silent parties.

> *"If [the communication] would be done smarter from the beginning, I personally hypothesize that in the long run, we would have significant benefits. We would finish faster and with fewer conflicts and differences in views. Communication is, in my opinion, the alpha and omega, and setting goals and such might sound simple, but they are damn complex in the end."*
> Senior Vice President of District Heating, EnergyCo

Interviewees also highlighted "accepting imperfection" as a crucial part of the learning process. Tackling tasks incrementally and acknowledging that, though a solution may not be perfect initially, it is worth progressing if it improves the current situation, allows for continuous integration of the DT into routine work and sets the stage for larger-scale developments.

Phase 3 (2024): Implementing the Digital Twin

The DT went live in Spring 2024, enabling real-time data flow between the district heating network and the cloud-based DT system. The DT has a simple 3D graphical user interface, which enables operators to monitor the network's current state (see Figure 3). The interface provides intuitive, mostly blue-to-red-colored graphs and time ramps, enabling users to quickly analyze past events and forecasts. Users can

Figure 4: The Digital Twin Visualizes all Relevant Parameters of Each Network Component in Real Time

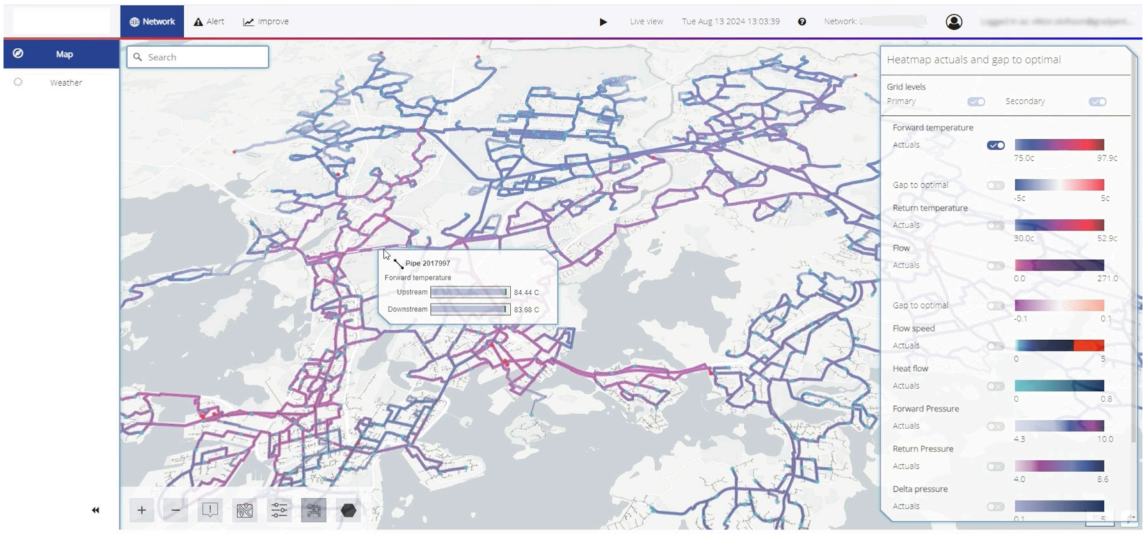

click on each part of the district heating network for more details. Overall, the user interface is designed to enable users to easily understand the complete network's current state and identify potential issues by prioritizing the most important parameters—a significant step up from traditional operating interfaces with numerous parameters.

The DT merges live sensor data from the grid and additional parameters, such as fuel prices and weather conditions, with an AI-assisted thermo-hydraulic model. This physics-based model uses machine learning to adapt to new situations and detect anomalies. It is trained using time-series data to accurately depict the physical network, including sources, hydraulic elements, pipelines and customers, and is constantly updated with live data to increase accuracy.

EnergyCo had previously used optimization software, typically to focus on a certain part of the district heating grid (e.g., production sites), providing a varying level of understanding of the system and optimization accuracy. In contrast, the DT provides a higher level of accuracy because it considers the entire end-to-end district heating network rather than just part of it. This is crucial in a complex CPS like EnergyCo's network, which comprises a growing number of consumers and an increasing number of decentralized heating sources and storage facilities, as well as other physical assets spread across different locations. Optimizing, for example, the temperatures of water heated from several sources in such an interconnected and complex system without understanding the whole system is very challenging.

Additionally, the thermo-hydraulic nature of district heating makes it even more necessary to understand the network fully, as it takes considerable time for temperature changes to reach consumers. By coupling a thermo-hydraulic engine with live data, the DT has enabled EnergyCo to visualize all relevant parameters in real time, including temperature drops in pipes (see Figure 4).

After 2024: Future Plans

EnergyCo's district heating unit aims to develop the DT's real-time optimization features for different parts of the network. In the future, the mental bandwidth of human operators will not be sufficient to comprehend all the different parameters to calibrate forecast models and align heat production with demand by adjusting settings manually. Future developments will address this issue: the AI-based DT will be able to recommend optimal temperatures, pressures and production scheduling, among other things. Eventually, the operator's role will change.

> *"As the system will outperform human capabilities, human intervention should focus on pushing the system further, training it and helping it tackle more challenging situations."* Senior Vice President of District Heating, EnergyCo

EnergyCo's growing emphasis is on developing the consumer side of the DT and district heating business model. However, this development has been delayed partly due to the high level of data protection required for data integration to comply with EU regulations. Once this obstacle is overcome, the DT will enable real-time usage optimization by providing a platform capable of managing thousands of consumers simultaneously with precision and efficiency. This means that operators will be able to identify customers whose consumption reduces grid efficiency and take proactive steps to mitigate the negative impact, either by changing consumers' consumption behaviors or directly limiting their usage.

> *"With 16,000 customers, each with unique preferences for flexibility in various situations, the potential for impactful management is significant. We're talking about more than just minor adjustments."* Senior Vice President of District Heating, EnergyCo

In the future, consumers might be encouraged to make choices about their energy consumption based on recommendations from the DT. Providing real-time information generated by the DT could help transform them from passive consumers into active participants in the energy management system.

> *"Pricing models need to be completely revamped to motivate customers [to make] large-scale adjustments. [This could] include all kinds of price signals [so] that you can also nudge the consumer in a way that's beneficial for total energy efficiency."* Product Market Lead, DigiTwinCo

In essence, this would mean that rather than making decisions reactively based on demand, proactive strategic decisions could be made based on sustainability and economic factors.

Outcomes and Strategic Values of EnergyCo's Digital Twin

EnergyCo implemented the DT as a part of its ongoing efforts to pursue its three key strategic priorities (moving toward carbon neutrality, dealing with fluctuating energy prices and achieving long-term flexibility). The company gained two primary benefits from the DT implementation that will help it achieve its strategic objectives. First, the DT solution provides the company with a comprehensive end-to-end, real-time representation of the entire district heating network. Second, the DT provides the district heating unit with integrated online and offline simulation capabilities and optimization solutions that take into account real-time data from the grid and various external parameters such as weather forecasts and fuel prices. Below, we describe how these two outcomes are instrumental in helping EnergyCo to achieve its strategic objectives.

Achieving Strategic Priority 1: Moving Toward Carbon Neutrality

Reducing and optimizing overall temperatures and pressures in the district heating grid is crucial to EnergyCo's pursuit of carbon neutrality. There are two ways that the DT can help EnergyCo achieve this objective.

1. The digital twin enables global management of the complex district heating network. Historically, EnergyCo used classic industrial control systems to optimize temperatures and pressures in the grid statically. However, this was cumbersome because such systems do not offer accurate, real-time insights for immediate decision-making. These systems become even less effective when intermittent thermal or electrical energy sources are integrated into the network. They fail to fully harness these sources' potential, resulting in inefficient energy use.

> *"The district heating network spans approximately 1,500 kilometers in [both] … directions. It's a complex network where heat can potentially travel through multiple routes, and the specific path it takes varies.*

  |

Currently, there's no comprehensive information available about how heat flows through the network. Operators rely on intuition to manage the network, but this approach inevitably leads to inefficiencies and heat loss." Senior Vice President of District Heating, EnergyCo

Compared to these traditional systems, which typically help with understanding just a certain part of the network, the DT provides a comprehensive representation of the entire network. Combining real-time data, prescriptive analytics and insights from a thermo-hydraulic model, the DT allows operators to capture the system's current state and analyze its future conditions (e.g., the network's hydraulic state, temperatures and production).

"That's really the foundation of the DT; it draws on that real-time data. But because it knows the topology of the district heating network, it can also fill in the gaps and estimate what is actually happening between different measurement points in the current real-time situation. In this way, it reduces the need to install sensors absolutely everywhere, as it helps to visualize what's happening both between and beyond the existing measurement points." AI Manager, EnergyCo

Moreover, compared to traditional modeling tools, the DT requires smaller datasets for training, takes less time and creates more detailed scenario models.

"The scenario models used previously [at EnergyCo] don't go into the level of detail achieved by the DT. Typically, they've been more about estimating broader network and main network interactions, focusing on aggregating district heating consumption needs rather than detailed analysis." Energy Systems Manager, EnergyCo

According to the interviewees, such a comprehensive and real-time understanding of the district heating network is crucial for optimizing the entire network and lowering overall temperatures and pressures.

2. The digital twin facilitates scenario-based decision-making. Another key DT feature is that it provides integrated simulation capabilities for scenario-based decision support, which is crucial for adjusting temperatures and pressures and forecasting their implications in the network. On the one hand, reducing the grid's forward temperature (i.e., when water is delivered to consumers) not only enables the integration of a wider array of renewable and low-temperature heat sources (e.g., geothermal energy) but also significantly reduces production costs and carbon dioxide emissions by regulating the outgoing water temperature at each production site. Lowering temperatures also reduces heat loss during transportation through the pipes, leading to direct environmental benefits.

"In addition to temperatures, the hydraulics of the system is also very important. Decreasing forward temperature normally increases flow, so operating pumps optimally is key." Solutions Director, DigiTwinCo

On the other hand, the growing complexity of the district heating system and changing flows require highly efficient temperature and pressure management to provide sufficient heat to even the farthest areas in the network, reduce hydraulic bottlenecks in the system and minimize pump costs. A major value of the DT is the ability to comprehend the physics of EnergyCo's complex district heating network and forecast heat demands at different times of the day and year.

"It's about thermodynamics and understanding how the fluid flows through the network. [The DT] also needs to account for various phenomena like turbulence, friction and other effects that occur inside the pipes." AI Manager, EnergyCo

Using real-time data, together with its integrated simulation capabilities, the DT allows operators to conduct online simulations and experiments to gain insight into the short-term (i.e., 24 to 48 hours) implications of various operational decisions (e.g., "What happens if we close a valve in the grid?"). Similarly, the DT can facilitate forecasting. Predicting sudden high or low demand or other changes allows time for operators to adjust the grid and pumps accordingly (e.g., by pre-charging the network by time and location), so that the temperatures and pressures stay within system limits.

Achieving Strategic Priority 2: Dealing With Fluctuating Energy Prices

The DT can help EnergyCo to achieve its objective of dealing with fluctuations in energy prices in two ways: by enabling exceptional operational flexibility and by enabling production forecasting.

1. The digital twin enables flexibility through real-time production optimization. While district heating is becoming more electrified, energy sources are becoming increasingly weather dependent, and electricity prices are increasingly volatile. District heating operators need to consider not only day-ahead electricity prices but, increasingly, intraday prices, requiring decisions every 15 minutes.[15] Thus, strategic management of heat and electricity production in combined heat and power plants, heat storage and electricity-powered sources—such as heat pumps and e-boilers—becomes crucial. Optimizing these district heating components requires a comprehensive and real-time understanding of many factors, including fuel costs, weather patterns, demand fluctuations and electricity prices.

> *"[EnergyCo] has a lot of combined heat and power and accumulators. At the moment of high electricity prices, [EnergyCo] would produce more electricity to sell it, and store the excess heat generated in the accumulator to use later when there is demand for it. This gives the system flexibility. But can an operator do this alone, using an Excel file? No, there are too many factors to consider."* Product Market Lead, DigiTwinCo

By providing a comprehensive view of the entire district heating network and additional parameters, the DT enables EnergyCo to deal with the increased complexity of the network and optimize the operation of combined heat and power plants with other sources. This means the company can strategically allocate energy production between heat and electricity and effectively use energy storage solutions, thus improving the system's performance. The increased flexibility enables EnergyCo to profit from volatile electricity markets and increases environmental sustainability by maximizing the use of renewable sources and reducing the use of peak boilers.

> *"The connection between electricity and heating is strong. As heating becomes more electrified, decisions about when to use electricity for heating and when to sell it on the wholesale market become crucial."* Energy Systems Manager, EnergyCo

The DT's ability to provide a comprehensive, real-time overview of the district heating network, together with its integrated simulation capabilities, enables EnergyCo to identify the most cost-effective energy production strategies in real time by considering fluctuating energy prices and consumer demands, multiple decentralized heat sources and accumulators, as well as the thermo-hydraulic state of the network.

2. The digital twin enables production forecasting. Production optimization requires scheduling because, due to thermo-hydraulic delays within the network, temperature changes take hours to reach consumers. By feeding real-time data into the AI-assisted thermo-hydraulic model and using future-oriented forecasting, the DT helps EnergyCo to understand the network's current state and optimize and eventually automate energy production and distribution. In addition, by factoring in heat propagation and travel time, scheduling optimizes the allocation of energy sources, ensuring timely heat delivery while minimizing disruptions and hydraulic bottlenecks in the grid's operation. The key added value of the DT in this respect is its ability to adapt to live data, offering a real-time estimation that effectively guides operational strategies and mitigates potential risks or disruptions. This is a unique strategic value of the DT because real-time optimization and simulation programs usually work separately, without data flowing dynamically between them.

Achieving Strategic Priority 3: Achieving Long-Term Profitability

Ensuring organizational efficiency and operational reliability of the district heating

15 Due to EU regulations, in October 2025, the European Electricity Exchange will move from hourly trading to 15-minute pricing intervals.

network is crucial for achieving long-term profitability. The DT can help EnergyCo achieve this goal in two ways: by enabling collaborative decision-making and by enhancing maintenance activities.

1. The digital twin enables collaborative decision-making. The DT is a key enabler for managing the entire value chain of district heating because it creates synergy between those involved in district heating operations and maintenance at different levels. By providing a global view of the entire district heating system, the DT enables real-time insights to be shared across divisions, which is crucial for cross-functional collaboration.

> *"Very few people have a comprehensive understanding of the entire chain from start to end. They see significant value in using the DT to enhance this understanding by making the processes more visible. The user-friendly interface of the DT could illustrate how adjustments made in one part of the system can impact another, effectively showing the interconnectedness of actions within the system."* Energy System Development Team Lead, EnergyCo

An example provided by an EnergyCo interviewee is of a heat pump station's control room staff focusing on a coefficient over 3.0 as an indicator of good performance. However, this somewhat isolated number does not reflect the investment, personnel and maintenance costs. Thus, while technically appearing efficient, operating the heat pump might consume substantial resources, thus negatively impacting profitability. This example highlights the DT's advantage: it enables real-time calculation of the broader picture by using more holistic, cross-functional parameters and eliminates the need to send data to the analytics team for recalculation whenever conditions such as geopolitical situations, prices and consumption change.

In addition, by providing integrated optimization solutions and simulation capabilities, the DT reduces the need for separate legacy software used by different divisions that do not synchronize well with each other. Thus, the DT enables multiple stakeholders to make better decisions collaboratively, simulate responses to various scenarios and improve their workflows without affecting the physical system.

2. The digital twin enhances maintenance activities. Maintaining a functional, long-lasting district heating network requires the operator to have the capability to ensure that pressures and temperatures do not build up to levels that may cause failure. The DT's capability in supporting global optimization of pressures and temperatures is therefore crucial. However, faults and failures will still occur due to the natural deterioration of assets and external factors, such as extremely low temperatures. Automating fault detection and identifying anomalies that indicate impending equipment failures allows EnergyCo to improve operational reliability by identifying and containing failures and potential bottlenecks as soon as possible. Malfunctioning district heating assets have a very high cost, so proactive maintenance practices help the company reduce costs and enhance efficiency in the long term.

> *"[We] could indicate whether a network segment is likely to remain stable for another five years or if it may soon experience issues based on recent pressure fluctuations and the network's age in that area. This predictive insight could identify potential failures before they occur, enhancing network management and maintenance planning."* Energy Systems Manager, EnergyCo

However, as the district heating grid constantly changes and becomes more complex, the impact of even minor maintenance activity on the overall system becomes harder to predict. Scenario analysis is essential for long-term maintenance decisions because it predicts the effects on the system. For emergency maintenance issues, conducting a what-if analysis is crucial for exploring potential solutions without causing disruptions. These types of decision-aiding analyses need to be run simultaneously with real-time optimization. EnergyCo representatives highlighted how the DT provides a comprehensive decision support system for such planning.

> *"The goal is to truly understand what is happening in the network in real time and to test scenarios, such as: What would happen*

in 10 minutes if I made this adjustment now? This allows for lightweight experiments to explore how different operating strategies would actually affect the system. Until now, [EnergyCo hasn't] had a tool that could do this." AI Manager, EnergyCo

In essence, the foresight provided by the DT's capabilities has enabled EnergyCo to shift from reactive maintenance approaches, which are often disruptive and costly, to proactive strategies that minimize system downtime, extend equipment lifespans and lower maintenance costs.

Recommendations for Implementing Digital Twins

The EnergyCo case highlights that the company's sociotechnical capabilities prepared it to successfully adopt a DT. Moreover, the company's approach to DT implementation and management not only enabled it to pursue its strategic priorities but also to navigate the challenges that emerged during the project. Reflecting on the lessons learned, we provide six recommendations for other companies aiming to implement real-time DTs. The first two recommendations are concerned with organizational capabilities and innovation readiness. The final four are concerned with developing, delivering and operating a DT solution.

1. Invest in Startups Developing Digital Twins Related to Broader Industry Trends to Build Collaborative Relationships With Vendors

EnergyCo's decision to implement a DT solution was driven by an "entrepreneurial" mindset encapsulated in an internal venture capital unit that scouted globally for new investment possibilities that were related to the larger energy landscape but did not specifically address the company's own needs. At EnergyCo, various teams (e.g., investments, business development and operations) are involved in identifying niche innovation opportunities and forward-looking ventures that focus on addressing long-term trends within the industry.

EnergyCo, together with two other investors, initially invested over €2 million in DigiTwinCo in 2020 before a specific business case was created, as that company's technology was promising and relevant to the energy sector. After a business case surfaced and DigiTwinCo was selected as the vendor, the prior investment helped to establish trust and collaboration. While most DT vendors are unwilling to reveal the inner workings of their AI models out of fear of exposing their intellectual property and compromising system security, DigiTwinCo implemented the DT solution in collaboration with EnergyCo's experts, enabling them to better understand the inner workings of the DT solution. DigiTwinCo also benefited from working closely with a forerunner client like EnergyCo, because it provided an interesting testing ground and generated lessons that could then be applied to other clients. EnergyCo continued to invest in DigiTwinCo through Series A and Series B funding rounds, enabling the DT supplier to raise over €40 million.

In addition to investing in DigiTwinCo, EnergyCo has invested in dozens of European startups engaged in developing forward-looking energy and climate technology solutions. Most recently, it has invested in an early-stage startup developing an AI-enabled cloud solution for the holistic optimization of electricity demand to seamlessly synchronize consumption across end users' devices and assets (e.g., electric vehicles, charging stations, solar inverters, batteries and smart thermostats) while accounting for fluctuating electricity prices. This further illustrates EnergyCo's financial commitment to long-term innovations.

2. Develop Robust Data Engineering and Management Capabilities to Facilitate Digital Twin Implementation

Developing in-house data engineering and management capabilities is crucial for DT implementation, especially for automating data integration from various sources into a format readable by the DT. As with many other complex CPSs, EnergyCo collects billions of individual data points or observations from thousands of metering points and various internal and external sources. This huge volume of heterogeneous data needs to flow seamlessly and automatically between the district heating grid and the DT platform, without being disrupted by maintenance activities. To achieve this level

  |

of integration, EnergyCo worked to standardize data formats. Our interviewees also emphasized the need to standardize data from the vendor's side because delays in working with and securing data pipelines could lead to client companies resorting to custom solutions under internal client pressures.

In addition to standardization, EnergyCo has invested in protecting data assets and data pipelines to ensure the information security and resilience of the entire district heating grid. Incidents and data breaches in a CPS can disrupt the supply chain and result in significant financial losses, the erosion of customer trust and reputational damage. However, given that CPSs are often composed of both modern and legacy systems, the standardization of data formats is a complicated task. EnergyCo's past investments in developing data engineering and management capabilities enabled it to build a solution within a comprehensive DT framework. Without such capabilities, companies risk fragmentation, which can diminish the overall value of a DT.

3. Select a Challenging Proof of Concept to Convince Skeptics and to Evaluate the Vendor's Capabilities

Some users will always be skeptical about the potential use and benefits of implementing a new technology like DTs. Despite EnergyCo's collaborative relationship and business partnership with DigiTwinCo, the district heating team approached the DT project with skepticism. To validate the feasibility and effectiveness of the DT solution, the team intentionally selected a challenging proof of concept. It tasked DigiTwinCo with modeling part of the district heating grid in a neighborhood far from the heat production sites, making the modeling particularly challenging. After this model proved to be accurate and provided a deeper understanding of that part of the grid, the district heating team's skepticism turned into optimism and even enthusiasm about the project.

In addition to convincing the skeptics, the challenging proof of concept allowed EnergyCo to validate the vendor's capabilities and refine the DT's functionalities through continuous experimentation and feedback. Teams from different parts of EnergyCo played a crucial role in such a learning-by-doing approach, as they had the domain knowledge necessary to comprehend these functionalities and capabilities. Our interviewees also stressed the importance of "accepting imperfection" as a crucial part of the learning process, allowing for continuous integration of the DT into routine work.

4. Optimize Decision-Making Globally, Not Locally, by Leveraging the Digital Twin's Capabilities to Facilitate Collaboration

The primary value of EnergyCo's DT stems from its ability to provide a comprehensive real-time view of the entire district energy grid rather than just part of it. This capability enables the company to break down silos and enhance collaborative decision-making to optimize operations globally. This capability is crucial in a complex CPS like EnergyCo's district heating network, which is managed by various teams and, at the same time, comprises a growing number of users and decentralized physical assets spread across different locations. Optimizing processes—for example, adjusting the temperature and flow of water heated with several sources—in such an interconnected and complex system without a comprehensive and real-time view and collaborative decision-making is very challenging.

The intuitive and functional user interface of EnergyCo's DT provides live monitoring tables and dashboards, helping operators to accurately forecast demand and set optimal temperatures, pressures and production schedules for the entire network. This global view shifts the focus from an expert's immediate field of expertise to understanding the broader impacts of operational decisions. Simplification and prioritization of the parameters shown to the users is beneficial in the long run because this can focus attention on the most relevant details and reveal the big picture, which may otherwise be hidden. In addition, the integration of the DT's simulation capabilities enables a comprehensive forecasting approach that allows, for example, a better understanding of how to prepare for cold periods and ensure adequate heat supply for everyone.

5. Do Not Delegate Decision-Making to an AI-Enabled Digital Twin Without Human Oversight, at Least Initially

Though AI-assisted DTs can support complex decision-making and automate operations, moving toward full automation is especially challenging in complex and critical CPSs, like EnergyCo's district heating network. While full automation may be a long-term goal, human operators remain essential in the process, and automation must advance gradually as operators build trust in the DT's value.

A rapid push toward AI-assisted decision-making risks employee resistance, particularly from operators whose engagement is vital for the successful implementation of a DT. In EnergyCo's case, employees did not express a significant fear of job loss, perhaps because the company avoided immediate automation, kept operators involved throughout the process and had already outsourced some functions while retaining a core emergency team. By positioning the DT as a tool to support, not replace, operators, the company fostered trust and collaboration. This incremental approach can also prepare operators for a future where some real-time operations will eventually be automated, while human expertise will continue to play a critical role in enhancing, training and extending the DT's capabilities. Clear communication from the outset, hands-on workshops and proof-of-concept demonstrations can further build trust, showing operators that their expertise is indispensable in developing the DT and ensuring that it augments rather than undermines their work.

6. Explore Additional Value Generation Opportunities to Fully Leverage the Digital Twin's Potential

The EnergyCo's case shows that DT implementation can help companies address immediate operational challenges while contributing to their long-term strategic objectives. As the challenges and strategic objectives change, so too can the DT—creating ongoing value for various stakeholders. For example, EnergyCo's initial priority was reducing and optimizing grid temperatures to support its carbon-neutrality goals. However, the company is already exploring ways to expand the DT's scope—for instance, by enhancing predictive maintenance capabilities and developing consumer-facing features that will engage consumers in decision-making and help them optimize their own energy use.

In the latter case, an intuitive user interface with well-designed data visualization—potentially via a mobile app—could provide consumers with real-time insights from the DT (e.g., on energy consumption), encouraging greater engagement, flexibility and proactive behavior. These developments would not only enhance consumer participation and transform them from passive into active participants in the energy management system, but could also reshape demand patterns and, ultimately, transform the district heating business model itself.

Concluding Comments

Implementing digital twins can enable significant sociotechnical transformations in organizations across domains. In the energy sector, the drive toward renewable energy sources, the decentralization of energy production, increased electrification of district heating and price fluctuations necessitate greater precision and flexibility and a comprehensive real-time understanding of district heating systems, all of which can be provided by a well-designed DT. By leveraging the advanced capabilities of DTs, companies can optimize operations, enable predictive maintenance and adapt to dynamic market demands, yielding economic and environmental benefits. To successfully capitalize on DTs, companies must not only cultivate an entrepreneurial mindset but must also develop robust data capabilities while engaging employees in DT implementation and adoption and following an incremental approach toward automation.

Appendix: Research Method

We conducted a five-month exploratory case study when EnergyCo's DT was implemented in the spring of 2024 (i.e., Phase 3 in Figure 2). Following a key informant methodology, we conducted nine semi-structured interviews with experts involved in the project in executive and management roles (see table below) who had deep insights and professional knowledge

List of Interviewees

Interviewee ID	Role/Responsibility	Company
Interviewee 1	Development Manager, Digital Twin	EnergyCo
Interviewee 2	Senior Vice President of District Heating	
Interviewee 3	AI Manager	
Interviewee 4	Digitalization Manager, District Heating	
Interviewee 5	Energy Systems Manager	
Interviewee 6	Vice President of Corporate Venture Capital	
Interviewee 7	Energy System Development Team Lead	
Interviewee 8	Solutions Director	DigiTwinCo
Interviewee 9	Product Market Lead	

about the project. Most interviews were 45 minutes to one hour long, except for the main contact from EnergyCo (Interviewee 1), whom we questioned on multiple occasions during the research. In addition, the second author, who was commissioned by EnergyCo during the study and was involved in the project, made field notes and observations. Finally, we relied on publicly available reports and information from both EnergyCo and DigiTwinCo, the vendor of the DT, as a complementary source of data to enhance our understanding of the case and to write the case narrative.

About the Authors

Hadi Ghanbari

Hadi Ghanbari (hadi.ghanbari@aalto.fi) is an assistant professor of information systems at Aalto University School of Business, Finland, and leads the urban analytics and data research theme at the FinEst Centre for Smart Cities, Estonia. His research focuses on the sociotechnical aspects of software development and digital innovation. He has published in top-tier periodicals, including *European Journal of Information Systems*, *Information and Organization*, *Journal of Systems and Software* and *Reliability Engineering & System Safety*. He has developed several executive education modules on cybersecurity and digital sustainability. Before his academic career, he worked in the IT industry for seven years.

Petter Nissinen

Petter Nissinen (petter.nissinen@gmail.com) is a product marketing and branding professional, specializing in climate technology and energy systems. His master's of social science degree from the School of Management, Tampere University, Finland, centered on systemic change, and his interdisciplinary master's of arts degree from Aalto University, Finland, combined computer science and design. He was studying at Aalto University when the research for this article was carried out. As an independent academic, his studies focus on the sociotechnical impact of digital twins on energy systems.