Position paper on the challenges posed by modern applications to cyber-physical systems theory

Frank Allgöwer, João Borges de Sousa, James Kapinski, Pieter Mosterman, Jens Oehlerking, Patrick Panciatici, Maria Prandini, Akshay Rajhans, Paulo Tabuada, Philipp Wenzelburger*

Abstract

Cyber-physical systems theory offers a powerful framework for modeling, analyzing, and designing real engineering systems integrating communication, control, and computation functionalities (the cyber part) within a natural and/or man-made system governed by the laws of physics (the physical part). New methodological developments in cyber-physical systems theory are required by traditional application domains such as manufacturing, transportation, and energy systems, which are currently experiencing significant and – to some extent – revolutionary changes to address the needs of our modern society. The goal of this position paper is to provide the cyber-physical systems community, and especially young researchers, a clear view on what are research directions worth pursuing motivated by the challenges posed by modern applications.

1 Introduction

Cyber-physical systems (CPS) are engineering systems characterized by the integration of communication, control, and computation within a natural and/or man-made system governed by the laws of physics. They are widely present in different domains, which include (but are not limited to) energy, transportation, and manufacturing that are fundamental pillars to the well functioning of our modern society. CPS theory has indeed the potential to make a transformational impact at global societal scale.

In the *smart energy* domain, CPS theory can help to redesign the power grids so as to increase their safety and reliability, as well as contribute at seamlessly integrating intermittent renewable energy sources such as wind energy by using predictive forecasting. In the *smart mobility* domain, there are complex CPS problems on land: connected and autonomous vehicles countering traffic congestion and fatalities; in the water: unmanned underwater vehicles and autonomous ships exploring the ocean; and in the air: unmanned aerial vehicles, colloquially referred to as drones, performing search and rescue operations. *Smart manufacturing* aims to leverage CPS to increase safety, reliability, and throughput in industrial production via the Industry 4.0 paradigm; the use of digital twins to deploy simulation counterparts of expensive physical assets to predict and detect failure, and take over the corresponding functionality from the physical counterpart; and extend our manufacturing capabilities to new and advanced problems[1].

Motivated by the breadth of the global societal-scale relevance of CPS, *our goal is to outline important research and development challenges in CPS theory that are key to their effective practical impact, possibly directing the efforts of the CPS community towards their study.*

To this purpose, the present paper brings together experts in the relevant application and methodological domains, who will offer some insight into the challenges posed by modern applications and provide their views on possible interesting theoretical topics to investigate.
The content of this paper reflects the personal view of its contributors.

*F. Allgöwer and P. Wenzelburger are with University of Stuttgart, Germany; J. Borges de Sousa is with Universidade do Porto, Portugal; J. Kapinski is with Toyota Research Institute of North America, US; P. Mosterman and A. Rajhans are with the MathWorks Inc, US; J. Oehlerking is with Robert Bosch GmbH, Germany; P. Panciatici is with RTE, France; M. Prandini is with Politecnico di Milano, Italy; P. Tabuada is with University of California at Los Angeles, US. Corresponding author: M. Prandini, e-mail: `maria.prandini@polimi.it`

[1] http://arminstitute.org/

The rest of the paper unfolds as follows.
We start by proposing a feature classification for CPS based on increasingly complex computation and communication capabilities, and outline exemplar challenges in each category in Section 2. We then consider smart manufacturing (Section 3), energy (Section 4), and automotive and marine transportation systems (Sections 5 and 6). The formulation of system requirements for these increasingly complex CPS applications becomes a challenge itself and offers new topics for research, as discussed in Section 7. Some final remarks are given in Section 8.

2 A Feature Classification for Cyber-Physical Systems

In this section, we propose a feature classification for CPS that leverages two celebrated frameworks on mental activity and agility in intelligent actions outlined in Fig. 1. The first framework shown in Fig. 1(a) is that of the levels of mental activities humans are capable of, as outlined by Marvin Minsky in the context of artificial intelligence [Min06]. The second framework shown in Fig. 1(b) depicts a modified drawing of the Observe-Orient-Decide-Act (OODA) Loop[2] developed by Colonel John Richard Boyd originally in the context of agility military strategy, but since applied to several technical domains such as agile software development [Ado06]. Building on these two frameworks, next we develop our classification of intelligent CPS as follows.

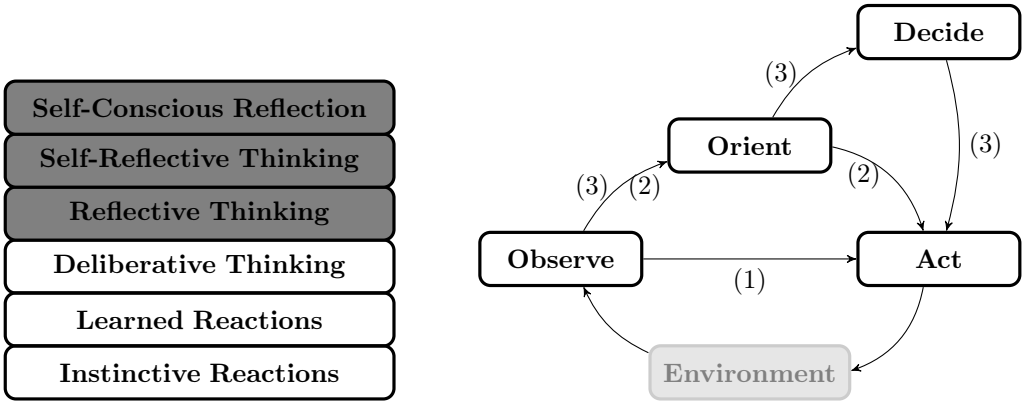

(a) Levels of Mental Activities.

(b) OODA Loop.

Figure 1: Frameworks from the literature applied to CPS intelligence hierarchy. On the left, in Fig. 1(a), a slight simplification omits the inputs feeding into the Levels of Mental Activities from the original drawing [Min06]. On the right, in Fig. 1(b), a slight modification of the Observe-Orient-Decide-Act (OODA) loop adds the gray Environment box to indicate closing the loop over the environment and shows simpler subset loops, namely the Observe-Act loop (1) and the Observe-Orient-Act loop (2) of the full OODA loop (3).

Computation and Communication Axes: Moore's Law and Metcalfe's Law. Over the past four decades, thanks to Moore's Law [Sch97], availability of fast, reliable, and cheap compute power has enabled us to embed increasing sophisticated intelligence onto our physical world. At the same time, advances in communication technology has interconnected and networked such smart systems where, following Metcalfe's Law[3], the value derived not just from itself, but from being connected to others in the network. Computation and communication dimensions form two axes in our classification scheme sketched in Fig. 2.

1. **Instinctive Reactions: Automatic and Distributed.** Instinctive Reactions is the simplest level of mental activity that humans and even animals are capable of. It involves our reflexes and our innate primal instincts that help us survive, e.g., looking towards a source of sound, or taking one's hand away on touching a hot object. Analogously, traditional *automatic* control systems like an inverted pendulum simply sense and actuate the environment, i.e., close the simplest Observe-Act loop on the environment. Ensembles of automatic

[2]https://en.wikipedia.org/wiki/OODA_loop
[3]https://en.wikipedia.org/wiki/Metcalfe%27s_law

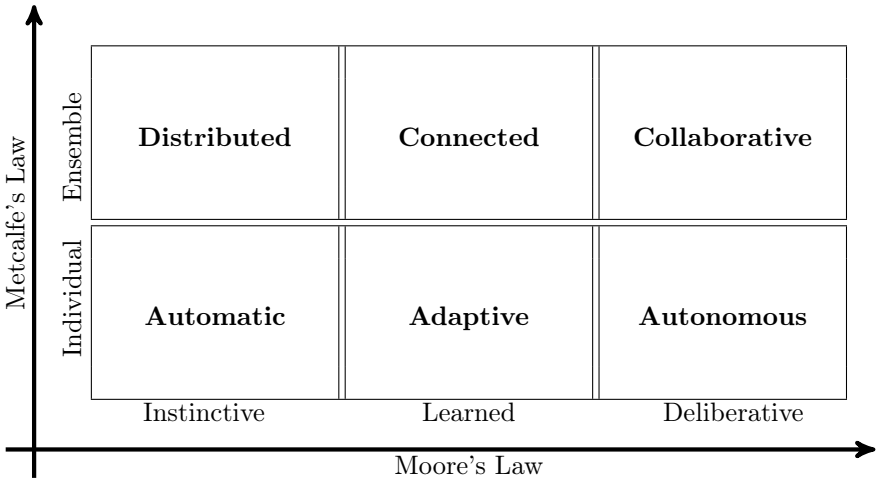

Figure 2: A CPS Feature Classification.

systems form *distributed* control systems [DD03, BCM09]. This class forms the baseline of intelligence in present day CPS.

2. **Learned Reactions: Adaptive and Connected Systems.** Learned Reactions is the next level of mental ability that involves learning patterns from historical information, e.g., identifying the shape of a ball or the sound of a car. *Adaptive* systems such as smart thermostats introduce such pattern-recognition ability into CPS, e.g., learning set points based on historical trends and occupancy. This adds the Orientation step in the Observe-Orient-Act loop. Other examples in this category include advanced driver assistance systems (ADAS) such as detecting a lane and steering to take a corrective action when a vehicle is about to leave it, or detecting the object such as a car or a pedestrian in front and its distance and applying emergency brakes to prevent a collision [LC06, PCSB16]. Ensembles of adaptive systems forms the class of *connected* systems in which, for example, a fleet of vehicles can share a better classification approach learned by one car with the entire fleet[4].

3. **Deliberative Thinking: Autonomous and Collaborative Systems.** Deliberative Thinking is a level of mental activity that humans and only some advanced mammals are capable of, e.g., planning a sequence of steps to execute a complex task, or figuring out what to do in a situation never encountered before. *Autonomous* systems such as the Mars rover, unmanned aerial/ground/underwater vehicles, fully-autonomous cars, etc. introduce a similar deliberative decision making step, effectively closing the entire OODA loop over the environment. *Collaborative* systems comprise of a number of collaborating autonomous systems that work together towards a common objective.

The more advanced uniquely-human levels of mental abilities that involve reflecting upon and action and on oneself, shown in a darker shade in Fig. 1(a), are an active area of research within computer science, e.g., the so-called artificial general intelligence (AGI) [GP07b], and as such are yet to make their way into CPS as of yet.

We next outline some domain-independent exemplar challenges as per our classification. We omit the baseline categories of automatic and distributed systems and focus on the future challenges in CPS brought about by adaptivity and autonomy.

1. **Adaptivity Challenges.** Because almost all CPS are safety critical, introducing learning into the mix must ensure that learning is limited to safe behavior. In terms of certification, testing or verifying a self-changing artifact will turn out to be a key research challenge in the near future. In case of connected ensembles, safe interpretation of data, especially in presence of possibly-malicious actors. Connected CPS must continue to operate safely regardless of loss of connectivity, which may mean architecting systems such as critical paths of operation do not depend on connectivity or tolerate intermittent connection.

[4]https://www.recode.net/2016/9/12/12889358/tesla-autopilot-data-fleet-learning

2. **Autonomy Challenges.** Assessing risk online and knowing whether a planned action is safe will be an important research challenge for autonomous CPS. Also, nontrivial interaction with humans for safely handing over control to a human operator during critical moments of operation will be of paramount importance for the utility of such systems [GKLB16]. For collaborative ensembles, ensuring safety of ad hoc rules in collaboration and gracefully enter/exit a collaboration will be important.

The following sections take a deeper dive into domain-specific challenges in manufacturing, automotive, underwater, and energy domains.

3 Manufacturing systems

The term "Industry 4.0" summarizes research topics for the development of the next level of industrial production. It was introduced by the German government, which started the strategic initiative Industry 4.0 in 2011 with the goal to prepare the German economy for the future [KHHW13]. Similar programs have been initiated in subsequent years by other governments under, partly, other names like for example "smart manufacturing" [LMF16].
In analogy to the first three major leaps in industrial production, namely mechanization through steam power, mass production through the introduction of the conveyor belt, and automation with programmable logic controllers and industrial robots, the growing computation power and most importantly the interconnection of industrial production units is considered the next industrial revolution. This time the revolution is predicted in advance allowing for a systematic development and a theoretic foundation of the new industrial environment. The goal of Industry 4.0 is to achieve the production of small lot sizes at the cost of mass production. This shall be achieved by real-time optimization and optimal coordination of the factory and its inventory, as well as the coordination with suppliers and customers. To this end, among other things intelligent and flexible production systems have to be developed that are able to integrate themselves into existing production systems and thereby react to changes in the production process [KZ15].
In contrast to the existing industrial machinery, the newly available machines have the capability to communicate in real-time and to process and distribute large amounts of sensor data, which they acquire through sensing and interaction with their environment. In combination with their computational capabilities, which are already present in today's machines and are steadily increasing, they fulfill all the criteria to be classified as Cyber-Physical Production Systems (CPPS) [LBK15]. Since the industrial processes being executed by the machines are mostly continuous and the decisions to be made are discrete changes of the state and the behavior of the machines, they can be described as hybrid systems. This already hints at the importance of the CPS community and the hybrid systems community for the development of Industry 4.0. In order to underline the relevance of cyber-physical and hybrid systems in this highly active research direction with great practical importance, we will briefly describe some important challenges and open problem that arise from the goals and visions of Industry 4.0 and the newly available technologies.

Challenges and open problems

The interaction of machines through communication allows for more flexibility of the production system and therefore enables faster reaction to the market, as well as self-reconfigurability in case of defective machines and robots. However, the new flexibility does not directly lead to improved efficiency and cannot simply be exploited, but comes with a set of challenges. The intelligent agents are thought to evaluate their possibilities autonomously. They are being developed to be increasingly self-aware and to make decisions on their own [LCK16]. This involves decision processes which have to be developed and are already challenging for a single autonomous system that has to adapt to a changing environment. Through the interconnection of multiple units, this task becomes even harder and might lead to undesired and unpredictable behaviors which have to be prevented by a profound analysis of the underlying systems. For a single CPPS, the flexibility of its hardware part can be used by processing its sensory information with the help of suitable software tools. Through the possibility of deploying different software, new possibilities for the overall system arise. When combining multiple CPPS, each with its own versatility, the degree of flexibility increases once again demanding a systematic approach in order to be manageable and requiring even more elaborate capabilities of the single systems.

The communication between the machines and possibly a central controlling unit brings further challenges. Communication links might be lossy and delayed, leading to the presence of outdated or missing information. For systems being dynamically coupled through a shared task or process, this can have devastating effects, if it is not considered appropriately [HNX07]. To this end, methods from the field of control of/over networks can be applied, by explicitly considering the communication medium.
Another widely discussed issue in Industry 4.0, which arises directly from its interconnected nature, is cyber security. Through the communication between agents, several (probably sensitive) data are exchanged. For a company it is crucial to keep this information secret in order to secure their know-how and possibly critical data about their customers and suppliers. On the other hand it is at least as important to prevent intruders from influencing the physical processes through malicious signals they might try to send to the machines. Otherwise, they might be able to stop the manufacturing process completely or even physically damage the production facilities. With the development of cloud control, attacks of this kind might become possible and defence strategies as well as resilient control techniques need to be developed and are already in the scope of the CPS community [ZSCC16] [KLH13].
During the execution of a manufacturing process, continuous dynamics will come into play, which can affect the discrete decisions to be made and render the overall set-up into a hybrid system. In classical manufacturing the process dynamics are controlled separately and are very well understood. However, through the combined control of the continuous process and the discrete decisions, the overall manufacturing system can be improved once again. This consideration of the hybrid system bears new challenges that have to be overcome. Here, the hybrid systems community has a great chance to contribute, even with existing methods and techniques. On the other hand new challenges for the hybrid systems community arise through the current industrial developments, especially through the connection of multiple systems, be it through a common continuous process to be handled, or through discrete events triggering subsequent processes.
An important topic on its own is human-machine interaction [GSLZ14]. It is generally accepted that humans are superior to their robotic counterparts when performing certain tasks and humans allow for more flexibility. Therefore, the vision of Industry 4.0 involves an effective interplay between humans and robots in order to exploit the strengths of both of them. There exists a wide range of methods and technologies that are being developed reaching from soft casings of robots and the increasing usage of additional sensors to new ways how humans and robots can communicate and interact [SECP13]. This moves the human into the control loop, either as additional controlling agent or as part of the process being controlled. In order to exploit the potential of this interplay, a model of the human behavior, or a description of the way a human interacts with a specific process, are possible approaches.

Concluding remarks

As a short conclusion one can state that Industry 4.0 is coined by the emergence of CPPS, which are triggered by discrete events and decisions, handle continuous processes, and offer great flexibility through their collaborative possibilities. At the boundary to different fields like robotics, control theory, and computer science, the cyber-physical systems community as well as the hybrid systems community can have a major contribution in this development.

4 Energy systems

Historical power system: emblematic example of system of systems

Electrification is considered by the National Academy of Engineering in USA as the greatest engineering achievements of the 20th Century[5]. This academy acknowledges that large power systems[6] are the most complex machines ever built by mankind. They are since their creation, the most emblematic examples of system of systems: thousands of large generating units interacting with millions of electrical loads through long distance connections (electrical grids).
The electricity as such was not easily storable and the balancing between electrical supply and demand had to be ensured almost instantaneously (in a time window of few seconds). This chal-

[5] http://www.greatachievements.org/?id=2949
[6] Power System = electrical grid with the generating units and the electrical consumption processes (loads)

lenging issue was elegantly solved using the synchronous generator[7] producing alternating electrical current and voltage. The voltage and current are sinusoidal signals oscillating at a certain frequency (f: nominal value in Europe is 50 Hz), this electrical frequency and the rotation speed of the synchronous generator are proportional. The speed of rotation is defined by the rotating mechanical equations [Kun94]

$$J\frac{d\omega}{dt} = T_m - T_e$$

where J is the moment of inertia, ω is the rotation speed, T_m is the mechanical torque, and T_e is the electrical torque.

The rotation speed is controllable by changing the mechanical torque (e.g., for turbo alternator, changing valve position in order to send more or less steam in the turbine). When a large population of synchronous generators is connected to an electrical grid, they are "magically" synchronized thanks to the structure of their interactions through this Alternating Current (AC) electrical grid. When an abrupt change occurs in the system, these synchronous generators are transiently not synchronized but they generally come back to synchronous operation[8] in less than a few seconds. Their rotation speed is the image of the grid frequency which has in steady state the same value all over the system. This frequency gives information about the system imbalance: if the frequency is larger than the nominal frequency that means that too much power is injected in the system; if frequency is lower than the nominal frequency that means that not enough power is injected. By measuring its rotation speed, an individual synchronous generator is aware of the system imbalance and knows if it must increase or decrease its generating power. This is why large power systems were able to ensure a reliable electrical supply without any complex telecommunication system and centralized control system, even if electricity was not storable (in fact a small storage of kinetic energy exists!). Moreover, this is a very resilient distributed control framework, all the synchronous generators of a vast area (e.g., continental Europe: from Portugal to Poland, from Denmark to Greece) participated in the control efforts. When an unplanned disconnection from the electrical grid of a large generator (1 GW) occurs, instantaneously (in few seconds) all the thousand other generators automatically compensated this lost of supply by generating individually a small additional amount of power (less than 1 MW).

This is typically what happens in system of systems when a lot of agents, physically connected (here through an electrical grid) contribute efficiently to a common objective (here supplying the electrical loads). We can propose an analogy with ants which could be helpful to understand the role of frequency similar to a "stigmergic" [9] channel. An ant finds food and builds a trail by leaving traces (pheromone) in the environment on its way back to its nest; other ants follow this trail and reinforce it by leaving more traces, optimizing the seeking of food. This is an efficient process to find food and to contribute to a common objective [10]. The frequency in a power system is similar to pheromone for the ants. This is one of the main amazing properties of our historical AC power systems, different projects in Europe and in USA propose to clarify these issues and to find new solutions.[11]

Ensuring an adequate level of reliability of a power system is not so easy while keeping costs under control. Finding the right balance between reliability, affordability and now sustainability is very challenging and involves political decisions in order to monetize the "reliability" and the "sustainability" [KW16]. In historical power systems, the "cyber" layer (information and communication technologies layer) had been used mostly for optimizing the system but not for ensuring safety or stability. Local protections and controls were very simple and implemented via analog devices. They have been digitalized but without changing the basic concepts and their functions. For optimizing the system, slow centralized controls have been installed using a minimal amount of remote information and actions, requiring only low performance telecommunication systems (low bandwidth, large latency and medium reliability). The system reliability was not affected by failures of this historical "cyber" layer. This is a description of the status of the "cyber" layer in transmission

[7] The beginning was in the 1880s based on principles discovered in 1831–1832 by Michael Faraday!

[8] There is a limit; a loss of synchronism could occur in case of too long lasting short circuit nearby a synchronous generator

[9] The biologist Grassé introduced the term "stigmergy" to describe the indirect information flow among the members of a termite colony when they coordinate their nest building activities

[10] inspired by the past European project **AMADEOS**: Architecture for Multi-criticality Agile Dependable Evolutionary Open System-of-Systems - http://amadeos-project.eu/

[11] For example, the European project Migrate, https://www.h2020-migrate.eu/ and in USA, SuNLaMP: https://www.energy.gov/eere/solar/project-profile-stabilizing-power-system-2035-and-beyond-evolving-grid-following-grid.

electrical grids (high voltage meshed grid, long distance connection) of the 20th century. Now, we are living a major evolution and perhaps a revolution in electrical grids. The term "smart grid" is used everywhere without a very precise definition. The concept of cyber-physical System of Systems seems a good framework to capture the essence of this (r)evolution.

(R)evolution of Power Systems: Cyber-Physical System of Systems[12]

The electrical grids and their management become more and more complex. This state of affairs has different causes that will not disappear in the near future.
The first reason is the massive integration of renewable but generally intermittent generation in the system. Power flows in the grid are created by differences in the location of sinks and sources. With a significant amount of intermittent generation, the predictability of the sources (location and amount of power injections) decreases and affects the predictability of the flows. Furthermore, some of these new generating units are small units (e.g. photovoltaic) connected to the distribution grid, changing the distribution grid into an active system. Moreover, Transmission System Operators (TSOs) have a poor observability of these power injections and have no control at all over them. Another factor is the inconsistency between the relatively short time to build new wind farms (2 or 3 years) and the time to go through all administrative procedures to build new lines (more than 5 years everywhere in Europe). In Europe, the best locations for wind farms are mostly along the coasts and offshore, while for photo-voltaic generation they are in the south of Europe. Since these locations do not generally match those of the large load centers, a transmission network is required and this network will have to cope with the variability of the flows induced by the stochastic nature of the novel generation subsystems.
The second main reason is that it is more difficult than ever to build new overhead lines because of low public acceptance and "Not In My BackYard" (NIMBY) attitude. People are more and more afraid of hypothetical electromagnetic effects or just don't like to see big towers in the landscape and in particular in protected areas which are more and more numerous around Europe.
The third reason is linked to the setup of electricity markets crossing the administrative and historical borders. Generators, retailers and consumers view the transmission system as a public resource to which they should have unlimited access. This approach has the desirable effect of pushing the system towards maximization of the social welfare and an optimal utilization of the assets. However, this optimization is constrained by security considerations because widespread service interruptions spanning over long periods of time are unacceptable in our modern societies due to their huge economic and social costs. The market players perceive reliability management by the TSOs as constraining their activities and reducing the European Social Welfare rather than as enablers of this large physical market place, as it would be the case if the grid was a unlimited copper plate.
The last reason is that the aging of grid assets needs increasing attention. A significant part of the European grids' assets are more than 50 years old. Asset management and maintenance in systems that cannot be stopped, are extremely challenging and need to be precisely anticipated when large numbers of assets are approaching simultaneously the end of their expected life times. To maintain the security of the supply in this context, TSOs have to change the architecture of the system by considering HVDC technologies and by optimizing the existing systems by adding more and more special devices such as Phase Shifting Transformers, Static Var Compensators and advanced controls and protection schemes, taking also advantage of the flexibility provided by HVDC links embedded in AC grids. At the same time, demand response or dispersed storage could offer new ways to control the system, even if business models and costs are still questionable. But in any case, how to use these flexibilities will require a rethinking of historical operating practices where grid operators made the assumption that the load is an uncontrollable exogenous stochastic variable.
All these evolutions are transforming power systems in cyber-physical system of systems (CPSoS). The "cyber" layer is going to play a key role in the system reliability. Indeed, more and more controls are embedded in subsystems which become "intelligent" and partially autonomous. The system behavior will be imposed by the interactions between these "intelligent" agents driven by local software (blue triangle in the figure 3) rather than by physical laws.
One example: the transformation of distribution grids (last miles grid at low voltage level; generally operated in a radial mode) in active distribution grids. More and more generating units are

[12] http://www.cpsos.eu/wp-content/uploads/2016/04/CPSoS-Brochure-LowRes.pdf

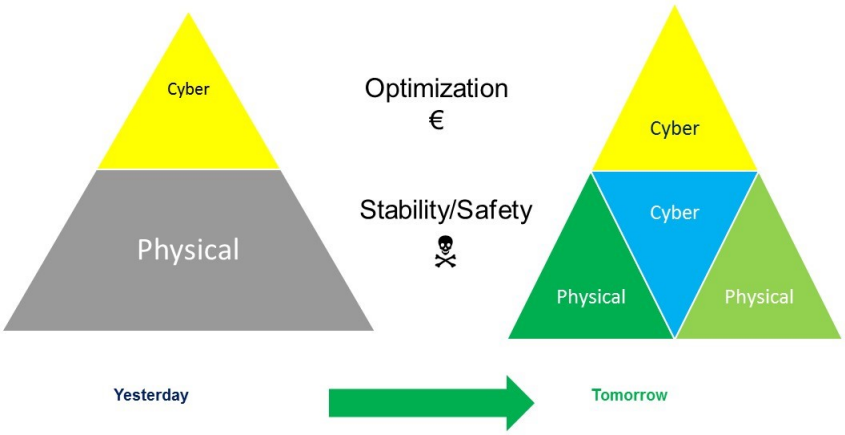

Figure 3: Digital transformation of Power Systems

connected to these distribution grids and possibly some electrical batteries. Moreover, demand response programs try to encourage consumers to adapt their consumption in order to "help" the system. More and more controls are installed in the distribution grids which transform them in active systems. In the past, distribution grids were "passive", in the sense that they were reacting to electrical signals following physical laws. The amount of consumed power was stochastic but for example the relationship between the consumed power (P) and the voltage magnitude (V) was captured sufficiently accurately by simple aggregated models. The typical behavior of a resistive load is $P = V^2/R$ where R is the resistance. In case of a sudden decrease of the voltage magnitude V, the consumed power P will decrease, this feedback helps the system and brings robustness. After this fast transient, the consumption process was changing very slowly the resistance (R) to keep for example a required temperature or will last a bit longer in order to consume a given energy. Tomorrow, the voltage magnitude in distribution grids will be controllable thanks to all the connected power electronics but if these controls are too fast, the consumed power will be constant and we will lose the helpful relationship between the consumed power and the voltage magnitude. A too local design of controls in distribution grid could have a disruptive impact on large power systems [AVC17].

The "cyber" layer ensuring the system reliability must be designed carefully using the CPSoS approach. We must pay attention; the emergence of possible negative impacts could occur when a certain critical level of penetration of new devices and processes will be reached. It is not easy, we must review our historical approach in order to specify behaviors which were yesterday imposed by physical laws and did not need any specification and which will be tomorrow defined by software in local controls.

We must take advantage of recent results and solutions proposed in the control, the computer science and the optimization communities. The first topic is related to approximations and to reductions of large complex systems, indeed they are mandatory steps in order to design control strategies for such large systems. Even classical reduction methods and approximation notions could be improved to deal with stochastic aspects ([GP07a], [JP09], [PGV14]) or to extend the validity domain of the reduced models (non-minimum phase behavior [KHR+16], ...). Moreover, these approximation methods must take into account the "cyber-physical" dimension: not only open loop physical systems but also all the embedded advanced controls including switching behaviors. Another important aspect is the control of large population of devices or agents with a partial autonomy; game theory is certainly a relevant framework to address this issue. For power system, aggregative game [Gra17] or mean field game [SMN18] seem interesting approaches. The definition of requirements going beyond classical stability requirements is also critical; Signal Temporal Logic and possible extensions to probabilistic STL [FMPS17] offer generic formal frameworks to improve the definition of these requirements. For security assessment (what is the largest perturbation that the system could endure?) and optimal control, we must deal with system non linearities. Since decades, energy based methods (Lyapunov like methods) have been proposed but using simplistic

unrealistic modeling assumptions which are not usable in practice. New ideas based on finite time reachability (occupation measure [HLS08]) seem more relevant and could propose guaranteed conservative assessments even if the scalability to tackle large systems is still questionable and should be associated to dedicated approximation methods.
Machine learning approaches are very fashionable nowadays and they could be useful to provide "optimal" operating points but generally these methods cannot provide any certificate. When we want to design a control strategy, we must be sure that the output-input relationships are really causal and not only reflecting spurious correlations which could be destroyed by a new control strategy [Pea00]. We must understand that we have to deal with two different questions; the first one: how to keep the physical quantities within acceptable ranges (feasible domain) or to bring back the system states in the feasible domain after a perturbation? and the second one: how to find an "optimal" operating point in the feasible domain? For the first question, we must have guarantees because the system cannot survive (explosion, destruction, cascading failures, ...) outside the feasible domain.
The ultimate solution should be a mixture of different approaches in order to tackle the increasing complexity of large power systems.

5 Automotive systems

Typically, a large portion of the development of classical embedded software functions in the automotive context can (in a simplified fashion) be seen as consisting of controller design, implementation and parametrization with the goal of meeting certain safety and performance requirements. Given a set of actuators and a set of sensors, and an execution platform with performance constraints, the goal typically is the design of a controller fulfilling various safety, stability and robustness constraints. Driven by the need to limit the number of sensors in a car and the computational cost, physical models are used as observers, taking the role of virtual sensors for quantities that are not physically observable, or inside model based controllers to improve control performance.
In Figure 1(b), this setting either corresponds to the Observe-Act loop, or to the Observe-Orient-Act loop if there is some adaptive component to the controller. Examples for such systems include engine control functions (e.g., airflow control via throttle valves, injection rail pressure control, engine torque control), vehicle stability functions (e.g., vehicle stability control, anti-lock control, steering control), electric network control (e.g., control of battery charge levels), among others.
With the strong push towards automated driving functions the scope of the problems to be solved are expanding considerably. It is generally accepted that the validation of driving functions with a high level of automation in complex environments is simply not feasible by collecting data in an unstructured fashion and then invoking statistical arguments. For instance, as stated in [KP17], to show that an automated vehicle causes 20 percent less fatalities than a human driver, with a confidence of 95 percent would require 8.8 billion miles to be driven without accident, and this would be true after every software changes. Therefore, the use of formal arguments to reduce this complexity is a necessity.
Formal arguments for this class of problems will have to relate to many classes of models. Machine learning-based perception is currently indispensable, and there is a general tendency to rely more strongly on data based (as opposed to model based) approaches. Also, there is the necessity to master complex interactions between ensembles of agents, and to interact with humans with all their unpredictability. This poses great challenges to CPS theory to expand beyond the classical embedded system into other domains.

Classical Control Problems

For classical control applications, system-level specifications typically represent some form of constraints restricting the admissible system traces over time. Furthermore, a single specification will only refer to a very limited number of (typically physical) quantities, on which a certain behavior is imposed. This may be a simple safety property ("The quantity must stay within x of a reference value"), a stability property ("The controller must not produce unstable oscillations"), or more complex temporal properties ("After a change in set speed, the actual velocity of the vehicle must reach the new set speed with tolerance of x within t seconds, and overshoot the new set speed by less than y"). Formal specification, for example in signal temporal logic [MN04], while often not straightforward, is generally possible. Formalisms geared towards this problem exist and can

be used. However, the challenge on how to come up with correct specifications, in the sense that they capture the intention behind the requirement correctly, is still open and discussed in detail in Section 7.
Classical control systems in the automotive domain can be seen as an instance of hybrid dynamical systems on some abstraction level, (i.e., leaving out implementation details like schedulers determining the exact computation times). The relevant physical effects are typically well understood and can be usually be modeled using some form of differential or algebraic equations (ODE, PDE, DAE), potentially with some form of switching. If a suitable physical model is unavailable, typically data based models (i.e., lookup tables) are supplemented. Also, effects like sensing disturbances are typically well enough understood to often be representable as stochastic distributions. Of course, some challenges still remain open, like efficient parametrization of these models from measurements or coming up with models that are simple enough to be amenable to formal methods and yet close enough to reality to be valid. However, it seems that hybrid dynamical systems are a suitable modeling framework for this class of problems.

Autonomy

For automated driving functions, this is only a part of the solution, however. For instance, specification for an automated driving function for urban contexts would at least require: formalization of (various) expected behaviors of other agents which our function assumes, formalization of the rules of interaction between traffic participants (this includes the actual traffic rules), protocols and rules of interaction for situations when these rules are incomplete, formalizations on the expected reliability of complex sensors (e.g. video with convolutional neural networks for object detection and classification), formalization of the road layout, and so on. This is true even if the property to be shown is only the global absence of collisions and not a complex performance requirement: here, the complexity of specification lies in the precondition, i.e., in the situation when we actually can guarantee that our vehicle can avoid a crash.
A specification formalism for this domain would at least require: 1) a temporal component, to describe the expected evolution of behavior over time, 2) a spatial component, to describe the locations of objects and their relation, as well as the shape of the roadway, 3) an epistemic component, to describe the knowledge about other agents and their intentions, 4) a predictive component, to describe possible future behaviors of physical objects, knowing their past behaviors, and 5) fault models, in particular for perception components, to describe the assumptions on faults under which guarantees can be made. While modeling paradigms for most of these components exist, combining them in one framework in such a way that the models are still natural and comprehensible is a difficult challenge. Furthermore, the question of how useful fault models for machine-learning based perception components should look like is still open at this point.

Conformance

In autonomous driving applications, the requirements with respect to model accuracy can be expected to increase as well. The sheer complexity of the verification problems means that very likely some aspects of the safety argument will be dealt with at runtime, by using potentially complex physical models. This means that the safety of the passengers hinges directly on the quality of the models. Key techniques here are online decision making and runtime monitoring based on hybrid systems models, as well as coverage based arguments about model validity.
Online decision making based on hybrid system models can for instance be based on the behavioral prediction of other actors (e.g., physical models, models of human behavior, blame in case of accidents), which can be obtained through the use of hybrid system models (e.g., using reachability analysis [AD14, LRH+17]). If the models that are used are correct, this would entail a correct-by-construction approach, significantly reducing testing effort. Since this correctness assumption still needs to be validated, the testing effort would actually be shifted toward model validation, which is hopefully a less complex problem.
Runtime monitoring (e.g., synthesized from theorem provers as in [MP16]) can then be used to check the validity of a model online, and reacting accordingly once the observations do not fit the model. For instance, it would be possible to then switch to more permissive predicted dynamics of other traffic participants, if the measurements do not conform to the original model, at the price of having to make more conservative control decisions. This would again reduce the effort required

for model validation.
Nevertheless, there is a strong need for methods to argue the quality of models based offline on a finite number of measurements as well. This is also true for the classical control domains, but the more prevalent use of models in safety critical contexts makes this even more crucial for autonomous systems. In fact, this step can be seen as a necessary data based first step in the overall verification chain which should be treated as formally as possible. Clearly, testing a model through measurements can never be truly exhaustive. However, the assumptions under which a set of measurements is considered sufficient to argue the model quality should be as formal and as explicit as possible.

Concluding remarks

In summary, the shift towards automation in the automotive industry poses a number of challenges for the cyber-physical system community not only on the practical but also on the theoretical side, as gaps between theories and research communities need to be bridge to arrive at formal frameworks that scale to this class of problems. In particular, these include: coming up with unified modeling and specification frameworks capturing both autonomy aspect and classical control, arguing about machine learning based system components in a formal fashion, and arguing about the quality of the models on which formal arguments are based. It seems that all of these point are a necessity in order to be able to argue the safety of highly automated driving.

6 Marine systems

Marine robots

Marine robots come in several shapes, sizes, configurations, and categories. Autonomous Underwater Vehicles (AUVs) are unmanned, untethered, submersibles. Autonomous Surface Vehicles (ASVs) and Unmanned Air Vehicles (UAV) are, respectively, their surface and air counterparts. Remotely Operated Vehicles (ROVs) are tethered submersibles, remotely controlled from a ship or shore by a skilled pilot.
Key technical specifications for marine robots include endurance, size, payload, range, communications, and navigation capabilities. Endurance is highly constrained by the limitations of current energy storage technologies. The size of the vehicle typically constrains the payload and energy storage. Power and size are the major limitations of the payload, as well as availability of sensor technologies. Range depends not only on endurance, but also on the mission profile. Communication and navigation capabilities determine the level of human intervention, the practical endurance, and the usefulness of the vehicle. Underwater navigation is very challenging because GPS is not available underwater. Communications are necessary for operating the vehicle and retrieving information from it. Above water communications typically rely on radio technology, with line-of-sight limitations imposed by the curvature of the Earth. Underwater communications have relied mostly on acoustic communications, in spite of recent advances in optical communications.
Heretofore, most robots have been automated but are not autonomous, as one could infer from the used acronyms. Autonomy basically means that decision-making takes place on-board without human intervention. In other words, in autonomous vehicles the typical Sense-Decide-Act cycle is intrinsic. This is in contrast to what happens with automated vehicles, where Sensing and Acting are mediated by scripted control procedures. In fact, full autonomy is still not feasible today; vehicles still lack the sensing and decision-making capabilities required for that purpose. This is partly why the concept of mixed initiative operation was introduced in the last decade. In this concept, human operators are part of the planning and control loops of the robot.

Ocean challenges

The oceans are, in fact, a "fragile", yet remote, thin layer of water, with an average depth of 4km in a planet with a 6,000+ km radius. This is why there is a pressing need for a sustained, persistent, and cost-effective presence in the oceans that will help us to understand and monitor how key issues such as climate change, ocean acidification, unsustainable fishing, pollution, waste, loss of habitats and biodiversity, shipping, security, and mining are affecting global ocean sustainability

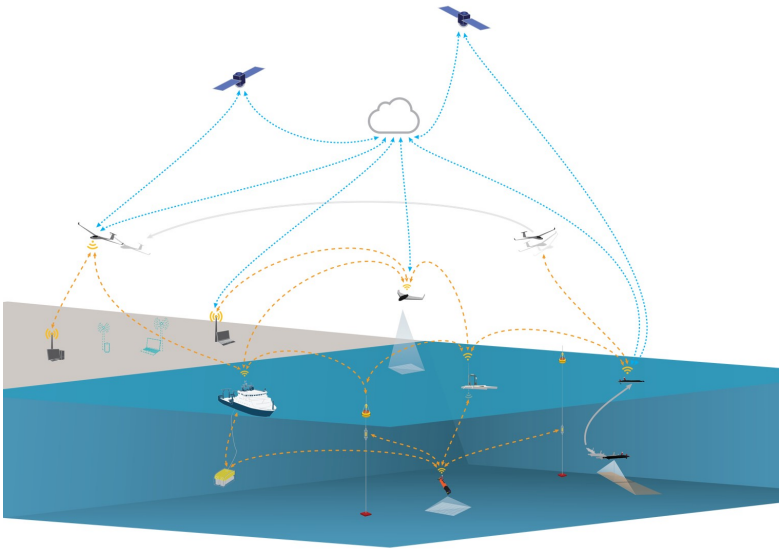

Figure 4: Networked marine systems.

and stewardship [IOC11]. Equally important is the development of systems and technologies that will allow us to explore and exploit the ocean in a sustainable manner.
Marine robotics holds the key to a sustained presence in the ocean [BR07]. But this is not a trivial task. It requires innovative approaches to systems development, operations, and management that can only be achieved with an incremental and multi-dimensional approach with the following characteristics: 1) the development of new robots with new sensing, communication, computation and energy storage capabilities, as well as intervention capabilities, in addition to high levels of autonomy; 2) the development and deployment of fleets of marine robots for persistent and/or specialized ocean operations; 3) the networking of existing systems, including manned assets, and new marine robots for large scale synergistic operations; 4) the development and deployment of interoperability standards with support for cyber-security to prevent unauthorized access to infrastructures; 5) the development of new concepts of operation to address different types of applications ranging from security and defense to environmental monitoring; and, 6) the development of new organizational frameworks to manage and coordinate the system(s) of systems, and the associated services, that will result from these networking trends.

Future marine operations

The future of marine operations will be significantly different from the current state-of-the-art. Marine robots are already delivering new capabilities, but this is just the beginning. Trends including miniaturization of sensors and computer systems, energy harvesting from atmospheric/ocean phenomena, power sources with increased energy-density, and increased subsystem standardization and modularity will have transformational effects in the future. Another significant trend is toward increased system autonomy via new command and control frameworks that facilitate integration of marine robots into higher-level marine systems.
Figure 4 presents a preview of future marine operations. These operations will typically involve multiple manned and unmanned assets, as well as human operators, interacting over interoperated underwater and above water communication networks, and taking place in remote and communications-challenged environments. Persistence will add another level of complexity. This is because some assets will come and go to recharge and exchange data at fixed or mobile docking stations. Finally, the number of assets involved in future maritime operations is also expected to increase significantly with respect to the current practice.
Future marine operations will be about dynamic networks of manned and unmanned assets exhibiting several distinctive features:

- Interactions among these assets will not be limited to information and commands. Other types of interactions will include: 1) Code migration to enable software updates for remote assets not in direct communication with control stations; 2) Energy transfer and refueling will be used to overcome endurance limitations; and, 3) Larger assets will be used to launch and recover smaller assets.
- Mobile connectivity and mobile locality. Examples of mobile locality arise when code migrates between computers or when smaller assets are transported by larger assets.
- Assets are coordinated to exhibit organization-like properties (e.g., some assets will be in charge of refueling and others will be in charge of data muling). These properties will be a function of the assets, of the communication networks, and of the interactions among these assets.
- Constant connectivity and full information assumptions will not hold because operations will take place in communications-challenged environments.
- The actions of some assets may affect the environment in which all assets evolve (e.g., high-power radios affect the quality of communications in the surrounding environment).

Dynamic networks of heterogeneous assets exhibiting such distinctive features pose new specification and control design challenges to CPS: from isolated systems to persistent dynamic networked systems [dSP].

Research challenges

Research in control engineering has not yet incorporated concepts such as link, interaction, and dynamic structure. Computer scientists, in part because of the pioneering work of Robin Milner [Mil96], were already making strides in this area in the early 90's. Milner's Turing award lecture is a must-read for this reason [Mil93]. The problem of mobile connectivity was addressed in the Π-calculus [Mil99], a calculus of communicating systems in which the components of a system may be arbitrarily linked and the communication over linked neighbors may carry information which changes that linkage. It was only one decade later, and in part because of the advent of ubiquitous mobile computing, that Milner and co-workers introduced the theory of Bi-graphical Reactive Systems (BRS's) [Mil09]. The theory is based on a graphical model of mobile computation with both mobile locality and connectivity. A bi-graph comprises a place graph, representing locations of computational nodes, and a link graph, representing interconnection of these nodes. Mobile connectivity and locality are expressed with BRS's by a set of reaction rules. In this model, systems of autonomous agents interact and move among each other, or within each other.

The BRS model is conceptually very powerful, but fails to be directly applicable to persistent networked systems exhibiting mobile connectivity and mobile locality. This is because the place and the link graphs fail to capture the true cyber-physical nature of these systems – physical and computational entities evolve and interact through coupled dynamics. This distinction between physical and computational entities is key to: 1) combine models of physical and computational dynamics; and, 2) map concepts of mobile connectivity and locality to the physical world and associated geographic constraints (e.g., communications between two physical entities are feasible only when these are within communication range of each other). Physical entities are governed by the laws of physics, computational entities by the laws of computation. Physical entities evolve in extended state-control spaces, may interact among themselves and with the environment, and can be "composed" to form other physical entities. For example, the state of one physical entity may include the computational and physical entities residing in it, as well as data; the controls available to a physical entity may affect the the environment (e.g., using electromagnetic signals to jam radio frequencies). Computational entities interact through communications. Computational entities may create other computational entities, and may be deleted as well. Some computational entities may be able migrate between physical entities over communication channels or by using physical entities as data "mules". The "composition" of computational entities is either local, with respect to the one physical entity in which they reside, or distributed, over communicating physical entities. The coupling of computational and physical entities enables mechanisms for self-awareness, for state propagation among unconnected entities, and for setting up controller structures in a networked system.

Further research in models of coupled physical and computational entities is needed to study behavioral equivalence and controller synthesis for the dynamic systems under consideration. Given a generic specification for the behavior of a system, the design problem consists of deriving a structure of computational entities which, when "composed" with the system, will satisfy the specification in some sense to be studied. Research in transition systems describing the co-evolution of physical and computational dynamics, and associated reachability concepts, may enable the specification of behaviors for networked vehicle systems in terms of traditional concepts from control engineering (e.g., invariance, reachability, and optimization) formulated in extended state-control spaces.

7 Requirements

As CPS applications become increasingly complex, improved engineering processes are needed to mitigate the concomitant increase in cost and intricacy of developing these new systems. New methodologies for design, synthesis, test, and verification allow improvements over traditional techniques. Despite this, many of the latest methodologies remain underutilized by industry, in part because the new techniques assume the existence of formalized system requirements, which can be challenging to create. Requirements engineering is the process of developing appropriate requirements for an application. Below we describe how requirements engineering fits within a CPS development process, identify the current technology gaps, and suggest directions for future research in this area.

Requirements for CPS Development

The purpose of system requirements is to provide documentation, act as a guide to develop design models, and to provide criteria against which test results may be evaluated. Requirements are vital for almost all aspects of system development. Requirements are created at the beginning of a development process, and they impact each subsequent phase of development.
Most CPS development processes maintain requirements in some form, even if they are not documented. Despite this, the requirements engineering process remains a pain point for many organizations developing CPS applications. In the sequel, we describe some of the ongoing challenges.

Requirements Engineering Gaps

When documented requirements are available, they are traditionally provided in the form of natural language, which can be ambiguous and can obscure inconsistencies. Formalized requirements, such as those given in the form of a specification language like a temporal logic [Pnu77, Koy90, MN04], offer a precise and unambiguous way to define expected system behaviors and are compatible with automated reasoning tools, such as theorem provers, that can check for consistency and completeness [FRNNS18, SMD+17]. Formalized requirements also allow recent development and analysis techniques, such as control synthesis [BVT18, FGP07] and formal verification methods [CÁS13, FLGD+11, DMVP15, PQ08]. Despite this, the use of formal requirements has been slow to gain traction for CPS domains. Formalized requirements are common for some application domains, such as aerospace, but even in those domains, the requirements usually define open-loop, software-only (i.e., not cyber-physical) behaviors. For other CPS domains, requirements are either not thorough or not formalized.

Research Challenges

One reason why formal requirements have not yet been adopted by industry is that they can be difficult to create and debug. Some groups are addressing this problem by developing tools that can automatically process a structured natural language (SNL) into a formalized representation. An SNL is a natural language, such as English, with restrictions on the allowed words and grammatical constructs [KGFP08]. An example of an SNL used to define requirements is the SADL Requirement Language (SRL), which is based on the Semantic Application Design Language (SADL). SRL takes requirements provided in an SNL format and processes them into an internal representations, for which properties can be proved using automated reasoning tools [SMD+17]. Other work attempts

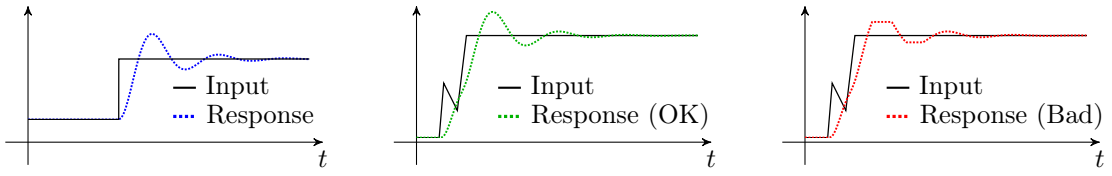

(a) (b) (c)

Figure 5: (a) Example of a canonical step-response Input and Response behaviors; (b) behaviors that are similar to the canonical step-response behavior and also demonstrate expected behavior; (c) response behavior that is too dissimilar to the canonical case to be considered acceptable.

to ameliorate the task of debugging requirements by providing methods to visualize a requirement given in a temporal logic language by automatically generating behavior examples that satisfy the requirement [DYHF17, RHM17, PLK]. A remaining challenge is in creating methods to visualize, or otherwise elucidate, the envelope (or complete set) of behaviors specified by the requirements.
Another reason why formal requirements have not yet found a footing in industry is that some types of expected system behaviors for CPS applications are subjective, or at least seemingly subjective. For example, consider a typical step-response behavior expectation for a CPS application, such as the behavior illustrated in Fig. 5-(a). The expectation is that when the system Input is a step function, the Response behavior will have an appropriate step-response shape, and the expectation may include constraints on specific aspects of the behavior, such as overshoot or settling time. In practice, step inputs are rare or nonexistent for real CPS applications, and yet the canonical step-response behavior informs the designer's expectations. In particular, designers expect that when a signal that is similar to a step is used to stimulate the system, then a response that is similar to the canonical step-response should be measured at the output. Figure 5-(b) illustrates an example of acceptable (OK) Input-Response behavior, based on the canonical expectation for the system, illustrated in Fig. 5-(a). The OK judgment is based on the similarity between the Response behaviors in Figs. 5-(a) and 5-(b), as compared to the similarity between their corresponding Input behaviors. This is in contrast to the behavior in Fig. 5-(c), which is deemed unacceptable (Bad). The Bad judgment is based on the qualitative difference between the Response behaviors in Figs. 5-(a) and 5-(c), as compared to the similarity between their corresponding Input behaviors. The Response behavior in Fig. 5-(c) exhibits some saturation qualities that are not present in the Response shown in Fig. 5-(a); this saturation behavior could indicate that the system has entered some unexpected operating regime, due to the particular input applied.
The notion of similarity used to make judgments like those illustrated in Fig. 5 is difficult to capture with existing formalisms. One approach to formally define these expected behaviors is to use appropriate metrics that capture distance between behaviors to reason about similarity to some canonical (or nominal) behavior [JBGN16, TN16, DMP17]. A remaining challenge is in providing better ways to understand these distance measures — including their corresponding units — in an engineering context so that designers may relate them to their expectations.

8 Conclusions

Cyber-physical systems, encompassing energy systems, mobility systems, manufacturing systems, are at the backbone of modern society. Ensuring a safe, reliable, but on the same time high performance operation of such systems is crucial from a societal point of view. This task is however becoming more and more challenging since "modern" cyber-physical systems typically involve many complex autonomous subsystems that interact with each other physically and/or exchange information via a communication network and possibly cooperate to achieve some common goal. For this kind of complex systems even the formulation of requirements is a challenge.

In this position paper, we pointed out that traditional methods need to be revisited, conceptually new operational paradigms need to be developed, and synergies between different fields of expertise (mainly control, computer science, optimization) have to be strengthened to address such a task. We also presented our views on what are interesting theoretical topics to investigate, motivated by the recent revolutionary changes in the manufacturing, energy, and transportation systems

domains.

It is perhaps worth mentioning that CPS also find application in the *smart health* domain, at the level of the single individual, where medical devices such as pacemakers, brain-computer interfaces, and robotic exoskeletons aim to improve quality of human life affected by disease, disorder, or injuries, and also at a societal scale, where NIST initiatives such as SmartAmerica[13] and Global City Teams Challenge[14] aim at creating a *smart society* by addressing problems of disaster recovery and readiness and more efficient water delivery networks. Smart health and further interesting application domains for CPS theory are however not discussed in this paper.

References

[AD14] Matthias Althoff and John M. Dolan. Online verification of automated road vehicles using reachability analysis. *IEEE Trans. Robotics*, 30(4):903–918, 2014.

[Ado06] S. Adolph. What lessons can the agile community learn from a maverick fighter pilot? In *AGILE 2006 (AGILE'06)*, pages 6 pp.–99, July 2006.

[AVC17] P. Aristidou, G. Valverde, and T. Van Cutsem. Contribution of distribution network control to voltage stability: A case study. *IEEE Transactions on Smart Grid*, 8(1):106–116, Jan 2017.

[BCM09] Francesco Bullo, Jorge Cortes, and Sonia Martinez. *Distributed control of robotic networks: a mathematical approach to motion coordination algorithms*. Princeton University Press, 2009.

[BR07] James Bellingham and Kanna Rajan. Robotics in remote and hostile environments. *Science*, 318:1098–1102, 2007.

[BVT18] Ayca Balkan, Moshe Vardi, and Paulo Tabuada. Mode-target games: Reactive synthesis for control applications. *IEEE Transactions on Automatic Control*, 63(1):196–202, Jan 2018.

[CÁS13] Xin Chen, Erika Ábrahám, and Sriram Sankaranarayanan. Flow*: An analyzer for non-linear hybrid systems. In *Computer Aided Verification*, pages 258–263, 2013.

[DD03] Raffaello D'Andrea and Geir E Dullerud. Distributed control design for spatially interconnected systems. *IEEE Transactions on automatic control*, 48(9):1478–1495, 2003.

[DMP17] Jyotirmoy V. Deshmukh, Rupak Majumdar, and Vinayak S. Prabhu. Quantifying conformance using the skorokhod metric. *Formal Methods in System Design*, 50(2-3):168–206, 2017.

[DMVP15] Parasara Sridhar Duggirala, Sayan Mitra, Mahesh Viswanathan, and Matthew Potok. C2E2: A verification tool for stateflow models. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 68–82, 2015.

[dSP] J. Borges de Sousa and F. Lobo Pereira. Coordination challenges in networked vehicle systems: are we missing something? In *Coordination Control of Distributed Systems. Lecture Notes in Control and Information Sciences, vol 456.*

[DYHF17] Adel Dokhanchi, Shakiba Yaghoubi, Bardh Hoxha, and Georgios E. Fainekos. Vacuity aware falsification for mtl request-response specifications. *2017 13th IEEE Conference on Automation Science and Engineering (CASE)*, pages 1332–1337, 2017.

[FGP07] Georgios E. Fainekos, Antoine Girard, and George J. Pappas. Hierarchical synthesis of hybrid controllers from temporal logic specifications. In *Hybrid Systems: Computation and Control*, pages 203–216, 2007.

[13] http://smartamerica.org/
[14] https://pages.nist.gov/GCTC/

[FLGD+11] Goran Frehse, Colas Le Guernic, Alexandre Donzé, Scott Cotton, Rajarshi Ray, Olivier Lebeltel, Rodolfo Ripado, Antoine Girard, Thao Dang, and Oded Maler. SpaceEx: Scalable verification of hybrid systems. In *Computer Aided Verification*, pages 379–395, 2011.

[FMPS17] S. S. Farahani, R. Majumdar, V. S. Prabhu, and S. E. Z. Soudjani. Shrinking horizon model predictive control with chance-constrained signal temporal logic specifications. In *2017 American Control Conference (ACC)*, pages 1740–1746, May 2017.

[FRNNS18] Predrag Filipovikj, Guillermo Rodriguez-Navas, Mattias Nyberg, and Cristina Seceleanu. Automated SMT-based consistency checking of industrial critical requirements. *SIGAPP Appl. Comput. Rev.*, 17(4):15–28, 2018.

[GKLB16] Christian Gold, Moritz Körber, David Lechner, and Klaus Bengler. Taking over control from highly automated vehicles in complex traffic situations: The role of traffic density. *Human Factors*, 58(4):642–652, 2016. PMID: 26984515.

[GP07a] A. Girard and G. J. Pappas. Approximation metrics for discrete and continuous systems. *IEEE Transactions on Automatic Control*, 52(5):782–798, May 2007.

[GP07b] Ben Goertzel and Cassio Pennachin. *Artificial general intelligence*, volume 2. Springer, 2007.

[Gra17] S. Grammatico. Dynamic control of agents playing aggregative games with coupling constraints. *IEEE Transactions on Automatic Control*, 62(9):4537–4548, Sept 2017.

[GSLZ14] D. Gorecky, M. Schmitt, M. Loskyll, and D. Zühlke. Human-machine-interaction in the industry 4.0 era. In *Proceedings of the 12th IEEE International Conference on Industrial Informatics (INDIN)*, pages 289–294, 2014.

[HLS08] D. Henrion, J. B. Lasserre, and C. Savorgnan. Nonlinear optimal control synthesis via occupation measures. In *2008 47th IEEE Conference on Decision and Control*, pages 4749–4754, Dec 2008.

[HNX07] J. P. Hespanha, P. Naghshtabrizi, and Y. Xu. A survey of recent results in networked control systems. In *Proceedings of the IEEE*, volume 95, pages 138–162, 2007.

[IOC11] FAO UNDP IOC/UNESCO, IMO. A blueprint for ocean and coastal sustainability. Technical report, Paris: IOC/UNESCO, 2011.

[JBGN16] Stefan Jaksic, Ezio Bartocci, Radu Grosu, and Dejan Nickovic. Quantitative monitoring of STL with edit distance. In *Runtime Verification - 16th International Conference, RV 2016, Madrid, Spain, September 23-30, 2016*, pages 201–218, 2016.

[JP09] A. Agung Julius and George J. Pappas. Approximations of stochastic hybrid systems. *IEEE Transactions on Automatic Control*, 54(6):1193–1203, June 2009.

[KGFP08] Hadas Kress-Gazit, Georgios E. Fainekos, and George J. Pappas. Translating structured english to robot controllers. *Advanced Robotics*, 22(12):1343–1359, 2008.

[KHHW13] H. Kagermann, J. Helbig, A. Hellinger, and W. Wahlster. *Umsetzungsempfehlungen für das Zukunftsprojekt Industrie 4.0: Deutschlands Zukunft als Produktionsstandort sichern; Abschlussbericht des Arbeitskreises Industrie 4.0.* Forschungsunion, 2013.

[KHR+16] K. Koorehdavoudi, M. Hatami, S. Roy, V. Venkatasubramanian, P. Panciatici, F. Xavier, and J. A. Torres. Input-output characteristics of the power transmission network's swing dynamics. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 1846–1852, Dec 2016.

[KLH13] C. Kwon, W. Liu, and I. Hwang. Security analysis for cyber-physical systems against stealthy deception attacks. In *Proceedings of the 2013 American Control Conference*, pages 3344–3349, 2013.

[Koy90] Ron Koymans. Specifying real-time properties with metric temporal logic. *Real-Time Syst.*, 2(4):255–299, 1990.

[KP17] Nidhi Kalra and Susan M. Paddock. Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability? Technical report, RAND Corporation, 2017.

[Kun94] Prabha Kundur. *Power System Stability and Control*. McGraw-Hill, 1 edition, 1994.

[KW16] E. Karangelos and L. Wehenkel. Probabilistic reliability management approach and criteria for power system real-time operation. In *2016 Power Systems Computation Conference (PSCC)*, pages 1–9, June 2016.

[KZ15] D. Kolberg and D. Zühlke. Lean automation enabled by industry 4.0 technologies. In *Proceedings of the 15th IFAC Symposium on Information Control Problems in Manufacturing*, pages 1870 – 1875, 2015.

[LBK15] J. Lee, B. Bagheri, and H.-A. Kao. A cyber-physical systems architecture for industry 4.0-based manufacturing systems. *Manufacturing Letters*, 3:18 – 23, 2015.

[LC06] Anders Lindgren and Fang Chen. State of the art analysis: An overview of advanced driver assistance systems (adas) and possible human factors issues. *Human factors and economics aspects on safety*, pages 38–50, 2006.

[LCK16] P. Leitão, A. W. Colombo, and S. Karnouskos. Industrial automation based on cyber-physical systems technologies: Prototype implementations and challenges. *Computers in Industry*, 81:11 – 25, 2016.

[LMF16] Y. Lu, K. C. Morris, and S. Frechette. Current standards landscape for smart manufacturing systems. *National Institute of Standards and Technology, NISTIR*, 8107:39, 2016.

[LRH+17] Stefan B. Liu, Hendrik Roehm, Christian Heinzemann, Ingo Lütkebohle, Jens Oehlerking, and Matthias Althoff. Provably safe motion of mobile robots in human environments. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2017, Vancouver, BC, Canada, September 24-28, 2017*, pages 1351–1357, 2017.

[Mil93] Robin Milner. Elements of interaction - turing award lecture. *Commun. ACM*, 36:78–89, 1993.

[Mil96] Robin Milner. Semantic ideas in computing. In *Computing tomorrow : future research directions in computer science*, pages 246–283. Cambridge University Press, 1996.

[Mil99] Robin Milner. *Communicating and mobile systems : the* Π*-calculus*. Cambridge University Press, 1999.

[Mil09] Robin Milner. *The Space and Motion of Communicating Agents*. Cambridge University Press, 2009.

[Min06] Marvin Minsky. *The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind*. Simon & Schuster, Inc., New York, NY, USA, 2006.

[MN04] Oded Maler and Dejan Nickovic. Monitoring temporal properties of continuous signals. In *Formal Modeling and Analysis of Timed Systems*, pages 152–166, 2004.

[MP16] Stefan Mitsch and André Platzer. Modelplex: verified runtime validation of verified cyber-physical system models. *Formal Methods in System Design*, 49(1-2):33–74, 2016.

[PCSB16] Aneesh Paul, Rohan Chauhan, Rituraj Srivastava, and Mriganka Baruah. Advanced driver assistance systems. Technical report, SAE Technical Paper, 2016.

[Pea00] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, NY, USA, 2000.

[PGV14] M. Prandini, S. Garatti, and R. Vignali. Performance assessment and design of abstracted models for stochastic hybrid systems through a randomized approach. *Automatica*, 50(11):2852–2860, 2014.

[PLK] Pavithra Prabhakar, Ratan Lal, and James Kapinski. Automatic trace generation for signal temporal logic. In *Submitted to 2018 Real-Time Systems Symposium.*

[Pnu77] Amir Pnueli. The temporal logic of programs. In *Symposium on Foundations of Computer Science*, pages 46–57, 1977.

[PQ08] André Platzer and Jan-David Quesel. KeYmaera: A hybrid theorem prover for hybrid systems. In *IJCAR*, volume 5195 of *LNCS*, pages 171–178, 2008.

[RHM17] Hendrik Roehm, Thomas Heinz, and Eva Charlotte Mayer. Stlinspector: STL validation with guarantees. In *Computer Aided Verification - 29th International Conference, CAV*, pages 225–232, 2017.

[Sch97] R. R. Schaller. Moore's law: past, present and future. *IEEE Spectrum*, 34(6):52–59, June 1997.

[SECP13] G. Schirner, D. Erdogmus, K. Chowdhury, and T. Padir. The future of human-in-the-loop cyber-physical systems. *Computer*, 46(1):36–45, 2013.

[SMD+17] Kit Siu, Abha Moitra, M B Durling, Andy Crapo, Meng Li, Han Yu, Heber Herencia-Zapana, Mauricio Castillo-Effen, Shiraj Sen, Craig McMillan, Daniel Russell, Sundeep Roy, and Panagiotis Manolios. Flight critical software and systems development using ASSERT^{TM}. *2017 IEEE/AIAA 36th Digital Avionics Systems Conference (DASC)*, pages 1–10, 2017.

[SMN18] R. Salhab, R. P. Malhamé, and J. Le Ny. A dynamic game model of collective choice in multiagent systems. *IEEE Transactions on Automatic Control*, 63(3):768–782, March 2018.

[TN16] Paulo Tabuada and Daniel Neider. Robust linear temporal logic. In Jean-Marc Talbot and Laurent Regnier, editors, *25th EACSL Annual Conference on Computer Science Logic*, number 10, pages 1–21, 2016.

[ZSCC16] H. Zhang, Y. Shu, P. Cheng, and J. Chen. Privacy and performance trade-off in cyber-physical systems. *IEEE Network*, 30(2):62–66, 2016.