Proceedings of the 2022 Winter Simulation Conference
B. Feng, G. Pedrielli, Y. Peng, S. Shashaani, E. Song, C.G. Corlu, L.H. Lee, E.P. Chew, T. Roeder, and P. Lendermann, eds.

DIGITAL TWIN AS AN AID FOR DECISION-MAKING IN THE FACE OF UNCERTAINTY

Vinay Kulkarni
Souvik Barat

Tata Consultancy Services Research
54B, Industrial Estate
Hadapsar, Pune, 411013, INDIA

Tony Clark

Aston University
Department of Computer Science
Birmingham, B4 7ET, UNITED KINGDOM

Balbir S. Barn

Middlesex University
The Burroughs
London, NW4 4BT, UNITED KINGDOM

ABSTRACT

You have heard extensively about digital twins and how they can be applied in a range of domains and for a range of different purposes. Do you need help in understanding how digital twins can be used for decision making requirements of in the context of uncertainty? This introductory tutorial provides you with a detailed understanding of why digital twins are an important technology for exploring the unique phenomena of complex socio-cyber-economic ecosystems. The tutorial presents a critical analysis of the current state of the art and practice of digital twins in decision making and includes a focus on our actor based language called Enterprise Simulation Language (ESL). This language is an example of an infrastructure for developing these digital twins for complex systems decision making. The tutorial helps you understand how ESL and related technologies can be used effectively by presenting multiple examples of real-world problems.

1 MOTIVATION

We live in a hyperconnected world where most of the business and social systems comprise a large number of interconnected socio-cyber-physical systems, i.e., a complex system of systems. In the context of an uncertain environment, the relevancy of such systems is determined by their ability to adapt or change as necessary. Changes can occur along multiple dimensions, such as regulatory or legal regimes, technology advance/obsolescence, mergers and acquisitions, and black swan events like COVID-19 pandemic. Introducing the right change in the right system at the right time is critical in order to benefit from an opportunity or to mitigate a threat. Deciding on an effective change within shortening time constraints is not easy. It requires a deep understanding of aspects, such as structural decomposition into systems and subsystems, relationships between these (sub) systems, and emergent behaviour of systems and subsystems. Further, the scale of system, socio-technical characteristics, and complex interaction with its operating environment make the decision making process a challenging endeavor.

978-1-6654-7661-4/22/$31.00 ©2022 IEEE

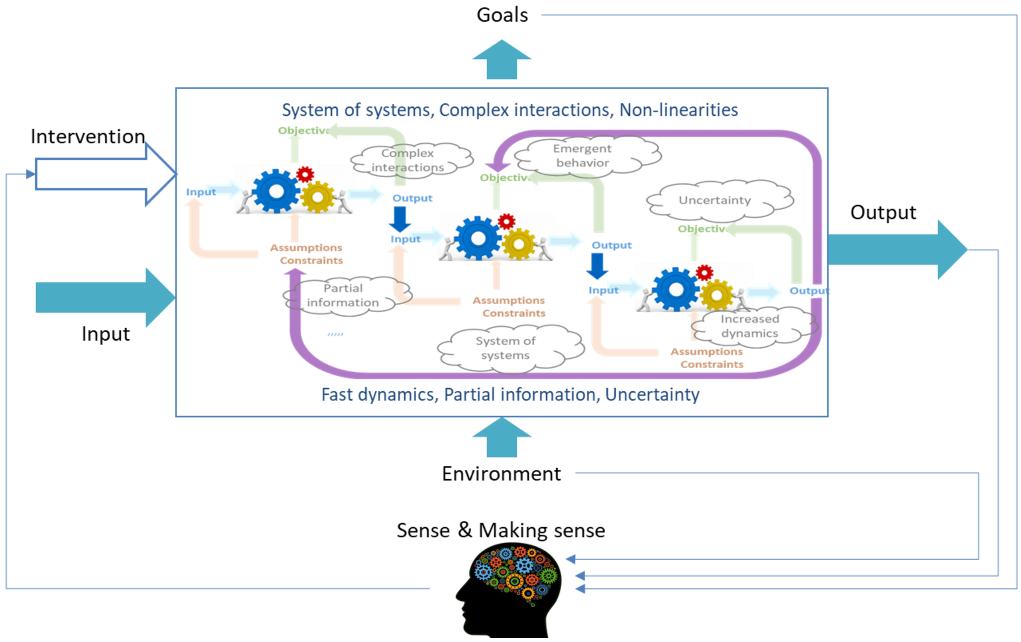

Figure 1. Decision-making in complex systems.

For decision-making, the system can be viewed as a transfer function from **Input** value space to **Output** value space as shown in Figure 1. A system exists to meet the stated **Goals** while operating in an environment which may constrain inputs and/or system behaviour. System goals are an objective function over an output value space, i.e. Measure, and may have temporal characteristics. Moreover, the goals can have a complex decomposition structure with inter-dependent and cross-cutting goals as shown in Figure 2. Designing a suitable transfer function is in itself a challenging task. The transfer function needs to be continually modified in response to changes in the environment and/or goals – the decision-making problem. This problem is further exacerbated when the available information about system is incomplete and distributed in fragments, and the environment is uncertain.

Current industry practice relies principally on human expertise to arrive at suitable corrective interventions. Typically, the decision maker operates with a fixed set of levers to nudge the system toward the desired state. This is an iterative process based largely on expertise and experience. Given the large size, complex goal decomposition, interfering goals, incomplete knowledge and inherent dynamism, this multi-criterion decision-making endeavor is typically an art. This intuition-based decision-making approach typically leads to iterative course corrections, and, more importantly, to suboptimal output or loss of business opportunities or dissatisfaction within beneficiaries or any combination of the three.

To address similar multi-dimensional decision-making problems, NASA explored the *Digital Twin* concept a decade ago for Apollo program. They crafted a digital twin of a flying vehicle using a multi-physics multi-scale probabilistic model, contextualized it with fleet history, and periodically synchronized using sensor updates to predict efficacy of a change in meeting the desired objectives. Use of digital twin helped them to understand the life of aircraft structure, ensure its integrity, design and maintain airframes, address the shortcomings of fleet management, and make the vehicle future ready (Shafto, et al., 2012).

While there are many definitions of Digital Twin, we adopt the notion of a digital twin as: a purposive virtual high-fidelity simulatable representation of the reality that is amenable to what-if and if-what scenario playing.

The digital twinning of physical systems, those governed by laws of physics, chemistry, thermodynamics etc. – for decision-making has seen good adoption. It is notable that adoption of digital twin concepts is less prevalent in other domains such as business systems and social-technical systems for

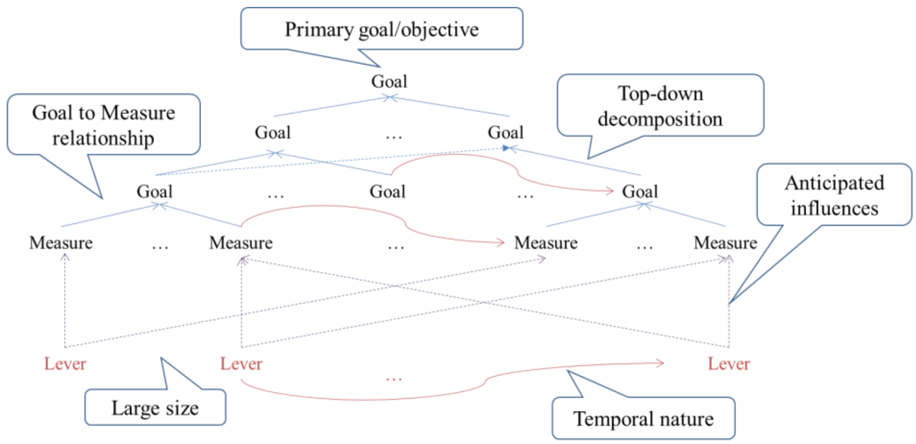

Figure 2. System goals, measures and levers.

which various reasons are offered such as: system behaviour not governed by well-defined laws, the emergent nature of system behaviour, the lack of expressiveness of the existing modelling techniques to faithfully capture system behaviour, and inability to comprehend dynamics of such systems (Sandkuhl, et al. 2016).

Our experience of applying digital twin technology for decision-making in business and social space (Kulkarni, et al. 2019; Barat, et al. 2020) also supports this observation. We believe that extending state-of-the-art modelling & simulation (M&S) techniques, and combining them meaningfully with proven ideas from Artificial Intelligence (AI) and Control Theory can comprehensively address this decision-making problem. In this tutorial, we will highlight the complexity of business and social systems, required extensions to the state-of-the-art M&S technology to make it relevant and our approach that combines M&S, AI and control theory. We will elaborate our ideas using four impact case studies from cyber-physical system, business system and social systems.

2 STATE OF THE ART AND CURRENT PRACTICE

Decision-making in the face of "deep" uncertainty (Marchau, et al. 2019) typically follows three key steps: decision framing, strategy evaluation, and trade-offs. Decision framing step focuses on precise definition of goals/objectives, goal decomposition structure, measures that indicate achievements of the specified goals, and potential options/changes that can be introduced to achieve goals (i.e., levers). Evaluation step aims to predict possible efficacy/consequences for potential options, i.e., understanding of measures. This step should consider all inherent complexities of the system under consideration that include uncertainties within the system and in environment, possible emergent behaviours, and all spatiotemporal influencing factors as highlighted in Figure 1. Finally, the Trade-off step analyzes predicted measures with respect to the goals to arrive at the best possible decision. The entire process needs to be repeated when a new option is available and/or there is a change in the system or environment or goal.

To date, industrial practice has adopted different form of human-centric and machine-assisted approach for the three decision-making steps outlined above. The business and social domain is primarily the preserve of intuition-centric approaches (Liebowitz et al. 2019). The key struggle for the pure intuition-centric approach is an inability to justify one option over other alternatives, i.e., trade-offs. Additionally, evaluation of options are based on intuitions,- thus, they are vulnerable to the law of bounded rationality. Machine assisted approaches, in contrast, make use of quantitative evaluation techniques and adopts statistical rigour for trade-offs. Evaluation of operations is broadly addressed using one of the three approaches: Optimization technique, Historical data-centric AI-based approach and , and Enterprise model-based approach.

Optimization: Expressing the desired behaviour as an optimization problem to be solved using rigorous mathematical techniques. This approach requires possible behaviour of the system and its operating environment to be known a-priori and expressed in pure analytical terms. However, this is possible only when the behaviour is governed by laws of nature, physics, thermodynamics, chemistry etc. For instance, controlling the boiler of a captive power plant (Biswas, et al. 2017), scheduling the crude arrival and mixing for a refinery (Wagle and Paranjape, 2020). Complexity and uncertainty arising from issues such as human behaviour, cyber-physical interactions, and communications with a changing environment, lead to an enterprise exhibiting emergent behaviour that is difficult to represent in terms of mathematical equations.

Historical data-centric AI models: Build AI/statistical model from the past data using machine learning techniques and use it to identify best change/option (i.e., lever) for achieving the stated goal of a system. In the light of recent advances in machine learning, this approach holds a lot of promise. However, the approach relies on two key conditions:

1. Past data must be available in order to learn a model and must be a representative of all possible behaviours of the system.
2. The future behaviour of interest must be a linearized extrapolation of the past as represented by the historic data.

Given these two conditions, machine learning can produce effective results (Parkhi, et al. 2020). However, the intrinsic uncertainty and incompleteness of historic data means that it is unlikely that the conditions are sufficiently met in order to apply AI-based learning techniques without additional approaches.

Enterprise model based approach: Enterprise Modelling community has approached analysis of system for the purpose of decision-making by adopting state of the art modelling and simulation techniques. Such an approach imitates a system using a (purposive) model, explores a range of levers by simulating a model that incorporates levers related changes, and develops simulation-led evidences for trade-offs. Existing modelling and simulation approach captures a system by adopting one of the broad paradigm: top-down or bottom-up (Thomas and McGarry, 1994). A top-down paradigm visualizes a system as a whole and often rely on reductionist view to decompose a large decision problem into smaller parts and address them in isolation. The enterprise model (EMs) approaches, such as ArchiMate (Jacob, et al. 2012), i* (Yu et al. 2006), BPMN (White, 2008), and System Dynamics (SD) (Meadows and Wright, 2008), are examples of top-down models that are used extensively for analyzing systems, such as Enterprises. The key concerns with top-down approaches are: they are not cognizant of individualistic characteristics of the constituent (micro) elements of the system and emergent behaviour. Construction of a system model using top-down approach expects information about whole enterprise, which is a difficult expectation for a large enterprise.

A bottom-up approach, in contrast, starts from the parts or micro-behaviours and arrives at a holistic view of a system through composition. The bottom up approach uses the agent and actor based technologies, such as Erlang (Armstrong, 1996), Akka (Allen, 2013), and Scala Actor (Haller and Odersky, 2009), for modelling and analyzing systems. They are capable of observing emergent behaviour but not capable of representing complex structure and not cognizant of uncertainties.

3 PROPOSED APPROACH

We envisage a line of attack that borrows proven ideas from modelling and simulation, control theory and artificial intelligence. We advance the state of the art by using these state of the art approaches in an integrated manner as shown in Figure 3 (a). At the heart of this line of attack is the concept of Enterprise Digital Twin (EDT) – a virtual hi-fidelity machine processable representation of system that can be contextualized for any business system or social system, such as a city. The novelty of our approach is innovative and integrated use of: a) a meta-model for decision-framing (GML structure), b) a simulatable language ESL (Clark, et al. 2017) that's capable of specifying the required techno-socio-economic aspects of systems, and c) digital-twin centric exploration where we adopt relevant techniques to construct and validate a simulatable model amenable to solution space exploration through what-if / if-what analysis. A GML structure precisely captures goals, goal decomposition structure, their relationships, measures and levers that symbolically represent (hypothetical) changes, interventions, and perturbations. A sufficiently high-fidelity close-to-real representation, i.e., a digital twin specified using ESL, provides an environment for "in silico" experimentation where experts subject the digital twin to a variety of perturbations, i.e.,

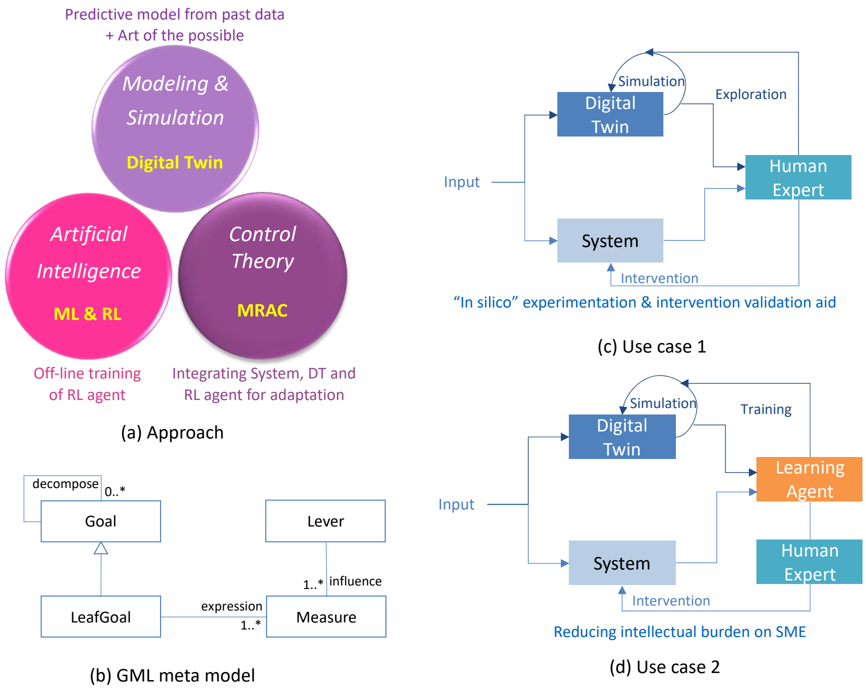

Figure 3. Digital twin centric approach to decision-making

experimentation through affectation of levers. As the digital twin is a hi-fidelity representation of the system, its Response to a perturbation, i.e., simulated estimation of Measures, is in the ballpark of actual system response. Experts can interpret multiple responses (simulation trends) in the light their experience and knowledge of the solution space to ascertain if the desired goals are met, and to arrive at a candidate set of suitable interventions i.e., affectation of levers. The set can be validated for correctness and efficacy using the digital twin itself by running multiple and appropriate simulations thus leading to identification of the most suitable lever to be effected. Thus, the digital twin considerably reduces the need for real life experimentation with the system and leads to significant savings in time, cost and effort.

Though highly useful as "in silico" experimentation aid, the digital twin does not reduce intellectual burden on human experts. To this end, we use an AI technique known as reinforcement learning. Basically, we use digital twin as "experience generator" from which the reinforcement learning agent (RL agent) learns *what action to perform when* to achieve the stated goal. We bring the system, digital twin and RL agent together in an adaptation architecture based on Model Reference Adaptive Control (MRAC) paradigm (Osburn et al, 1961). Our future research plans include ideas to extend this architecture to support dynamic adaptation where even the goals can change over time.

Our digital twins are more than traditional simulation models. There are many definitions of digital twins which have been framed by key characteristics such as rate of synchronization with a real world artifact, the context of the real world artifact, the purpose of the digital twin and the level of machine led adaptation possible. Our digital twin places these characteristics foremost in the design process. For example, adaptation is possible through the MRAC paradigm.

3.1 GML Structure – Decision Framing

The process by which decisions are made starts with a decision framing exercise, where experts define the scope by capturing **goals** and sub-goals, **measures** - qualitative/quantitative criteria that indicate the reachability of goals and help to rank levers, and **levers** – potential changes/perturbations. We use a meta-model as described in Figure 3 (b) to capture Goal, Measure and Lever (GML) in a precise form – it guides

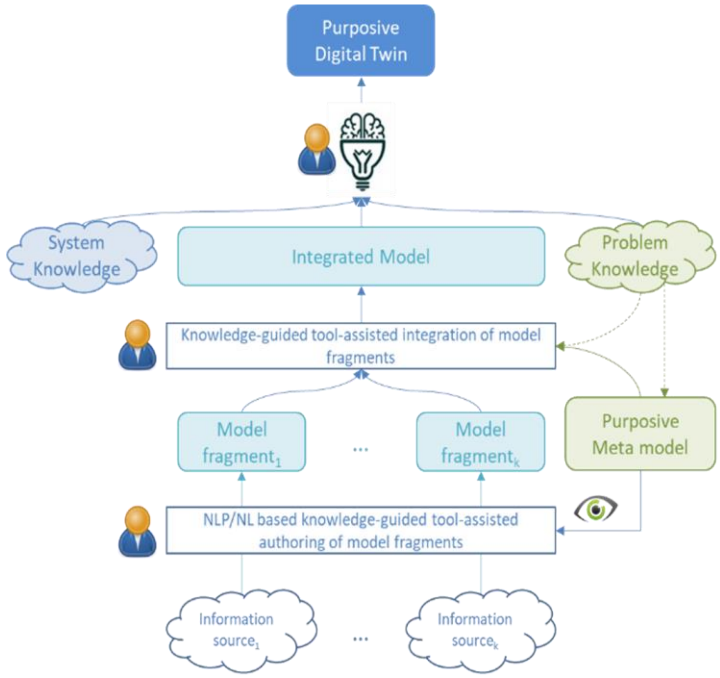

Figure 4. Knowledge-guided tool-assisted construction of digital twin

the decision space exploration and helps to capture the experience in a systematic manner. GML is an extension of i* specification (Yu et al. 2006).

3.2 ESL - Technology Infrastructure

We have developed an actor based language, called ESL (Clark, et al. 2017), by extending the canonical concept of "actor" (Agha, et al. 1997) to specify digital twins of a techno-socio-economic system (of systems). It helps represent a system as a set of intentional (i.e. there is a well-defined goal to be achieved), autonomous (i.e. capability to achieve the goal in a pro-active manner) and composable (i.e. an actor can be realized in terms of a set of fine-grained interacting actors) actors. An actor tries to achieve its stated goal by responding suitably to the events of interest and by exchanging messages with other actors. We extend the canonical "Actor" abstraction to include stochastic behaviour, i.e. a probability distribution of actions associated with an event, and the notion of "time" to capture temporal relations. We translate ESL to Java – and in fact have an implementation in the form of Java library so that digital twin can be implemented in pure java making use of the aforesaid Java library. We have also developed a simulator for ESL.

3.3 "In silico" Experimentation Aid

An Enterprise Digital Twin (EDT) is a virtual, high fidelity representation of a complex system of systems that is amenable to rigorous quantitative analysis through what-if and if-what scenario playing using real data to facilitate local optimality, global robustness and continuous learning. A decision-making problem starts with a precise GML structure. It is constructed using a manual knowledge-guided top-down approach considering high-level goals as a starting point. Domain experts elaborate the goals, define appropriate measures, and list potential levers to sufficiently frame a decision problem.

Construction: Construction of purposive EDT for a GML follows a bottom-up paradigm where constituent elements and sub-systems are identified by reflecting on the domain under consideration and captured using ESL actors. Actors represent domain concepts and compose/decompose to closely represent the notion of system of systems. Domain uncertainties are captured as stochastic behaviours within the actor

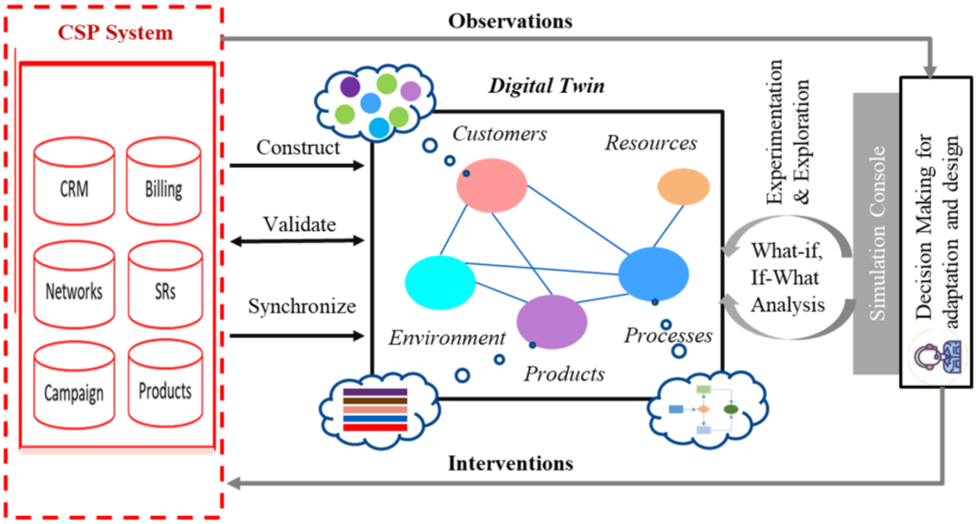

Figure 5. Digital twin for maximizing CLV in telecom

specification. The required domain information to be captured in a purposive DT comes from a wide spectrum of semi-/un-/structured sources including data bases, execution logs, standard operational procedure notes, policy documents and understanding of domain experts. As a result, wide-ranging fields of expertise are required to manually create a digital twin model – clearly a time-, cost-, effort- and intellect-intensive endeavor. The purposive meta model serves as a lens to mine or author appropriate model fragments, corresponding to a view of the purposive meta model, from these information sources. The meta model also serves to integrate these model fragments in ensuring correctness and internal consistency. Knowledge of the problem domain and system helps construct the purposive digital twin from the integrated model. Figure 4 presents a pictorial overview of the framework we have developed for accelerated creation of digital twin models from information available in semi- / un-/ structured form. It comprises of automation aids based on: (i) Natural Language Processing (NLP) and Machine Learning (ML) techniques for gathering the desired information from a given information source, (ii) Meta model driven techniques for integration and reconciliation of the model fragments, and (iii) Model validation based techniques for identifying the missing model fragments.

Validation: The utility and efficacy of a constructed digital is largely dependent on comprehensiveness and closeness of the digital twin with respect to the real system. We provide two established ways of validating a simulation model – conceptual validity and operational validity (Sargent, 2004). In conceptual validity, the domain experts certify comprehensiveness of the constructed actor-topology from a domain perspective. We ensure operational validity through simulation wherein the constructed digital twin (suitably initialized) is subjected to known past events leading to a simulation trace which is then examined to ascertain the resultant behaviors are identical to the ones from the past. Our simulation engine generates rich execution traces containing the detailed information necessary for analysis. We have developed a pattern language to specify the desired behaviour and a pattern matching engine to look for the specified patterns in the simulation trace (Clark, et al., 2017). This generic solution to ascertain correctness can be further augmented by manual validation of the input and output and control variables of the simulation.

A validated EDT supports knowledge-guided simulation-led evidence-based decision-making as shown in Figure 3(c) and (d). The key tenets of the approach are:

Analysis: Data-backed explanation of why things are the way they are i.e. "holding a mirror".

Design: Exploring the solution space in a evidence-backed manner i.e., "art of the possible".

Control: Bring the system back to the desired state in response to changes in the environment or perturbations to input.

Adaptation: Similar to 'Control' but in a significantly highly dynamic setting wherein goals may also change.

Transformation: A sequence of interventions to nudge the system from as-is state to the desired to-be state in a data-driven justification-backed manner.

Simulation of a digital twin of an as-is system helps *analysis.* Simulation of the digital twins of a hypothetical to-be system and possible design alternatives enables design space exploration. Control, adaptation and transformation can be experimented using a digital twin of an as-is system with the lever specification incorporated in the model. Sufficient in-silico experimentation using digital twin considerably reduces the need for real life experimentation with the system and leads to significant savings in time, cost and effort (as shown in Figure 3 (c)).

4 REAL WORLD INDUSTRY SCALE USE CASES

The approach and supporting tools have been validated for utility and efficacy by constructing applications that address real world problems spanning a wide spectrum of business verticals. This section provides case studies based on these applications. Details of the applications and their associated research outcomes have been described in detail elsewhere (Barat, et al., 2019; Barat, et al., 2020; Barat, et al., 2021; Ghosh, et al., 2021). A representative sample is presented below.

4.1 Customer Lifetime Value Optimization in Telecom

The Telecom Space can be viewed along two broad dimensions: physical systems of telecom network infrastructure, and business systems that offer products to customers through various processes. This case study focuses on the business system where the key objective is to assist Sales, Product and Customer Care heads to fulfil their respective goals using digital twin technology.

The principal goal of the Head of Sales is to acquire as many valid customers as possible given the current products portfolio and customer care processes. The principal objective of the Head of Products is to design (and/or personalize) products that will best meet the communication needs of prospective (and/or existing) customers given the current customer base and customer care processes. The principal objective of the Head of Processes is to help acquire new customers and retain existing customers at minimum cost given the current customer base and products portfolio.

Communication Services Providers (CSP) operate in arguably the most dynamic of business verticals. Current practice is for the three heads to operate in silos to arrive at strategies aimed at achieving their stated goals in an independent manner. These strategies are evaluated in a sandbox environment using an A/B technique. This is a time-, cost-, and effort-intensive endeavor. While this strategy could be the most suitable for achieving the local objective (i.e. Sales, Product or Process), there is no way to ascertain the ripple effect of its introduction on the other two aspects. As a result, CSPs end up operating in a reactive (or catch-up) mode for most of the time with little a-priori assurance of correctness. Therefore, this domain is characterized by high customer churn, high time to market for products, expensive product launches, long tail of inactive products, arguably the lowest customer satisfaction amongst all business verticals, and poor Customer Lifetime Value (CLV).

Additionally, there is a greater need to personalize the overall user experience for individual customers. Advances in technology such as 5G and 6G are increasingly making such personalization possible and the current practice of bucketing customers into problem-specific clusters is falling well short of expectations.

We addressed the problem of CLV optimization for individual customers of several leading CSPs worldwide. Typically, a CSP major has tens of millions of customers, hundreds of active products, and hundreds of processes. As a result, the scale of the problem is very large and associated with high dynamics.

We constructed a purposive digital twin (Barat, et al., 2020) that encapsulates the knowledge of Products, Processes, and Individual customers (e.g., age, gender, station of life, early adaptor/laggard etc.) leading to a good estimate of product and service expectations. A simulation-based iterative process helped us arrive at the right configuration of Products (features, launch strategy, price point etc.), Processes (SLAs, technology investment, price point etc.), and Customers (who will opt for which product at what price point, what's the best way to sell/upsell/cross-sell, how best to retain at what cost etc.) as shown in Figure 5. The

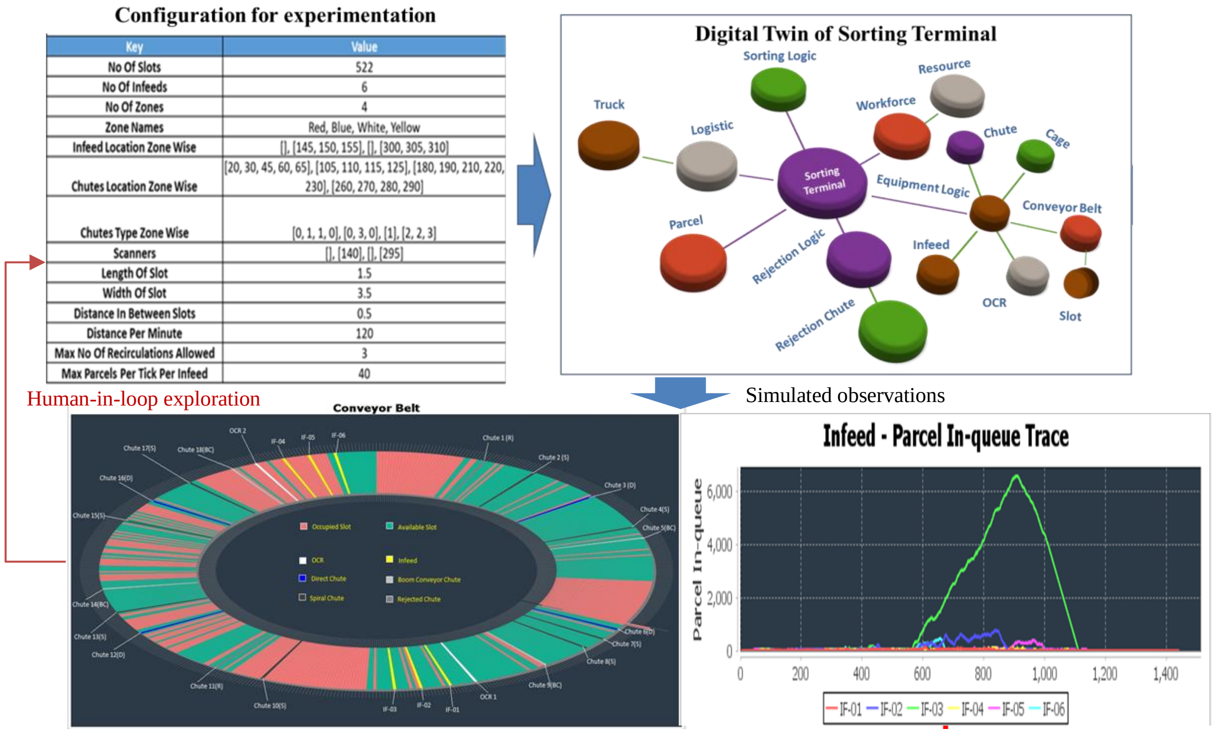

Figure 6. Digital twin for optimizing throughput of a sorting terminal.

measures we focused on were: Average monthly revenue, Average cost of servicing customers, Number of new customers added, and Number of customers left.

This DT was used to explore and define product and its launch time (i.e., levers) for a large US telecom organization. Multiple explorations helped them to define a segment-based product launch strategy (i.e., a set of products and their launch time) with expected 2.1% Take rate and 17% Retention rate (simulated measures) from second month of the product launch. In reality, Take rate was 2.5% and Retention rate was 16%. Furthermore, DT was used to fine tune the product launch strategy thus delivering 2x improvement.

Other problems in the telecom space that we have addressed include minimizing customer churn, optimizing customer care processes, and designing product launches. We could also repurpose this DT for the Banking domain to address the problem of minimization of Non-Performing Assets.

4.2 Maximizing throughput of Sorting Terminals

Package sorting is a critical activity in the package delivery industry. Sorting Terminal is a machine-led human in the loop cyber-physical system that aims to maximize throughput for efficiently routing packages to appropriate destinations, ensuring faster delivery and improving utilization of physical infrastructure. Figure 6 depicts schematic of a typical sorting terminal comprising of: Conveyer belt (called Sorter) that carries the packages to be sorted, Infeed to introduce packages from carriers to Conveyer belt, Scanner for identifying the destination delivery zone from address written on the package, Chutes to collect the packages, and Robotic arm (in front of every chute) to push the package into chute.

Each chute is assigned a team that collects the packages to be taken to designated loading stations. The chute can be configured along several parameters, such as number and types of operational chutes, sorter speed, placement of scanner, assignment of destination delivery zones to a chute, assignment of collecting team to a chute, and possible schedule of collecting-team-to-chute assignment. Some of these parameters are more static than others e.g., placement of scanner. For maximal throughput of the sorting terminal, the following conditions must hold (i.e., measures to be satisfied):

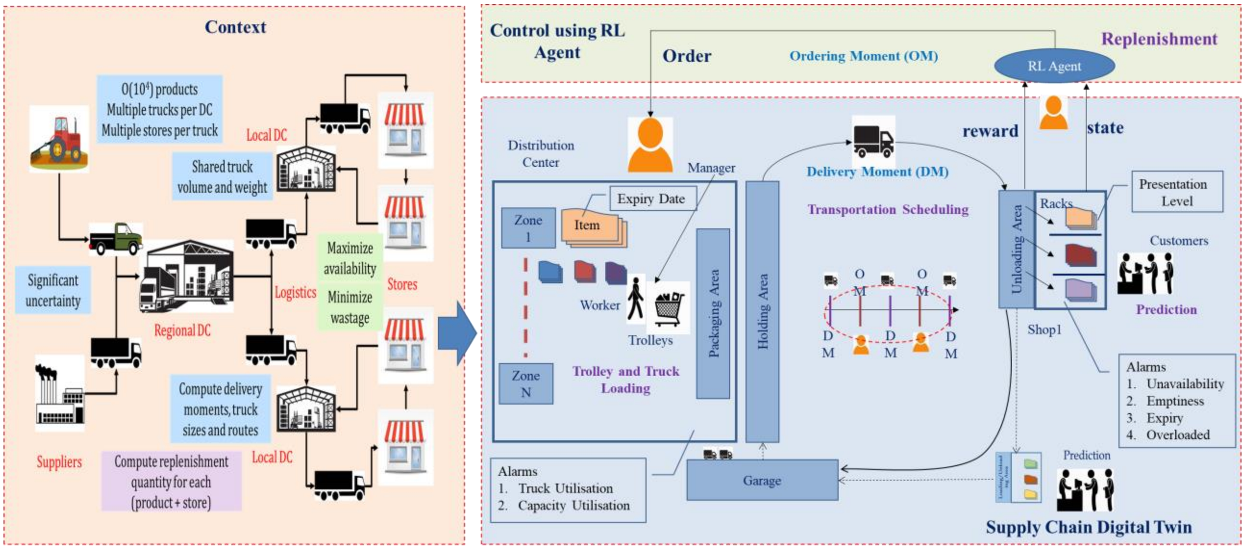

Figure 7. Digital twin for responsive and optimal shop stock replenishment

(i) A package should minimize the time spent on the sorter,
(ii) A package should get collected in the right chute,
(iii) Chutes should get emptied as quickly as possible,
(iv) No chute should ever be without a collecting team,
(v) No package should remain unsorted and thereby require manual handling.

Current practice is to predict the workload for a shift from past workloads using a statistical prediction algorithm. Knowledge of workload is then used to define suitable sorting terminal configuration that includes a sorting plan (package to chute assignment) and the number of active chutes and resources allocated to the chute. This step is part estimate part art. Significant uncertainty in shift workloads, high variance in the temporal order of the workload, and varying skills of collecting team are the principal sources for deviation from business-as-usual operation during the shift. These outliers are addressed by the shift manager in an ad hoc manner and, as a result, there is little a-priori assurance regarding the throughput of the sorting terminal for a given shift.

To overcome this problem, we constructed a digital twin of the sorting terminal (Ghosh, et al., 2021) as shown in Figure 6. We experimented with a wide range of what-if simulations to arrive at the right configuration and also to be prepared for possible outlier conditions. The parameters we focused on were: Average time to chute, Average chute clearing time, Number of unassigned packages, and Number of corrective actions introduced. We also ran the digital twin in parallel with the actual sorting terminal in a shift to serve as an early warning system. Moreover, plausible solutions to mitigate the outlier condition were worked out "in silico" – the proverbial "forewarned is forearmed" situation.

4.3 Optimizing Shop Stock Replenishment for a Retail Chain

Optimum stock replenishment is a critical requirement for the retail industry. We addressed this need for a grocery retailer with a network of stores and warehouses served by a fleet of trucks for transporting products. The goal of replenishment is to regulate the availability of the entire product range in each store subject to the spatiotemporal constraints imposed by (i) available stocks in the warehouses, (ii) labour capacity for picking and packaging products in the warehouses, (iii) the volume and weight carrying capacity of the trucks, (iv) the transportation times between warehouses and stores, (v) the product receiving capacity of each store, and (vi) available shelf space for each product in each store. A schematic of the product flow is shown in Figure 8. The macro-level uncertainties that emerge in the whole replenishment process are due to the probabilistic behaviours of the individual elements. For example: unavailability and

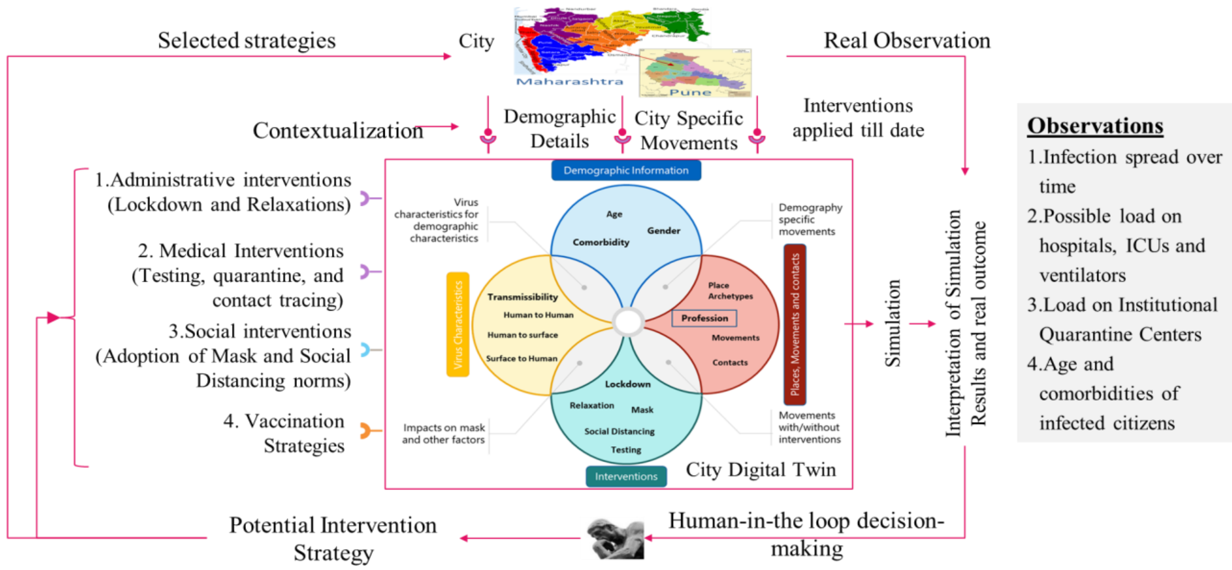

Figure 8. City digital twin for prediction and control of COVID-19 pandemic

varying productivity of the resources, sub-optimal machine throughput and unavailability and unaccounted delays of the trucks. Trucks are constrained by the volume and weight capacities, often they are suited for specific types of products and each of them has probabilistic characteristics, such as: propensity for transportation delay and break down.

Current practice is to construct a predictive model based on the past data pertaining to shop stock replenishment with above details. Given the scale of the problem i.e., number of warehouses, number of stores, number of trucks, number of product types, and number of products, purely analytical specification gets too large and hence vulnerable to the errors of omission and commission. As a result, an aggregated lumped up model is used to train a Reinforcement Learning (RL) agent to learn a policy i.e., a sequence of actions to be performed by the various stakeholders for shop stock replenishment. The coarseness of model leads to the RL agent learning a sub-optimal policy which turns out to be satisfactory.

We constructed a fine-grained model of the supply chain (Barat, et al., 2019) where each warehouse, shop, truck, container, shelf, product, customer etc. is modelled as an actor having well-defined (though probabilistic) behaviour. We are able to model the individual characteristics such as buying propensity of a customer, breakdown vulnerability of a truck, packing errors of a packer etc. at the finest level of detail. This fine-grained model of the supply chain is used to train RL algorithm to learn a policy for shop stock replenishment. Our approach led to fine-grained learning with faster convergence and the resulting digital twin can be easily repurposed to address similar problems for other supply chains.

4.4 Prediction and Control of COVID-19 Pandemic in a City

In the midst of a pandemic like COVID-19, one of the key priorities of the a public health administration is to understand the dynamics of the transmission of virus (World Health Organization 2020;) and use that knowledge to design effective control measures to keep its impact on public health within manageable and tolerable limits. In the case of COVID-19, while the characteristics of the virus (i.e., mode of transmission and the typical trajectory of infection in an individual) are known to an extent from the existing research (World Health Organization 2020; Asadi et al. 2020; Cai et al. 2020; Wang et al. 2020; He et al. 2020), the dynamics of its transmission and spread in a heterogenous population is not fully understood. It is known, though, that the spread of infection is related to people's movement, the nature of the area where people congregate (open-air versus closed), and number and frequency of proximal contacts. It is also known that demographic factors and comorbidity play a role in the lethality. Therefore, the primary non-pharmaceutical intervention (NPI) of the public health authorities has been to restrict people's movement to varying degrees

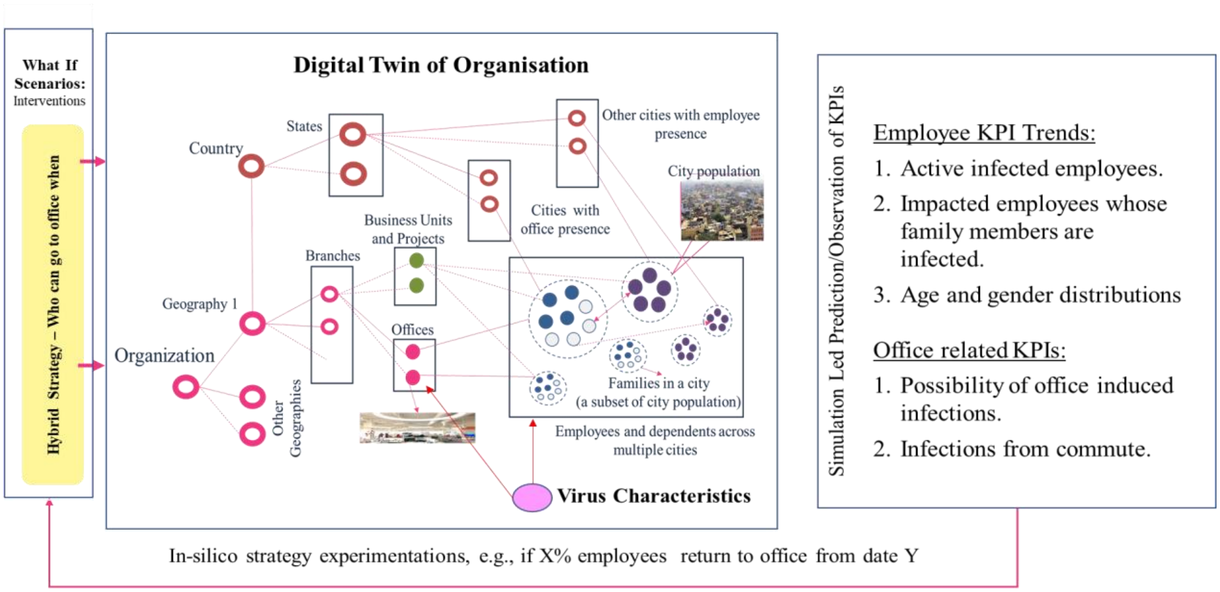

Figure 9. Organisation digital twin for WFH to WFO transition.

through the so-called lockdowns. In addition to saving lives, lockdowns have been the primary instruments for managing the load on local healthcare systems.

The economic impact of the lockdowns imposed in 2020 has been recorded as being amongst the most adverse phenomena to impact the world economy (Fernandes 2020). Until the pandemic is brought under control through largescale availability of medication or vaccines, the administrators need to decide whether or not lockdowns are needed, and their nature and duration. As such, there is no universal formula for answering these questions because the dynamics of the spread of COVID-19 depend heavily on individual localities: their demographic profile, the prevalent social etiquette, the capacity of their healthcare systems, whether or not people comply with the administrative recommendations, etc. Therefore, devising effective tools and models (possibly on a continuously changing basis) to help administrators take decisions at a local level is an urgent requirement in the midst of the pandemic.

Use of statistical and mathematical models to understand the spread of a virus and to explore effective control measures is a well-established decision-making aid (Hethcote 1989; Marathe and Vullikanti 2013). A wide range of modeling, data visualization, and interpretation techniques have been developed to predict the spread of COVID-19 and to explore the efficacy of NPIs (Wynants et al. 2020). While some models have been found to be useful for exploring NPIs in a specific geography, others have been found wanting for their accuracy of prediction (Holmdahl and Buckee 2020).

We believe that a universal model to predict the efficacy of NPIs for all geographies, countries, and cities across the world is a difficult proposition. Instead, a purpose-specific, locality-based, fine-grained model addressing a set of relevant aspects of interest can play a crucial role in decision-making for controlling the pandemic. Therefore, we developed a novel agent-based digital twin of a city (Barat, et al., 2021) to predict and control the COVID-19 epidemic as shown in Figure 8. The defining characteristic of the city digital twin is a set of suitable agent types necessary to capture heterogeneity in terms of people, places, transport infrastructure, health care infrastructure, leading to a fine-grained model of the city that is amenable to what-if and if-what scenario playing. We populated the city digital twin using data from the city administration together with suitable augmentation. The fine-grained nature of the digital twin enabled us to address the critical concerns such as the rate and the extent of the spread of the epidemic, demographic, and comorbidity characteristics of the infected people, load on the healthcare infrastructure in terms of specific needs such as number of admissions requiring critical care (supplementary oxygen, ventilator support, intensive care, etc.), load on institutional quarantine centers, and so on. We set up appropriate what-if scenarios to identify the most effective intervention from the candidate set to control epidemic as well as bring back normalcy. We vetted the simulation results against epidemic-related data released by an Indian city.

4.5 Helping Organizations Transition from Work From Home to Work From Office mode

Organizations are struggling to ensure business continuity without compromising on delivery excellence in the face of COVID-19 pandemic related uncertainties. The uncertainty exists along multiple dimensions such as virus mutations, infectivity and severity of new mutants, efficacy of vaccines against new mutants, waning of vaccine induced immunity over time, and lockdown/opening-up policies effected by city authorities. Moreover, this uncertainty plays out in a non-uniform manner across nations, states, cities, and even within the cities thus leading to highly heterogeneous evolution of pandemic. While Work From Home (WFH) strategy has served well to meet ever-increasing business demands without compromising on individual health safety, there has been an undeniable reduction in social capital. With COVID-19 showing definite waning trends and employees beginning to miss the office environment, several organizations are considering the possibility of safe transition from WFH to Work From Office (WFO) or a hybrid mode of operation. An effective strategy needs to score equally well on possibly interfering dimensions such as risk of infection, project delivery, and employee (and their dependents) wellness. As large organizations will typically have a large number of offices spread across a geography, the problem of arriving at office-specific strategies becomes non-trivial. Moreover, the strategies need to adapt over time to changes that cannot be deduced upfront. This calls for an approach that’s amenable to quick and easy adaptation.

We developed a Digital Twin centric approach (Figure 9) that: (i) Leverages pure data-centric statistical model, coarse-grained system dynamic model, and fine-grained agent-based model, (ii) Helps human experts arrive at pragmatic strategies to effect WFH to WFO transition keeping the key stakeholders satisfied, and (iii) Easily adapt the strategies over time. We have validated the approach with a large organization and the results are encouraging.

5 SUMMARY AND FUTURE WORK

We live in a world comprised of complex systems of systems that change along multiple dimensions in manner that cannot be deduced upfront. In this context, decision-making in the face of uncertainty is a critical need. We argue that state of the art and current practice does not address this need to the desired level of satisfaction and sophistication.

We propose an innovative simulation-based approach that integrates and builds upon proven ideas from M&S, Artificial Intelligence, and Control Theory. At the heart of the approach is the concept of Digital Twin as a purposive virtual high-fidelity virtual simulatable representation of reality that is amenable to what-if and if-what scenario playing. We described the core technology infrastructure necessary to implement an approach that places “human in the loop” as key requirement. We establish the utility and efficacy of the approach in terms of real world industry scale applications spanning business (telecom case study and retail case study), cyber-physical (sorting terminal case study), and societal (COVID-19 case studies) domains. Almost everywhere the proposed approach has fared better than current practice.

A particular insight to highlight is that we consider the city digital twin can be repurposed to address emerging socio-techno-economic challenges, such as healthcare, sustainable enterprise and smart city. While our broad trajectory is well-directed our future research road map indicates that there is a lot that needs to be done as regards: multi-paradigm digital twins, multi-objective reinforcement learning, adaptive digital twins, method support for construction and use of digital twins, and domain-specific knowledge. We also see possibility of taking the idea of digital twin to the software systems to support dynamic adaptation in the face of uncertainty. The combined ability of arriving at the right decisions and effecting them in an efficacious way will help realize an adaptive enterprise where adaptations are justification-backed.

REFERENCES

Allen, J. 2013. *Effective Akka*. Sebastopol, California: O'Reilly Media, Inc.

Agha, G. A., Mason, I. A., Smith, S. F., and Talcott, C. L. 1997. "A Foundation for Actor Computation". *Journal of Functional Programming*, 7(1), 1-72.

Armstrong, J. 1996. "Erlang - a Survey of the Language and its Industrial Applications". In *Proceedings of the Symposium on Industrial Applications of Prolog (INAP)*, October 16th – 18th, Hino, Tokyo, 1-8.

Asadi, S., Bouvier, N., Wexler, A.S. and Ristenpart, W.D. 2020. "The Coronavirus Pandemic and Aerosols: Does COVID-19 Transmit via Expiratory Particles?". In *Proceedings of Aerosol Science and Technology*, Taylor & Francis, June 2nd, 54(6):635-638.

Barat, S., Khadilkar, H., Meisheri, H., Kulkarni, V., Baniwal, V., Kumar, P., and Gajrani, M. 2019. "Actor Based Simulation for Closed Loop Control of Supply Chain using Reinforcement Learning". In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems*, May 13th – 17th, Montreal, Canada, 1802-1804.

Barat, S., Kulkarni, V., Kumar, P., Bhattacharya, K., Natarajan, S., and Viswanathan, S. 2020. "Towards Effective Design and Adaptation of CSP Using Modelling and Simulation Based Digital Twin Approach". In *Proceedings of the 2020 Summer Simulation Conference*, July 20th – 22th, Virtual, Spain, 1-12.

Barat, S., Parchure, R., Darak, S., Kulkarni, V., Paranjape, A., Gajrani, M., and Yadav, A. 2021. "An Agent-Based Digital Twin for Exploring Localized Non-pharmaceutical Interventions to Control COVID-19 Pandemic". *Transactions of the Indian National Academy of Engineering*, 6(2), 323-353.

Biswas, J., Kumar, R., Mynam, M., Nistala, S., Panda, A., Pandya, R., Rathore, R., and Runkana, V. 2017. "Method and System for Data Based Optimization of Performance Indicators in Process and Manufacturing Industries". Patent No. US10636007B2.

Cai J., Sun W., Huang J., Gamber M., Wu J., and He G. 2020. "Indirect Virus Transmission in Cluster of COVID-19 Cases", *Emerging Infectious Diseases* 26(6):1343–1345

Clark, T., Kulkarni, V., Barat, S., and Barn, B. 2017. "ESL: An Actor-based Platform for Developing Emergent Behaviour Organisation Simulations". In *International Conference on Practical Applications of Agents and Multi-agent Systems*, June 21st -23rd Porto, Portugal, 311-315.

Clark, T., Barn, B., Kulkarni, V., and Barat, S. 2017. "Querying Histories of Organisation Simulations". *In Proceedings of the Information Systems Development*, September 6th – 8th, Larnaca, Cyprus, 1-12.

Fernandes, Nuno. 2020. Economic Effects of Coronavirus Outbreak (COVID-19) on the World Economy. IESE Business School Working Paper No. WP-1240-E, https://ssrn.com/abstract=3557504, accessed 27th Feb 2022.

Haller, P., and Odersky, M. 2009. Scala Actors: Unifying Thread-based and Event-based Programming. *Theoretical Computer Science*, 410(2-3): 202-220.

Ghosh, S., Pal, A., Kumar, P., Ojha, A., Paranjape, A.A., Barat, S., and Khadilkar, H. 2021. "A Simulation Driven Optimization Algorithm for Scheduling Sorting Center Operations". In *Proceedings of the 2021 Winter Simulation Conference*, edited by S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Loper, 1-12. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Hethcote, H.W. 1989. "Three Basic Epidemiological Models". In *Applied Mathematical Ecology,* edited by S.A. Levin, T.G. Hallam, and L. J. Gross, 119-144. Heidelberg: Springer.

He, S., Peng, Y., and Sun, K. 2020. "SEIR Modeling of the COVID-19 and Its Dynamics". Nonlinear Dynamics 101(3):1667-1680.

Holmdahl, I., and Buckee, C. 2020. "Wrong But Useful – What Covid-19 Epidemiologic Models Can and Cannot Tell Us". New England Journal of Medicine 383(4):303-305.

Iacob, M.E., Jonkers, H., Lankhorst, M.M., Proper, E., and Quartel, D. 2012. Archimate 2.0 Specification: The Open Group. Netherlands: Van Haren Publishing.

Kulkarni, V., Barat, S., and Clark, T. 2019. "Towards Adaptive Enterprises Using Digital Twins". In *Proceedings of the 2019 Winter Simulation Conference*, edited by N. Mustafee, M. Rabe, K. G. Bae, C. Szabo, S. Lazarova-Molnar, P. Haas, and Y. Son, 60-74. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Liebowitz, J., Chan, Y., Jenkin, T., Spicker, D., Paliszkiewicz, J., and Babiloni, F. 2019. "If Numbers could "Feel": How Well Do Executives Trust Their Intuition?". *VINE Journal of Information and Knowledge Management Systems* 49(4): 531:545.

Marchau, V.A., Walker, W.E., Bloemen, P.J., and Popper, S.W. 2019. *Decision Making Under Deep Uncertainty: from Theory to Practice*. Switzerland: Springer Nature.

Marathe, M., and Vullikanti, A.K.S. 2013. "Computational Epidemiology". *Communications of the ACM* 56(7): 88-96.

Meadows, D.H. 2008. *Thinking in systems: A primer*. USA: Chelsea Green Publishing.

Osburn, P. V., Whitaker, H. P., and Kezer, A. 1961. "New Developments in the Design of Model Reference Adaptive Control Systems". USA: Institute of Aeronautical Sciences.

Parkhi, O.M., Vedaldi, A., and Zisserman, A. 2015. *Deep Face Recognition*. https://web.archive.org/web/20161213020207id_/http://www.robots.ox.ac.uk:80/~vgg/publications/2015/Parkhi15/parkhi15.pdf, accessed 27th Feb 2022.

Sandkuhl, K., Fill, H.G., Hoppenbrouwers, S., Krogstie, J., Leue, A., Matthes, F., Opdahl, A.L., Schwabe, G., Uludag, Ö., and Winter, R. 2016. "Enterprise Modelling for the Masses – from Elitist Discipline to Common Practice". In *Proceedings of the*

IFIP Working Conference on The Practice of Enterprise Modeling edited by J. Horkoff, M. Jeusfeld, and A. Persson, 225-240. Springer, Cham.

Sargent, R.G. 1992. "Validation and Verification of Simulation Models". In *Proceedings of the 1992 Winter Simulation Conference,* edited by J.J. Swain, D. Goldsman, J.R. Wilson, and R. C. Crain, 104-114. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Shafto, M., Conroy, M., Doyle, R., Glaessgen, E., Kemp, C., LeMoigne, J., and Wang, L. 2012. "Modeling, Simulation, Information Technology & Processing Roadmap". *National Aeronautics and Space Administration*, 32(2012): 1-38.

Simon, H.A. 1990. "Bounded Rationality". In *Utility and Probability* edited by J. Eatwell, M. Milgate, and P. Newman, 15-18. London: Palgrave Macmillan.

Thomas, M., and McGarry, F. 1994. "Top-down vs. Bottom-up Process Improvement". *IEEE Software*, 11(4): 12-13.

Wagle, S., and Paranjape, A.A., 2020. "Use of Simulation-aided Reinforcement Learning for Optimal Scheduling of Operations in Industrial Plants". In *Proceedings of 2020 Winter Simulation Conference*, edited by K.G. Bae, S. Lazarova-Molnar, Z. Zheng, B. Feng, S. Kim, T. Roeder, and R. Thiesing, 572-583. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Wang, Y., Wang, Y., Chen, Y., and Qin, Q. 2020. "Unique Epidemiological and Clinical Features of the Emerging 2019 Novel Coronavirus Pneumonia (COVID-19) Implicate Special Control Measures". *Journal of Medical Virology* 92(6): 568-576.

World Health Organization. 2020. Modes of Transmission of Virus Causing COVID-19: Implications for IPC Precaution Recommendations: Scientific Brief, 29 March 2020 (No. WHO/2019-nCoV/Sci_Brief/Transmission_modes/2020.2).

White, S.A., 2004. *Introduction to BPMN*. http://yoann.nogues.free.fr/IMG/pdf/07-04_WP_Intro_to_BPMN_-_White-2.pdf, accessed 27th Feb 2022.

Wynants, L., Van Calster, B., Collins, G.S., Riley, R.D., Heinze, G., Schuit, E., Bonten, M.M., Dahly, D.L., Damen, J.A., Debray, T.P., and De Jong, V.M. 2020. "Prediction Models for Diagnosis and Prognosis of Covid-19: Systematic Review and Critical Appraisal". *British Medical Journal* 369 (2020).

Yu, E., Strohmaier, M., and Deng, X. 2006. "Exploring Intentional Modeling and Analysis for Enterprise Architecture". In *Proceedings of 10th IEEE International Enterprise Distributed Object Computing Conference Workshops (EDOCW'06)*, 32-32. IEEE.

AUTHOR BIOGRAPHIES

VINAY KULKARNI is a Distinguished Chief Scientist at Tata Consultancy Services Research. An alumnus of Indian Institute of Technology Madras, Vinay is a Fellow of Indian National Academy of Engineering. He is a Visiting Professor at Middlesex University London, and Indian Institute of Technology Jodhpur. His research interests include Digital Twins, Adaptive Enterprises, Model Driven Engineering, Artifical Intelligence, and Software Engineering. His e-mail address is vinay.vkulakrni@tcs.com.

SOUVIK BARAT is a Principal Scientist at Tata Consultancy Services Research. He holds a PhD Computer Science and his research interests include digital twin technology, modelling and simulation of complex systems, agent and actor technology, enterprise modeling, and business process modelling. His email address is souvik.barat@tcs.com.

TONY CLARK is a Tony Clark is a Professor of Computer Science and Deputy Dean (Education) in the College of Engineering and Physical Science at Aston University. He has experience of working in both Academia and Industry on a range of software projects and consultancies. His research interests include Digital Twins and Software Engineering, Languages and Modelling for Agent-based Adaptive Systems. His e-mail address is tony.clark@aston.ac.uk.

BALBIR BARN is Professor of Software Engineering and Academic Dean of the Faculty of Science and Technology at Middlesex University. Balbir has extensive commercial research experience working in research centres at Texas Instruments and Sterling Software as well as leading on academic funded research (Over £2.5 million). His research interests include socio-technical digital twins, value sensitive design and model driven software engineering. His email address: b.barn@mdx.ac.uk.