

Document downloaded from:

<http://hdl.handle.net/10251/199518>

This paper must be cited as:

Onggo, B.; Corlu, CG.; Juan-Pérez, ÁA.; Monks, T.; Torre-Martínez, MRDL. (2021).  
Combining symbiotic simulation systems with enterprise data storage systems for real-time  
decision-making. *Enterprise Information Systems*. 15(2):230-247.  
<https://doi.org/10.1080/17517575.2020.1777587>



The final publication is available at

<https://doi.org/10.1080/17517575.2020.1777587>

Copyright Taylor & Francis

Additional Information

# Combining Symbiotic Simulation Systems with Enterprise Data Storage Systems for Real-Time Decision Making

B. S. Onggo<sup>a</sup>, Canan G. Corlu<sup>b</sup>, Angel A. Juan<sup>c,\*</sup>, Thomas Monks<sup>d</sup>, and Rocio de la Torre<sup>e</sup>

<sup>a</sup>CORMSIS, Southampton Business School, University of Southampton, Southampton, United Kingdom (B.S.S.Onggo@soton.ac.uk); <sup>b</sup>Metropolitan College, Boston University, Boston, MA, United States (canan@bu.edu); <sup>c</sup>IN3–Computer Science Dept., Universitat Oberta de Catalunya, Barcelona, Spain (ajuanp@uoc.edu); <sup>d</sup>University of Exeter Medical School, University of Exeter, Exeter, United Kingdom (T.M.W.Monks@exeter.ac.uk); <sup>e</sup>INARBE Institute, Public University of Navarre, Pamplona, Spain (rocio.delatorre@unavarra.es)

## ARTICLE HISTORY

Compiled November 9, 2023

## ABSTRACT

A symbiotic simulation system (S3), sometimes also called a ‘digital twin’, enables interactions between a physical system and its computational model representation. With the goal of supporting operational decisions, an S3 uses real-time data from the physical system, which is gathered via sensors. This real-time data is also saved in an enterprise data storage system (EDSS), so it can be used as historical data for future use. Both real-time and historical data are then used as inputs to the different components of an S3, which typically comprises several modules: data acquisition, simulation, optimisation, machine learning, and an ‘actuator’. The latter is needed when there is not a human agent between the S3 and the system. Given the amount of data generated by today’s smart systems, an S3 needs to be coupled with an EDSS. Furthermore, the S3 may produce a large amount of output data that needs to be stored, since it might be re-used by the machine learning module to make the S3 adaptive in dynamic scenarios. With the goal of supporting real-time operational decision-making – specially in Industry 4.0 applications such as smart cities, smart factories, intelligent transportation systems, and digital supply chains –, this paper proposes a generic system architecture for an S3 and discusses its integration within EDSSs. Moreover, the paper reviews the state-of-the-art in S3, and analyses how these systems can interact with EDSSs to make real-time decision making a reality. Finally, the paper also points out several research challenges in S3.

## KEYWORDS


Symbiotic simulation systems; enterprise data storage systems; online simulation; digital twin; real-time decision; big data.

## 1. Introduction

Managerial decisions can be classified into three categories according to their associated planning horizon: strategic (long term, typically above one year), tactical (medium term, typically several months), and operational (short term, typically days

or weeks). In the current context of Internet of Things (IoT), Industry 4.0, and smart sustainable cities, *enterprise data storage system* (EDSS) plays a relevant role in order to take data-driven and fast, or even real-time, operational decisions. One of the first works proposing simulation as a real-time / online decision support tool is due to Rogers and Gordon (1993), who used the term ‘online simulation’ to describe their real-time simulation-based decision-support tool for manufacturing systems. Likewise, Davis (2007) discusses how online simulation can be used to control any physical system. Darema (2004) uses the term ‘dynamic data-driven simulation’ to emphasise the ability of a simulation model to: (i) react to additional data from the physical system – and / or from the EDSS – while the simulation is running; and (ii) control the physical system in (near) real time.

The idea of using simulation to control a physical system leads to the term *symbiotic simulation* (Fujimoto et al. 2002). Aydt et al. (2008) propose a more general definition of symbiotic simulation, highlighting the benefits generated by the association between an S3 and a physical system. This close association is enabled by IoT, which allows the simulation system and the physical (or real-life) system to interact in real time. These *symbiotic simulation systems* (S3) need an EDSS that can manage the large volume of data gathered and transmitted via IoT-related technologies. For the case of an airport system, Figure 1 illustrates the relationship among these concepts, while Onggo (2019) provides a review of various terms used to describe a similar simulation system.



**Figure 1.** Interactions among a real-life system, an enterprise data storage system, and a symbiotic simulation system.

Similarly, in the context of smart cities, logistics activities such as waste collection and urban traffic management can significantly benefit from the S3 concept: sensors and other devices can be located across the streets, on board of unmanned aerial vehicles, or even inside waste containers to gather data in real time; this data can then be sent to the S3 in order to predict the evolution of the traffic conditions and support policy makers in charge of designing the best strategy to be applied (Suh 2019). In the industry, the term ‘digital twin’ has been sometimes used to refer to an S3. Many simulation software vendors have marketed their products and services to provide solutions for typical Industry 4.0 applications, such as: smart cities, smart factories, and intelligent transportation systems. Increasingly, the term has also been used in academia (Grieves 2014). This shows that both academia and industry consider

S3 as an important development in real-time decision support systems.


As conceptualised in Figure 1, the interaction among an EDSS, a physical system, and an S3 might provide benefits that would otherwise be unavailable if they are used separately. This resembles what we know as hybrid systems (Mustafee et al. 2017; Eldabi et al. 2016). In a recent review, Brailsford et al. (2019) show that the use of hybrid systems to support complex decision making is gaining popularity among simulation researchers and practitioners. This is partly because as systems become more complex and multi-faceted, there is a need to use a combination of methods (e.g., data analysis, intelligent algorithms, optimisation, simulation, machine learning, etc.) and technologies (e.g., IoT, EDSS, etc.) to better represent and manage the system of interest. In fact, the increasing popularity of S3 cannot be separated from developments in Industry 4.0, whose applications put emphasis on fast or real-time situational awareness to address the increasing complexity of real-world systems – e.g., manufacturing and transportation systems. According to Szozda (2017), these new concepts are enabled by four groups of technologies: *(i) data and communication*, including enterprise data storage systems and IoT; *(ii) big data, cloud, and advanced analytics*; *(iii) advanced man-machine interfaces*, including augmented reality; and *(iv) advanced actuators*, including robotics, 3D printers, and autonomous vehicles.

This paper discusses S3 and its integration with EDSS. An S3 is typically composed of analytics and machine learning modules, an optimisation module, a data storage / acquisition module, and an ‘actuator’ module. An actuator is the software / hardware that is responsible from triggering actions on the physical system based on the results from its S3 or digital twin. For example, taking the case of manufacturing, a conveyor belt in a factory (physical system) may be slowed down after receiving a signal from its S3. Similarly, when considering the services sector, applying modifications to the real system might require the authorisation of a manager. Thus, for example, even when an S3 might suggest re-allocating stuff in a hospital, this decision might require the final endorsement of the hospital manager.

Our work makes a novel contribution to the existing literature in the following ways: *(i)* it presents the architecture of an integrated S3 and EDSS; *(ii)* it defines and discusses the components of a typical S3, with a reference to existing research and software implementation; and *(iii)* it identifies future work to advance research in S3. The remainder of this paper is organised as follows. Section 2 discusses the S3 architecture and describes the components of an S3. Section 3 provides a review on the use of EDSS in simulation. Section 4 discusses the applications of S3 in operations management. Section 5 identifies the research challenges in S3 and highlights future research directions. Finally, Section 6 concludes by summarising the findings of our study.

## 2. Architecture of a Symbiotic Simulation System

Authors such as Aydt et al. (2008) classify S3 into three types, which are based on the communication between an S3 and the corresponding physical system. The three types are shown in Figure 2. **As shown in the figure, the role of EDSS is important in S3. This is because S3 may receive a large amount of data not only from sensors in today’s smart systems, but also the output data from S3 that needs to be stored to make S3 adaptive in dynamic scenarios. This section explains the components of an S3 in detail, including a review of the latest developments. We will explain the role of EDSS in S3 separately in Section 3.**




**Figure 2.** Three types of symbiotic simulation systems.

- Type 1 (decision support S3): The enterprise data storage system is used to store data from the physical system. The symbiotic simulation reads real-time data from the physical system and historical data from the enterprise data storage system. Then, the outcomes from the S3 are considered by a manager to introduce changes in the real system, i.e.: the symbiotic simulation indirectly controls the physical system via a decision maker.
- Type 2 (control symbiotic simulation system): This is similar to Type 1 except that the symbiotic simulation directly controls the physical system through an actuator – i.e., without any human intervention.
- Type 3 (open loop symbiotic simulation system): While the first and second types are characterised by the closed loop between a physical system and its corresponding S3, the third type is an open-loop S3. Now, the S3 reads all data from the real system (e.g., input data and key performance indicators), and uses it to identify strange patterns in the real system or to forecast its future performance. Notice that this is analogous to non-symbiotic simulation with the incorporation of real-time data. Although the S3 is not employed here to control the real system, the output generated by the S3 can be utilised by third systems that are in charge of controlling the physical one.

Earlier works (Onggo et al. 2018; Onggo 2019) have identified the components of an S3, which includes a simulation module, a data acquisition component, an optimisation module, a machine-learning component, a data analytics module, and a scenario manager. However, the role of EDSS was not considered. As discussed before, in many industries it is not possible to have an S3 without an appropriate EDSS. Our paper extends the earlier work from Onggo et al. (2018) and Onggo (2019) by extending the architecture of S3 to include the interaction with an EDSS. We also add the latest developments in each of the S3 components and in particular we significantly extend the discussion on the machine learning component.

Figure 3 shows the interaction between a physical system (real-life system), its S3, and the EDSS. The figure also illustrates the main components of an S3 and their interactions. The data from the physical system and the enterprise data storage system is collected by a data acquisition module. The use of appropriate data analytics approaches (e.g., forecasting, regression models, etc) and machine learning (ML) methods allow the system to combine historical data with real-time one. The combined use analytics models and historical data might be helpful, for instance, to adjust the parameters of the simulation, scenario manager, or optimisation modules. Notice that a scenario manager considers several ‘what-if’ analyses based on the outcome provided by the simulation model. Alternatively, an optimisation module can be used to produce a near-optimal solution. Machine learning methods can be used to adapt the

scenario manager or optimisation model, the simulation model, and the data analysis. The results from the scenario manager / optimisation module are shared with an actuator or a decision maker, who decides the nature of the changes to be implemented in the physical system. The unifying objective of these systems is to make fast (or even real-time) adjustments in the physical system, so that its performance does not deteriorate. **The detailed explanation of the S3 components is as follows.**



**Figure 3.** Symbiotic system formed by an enterprise data storage system, a physical system, and its S3.

### 2.1. *Data Acquisition*

Using the data acquisition component (Figure 3), S3 extracts, transforms, and loads historical data from the enterprise data storage system and real-time data from a physical system. The data can be read online or offline. In the former case, direct communication with the sensors is used, while in the latter case the S3 employs historical data available from the EDSS. The loaded data is then analysed using appropriate analytics methods. The objective of these methods is to select the best way to use the combination of historical data and new data, which is only available when the model is running. The length of a time window in which the new data becomes historical data is application specific (e.g., one year may be reasonable for a system with seasonal data). The information extracted by analytical methods may be used to update the scenario manager, the optimisation module, or the symbiotic simulation model. ML methods can also be used to adapt the previous components and, as more data is collected, increase their performance.

Notice that the data acquisition module could be deployed as a web service or even as a mobile app, and receive data from different information systems (Marmor et al. 2009). Espinoza et al. (2014) use the data stored in a workflow management system as an input to their real-time simulation model – which was built to improve the daily operations at an emergency room. Mustafee et al. (2017) discuss data specification and data acquisition for a platform that displays real-time data on urgent-care waiting times from different medical centres. Rossmann et al. (2011) use data from remote sensing for their virtual forest simulation system. In a similar way, other authors have used S3 in the airline management (Rhodes-Leader et al. 2017).

## 2.2. *Symbiotic Simulation Model*

The core model of an S3 is the symbiotic simulation model. A symbiotic simulation model is capable of responding to new data while the simulation is running. Often, data analytics or ML methods are used to investigate the best way of using various data sources to update the appropriate parts of the S3. The new data may indicate a change in the physical system. Hence, symbiotic simulations should be able to make an appropriate response – as specified by the analysts – to the new data. The response can be in several forms: from re-initialising the simulation (i.e., the system states) to making adjustments in the simulation model structure. Adjusting the input distributions and the remaining time for entities that are currently being simulated are also among the responses that a symbiotic simulation model can perform with the arrival of new data.

In a typical non-symbiotic simulation, the system states (such as the number of entities in the system or the queue length at a certain time) are initialised at the start of the simulation. In a symbiotic simulation, however, the system states may need to be re-initialised based on the new information available. Fortunately, most simulation software allows this functionality. Re-initialisation in symbiotic simulation affects the recording of times in the simulation, i.e.: the service time of an existing entity that is being served when the simulation is initialised should be sampled based on the time already spent being serviced. Otherwise, the output statistics might be biased. The key here is to simply use a conditional probability distribution function (i.e., the sampling of a service time takes into account the time already spent in the service). For example, Oakley et al. (2019) derive conditional length-of-stay distributions to be used in their symbiotic simulation model for hospital bed management. They also note that most commercial simulation software allow to specify user-defined distributions, including conditional probability ones.

Several simulation tools offer a functionality to adjust the simulation parameters while the simulation is running. One input needed from the modeller is a condition that triggers the parameter adjustment. If the trigger is too sensitive, the physical system may become unstable. On the other hand, if the trigger is not sensitive enough, the physical system may not be responsive enough to the changes in the real world. Another issue is that the symbiotic simulation model need to be re-validated after the parameter adjustment. If the parameter adjustment is done frequently, then an auto-validation mechanism would help. Onggo and Karatas (2016) discuss a validation suite in test-driven simulation modelling, which could be a viable approach for the auto-validation of a symbiotic simulation model.

The most challenging response to the new data is perhaps the ability to make adjustments to the structure of the simulation model while it is running. In this regard, Mitchell and Yilmaz (2009) use a genetic algorithm to configure and coordinate a series of simulation models acting in parallel. Authors such as Zhang and Zhao (2012) use adaptive simulation models that are able to modify their logic as new data is received from the physical system. A different possibility is to make the approach adaptive by adding a machine learning model. For example, process mining methods can be used to learn about changes to real process from process log data. This might ensure that the simulation model accurately reflects the current process (Mesabbah et al. 2019). If there is a drift from the real process, the model can be modified accordingly. In these examples, the key again is to define a specific condition that triggers the change in the structure of the model. Also, the auto-validation of the model gains more importance because of the structural change in the simulation model. The auto-verification and

auto-validation methods described in Onggo and Karatas (2016) are especially relevant here.

Another characteristic that differentiates symbiotic simulations from non-symbiotic simulations is that they are designed to support short-term operational management decisions. Therefore, they are expected to quickly find a solution. This requires that they run fast. Finally, another requirement is that symbiotic simulations need an enterprise storage system with the capability of efficiently handling large volumes of data.

### 2.3. *Scenario Manager and Optimisation Model*

The scenario manager performs several ‘what-if’ analyses, which goal is to compare different strategies or policies. As in the case of a non-symbiotic simulation system, the main role of the scenario manager module is to define and execute the simulation models to be run. Typical analysis, such as sensitivity analysis and experiment design, may be used here.

Also, an optimisation model may be employed instead of (or in combination with) a scenario manager, especially when it is impractical to define a set of reduced scenarios. In this case, the optimisation model identifies the optimal solution based on a pre-defined objective function, which is estimated by running the simulation model. The combination of simulation and optimisation models in this manner is referred to as simulation-optimisation (Fu et al. 2015). Many complex short-term operational decisions can be formulated as combinatorial optimisation problems (COPs). A typical COP encompasses a vast number of possible solutions to choose from. Examples include routing vehicles in a city (Calvet et al. 2016; Gruler et al. 2017), staffing decisions to minimise waiting time, scheduling jobs to maximise throughput (Gonzalez-Neira et al. 2017), inventory management to minimise stock-outs (Gruler et al. 2020), etc. Whenever the optimisation problem is *NP-hard*, finding an optimal or near-optimal solution to a large-sized COP using traditional methods is likely to take a very long time. Hence, alternative methods are needed to find good-enough solutions that can be achieved over a short planning horizon. The alternative methods include ‘simheuristics’ (Juan et al. 2018), multi-fidelity modelling (Rhodes-Leader et al. 2018), or parallel computing (Panadero et al. 2018).

Simheuristics constitute an example of a simulation-optimisation method that can be used for symbiotic simulation-optimisation. In that regard, simheuristics have been used to solve stochastic optimisation problems – employing short computational times – in several fields, which include: transportation and logistics, supply chain management, healthcare, etc. One advantage of these algorithms is that they can further be coupled with parallel and distributed computing, which makes the algorithm more scalable. Thus, for instance, De Armas et al. (2017) consider an uncapacitated facility location problem with stochastic demands and propose a fast heuristic algorithm for the deterministic version of the problem. This algorithm is then extended into a simheuristic framework, where the feedback from simulation is shown to eliminate non-promising solutions and hence further decrease the computational time. Similarly, Pagès-Bernaus et al. (2019) consider a capacitated facility location problem in the context of e-commerce operations. These authors propose a two-stage stochastic programming approach, and show that this approach can only solve small-sized instances of the problem. A simheuristic approach, on the contrary, is able to solve large-sized instances in reasonable computing times. These examples show that simheuristics con-

stitute a promising simulation-optimisation method for S3.

#### 2.4. Data Analytics and Machine Learning

A key challenge for S3 and simulation-optimisation is model fidelity (Xu et al. 2014). Analytics and ML offer exciting research opportunities to tackle these challenges. To illustrate the problem, assume that a simulation model uses  $m$  sensors to infer the physical state of the system in time  $t$ ,  $S_t$ . In simulation-optimisation, the task is to select the best design(s),  $I^*$ , from a set of  $k$  competing designs  $I = \{1, 2, \dots, k\}$ . Suppose that the optimisation procedure commences at time  $t$  and requires a budget of  $h$  time periods to complete. When  $h$  is large, in an online optimisation setting, it may be problematic to assume that  $S_t = S_{t+h}$  and  $I_t^* = I_{t+h}^*$ , i.e., that the optimal design(s) are the same if the state of the system has changed. There are two approaches to cope with this challenge: (i) to predict the system state  $h$  steps ahead; and (ii) to reduce the  $h$  required to select  $I^*$ .

Each input into the simulation model collected via sensors is a time series collected at a given sensor frequency. For sensor  $i$ , the mean  $h$ -step forecast is  $\hat{y}_{T+h|T}^i$ , where  $T$  is all of the observations up to time  $T$  (or, alternatively, a sliding window of these observations). A forecast of  $\hat{S}_{T+h|T}$  is therefore the set of mean point forecasts for all data inputs that are used to set the parameters of the simulation model. In a statistical forecasting paradigm, these point forecasts are part of a forecast distribution, and can be used to inform the uncertainty in the model parameters. It is typical that, as  $h$  increases, so does the width of a point forecast prediction intervals at a given level of confidence. The result is that the more simulation run-time budget that is needed by the simulation-optimisation approach, the longer the forecast horizon and the more uncertain the parameter estimates. For high fidelity simulation, the number of sensor inputs,  $m$ , could be very large and possibly count in the hundreds of time series. In these circumstances, modellers face a large-scale forecasting problem, where automatic forecasting and model selection procedures are needed. Two powerful and commonly used automatic procedures are ETS and autoARIM, which are provided by the R’s Tidyverse fable package (Hyndman and Athanasopoulos 2018). However, in S3 sensor data may be high frequency and / or have complex periodicity. In these instances, researchers will need to look at more flexible automatic methods, such as TBATS (De Livera et al. 2011) or the Facebook’s Prophet library (Taylor and Letham 2018).

To date, forecasting research in the context of symbiotic simulation has been barely explored. Historically, machine learning methods, such as Gaussian processes or long short-term memory processes (LSTM) perform poorly relative to statistical procedures for time series forecasting (Makridakis et al. 2018). Statistical procedures also have the benefit of natively providing prediction intervals, while methods based on neural networks – such as LSTM – must be adapted (potentially at great computational expense) to provide an estimate of uncertainty. Cutting edge research might include hybridising the structural benefits of time series and the flexibility of ML methods. Such a method won the recent M4 forecasting competition (Makridakis et al. 2020; Smyl 2020), where it was the most accurate at forecasting 100,000 time series.

A complementary approach to forecasting  $h$ -steps ahead is to reduce  $h$  itself – the run-time budget needed to find an optimal solution. A classic approach is optimisation by meta-model (Barton 2015). Meta-models are a ‘model of a model’, and offer substantially reduced run-times relative to a simulation model. Examples are first or second order polynomials or stochastic kriging (Ankenman et al. 2010). Meta-models

are fitted offline to the inputs and responses of a series computer experiments (following a design such as a Latin hyper-cube). Meta-models are then used online in place of the simulation model and interpolate between the design points. A complication is that design of experiments for meta-models is often difficult. This is specially the case if there are a large number of sensor inputs, states of the physical system, and options – as in an S3. It may also be foolish to assume that we can know all states of the physical system in advance. As such, there may be regions of the solution space that have limited (or no) coverage, or where the model has a poor fit. In these instances, the meta-model might be thought of as a tool to narrow down the solution space. Full computer simulation would still be needed to select the best design, albeit from a much more limited set of candidates.

An exciting, but challenging area of research that holds great potential for S3 is reinforcement learning (RL). When RL is implemented using neural network approaches, it is called deep RL. In RL, an agent takes actions within an environment, and learns the value of its actions via (delayed) feedback, which may consist of observations and a reward. RL approaches, such as tabular Q-learning and deep Q-networks (DQN) are suitable for stochastic environments that would be found in S3, with DQN being suitable for problems with more states. Unlike in a meta-modelling approach, where design of experiments is employed, an agent starts with no information and steps forward in time. The agent takes actions that balance exploitation of promising actions already tested,  $S_t$ , and exploration of actions where it has limited experience. RL agents are often initially trained in simulated environments – making it ideal for S3 – and then used with the physical system. Over time, an agent would also receive feedback from the physical system, refining its internal estimates of the value of actions given a state. For instance, in Q-Learning this is done by blending old and new values of the quality of an action in a given state. A particular challenge in RL is defining reward values. In an example of a production facility, Creighton and Nahavandi (2002) defined a reward based on inventory storage, set-up, and production costs. RL is a highly active research area within artificial intelligence, and there is great opportunity to transfer the gains across to S3.

Relate to symbiotic simulation, the combination of heuristic optimisation and machine learning methods (Calvet et al. 2017) can be a valuable methodology for predicting the future performance of a physical system and adjust its parameters accordingly. Thus, for instance, Calvet et al. (2016) use such a ‘learnheuristic’ method to optimise a dynamic vehicle routing problem, which might take into account variations in the system inputs.

### 3. Use of Enterprise Data Storage Systems in Simulation

Data storage systems have been researched for a long time and their application areas are very wide (Al-Awami and Hassanein 2016). Examples include Industry 4.0, smart sustainable cities or service industry. As discussed in previous sections, S3 requires an EDSS that can manage a large amount of data in the form of historical data and real-time data from sensors (some of them come in high frequency). In big data analytics literature, these two dimensions are called volume and velocity, respectively. S3 may also receive data in the form of images, videos and sounds (also known as the variety dimension in big data). Finally, with the increasing number of sensors used in smart systems, the accuracy and the quality of the data received by S3 from the sensors vary. This is called the veracity dimension in big data. Hence, to support S3, it is important

that the EDSS contains capabilities to handle data with such characteristics. This section provides a broad overview of the existing research on EDSSs and how they relate to S3 and data storage technologies that help dealing with a variety of data characteristics.

In the field of sensors, Al-Awami and Hassanein (2016) present a study for designing data survivability schemes by using wireless sensor network (WSN). These types of networks are very useful in intelligent transportation systems, smart grids, and IoT. In order to protect the sensed data, WSNs benefit from the distributed data storage systems technology. To do so, authors propose a data storage system based on decentralised erasure codes that are able to simplify and decentralise the construction of the target code. Storage systems are critical for S3 to deal with the volume and veracity of the data from the sensors installed inside the physical system.

In order to procure their financial sustainability financially, data storage centres are changing their methods from simple data replication to more powerful erasure codes (Andriolli et al. 2009; Kim et al. 2014). These can ensure the same level of reliability than those provided by the replication methodology, but incurring in significantly lower storage costs. Lee et al. (2017) provide new data (both lower and upper bounds) storage systems based on maximum distance-separable codes. For real-world implementation of S3, cost is a factor that must be considered. Others, like Cai et al. (2017) depict a method for designing systematic codes that permit create efficient spectral shaping codes with the maximum run length possible. These codes can be very useful in high-density data storage systems. An enterprise storage system with this method can deal with the volume and velocity of data that feeds into S3.

Cao et al. (2019) propose a new approach for decoding fixed-length and variable-length constrained sequence codes, which permit a more efficient and reliable transmission of coded symbols used in the frontiers of wireless communication systems and data storage systems. This method can further help an EDSS to deal with veracity issues that may arise and boost its use in S3. Associated with the popularity and price of the data storage systems comes the loss of data compression, which leaves a lot of files stored uncompressed, despite there are losslessly compressible. In order to exploit the lossless compressibility, Xie et al. (2010) propose to boost storage systems metrics, like energy efficiency and access speed, rather than trying to save storage space in a conventional way. By doing so, finite word length configurations can be figured out, and the configurations that ensure no performance gain, even considering a possible increment in the finite word length can be chosen. This is useful as for the application of S3, the cost and the variety in the structured data need to be considered. For improving phase-change memory, Wu et al. (2012) show three design methodologies (i.e., time-aware variable-strength error correction code decoding, time-aware partial rewrite, and time-aware read-&-refresh). These methodologies can enhance the data retention limit by a few orders of magnitude, and can enable considerable energy savings, which can be beneficial for an S3.

#### **4. Applications of S3 in Operations Management**

The idea of using S3 for control has been around since the very beginning of the S3 concept. The way S3 is used to control a physical system can take several forms. First, S3 can create a reference model of a physical system. In the manufacturing context, for example, Katz and Manivannan (1993) developed an S3 to detect discrepancies between the expected behaviour of the performance of a shop-floor and its actual performance.

Secondly, S3 can be used as an early-warning system if the future performance of a system is not within an acceptable range. For instance, Hetu et al. (2018) designed a framework using virtual reality as a part of the symbiotic system. This approach helps to deal with crisis scenarios and to add value to the current expert advice process, by creating a symbiotic feedback loop that helps during crisis management. Similarly, Oakley et al. (2019) develop an early-warning system to alert hospital managers when there is a high probability of potential overcrowding in the next few days. Thirdly, S3 can be used to help decision-makers compare different scenarios. It could be done using the scenario manager in Figure 3. For example, Teixeira et al. (2019) use symbiotic simulation for decision support in a remanufacturing system. Authors demonstrate that the application of symbiotic simulation helps dealing with uncertainty, especially in forecasting the outcomes and the performance of the remanufacturing process and the obtained products, and the managerial tasks.

S3 have been applied to different fields, and in particular they have been applied in manufacturing processes (Low et al. 2005; Kück et al. 2016), which is partially due to the fact that real-life data can be obtained from sensors in many modern factories. Obviously, one can think of other ecosystems when S3 can have a significant number of applications, e.g., smart cities, airports, hospitals, etc.

Although the concept of S3 – or digital twins – has received significant attention in the manufacturing industry, its implementation in the service industry is relatively new. Service industries that use S3 include: *(i)* product services; *(ii)* healthcare services; and *(iii)* airline services. Regarding product services, most of the research has focused on the production itself, ignoring the product service component in the life cycle management of a product. However, for some products, the majority of the margin comes from the after-sale services (Cohen and Whang 1997). In a recent survey paper, Zhang et al. (2019) investigate the literature on the role of digital twins in product services, and find only a few relevant publications. Among those, Abramovici et al. (2018) is perhaps the most noticeable study. It conceptually discusses the potential use of digital twins for after-sale services of products. Similarly, El Saddik (2018) is among the first authors who discusses the potential use of digital twins in healthcare broadly. The application areas discussed in the previous work include: predicting the occurrence of an illness, providing customised recommendations to patients on how to improve their health conditions, and offering advice on well-being. Laaki et al. (2019) discuss the use of digital twins for a remote surgery operation. Liu et al. (2019) focus on elderly healthcare services and develop a novel digital twin healthcare to monitoring, diagnosing, and predicting health conditions of the individuals. At the operational decision-making level in healthcare, Marmor et al. (2009), Bahrani et al. (2013), Tan et al. (2013), and Espinoza et al. (2014) focus on real-time decision making in the operations of emergency departments. More recently, Oakley et al. (2019) develop a symbiotic simulation model in a hospital for the operational management of inpatient beds. Rhodes-Leader et al. (2017) develop a symbiotic simulation for the stochastic aircraft recovery problem. The simulation, which is used to compare operational decisions (e.g., cancelling the flight or replacing the aircraft with another aircraft), uses data from a small airline company.

Conceptually, a combination of an S3 and an EDSS can provide a tool to support the four levels of managerial capacities in operations management (Moeuf et al. 2018): monitoring, control, optimisation, and automation (Figure 4). The upper level in Figure 4 receives data or components from the lower level. Hence, IoT enables us to monitor the various parts of a physical system and also to provide the data needed for the upper levels. Based on ‘monitoring’ data, it is possible to define the ‘standard’

behaviour of the physical system, which then could be used to control the behaviour of the physical system. The capacity to control the physical system allows us to incorporate a simulation-optimisation model to find the most optimal decision. Finally, machine learning method can be used to make the simulation-optimisation model learn from past system performance and adjust the system. Hence, it creates an autonomic system. S3 can be used to support control, optimisation, and autonomy.




Figure 4. S3 supports the top three levels of managerial capacities.

## 5. Open Research Challenges in S3

The following research challenges and opportunities have been identified regarding S3:

- The first challenge for S3 is to integrate its different technologies, i.e., data acquisition, analytics, and machine learning. Hence, in order to reach common system-level objectives, and make S3 components inter-operable, an integrative framework must be designed. In this regard, Onggo et al. (2018) have detected the existing challenge in integrating various analytics models with ML methods in an S3. Apart from the aforementioned requirements, the new framework should also consider the communication between the virtual system and the physical one. Further, this integrative approach should not stop at one symbiotic system, but it should include all the existing symbiotic systems across the entire value chain. This is particularly true for Industry 4.0 and the modern services industry, since they also involve end-to-end digital integration across organisations in a value chain including end-customers.
- The second challenge is the standardisation, as a key aspect, of an integrated framework. In terms of Industry 4.0, the *industrial Internet reference architecture* and the *reference architecture model for Industry 4.0* provide a good starting point for the development of standards for the real-world implementation of S3. According to Yli-Ojanperä et al. (2019), the preference of one or more architectures, the related standards, and the standard-compliant products and services are being established in the nations currently boosting the Industry 4.0 revolution. However, since Industry 4.0 technology and standards are constantly

changing, its application is not simple for researchers and practitioners, hence there is still a lack of standardisation for enabling inter-operability (Pereira and Romero 2017).

- The third challenge is related to scalability issues, which may arise as a consequence of providing more data sources for S3 (mainly because computing and communication capabilities are embedded in more devices). In order to address these enormous managerial challenges, it is mandatory to focus the current research into how S3 can efficaciously handle the amount of data that enters at high frequency in various forms (for example, text, numbers, images, etc.), and which may contain noise. Furthermore, in order to enhance the best use of such data, methods suitable for data analysis are required. In a preliminary proposal, Liu and Yang (2016) present a two-tier replication framework that assures good scalability for replication across multiple data centres in the context of big data storage.
- The fourth challenge concerns security and data-privacy issues as one of the main difficulties against the implementation of Industry 4.0, smart sustainable cities, and other smart services (Fan et al. 2015). In fact, as virtual and physical systems become more integrated, the risk of a physical system being attacked or hacked also increases. Further, data-privacy can also arise due to the sensitive nature of the data, that is the case of end-to-end digital integration within the value-chain. As a consequence, we believe that research has to focus on the security and data-privacy aspect of S3.
- Finally, apart from the aforementioned necessities regarding the integration of frameworks, there are also other methodological challenges for S3 (Onggo 2019), among them the following ones: *(i)* how to deal with highly dynamic physical systems; *(ii)* the need for suitable algorithms (e.g. optimisation, machine learning) in a short-term decision-making environment; and *(iii)* how to create simulation models that are also adaptive in order to reflect changes in the physical system.

## 6. Conclusions

This paper has discussed a decision support tool called symbiotic simulation system (S3) or digital twin. These S3 are different from non-symbiotic simulations, mainly because they are designed for fast or even real-time operational decision making. The real-time decision making aspect of an S3 is very beneficial to be used as a decision support tool, especially in Industry 4.0 applications, such as smart factories, smart cities, digital supply chains, intelligent transportation systems, etc. However, making real-time decision making a reality bears several challenges.

To help us understand the different methods, models, and technologies that form the S3 and how these components work together, we have proposed a generic system architecture. The components include an analytics module, a machine learning module, a simulation module, an optimisation module, a scenario manager, an enterprise data storage system, a data acquisition module, and an actuator. The paper also provides a review on the latest developments in each of these components.

Given the amount of data generated by the Internet of Things and by simulation runs, we also discuss the role of EDSSs in S3. Moreover, a review of S3 applications to different fields is provided, and these applications are categorised based on the supported managerial capacities. Although S3 is promising, we acknowledge that many

non-trivial theoretical and practical challenges exist. Hence, we have summarised these challenges and highlighted potential research areas in S3. Among these challenges, the paper has identified and described the following ones: *(i)* technological integration of different modules in an S3; *(ii)* standardization of the architecture and its modular composition, so that component inter-operability is promoted; *(iii)* data-scalability of an S3 components; and *(iv)* security and data privacy in an S3.

All in all, integrated frameworks combining a symbiotic simulation system with an enterprise data storage system offer the possibility of ‘learning’ from the past (historical data), the present (real-time data), and the future (simulation data), and thus constitute a major player in many digital industries, from manufacturing to logistics and transportation services.

## Acknowledgements

This work has been partially funded by the Staff Mobility programme from the Erasmus+ (2020-2021).

## References

- Abramovici, M., Savarino, P., Göbel, J. C., Adwernat, S., and Gebus, P. (2018). Systematization of virtual product twin models in the context of smart product reconfiguration during the product use phase. *Procedia CIRP*, 69:734–739.
- Al-Awami, L. and Hassanein, H. S. (2016). Distributed data storage systems for data survivability in wireless sensor networks using decentralized erasure codes. *Computer Networks*, 97:113–127.
- Andriolli, N., Scaffardi, M., Barman, A. D., Castoldi, P., Potì, L., and Bogoni, A. (2009). All-optical packet-switched interconnection network based on modular photonic digital processing. *IET Communications*, 3(3):477–486.
- Ankenman, B., Nelson, B. L., and Staum, J. (2010). Stochastic kriging for simulation meta-modeling. *Operations Research*, 58(2):371–832.
- Aydt, H., Turner, S. J., Cai, W., and Low, M. Y. H. (2008). Symbiotic simulation systems: An extended definition motivated by symbiosis in biology. In *Proceedings of the 2008 Workshop on Principles of Advanced and Distributed Simulation*, pages 109–116. IEEE.
- Bahrani, S., Tchemeube, R. B., Mouttham, A., and Amyot, D. (2013). Real-time simulations to support operational decision making in healthcare. In *Proceedings of the 2013 Summer Computer Simulation Conference*, pages 1–7.
- Barton, R. (2015). Tutorial: Simulation metamodeling. In *Proceedings of the 2015 Winter Simulation Conference*, pages 1765–1779.
- Brailsford, S. C., Eldabi, T., Kunc, M., Mustafee, N., and Osorio, A. F. (2019). Hybrid simulation modelling in operational research: A state-of-the-art review. *European Journal of Operational Research*, 278(3):721 – 737.
- Cai, K., Immink, K. A. S., Zhang, M., and Zhao, R. (2017). On the design of spectrum shaping codes for high-density data storage. *IEEE Transactions on Consumer Electronics*, 63(4):477–482.
- Calvet, L., de Armas, J., Masip, D., and Juan, A. A. (2017). Learnheuristics: hybridizing metaheuristics with machine learning for optimization with dynamic inputs. *Open Mathematics*, 15(1):261–280.
- Calvet, L., Ferrer, A., Gomes, M. I., Juan, A. A., and Masip, D. (2016). Combining statistical learning with metaheuristics for the multi-depot vehicle routing problem with market segmentation. *Computers & Industrial Engineering*, 94:93–104.

- Cao, C., Li, D., and Fair, I. (2019). Deep learning-based decoding of constrained sequence codes. *IEEE Journal on Selected Areas in Communications*, 37(11):2532–2543.
- Cohen, M. A. and Whang, S. (1997). Competing in product and service: a product life-cycle model. *Management Science*, 43(4):535–545.
- Creighton, D. and Nahavandi, S. (2002). Optimising discrete event simulation models using a reinforcement learning agent. In *Proceedings of the 2002 Winter Simulation Conference*, pages 1945–1950.
- Darema, F. (2004). Dynamic data driven applications systems: A new paradigm for application simulations and measurements. In *Proceedings of the International Conference on Computational Science*, pages 662–669. Springer.
- Davis, W. J. (2007). *On-Line Simulation: Need and Evolving Research Requirements*, chapter 13, pages 465–516. John Wiley & Sons, Ltd.
- De Armas, J., Juan, A. A., Marquès, J. M., and Pedroso, J. P. (2017). Solving the deterministic and stochastic uncapacitated facility location problem: from a heuristic to a simheuristic. *Journal of the Operational Research Society*, 68(10):1161–1176.
- De Livera, A., Hyndman, R., and Snyder, R. D. (2011). Forecasting time series with complex seasonal patterns using exponential smoothing. *Journal of the American Statistical Association*, 106(496):1513–1527.
- El Saddik, A. (2018). Digital twins: The convergence of multimedia technologies. *IEEE MultiMedia*, 25(2):87–92.
- Eldabi, T., Balaban, M., Brailsford, S., Mustafee, N., Nance, R. E., Onggo, B. S., and Sargent, R. G. (2016). Hybrid simulation: Historical lessons, present challenges and futures. In *Proceedings of the 2016 Winter Simulation Conference*, pages 1388–1403. IEEE.
- Espinoza, C., Pascual, J., Ramis, F., Bórquez, D., and Sepúlveda, J. A. (2014). Real-time simulation as a way to improve daily operations in an emergency room. In *Proceedings of the 2014 Winter Simulation Conference*, pages 1445–1456. IEEE.
- Fan, X., Yang, G., Mu, Y., and Yu, Y. (2015). On indistinguishability in remote data integrity checking. *The Computer Journal*, 58(4):823–830.
- Fu, M. C. et al. (2015). *Handbook of Simulation Optimization*, volume 216. Springer.
- Fujimoto, R., Lunceford, D., Page, E., and Uhrmacher, A. M. (2002). Grand challenges for modeling and simulation. *Schloss Dagstuhl*, 350.
- Gonzalez-Neira, E. M., Ferone, D., Hatami, S., and Juan, A. A. (2017). A biased-randomized simheuristic for the distributed assembly permutation flowshop problem with stochastic processing times. *Simulation Modelling Practice and Theory*, 79:23–36.
- Grieves, M. (2014). Digital twin: manufacturing excellence through virtual factory replication. *White paper*, 1:1–7.
- Gruler, A., Fikar, C., Juan, A. A., Hirsch, P., and Contreras-Bolton, C. (2017). Supporting multi-depot and stochastic waste collection management in clustered urban areas via simulation–optimization. *Journal of Simulation*, 11(1):11–19.
- Gruler, A., Panadero, J., de Armas, J., Moreno, J. A., and Juan, A. A. (2020). A variable neighborhood search simheuristic for the multiperiod inventory routing problem with stochastic demands. *International Transactions in Operational Research*, 27(1):314–335.
- Hetu, S. N., Gupta, S., Vu, V.-A., and Tan, G. (2018). A simulation framework for crisis management: Design and use. *Simulation Modelling Practice and Theory*, 85:15–32.
- Hyndman, R. J. and Athanasopoulos, G. (2018). *Forecasting: principles and practice*. OTexts.
- Juan, A. A., Kelton, W. D., Currie, C. S., and Faulin, J. (2018). Simheuristics applications: dealing with uncertainty in logistics, transportation, and other supply chain areas. In *Proceedings of the 2018 Winter Simulation Conference*, pages 3048–3059. IEEE.
- Katz, D. and Manivannan, S. (1993). Exception management on a shop floor using online simulation. In *Proceedings of the 1993 Winter simulation Conference*, pages 888–896.
- Kim, D.-H., Jeon, S., Park, N.-C., and Park, K.-S. (2014). Iterative design method for an image filter to improve the bit error rate in holographic data storage systems. *Microsystem Technologies*, 20(8-9):1661–1669.
- Kück, M., Ehm, J., Hildebrandt, T., Freitag, M., and Frazzon, E. M. (2016). Potential of

- data-driven simulation-based optimization for adaptive scheduling and control of dynamic manufacturing systems. In *Proceedings of the 2016 Winter Simulation Conference*, pages 2820–2831. IEEE.
- Laaki, H., Miche, Y., and Tammi, K. (2019). Prototyping a digital twin for real time remote control over mobile networks: Application of remote surgery. *IEEE Access*, 7:20325–20336.
- Lee, K., Shah, N. B., Huang, L., and Ramchandran, K. (2017). The mds queue: Analysing the latency performance of erasure codes. *IEEE Transactions on Information Theory*, 63(5):2822–2842.
- Liu, F. and Yang, Y. (2016). A two-tier replication framework and its consensus protocol for big data storage of internet of things. *Journal of Internet Technology*, 17(7):1453–1460.
- Liu, Y., Zhang, L., Yang, Y., Zhou, L., Ren, L., Wang, F., Liu, R., Pang, Z., and Deen, M. J. (2019). A novel cloud-based framework for the elderly healthcare services using digital twin. *IEEE Access*, 7:49088–49101.
- Low, M. Y. H., Lye, K. W., Lendermann, P., Turner, S. J., Chim, R. T. W., and Leo, S. H. (2005). An agent-based approach for managing symbiotic simulation of semiconductor assembly and test operation. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 85–92.
- Makridakis, S., Spiliotis, E., and Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. *PLoS ONE*, 13(3).
- Makridakis, S., Spiliotis, E., and Assimakopoulos, V. (2020). The m4 competition: Results, findings, conclusion and way forward. *International Journal of Forecasting*, 34(4):802–808.
- Marmor, Y. N., Wasserkrug, S., Zeltyn, S., Mesika, Y., Greenshpan, O., Carmeli, B., Shtub, A., and Mandelbaum, A. (2009). Toward simulation-based real-time decision-support systems for emergency departments. In *Proceedings of the 2009 Winter Simulation Conference*, pages 2042–2053. IEEE.
- Mesabbah, M., Abo-Hamad, W., and McKeever, S. (2019). A hybrid process mining framework for automated simulation modelling for healthcare. In *Proceedings of the 2019 Winter Simulation Conference*, pages 1091–1102. IEEE.
- Mitchell, B. and Yilmaz, L. (2009). Symbiotic adaptive multisimulation: An autonomic simulation framework for real-time decision support under uncertainty. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 19(1):1–31.
- Moeuf, A., Pellerin, R., Lamouri, S., Tamayo-Giraldo, S., and Barbaray, R. (2018). The industrial management of smes in the era of industry 4.0. *International Journal of Production Research*, 56(3):1118–1136.
- Mustafee, N., Brailsford, S., Djanatliev, A., Eldabi, T., Kunc, M., and Tolk, A. (2017). Purpose and benefits of hybrid simulation: contributing to the convergence of its definition. In *Proceedings of the 2017 Winter Simulation Conference*, pages 1631–1645. IEEE.
- Oakley, D., Onggo, B. S., and Worthington, D. (2019). Symbiotic simulation for the operational management of inpatient beds: model development and validation using  $\delta$ -method. *Health Care Management Science*, pages 1–17.
- Onggo, B. S. (2019). Symbiotic simulation system (s3) for industry 4.0. In *Simulation for Industry 4.0*, pages 153–165. Springer.
- Onggo, B. S. and Karatas, M. (2016). Test-driven simulation modelling: A case study using agent-based maritime search-operation simulation. *European Journal of Operational Research*, 254(2):517–531.
- Onggo, B. S., Mustafee, N., Smart, A., Juan, A. A., and Molloy, O. (2018). Symbiotic simulation system: hybrid systems model meets big data analytics. In *Proceedings of the 2018 Winter Simulation Conference*, pages 1358–1369. IEEE.
- Pagès-Bernaus, A., Ramalhinho, H., Juan, A. A., and Calvet, L. (2019). Designing e-commerce supply chains: a stochastic facility–location approach. *International Transactions in Operational Research*, 26(2):507–528.
- Panadero, J., Juan, A. A., Mozos, J. M., Corlu, C. G., and Onggo, B. S. (2018). Agent-based simheuristics: extending simulation-optimization algorithms via distributed and parallel computing. In *Proceedings of the 2018 Winter Simulation Conference*, pages 869–880.

- IEEE.
- Pereira, A. and Romero, F. (2017). A review of the meanings and the implications of the industry 4.0 concept. *Procedia Manufacturing*, 13:1206–1214.
- Rhodes-Leader, L., Onggo, B. S., Worthington, D. J., and Nelson, B. L. (2017). Airline disruption recovery using symbiotic simulation and multi-fidelity modelling. In *In Proceedings of the 2017 Winter Simulation Conference*, pages 4588–4589. IEEE.
- Rhodes-Leader, L., Onggo, B. S., Worthington, D. J., and Nelson, B. L. (2018). Multi-fidelity simulation optimisation for airline disruption management. In *Proceedings of the 2018 Winter Simulation Conference*, pages 2179–2190. IEEE.
- Rogers, P. and Gordon, R. J. (1993). Simulation for real-time decision making in manufacturing systems. In *Proceedings of the 1993 Winter Simulation Conference*, pages 866–874.
- Rossmann, J., Schluse, M., Waspe, R., and Moshhammer, R. (2011). Simulation in the woods: From remote sensing based data acquisition and processing to various simulation applications. In *Proceedings of the 2011 Winter Simulation Conference*, pages 984–996. IEEE.
- Smyl, S. (2020). A hybrid method of exponential smoothing and recurrent neural networks for time series forecasting. *International Journal of Forecasting*, 36(1):75–85.
- Suh, W. (2019). Application of real-time sensor data to evaluation of performance of ad hoc distributed traffic simulation. *Sensors and Materials*, 31(5):1477–1488.
- Szozda, N. (2017). Industry 4.0 and its impact on the functioning of supply chains. *LogForum*, 13.
- Tan, K. W., Tan, W. H., and Lau, H. C. (2013). Improving patient length-of-stay in emergency department through dynamic resource allocation policies. In *2013 IEEE International Conference on Automation Science and Engineering (CASE)*, pages 984–989. IEEE.
- Taylor, S. J. and Letham, B. (2018). Forecasting at scale. *The American Statistician*, 72(1):37–45.
- Teixeira, E. L. S., Tjahjono, B., Alfaro, S. C. A., and Wilding, R. (2019). Extending the decision-making capabilities in remanufacturing service contracts by using symbiotic simulation. *Computers in Industry*, 111:26–40.
- Wu, Q., Sun, F., Xu, W., and Zhang, T. (2012). Using multilevel phase change memory to build data storage: A time-aware system design perspective. *IEEE Transactions on Computers*, 62(10):2083–2095.
- Xie, N., Dong, G., and Zhang, T. (2010). Using lossless data compression in data storage systems: Not for saving space. *IEEE Transactions on Computers*, 60(3):335–345.
- Xu, J., Zhang, S., Huang, E., Chen, C.-H., Lee, L. H., and Celik, N. (2014). Efficient multi-fidelity simulation optimization. In *Proceedings of the 2014 winter simulation conference*, pages 3940–3951.
- Yli-Ojanperä, M., Sierla, S., Papakonstantinou, N., and Vyatkin, V. (2019). Adapting an agile manufacturing concept to the reference architecture model industry 4.0: A survey and case study. *Journal of Industrial Information Integration*, 15:147–160.
- Zhang, F. and Zhao, Q. (2012). Runtime reconstruction of simulation models for dynamic structure systems. In *Asian Simulation Conference*, pages 27–36. Springer.
- Zhang, H., Ma, L., Sun, J., Lin, H., and Thürer, M. (2019). Digital twin in services and industrial product service systems: Review and analysis. *Procedia CIRP*, 83:57–60.