

# Modeling Cyber–Physical Systems

*Developing cyber–physical systems (CPS) models is the main concern of this paper which provides a comprehensive overview of their current state and future directions.*

By PATRICIA DERLER, *Member IEEE*, EDWARD A. LEE, *Fellow IEEE*, AND ALBERTO SANGIOVANNI VINCENTELLI, *Fellow IEEE*

**ABSTRACT** | This paper focuses on the challenges of modeling cyber-physical systems (CPSs) that arise from the intrinsic heterogeneity, concurrency, and sensitivity to timing of such systems. It uses a portion of an aircraft vehicle management system (VMS), specifically the fuel management subsystem, to illustrate the challenges, and then discusses technologies that at least partially address the challenges. Specific technologies described include hybrid system modeling and simulation, concurrent and heterogeneous models of computation, the use of domain-specific ontologies to enhance modularity, and the joint modeling of functionality and implementation architectures.

**KEYWORDS** | Computers and information processing; computer simulation; continuous time systems; discrete time systems; embedded software; modeling; simulation; software; system analysis and design; systems engineering and theory; time factors

## I. INTRODUCTION

Cyber–physical systems (CPSs) are integrations of computation and physical processes. Embedded computers and networks monitor and control the physical processes,

usually with feedback loops where physical processes affect computations and *vice versa*. The design of such systems, therefore, requires understanding the joint dynamics of computers, software, networks, and physical processes. It is this study of *joint* dynamics that sets this discipline apart.<sup>1</sup>

When studying CPS, certain key problems emerge that are rare in so-called general-purpose computing. For example, in general-purpose software, the time it takes to perform a task is an issue of *performance*, not *correctness*. It is not incorrect to take longer to perform a task. It is merely less convenient and therefore less valuable. In CPS, the time it takes to perform a task may be critical to correct functioning of the system.

In CPS, moreover, many things happen at once. Physical processes are compositions of many things occurring at the same time, unlike software processes, which are deeply rooted in sequential steps. Abelson and Sussman [1] describe computer science as “procedural epistemology,” knowledge through procedure. In the physical world, by contrast, processes are rarely procedural.

Physical processes are compositions of many parallel processes. Measuring and controlling the dynamics of these processes by orchestrating actions that influence the processes are the main tasks of embedded systems. Consequently, concurrency is intrinsic in CPS. Many of the technical challenges in designing and analyzing embedded software stem from the need to bridge an inherently sequential semantics with an intrinsically concurrent physical world.

The goal of this paper is to highlight progress that has been made, but mostly still not yet adopted by the

Manuscript received March 1, 2011; revised June 17, 2011; accepted June 23, 2011. Date of publication August 22, 2011; date of current version December 21, 2011.

This work was supported in part by the Center for Hybrid and Embedded Software Systems (CHESS), University of California Berkeley, which receives support from the National Science Foundation [NSF awards #0720882 (CSR-EHS: PRET), #0931843 (ActionWebs), and #1035672 (CSR-CPS Ptides)], the U.S. Army Research Office (ARO #W911NF-07-2-0019), the U.S. Air Force Office of Scientific Research (MURI #FA9550-06-0312), the Multiscale Systems Center (MuSyC), one of six research centers funded under the Focus Center Research Program, a Semiconductor Research Corporation program, and the following companies: Bosch, National Instruments, Thales, and Toyota.

The authors are with the Department Electrical Engineering and Computer Science, University of California Berkeley, Berkeley, CA 94720 USA (e-mail: pd@eecs.berkeley.edu; eal@eecs.berkeley.edu; alberto@eecs.berkeley.edu).

Digital Object Identifier: 10.1109/JPROC.2011.2160929

<sup>1</sup>As an experiment, the models discussed in this paper are available in executable form using Web Start at <http://ptolemy.org/modelingCPS>.

community, on an enormously difficult and complex problem, the design of CPSs. The major theme of this paper is on *models* and their relationship to realizations of CPSs. The models we study are primarily about dynamics, the evolution of a system state in time. We do not address structural models, which represent static information about the construction of a system, although these too are important to system design.

Working with models has a major advantage. Models can have formal properties. We can say definitive things about models. For example, we can assert that a model is deterministic, meaning that given the same inputs it will always produce the same outputs. No such absolute assertion is possible with any physical realization of a system. If our model is a good abstraction of the physical system (here, “good abstraction” means that it omits only inessential details), then the definitive assertion about the model gives us confidence in the physical realization. Such confidence is hugely valuable, particularly for embedded systems where malfunctions can threaten human lives. Studying models of systems gives us insight into how those systems will behave in the physical world.

In this paper, we begin in Section II by considering a vehicle management system (VMS) challenge problem provided to us by the United Technologies Corporation as support to our MuSyC research effort (Multiscale System Center). We focus on the fuel management subsystem, using it to illustrate modeling challenges that arise from the intrinsic heterogeneity, concurrency, and sensitivity to timing of such systems. In Section III, we survey a few existing techniques that at least partially address these challenges, and we identify open research challenges. In Section IV, we analyze the state of the art in existing tools and methods.

## II. MODELING CHALLENGES

In model-based design [73] and model-driven development [72], models play an essential role in the design process. They form the specifications for systems and reflect the evolution of the system design. They enable simulation and analysis, both of which can result in earlier identification of design defects than prototyping. Automated or semiautomated processes can, under certain circumstances, synthesize implementations from models. But the intrinsic heterogeneity and complexity of CPS stresses all existing modeling languages and frameworks.

A model of a CPS comprises models of physical processes as well as models of the software, computation platforms, and networks. The feedback loop between physical processes and computations encompasses sensors, actuators, physical dynamics, computation, software scheduling, and networks with contention and communication delays. Modeling such systems with reasonable fidelity is challenging, requiring inclusion of control engineering, software engineering, sensor networks, etc. Addi-

tionally, the models typically involve a large number of heterogeneous components. Composition semantics becomes central.

We will illustrate modeling challenges in CPS using a portion of an aircraft VMS, specifically the fuel management subsystem. We give small illustrative models constructed in Ptolemy II [19], a modeling and simulation environment for heterogeneous systems. Ptolemy uses an actor-oriented design approach [44] to model components that communicate via ports. Actors can execute and communicate with other actors. The rules of interaction between actors are defined by the model of computation (MoC).

In Ptolemy, an MoC is indicated and implemented by a *director*, which is a component in a model or submodel. A multiplicity of directors realizing distinct MoCs can be combined in a single hierarchical model. A region of a model governed by a single director is called a *domain*. The open-source Ptolemy software distribution includes directors that implement a variety of both mature and experimental MoCs, including discrete events (DEs) [35], continuous time (CT) [49], finite state machines (FSMs) [47], synchronous reactive (SR) [17], process networks (PNs) [45], and several varieties of dataflow [43]. Hierarchical compositions that combine CT models with discrete domains such as FSM or DE can be used to model hybrid systems.

Modern aircraft are typically equipped with several fuel tanks and a variety of valves, pumps, probes, sensors, and switches. The primary purpose of the fuel system is to reliably supply engines with fuel. Secondary functions of the fuel system include providing an engine coolant and distributing weight on an aircraft to maintain balance and optimal performance. Fuel is transferred from collector tanks to engine feed tanks, and, to maintain the center of gravity of the vehicle, fuel is transferred between storage tanks. Fuel might also be used as a heat sink for heat generated during the flight. For example, fuel is used to cool the engine oil. Raising the temperature of the fuel can also make engines operate more efficiently. Fuel venting systems are required to adjust ullage space in tanks during aircraft climbs and descends. Without such venting systems, high forces on the tanks endanger the structural integrity. Fuel pressurization is required to perform efficient refueling and to transfer fuel between tanks. The system is far from simple.

Several physical processes affect the fuel control system. Weight and density of fuel changes with the temperature of the fuel, and different fuel types have different characteristics. Measuring the fuel level is complex because one has to consider the shape of the fuel tank, the motion of the aircraft, the temperature, and the pressure. In order to measure fuel level during flight operations, a large number of sensors, including float level, diodes, capacitance sensors, and ultrasonic sensors, are used. The fuel control system monitors the values retrieved from sensor tanks, computes and displays the fuel levels and

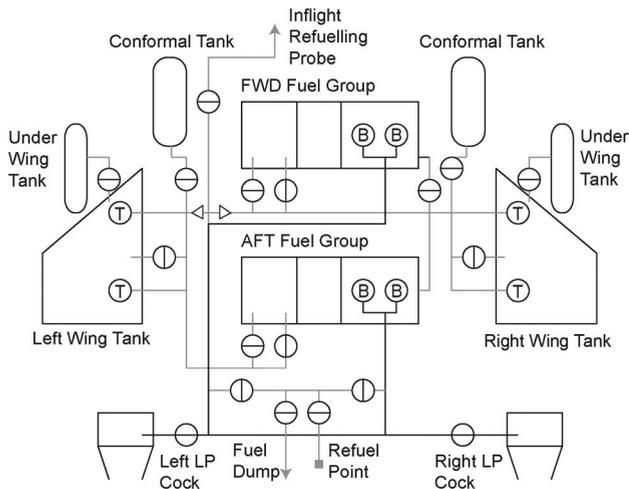


Fig. 1. Military aircraft system (after [57]).

operates pumps and valves. Fig. 1 shows a schematic view of a typical aircraft fuel system for a military jet aircraft.

The importance of a fuel management system that supports flight crew decisions is underscored by the experience with Air Transat Flight 236 on August 23–24, 2001 [14] (see also [76]). On that flight, a fuel leak was discovered too late and led to an emergency landing. Initially, the only indication of a problem was that the flight control system reported a low oil temperature and a high oil pressure. This resulted in an increased fuel flow caused by a leak in the engine. Fuel is used as a heat sink, and the increased fuel flow cooled the engine oil more than usual. The decreased oil temperature increased its viscosity, resulting in higher oil pressure. Later, the flight control system reported a fuel imbalance between the tanks. The flight crew corrected this imbalance by transferring fuel into the leaking tank. A smarter flight control system that integrated models could have helped the flight crew detect the leak sooner, for example, by using model-integrated computing [73], where the behavior of a model can be compared at run time to the measured behavior of the deployed system.

Using the fuel system as an illustration, we next discuss a few specific modeling challenges.

**Challenge 1—Models With Solver-Dependent, Nondeterminate, or Zeno Behavior:** A CPS may be modeled as a hybrid system where physical processes are represented as continuous-time models of dynamics and computations are described using state machines, dataflow models, synchronous/reactive models, and/or DE models. Continuous-time models work with solvers that numerically approximate the solutions to differential equations. Design of such solvers is an established, but far-from-trivial art. Integrating such solvers with discrete models is a newer problem, and problems persist in many available tools.

One of the problems that can arise is that the behavior defined by a model may be nondeterminate even when the underlying system being modeled is determinate. This means that the model defines a multiplicity of behaviors, rather than a single behavior. This can occur, for example, when DEs are simultaneous and the semantics of the modeling language fails to specify a single behavior. Such nondeterminism can surprise system designers. Another problem that can arise is that numerical solvers typically dynamically adjust the step size that they use to increment time, and the behavior of the model can depend on the selected step sizes. Again, this will surprise designers, who rarely have insight into the basis for step-size selection. A third problem that can arise is that models may exhibit Zeno behavior, where infinitely many events occur in a finite time interval. Such behavior from a model may reflect physical phenomena, such as chattering, but Zeno behavior can also arise as an artifact of modeling.

We can illustrate some of these difficulties using a model of a fuel tank. One challenging part of such a model is to accurately specify how the behavior changes when the tank becomes full or empty. It is reasonable to model “becoming full” and “becoming empty” as DEs, even though the physical world is not so discrete. We do not want to force modelers to use overly detailed models of physical phenomena that are not central to their design problem, and such discrete models provide adequate abstractions.

Fig. 2 depicts one possible Ptolemy II model of a single fuel tank. This model is constructed using a continuous-time director with semantics similar to popular continuous-time simulation tools such as Simulink (from The MathWorks) and the LabVIEW Control Design and Simulation Module (from National Instruments). (For example, Simulink is used in [31] to construct similar models of aircraft fuel systems.) In this MoC, components (actors) have input and output ports, in contrast to equational modeling tools such as Modelica [23], where ports are neither inputs nor outputs. In Modelica ports represent continuous-time signals, and actors and connections between ports represent declarative relationships among these signals.

While it may seem intuitive that a fuel tank is naturally modeled with inputs and outputs—inputs provide fuel for the tank, and outputs withdraw fuel—this is actually not the case. An input to a fuel tank cannot insert fuel if the tank is full, and an output cannot withdraw fuel if the destination is full. Hence, if we want to model, for example, the transfer of fuel from one tank to another, the amount of fuel transferred per unit time must be determined collaboratively by the two tanks and a model of the intervening pumps and valves.

Fig. 2 shows one possible resolution to this dilemma. Two of the input ports of the tank represent *desired input* and *desired output* flows, rather than actual flows. A third input represents the *actual output* flow. The actual output flow has to be an input to the tank model because fuel

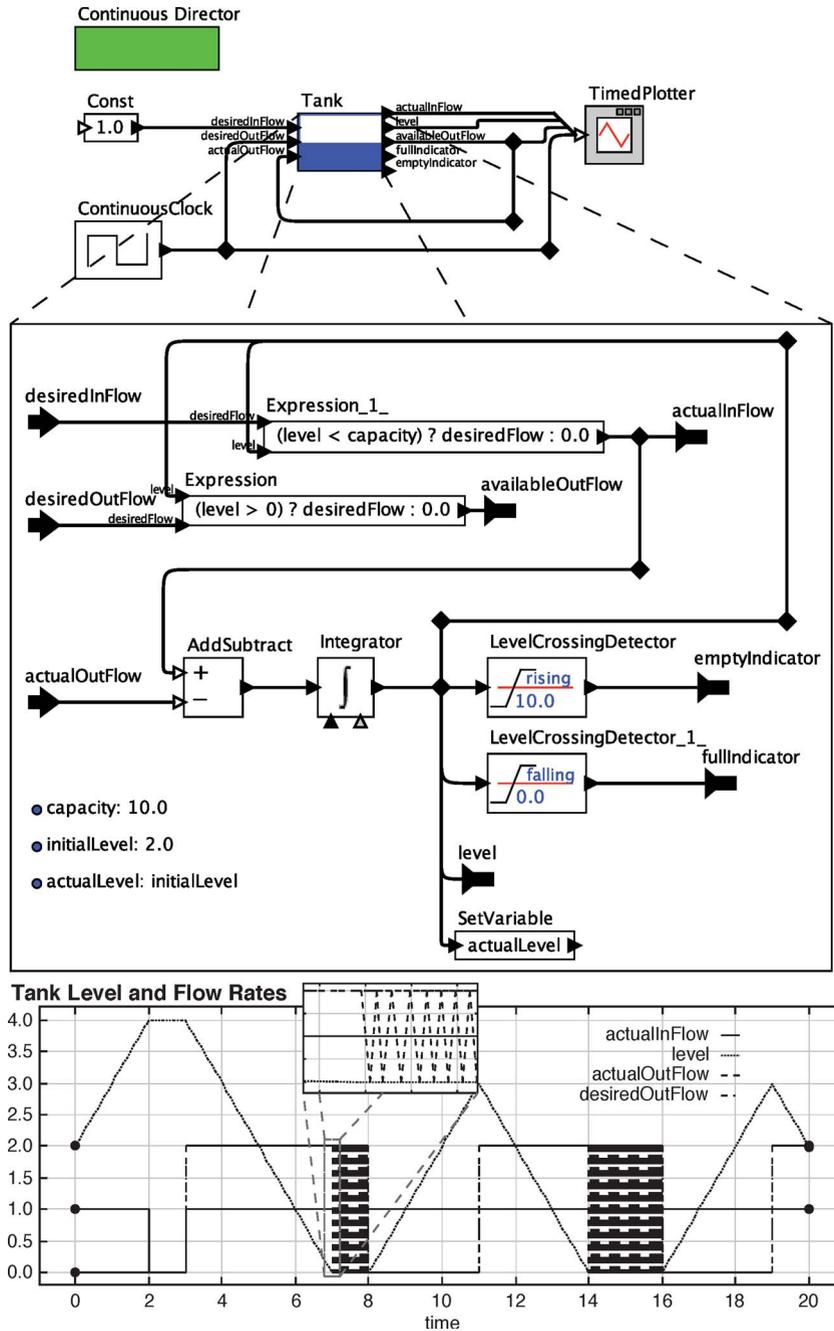


Fig. 2. Naive model of a single fuel tank in a simple test scenario.

cannot flow out of a tank unless it has somewhere to go. One of the output ports represents *available output* flow, thus enabling models where the destination for the fuel refuses to accept some or all of the output flow. At the top level in Fig. 2, the *available output* flow output is fed back to the *actual output* flow input, which models the situation where there is always somewhere for the fuel to go. The destination can accept any amount of fuel.

Another output port represents the *actual input* flow. If the tank is not full, then, ideally, the actual input flow will

closely match the desired input flow. A true match between flows is desired but in a real system many factors such as the shape of the pipe might influence the flows and therefore introduce slight variations. But if the tank is full, then the actual input flow will significantly differ from the desired flow.

The model in Fig. 2 also shows output ports that provide the fuel level and full and empty indicators. The latter indicators are DEs provided by LevelCrossingDetector actors.

The particular model in Fig. 2 uses two expression actors to implement what may seem to be an intuitively accurate description of how a fuel tank behaves. In this model, the tank ignores the desired input flow to the tank when the tank is full, setting the actual input flow to zero. It also stops the output flow when the tank is empty.

The behavior of the model is problematic, however. In particular, Fig. 2 embeds the tank model in a simple test scenario where the desired input flow is a constant 1.0 (in some units, such as liters per second), and the desired output flow is 0.0 for three time units, followed by 2.0 subsequently. As shown in the plot in the figure, when the tank becomes empty, the model begins chattering. This behavior exposes one flaw in the model that results in solver-dependent behavior and/or the possibility of a Zeno condition. Specifically, when the tank becomes empty, actual output flow goes to zero. But an infinitesimal amount of time later, the tank is no longer empty, so the output flow can resume. But then it becomes empty again, then resumes, etc. In simulation, it appears as if time stops advancing. Only close examination of the output signals reveals the chattering behavior.

This model has several features that challenge simulation technology. First, the expression actors have input-output relations that are not continuous. The lower one in Fig. 2, for example, has the expression

$$(\text{level} > 0) \ ? \ \text{desiredFlow} : 0.0.$$

The output of this actor is equal to the input if the level is greater than zero, and otherwise the output is zero. When the level hits zero, the actual behavior of this actor will depend on numerical precision and solver-chosen step sizes. Giving a precise semantics to the time at which the level hits zero is far from trivial. And controlling the behavior at the discontinuity requires collaboration between the solver and the expression actor, something not easy to achieve in a modular way.

In addition, the LevelCrossingDetector should produce a DE at the precise time at which the tank becomes full or empty. This also requires careful collaboration with the solver to ensure that step sizes in time are chosen to include the time point at which the event occurs.

The chattering exhibited by the model in Fig. 2 may in fact accurately represent the physical realization. One could design a fuel tank controller that closes an output valve when the tank becomes empty. Such a design could, in fact, chatter. A better design would limit the input and output flows when the tank becomes either full or empty. When the tank becomes full, the input flow will be limited to be no greater than the output flow, and when the tank becomes empty, the output flow will be no greater than the input flow.

Below, we will give such a model, showing how it can leverage a rigorous hybrid-systems semantics to provide a determinate and intuitive model. But before we develop the better model, we address a second challenge that concerns how to evolve and maintain a multiplicity of models and keep them consistent.

*Challenge 2—Keeping Model Components Consistent:* In the previous section, we developed a simple tank model and tested it in a scenario that does not reflect any real-world scenario, but is useful nonetheless as a regression test for the tank model. In practice, unfortunately, such test scenario models are usually lost. They evolve into much more complex and realistic models, so the simple regression test disappears. If the model is kept as a regression test, then there is little assurance that the test model and the final design match. The final design is likely to be evolved from a copy of the test model, and the copy will inevitably diverge from the test model. Worse, when it comes time to design the embedded software, even if code generators are used to synthesize software from models, the software will often evolve separately from the model, increasing divergence.

The same problem arises as a simple model evolves into a complex one, where a single component in the simple model becomes multiple components in the complex one. How can we ensure that the multiple components evolve together? Until now we have only considered models with one fuel tank. Modern aircraft typically have several fuel tanks, and fuel is transferred between these tanks.

In this section, we consider the problem of evolving multiple models with multiple variants of components, all the while ensuring some measure of consistency across the models.

In the model in Fig. 3, we add functionality for moving fuel between tanks. Whenever the fuel level of Tank1 falls below one fourth of its capacity, a FuelMover component transfers fuel from Tank2 to Tank1. This component uses a simple modal model to change the desired output flow of Tank2 from zero in the initial `idle` mode to `moveRate` (a parameter) in the `moving` state. The FuelMover controller specifies the desired output flow of Tank2, which in turn provides an available output flow to the desired input flow port of Tank1. Tank1 responds with an actual input flow, which in turn becomes the actual output flow of Tank2. Thus, if either tank is full or empty, the flows are appropriately limited.

In Fig. 3, the FuelMover and the two tanks are embedded in a simple test scenario where Tank1 has a constant desired output flow of 1. This could, for example, model the flow of fuel to an engine. The desired input flow of Tank1 is regulated by Tank2. Tank2 does not receive any fuel (its desired input flow is 0) and its desired output flow is regulated by the FuelMover. This component is a modal model that receives information about the fuel level of Tank1 and regulates the desired output flow of Tank2. The

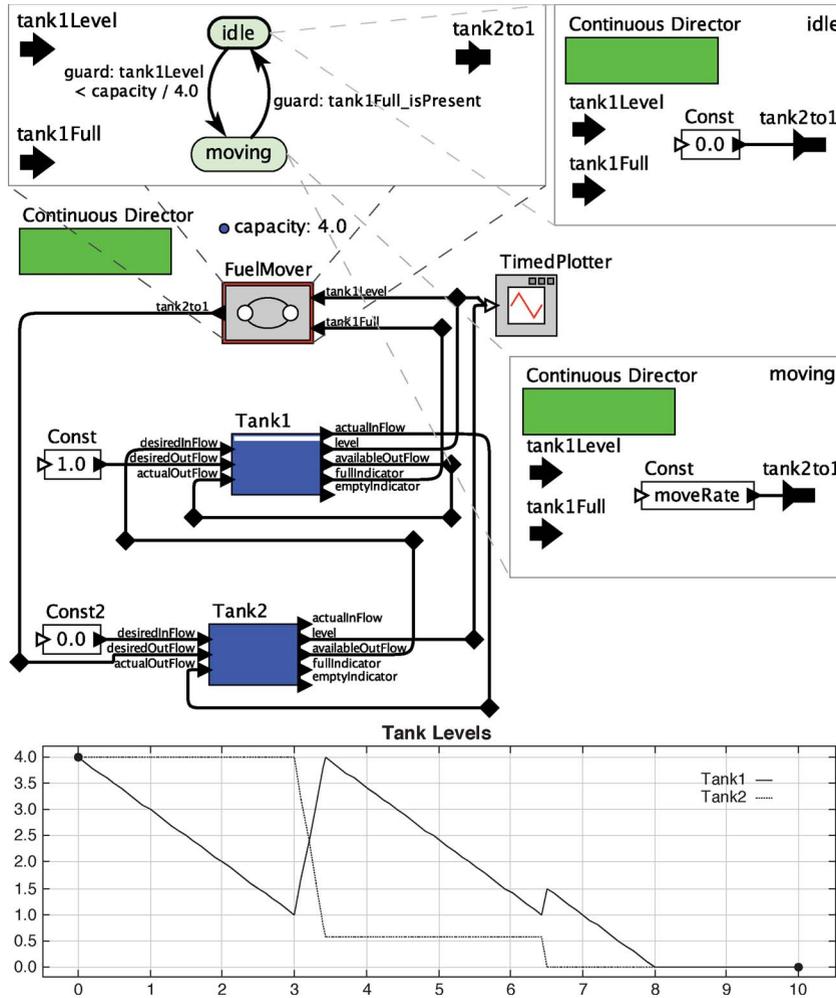


Fig. 3. Fuel system with two tanks.

fuel mover can be interpreted as the cyber part of this CPS, while the Tank models the physical processes.

In a modeling environment, the tank model can be copied and reused in various parts of the model. However, if later a change in the tank model becomes necessary, the same change has to be applied to all other tank models that were copied. This procedure is error prone, because there is no way to ensure that all copies are updated accordingly.

Below, we explain how object-oriented design principles can be adapted to CPS models using the notion of actor-oriented classes [42], and how semantics-preserving transformation and code generation can reduce divergence in evolving designs.

*Challenge 3—Preventing Misconnected Model Components:* The bigger a model becomes, the harder it is to check for correctness of connections between components. Typically model components are highly interconnected and the possibility of errors increases. We identify three types of errors: 1) unit errors, 2) semantic errors, and 3) transposition

errors, illustrated in Fig. 4. Unit errors occur when a port provides data with different units than those expected at the receiving port. For example, in Fig. 4, the output port of the left actor provides the fuel level in gallons, whereas the input port of the right actor expects a value in kilograms. An example of a semantic error 2) is shown, where

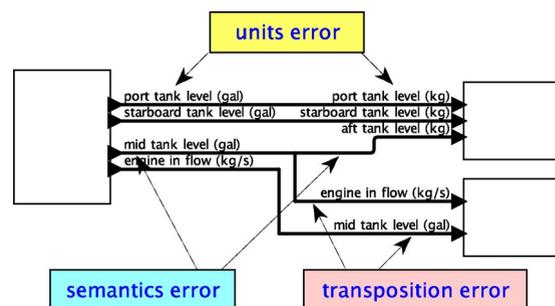


Fig. 4. Misconnected model components.

a midtank level is provided to a port that expects an aft tank level. A transposition error 3) occurs when connections are reversed, as in the example where the engine input flow and a tank level are reversed. None of these errors would be detected by a type system. Below we will explain how expressing knowledge formally in an ontology [51] can encode domain knowledge and enable measures to automatically check for such errors.

*Challenge 4—Modeling Interactions of Functionality and Implementation:* The models discussed above describe functionality of computations and physical processes. All of the above models ignore implementation details such as the underlying platform and the communication networks. These models implicitly assume that data are computed and transmitted in zero time, so the dynamics of the software and networks have no effect on system behavior. However, computation and communication takes time. We argue that, in order to evaluate a CPS model, it is necessary to model the dynamics of software and networks.

Implementation is largely orthogonal to functionality and should therefore not be an integral part of a model of functionality. Instead, it should be possible to conjoin a functional model with an implementation model. The latter allows for design space exploration, while the former supports the design of control strategies. The conjoined models enable evaluation of interactions across these domains.

One strategy for conjoining distinct and separately maintained models uses the concept of aspect-oriented programming (AOP) [34]. In AOP, distinct models are maintained and developed, but they can be “woven” together to conjoin their behaviors. AOP has been developed primarily as a software engineering technique, but the concept has also been applied to conjoining models of functionality and implementation. Specifically, below we explain how to use the concept of quantity managers, introduced in Metropolis [3], to associate functional models with implementation models.

*Challenge 5—Modeling Distributed Behaviors:* The distributed nature of CPS requires strategies that facilitate composition of components that are separated in space. Modeling distributed systems adds to the complexity of modeling CPS by introducing issues such as disparities in measurements of time, network delays, imperfect communication, consistency of views of system state, and distributed consensus. Below, we will describe a distributed programming model called PTIDES [78] that addresses many of these issues.

*Challenge 6—System Heterogeneity:* An undercurrent throughout the above challenges arises from the intrinsic heterogeneity and complexity of CPSs. By definition, CPS integrates diverse subsystems. Often these subsystem have domain-specific, and sometimes quite sophisticated mod-

eling techniques. Integrating the resulting models to develop holistic views of the system becomes very challenging.

### III. ADDRESSING THE CHALLENGES

Over the last 20 years, we at the University of California Berkeley have been researching model-based design techniques that we believe have sufficiently well-defined semantics to provide an effective basis for platform-based design and engineering. Components can be designed to operate with a model, and when deployed, will operate in predictable ways with the deployed system. The rigorous foundations of the models [46] provide a foundation for integration across design domains, design adaptation and evolution, and analysis and verification. Our work has been demonstrated in the open-source software frameworks Ptolemy Classic [9], Ptolemy II [19], Polis [2], Metropolis [3], and MetroII [15]. Many of the techniques that we developed have been deployed in a wide range of domain-specific applications, including hardware and field-programmable gate array (FPGA) synthesis, signal processing, automotive system design, computer architecture design and evaluation, instrumentation, wireless system design, mixed signal circuit design, network simulation and design, building system simulation and design, financial engineering, and scientific workflows.

In the remainder of this section, we review some of the key technologies that can help. In the following section, we outline the considerable work remaining to be done.

#### A. Models of Computation

Several of the challenges arise fundamentally from the complexity and heterogeneity of CPS applications. Complexity of models is mitigated by using more specialized, domain-specific models and by using modeling languages with clear, well-defined semantics. Heterogeneity necessitates combining a multiplicity of models (Challenge 6).

To address the first of these issues, models should be built using well-defined MoCs [16]. An MoC gives semantics to concurrency in the model, defining, for example, whether components in the model execute simultaneously, whether they share a notion of time, and whether and how they share state. An MoC also defines the communication semantics, specifying, for example, whether data are exchanged using publish-and-subscribe protocols, synchronous or asynchronous message transmission, or time-stamped events. We have provided a formal framework within which the semantics of a variety of MoCs can be understood and compared, and heterogeneous interactions across MoCs can be defined [46]. This formal foundation has been elaborated and applied to multiclock (latency insensitive) systems [12], globally asynchronous, locally synchronous (GALS) designs [4], and to timed MoCs capable of reflecting real-time dynamics [54]. The challenge is to define MoCs that are sufficiently expressive and have strong formal properties that enable systematic validation

of designs and correct-by-construction synthesis of implementations. A second challenge is to identify which of the many MoCs and variants are actually needed, and to figure out how to educate the community to use them.

## B. Abstract Semantics

In many situations, using a single general MoC for an entire design requires giving up any possibility of property checking except by extensive simulation. More restricted (less expressive) MoCs yield better to analysis, enabling systematic exploration of properties of the design, often including formal verification. But less expressive MoCs cannot capture the richness and diversity of complex designs. The solution is heterogeneous mixtures of MoCs. Indeed, the heterogeneous nature of most CPS applications makes multiple MoCs a necessity (Challenge 6).

In addition, during the design process, the abstraction level, detail, and specificity in different parts of the design vary. The skill sets and design styles that engineers use on the project are likely to differ. The net result is that, during the design process, many different specification and modeling techniques will be used. The challenge is how to combine heterogeneous MoCs and determine what the composition's behavior is. Unfortunately, the semantics of different MoCs are typically not directly compatible.

A way to solve this problem is to embed the detailed models into a framework that can understand the models being composed. A theoretical approach to this view, which is well beyond the scope of this paper, can be found in [10], where an abstract algebra defines the interactions among incompatible models. In some sense, we are looking at an abstraction of the MoC concept that can be refined into any of the MoCs of interest. We call this abstraction an *abstract semantics*, first introduced in [37] and [44].

The inspiration for how to define the abstract semantics comes from the consideration that MoCs are built by combining three largely orthogonal aspects: sequential behavior, concurrency, and communication. Similar to the way that an MoC abstracts a class of behavior, abstract semantics abstract the semantics. The concept is called a “semantics meta-model” in [70], but since the term “meta-model” is more widely used in software engineering to refer instead to models of the structure of models (see [59] and <http://www.omg.org/mof/>), we prefer to use the term “abstract semantics” here. The concept of abstract semantics is leveraged in Ptolemy II [19], Metropolis [3], and Ptolemy Classic [9] to achieve heterogeneous mixtures of MoCs with well-defined interactions.

## C. Actor-Oriented Models

Model-integrated development for embedded systems [30], [33] commonly uses actor-oriented software component models [36], [44]. In such models, software components (called actors) execute concurrently and communicate by sending messages via interconnected ports.

Examples that support such designs include Simulink, from The MathWorks, LabVIEW, from National Instruments, SystemC, component and activity diagrams in SysML and UML 2 [5], [60], and a number of research tools such as ModHel’X [26], TDL [66], HetSC [28], ForSyDe [67], Metropolis [24], and Ptolemy II [19]. The fuel tank models given above are all actor-oriented models.

The key challenge is to provide well-defined actor-oriented MoCs with well-defined semantics. All too often the semantics emerge accidentally from the software implementation rather than being built-in from the start. One of the key challenges is to integrate actor-oriented models with practical and realistic notions of time. To address Challenge 5, for example, modeling distributed behaviors, it is essential to provide multiform models of time. Modeling frameworks that include a semantic notion of time, such as Simulink and Modelica, assume that time is homogeneous in the sense that it advances uniformly across the entire system. In practical distributed systems, even those as small as a single chip, no such homogeneous notion of time is observable. In a distributed system, even when using network time synchronization protocols (such as IEEE 1588 [29]), local notions of time will differ, and failing to model such differences could introduce artifacts in the design.

One interesting project that directly confronts the multiform nature of time in distributed systems is the PTIDES project [18]. PTIDES is a programming model for distributed systems that relies on time synchronization, but recognizes imperfections. Simulations of PTIDES systems can simultaneously have many time lines, with events that are logically or physically placed on these time lines. Despite the multiplicity of time lines, there is a well-defined and deterministic semantics for the interactions between events.

## D. Hybrid Systems

CPSs integrate computation, networking, and physical dynamics. As a consequence, modeling techniques that address only the concerns of software are inadequate [38], [40]. Integrations of continuous physical dynamics expressed with ordinary differential equations with the discrete behaviors expressed using finite automata are known as *hybrid systems* [56]. There are many challenges in defining good hybrid systems modeling languages, a few of which are described in Challenge 1. These challenges include ensuring that models of deterministic systems are indeed determinate, not having behavior that, for example, depends on arbitrary decisions made by a numerical differential equation solver. It is also challenging to accurately represent distinct events that are causally related but occur at the same time.

Tools supporting hybrid systems modeling and simulation are compared in [11]. A rigorous MoC that provides determinate semantics to such hybrid systems is described in [49], [52], and [55]. This work has influenced

development of commercial tools such as Simulink and LabVIEW and has been realized in the open-source tool HyVisual [7]. Moreover, we have leveraged the notation of abstract semantics to integrate such hybrid systems with other MoCs such as synchronous/reactive and DE models [50]. This integration enables heterogeneous models that capture the interactions of software and networks with continuous physical processes.

We can use the rigorous hybrid systems MoC implemented in Ptolemy II to improve the tank model described in Challenge 1 above. Such a model is shown in Fig. 5, which uses a modal model [47] to describe the different modes of operation based on the state of the tank. In this model, the net flow is equal to the difference between the actual input flow and the actual output flow. The model has three modes. In the `normal` mode, the tank is neither full nor empty, and the actual input flow equals the desired flow, and the available output flow equals the desired output flow. In the `full` and `empty` modes, either the actual input flow (if the tank is full) or the available output flow (if the tank is empty) is limited to match the other. The transitions between modes are triggered by the guards, which are predicates shown adjacent to the transitions in the figure. This model does not chatter and is determinate.

Note that in a manner very similar to Fig. 5, modal models can be used to encode adaptation strategies that all CPSs need to deal with uncertain environments and with faults. For example, the state in which a tank that feeds an engine becomes empty might be considered an error condition if there remains fuel in other tanks. A modal model could specify a strategy for adapting to that eventuality. Also, analysis of the dynamics of the model could identify conditions under which such an error condition occurs.

The hybrid systems semantics of Ptolemy II handles rigorously the discontinuities of the mode transitions and the event detection of the `LevelCrossingDetector` actors. The semantics ensures determinism in the model, helps avoid solver-dependent behavior, and helps to identify Zeno conditions. This helps in developing confidence in the models, and hence confidence in the system.

At the bottom of Fig. 5, the plot shows the behavior of the system in a test scenario where the desired input flow is a constant 1.0, and the desired output flow alternates between zero (for three time units) and 2.0 (for five time units). Such a scenario provides a reasonable regression test for behavior of the tank model that includes full and empty conditions.

## E. Heterogeneity

Integrating multiple MoCs such that they can interoperate, which is far from trivial, has been called “multi-modeling” [8], [22], [58]. Many previous efforts have focused on tool integration, where tools from multiple vendors are made to interoperate [25], [32], [53]. This approach is challenging, however, and yields fragile tool chains. Many tools do not have adequate published ex-

tension points, and maintaining such integration requires considerable effort. A better approach is to focus on the semantics of interoperation, rather than the software problems of tool integration. Good software architectures for interoperation will emerge only from a good understanding of the semantics of interoperation.

Key to heterogeneity is to focus on the interfaces between MoCs. We have built a variety of modeling, analysis, and simulation tools based on different MoCs [16], and have shown how such interfaces can facilitate more robust interoperability. These include DE [35] (useful for modeling networks, hardware architecture, and real-time systems), synchronous-reactive [17] (useful for modeling and designing safety-critical concurrent software), dataflow [43] (useful for signal processing), process networks [63] (useful for asynchronous distributed systems), and continuous-time models [50] (useful for physical dynamics).

For nearly all of these MoCs, the emphasis in our design has been on providing determinate behavior (where the same inputs always result in the same outputs), and introducing nondeterminacy only where it is needed by the application (for example, to model faults). The result is a family of far better concurrency models than the widely used thread-based models that dominate software engineering [39].

The hybrid systems MoC of Ptolemy II is, in fact, built using such interfaces between MoCs. The modal model in Fig. 5 is constructed using a Ptolemy director that has no particular knowledge of continuous-time modeling, yet it interoperates with a director that realizes continuous-time modeling. It also interoperates with other Ptolemy directors, enabling modal models to be used anywhere in a design with a wide variety of modeling styles. This stands in stark contrast to, say, Statecharts techniques, where the state machine modeling framework is designed to specifically work with exactly one concurrent MoC, typically synchronous/reactive.

Systematic support for heterogeneity can also help with Challenge 2, keeping model components consistent, because models that interoperate are less likely to diverge than models that do not. It also helps with Challenge 4, modeling interactions of functionality and implementation, because different pieces of a system can evolve nonuniformly from abstract models to detailed models of realizations. It also helps with Challenge 5, modeling distributed behaviors, because models of networks (particularly wireless networks) benefit enormously from domain-specific modeling and simulation techniques created particularly for networks. If heterogeneity is supported, then models of networks interoperate with models of the distributed components, be they mechanical, chemical, or software components.

Influenced in part by our work, SystemC, a widely used language in hardware design, is capable of realizing multiple MoCs [28], [65], although less attention in that community has been given to interoperability.

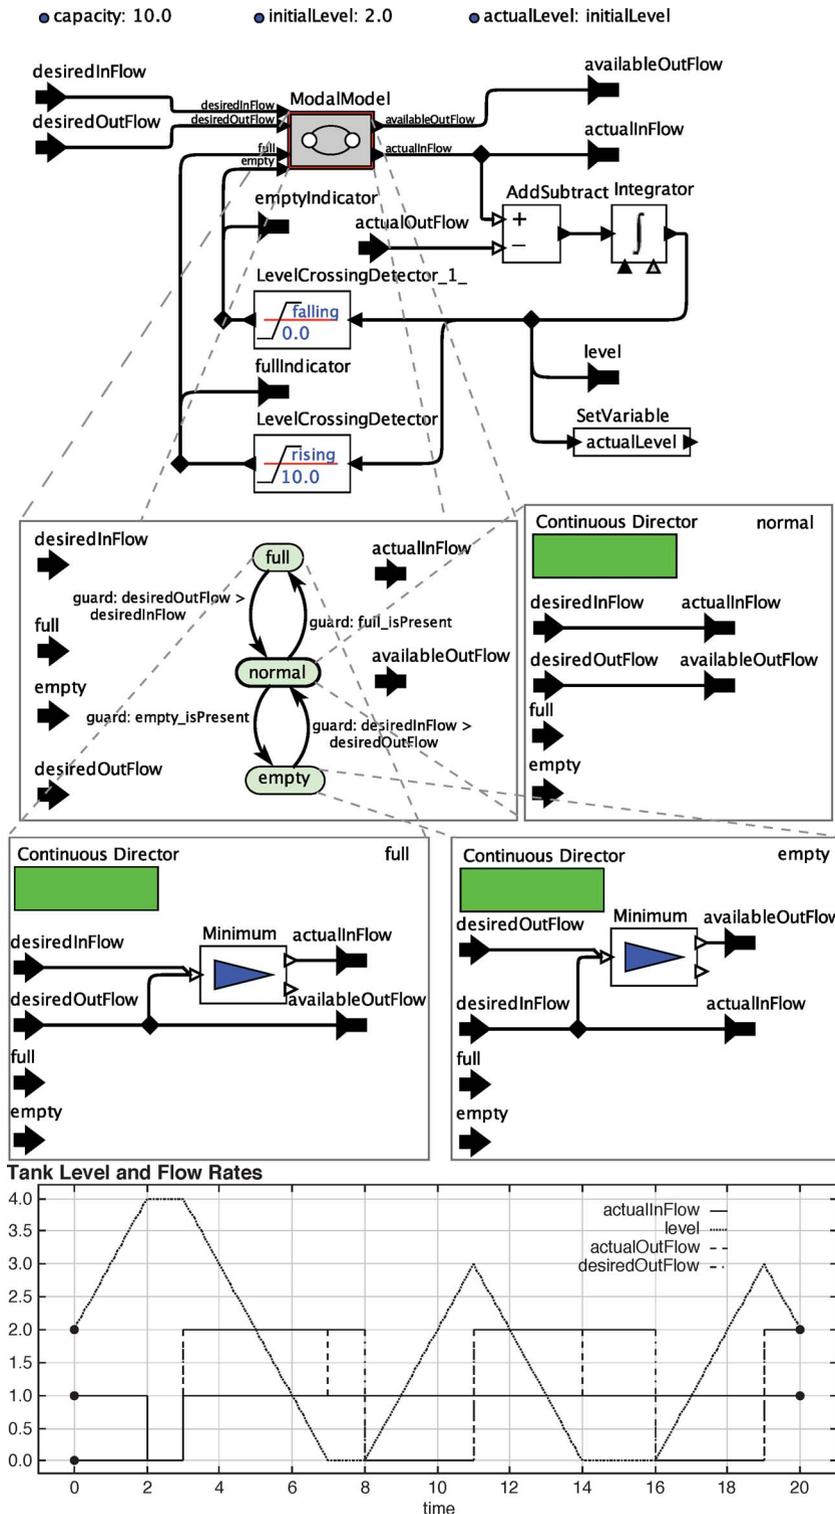


Fig. 5. ModalTank: Improved model of a fuel tank using modal models.

F. Modularity

Key to effective design of complex systems is modular design, where modules have well-defined interfaces, and composition of modules can be checked for compatibility.

A key challenge is to ensure that modules are correctly composed (Challenge 3 above). It is often easy to misinterpret data provided by one module in another module. Such misinterpretation can be catastrophic. Type systems

provide rudimentary protection, but much stronger protection is possible. Another key challenge is to enable reuse and customization of modules, as, for example, is supported by object-oriented concepts such as classes, inheritance, and polymorphism.

Like some other modeling frameworks (Modelica, SystemC, SysML, and others), Ptolemy II supports object oriented design principles in modeling, although the approach in Ptolemy II is unique [42]. In particular, we use classes to describe interface and behavior of an actor. Inheritance is used to describe relationships between objects. Instances of actors implement the behavior specified by the actor class definition. Subclasses inherit the class definitions and can be extended with additional objects as well as override parameters of the superclass. And updates to classes propagate to instances even during interactive model development.

In the fuel system models we have shown, the Tank component is defined as a class and instantiated in different parts of a model and in different models. Fig. 3 displays a model with two tanks. The class definition of the tank is maintained in a separate XML file, and Tank1 and Tank2 are instances of this class. Editing the class definition results in changes in all the instances. This strategy also facilitates maintenance of regression tests that can help assure that behavior remains reasonable as the tank model evolves. The regression tests are themselves models that use instances of the same class. This ensures that the model being tested and the model being developed remain consistent (Challenge 2).

Fig. 6 shows a possible UML class hierarchy for the Tank models we have implemented so far. The AbstractTank class models the interface and the functionality of computing the tank level. The SimpleTank class extends the AbstractTank with the functionality shown in Fig. 2. The ModalTank class represents the tank model in Fig. 5. If all models, including regression tests, are developed using these class definitions, then the tank model can

evolve to become more sophisticated, regression tests can check whether the evolution has changed expected behavior, and a multiplicity of instances of tank components across several models are assured to be consistent.

We have shown that object-oriented concepts such as classes, inheritance, and polymorphism can be adapted to concurrent, actor-oriented components [42], making them more suitable for CPS. But there are several more techniques that can be applied in modeling that enhance modularity. In particular, we have also developed advanced type systems for component compositions, enabling type inference and type checking across large models with polymorphic components [77]. We have also adapted such type systems to capture domain-specific ontology information, checking for correct usage and correct interpretation of shared data [51]. And we have shown how to check for compatibility of protocols in compositions [48] and to synthesize interface adapters for separately defined components [64].

We illustrate how a domain-specific ontology can address Challenge 3 discussed above using the techniques given in [51]. Fig. 7 shows at the upper right a simple ontology for dimensions used in the fuel system with the key concepts Level and Flow. The goal is to prevent model builders from accidentally misinterpreting a flow as a level or vice versa.

A model can be annotated with constraints that represent the model builder's domain knowledge. Such constraints are shown at the bottom of Fig. 7, where the parameters *tank2InFlow* and *tank1OutFlow* are associated with the concept *Flow*, and *capacity* is associated with the concept *Level*.

The system described in [51] can infer concept associations throughout a model from just a few annotations. To do this, actors may impose particular constraints. In the simple ontology shown in the figure, we have asserted that an Integrator actor converts a Flow to a Level. That is, if its output is a Level, then its input must be a Flow, and vice versa. Given this and similar constraints and the associations on the parameters, the ontology solver infers associations throughout the model. The result is shown as annotations close to the ports of the actors in the model, which are also color-coded according to the color scheme in the ontology. The component in the model labeled *FuelDimensionSystemSolver* handles executing the inference algorithm. With an appropriately designed ontology, errors like those in Fig. 4 will be reported by the solver.

Constraints on model components such as actors impose constraints on proper relationships. Checkers such as the type system infrastructure in Ptolemy can infer these properties and detect errors. By allowing users to specify ontologies graphically and describing constraints on model elements, the domain knowledge is built up. Ontological information about the overall system can be inferred automatically from the relatively small constraint set. Inconsistencies can be detected and reported automatically.

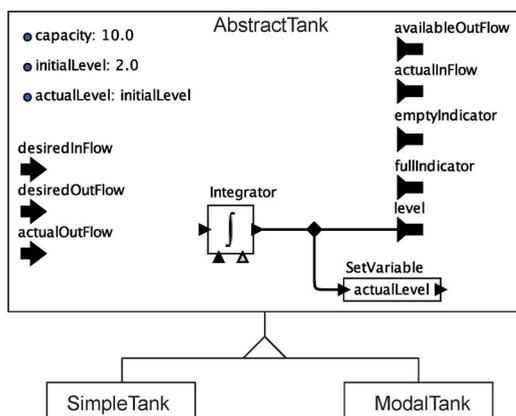


Fig. 6. UML diagram for abstract and concrete tank classes.

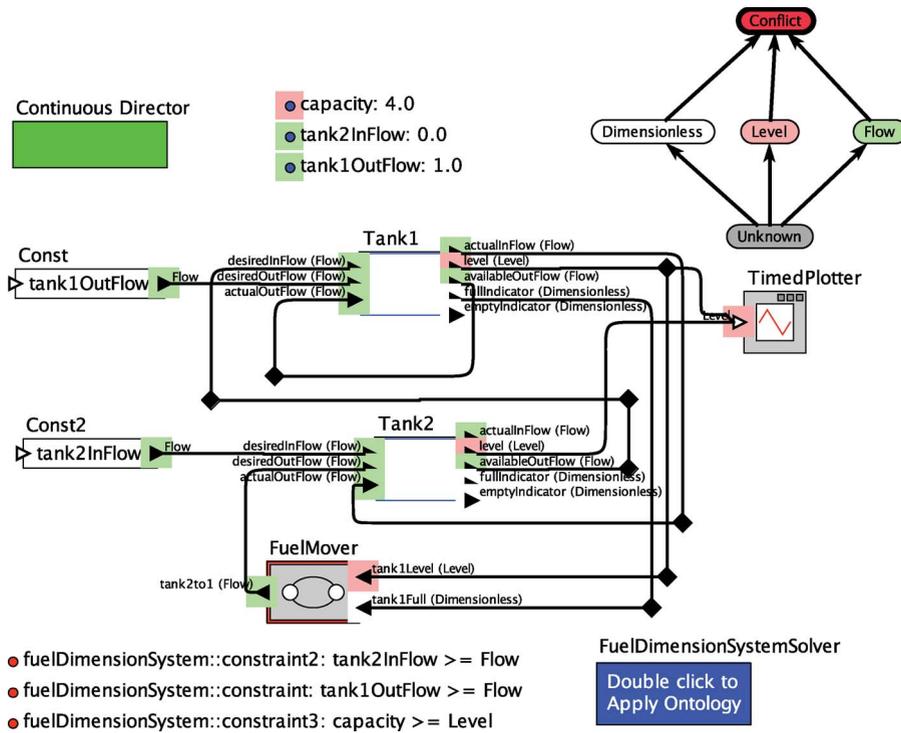


Fig. 7. Domain-specific ontology and solver ensure that flows are not accidentally misinterpreted as levels and vice versa.

### G. Linking Behavior to Implementation: Quantity Managers

To support evaluation of design choices, modeling frameworks need to enable weaving together a multiplicity of models that cover different aspects of a system (Challenge 4 above). For example, a choice of networking fabric will affect temporal behavior, power usage, and vulnerability to faults. A key challenge is that this requires bridges between models that are modeling very different aspects of the same system.

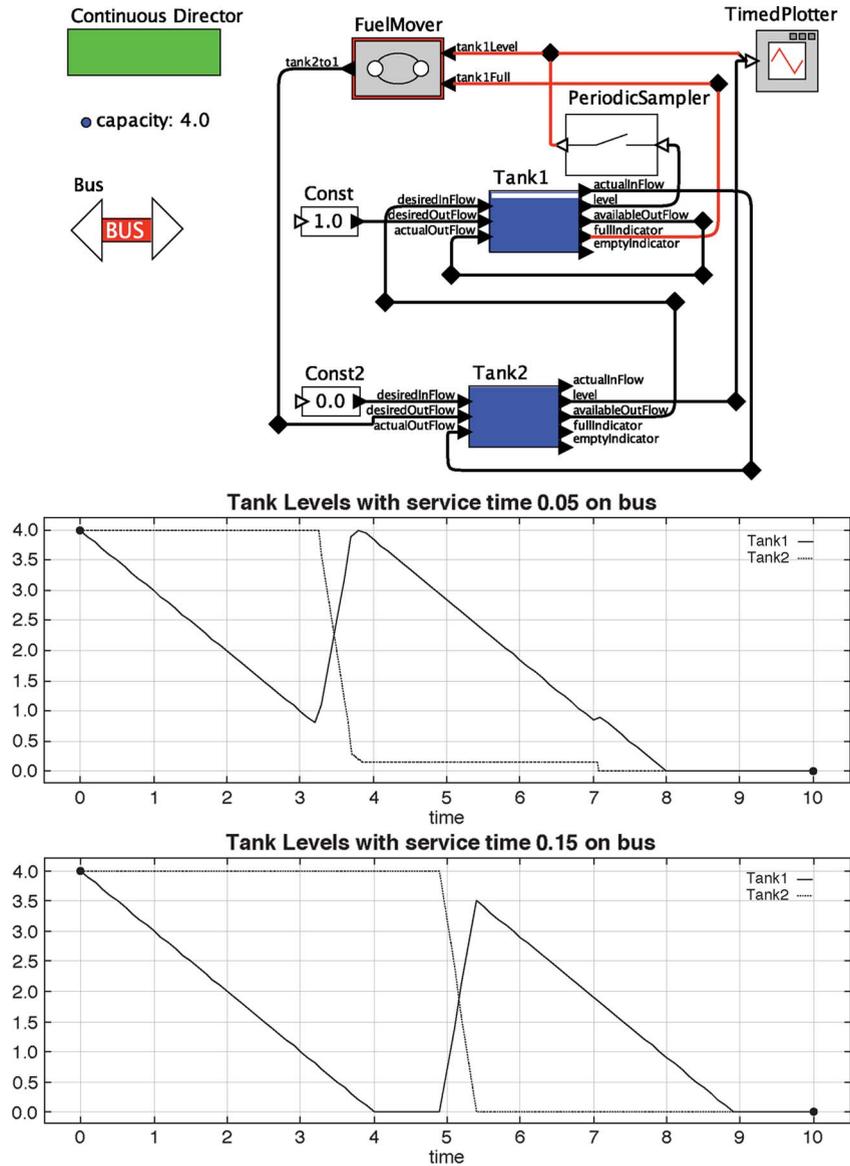
The Metropolis project [3], [15] introduced the notion of “quantity manager,” a component of a model that functions as a gateway to another model. For example, a purely functional model that describes only idealized behavioral properties of a flight control system could be endowed with a quantity manager that binds that functional model to a model of a distributed hardware architecture using a particular network fabric. By binding these two models, designers can evaluate how properties of the hardware implementation affect the functional behavior of the system. For example, how does a time-triggered bus protocol affect timing in a distributed control system, and how do the timing effects change the dynamics of the system? Similarly, a functional model could be endowed with a quantity manager that measures power and identifies potential overloads that may result from unexpectedly synchronized interactions across a distributed system.

Quantity managers bring a capability analogous to aspect-oriented programming in software engineering [34]. Separately designed models can be woven together using quantity managers in a manner similar to the weaving of separately designed classes in aspect-oriented design.

Fig. 8 shows a variant of the model of Fig. 3 that has been annotated with information about the implementation. In particular, the two inputs to the fuel mover component are implemented on a network bus, which introduces a communication delay. The plot at the bottom of Fig. 8 compares the fuel levels of the tanks in two cases. In the first case the service time of the bus is 0.05 and in the second case the service time is 0.15. The latter shows a clear degradation of the behavior. Between time 4 and time 5, Tank1 does not receive any fuel from Tank2 although it runs out of fuel at time 4. In a real system, this might cause the plane to crash or at least introduce severe system instability.

### H. Semantics-Preserving Transformation and Implementation

Effective use of models requires well-defined relationships between the models and systems being modeled. In many cases, models can be used as specifications, and implementations can be synthesized from these specifications. The key challenge is that such synthesis must preserve the semantics of the implementation.



**Fig. 8.** Fuel system model that includes some implementation detail, where the two signals into the FuelMover are transported over a single shared bus, and the transport takes time.

Code generation from models is a very active area of work. UML tools such as Rational Rose can generate code from executable models and templates from nonexecutable models. Code generators from actor modeling languages such as Simulink exist (TargetLink from dSPACE and RealTimeWorkshop from The MathWorks, for example). LabVIEW Embedded and LabVIEW FPGA are able to generate code and hardware designs, respectively, from dataflow models.

Despite considerable progress in this area, in our opinion, code generation from models is still in its infancy. The research investment so far is considerably smaller than the investment in compiler technology, despite the fact that the technical challenges are at least as great. And

many problems remain only partially solved, such as modular code generation [74] and model transformation [21].

#### IV. THE STATE OF THE ART

Despite considerable progress in languages, notations, and tools, major problems persist. In practice, system integration, adaptation of existing designs, and interoperation of heterogeneous subsystems remain major stumbling blocks that cause project failures. We believe that model-based design, as widely practiced today, largely fails to benefit from the principles of platform-based design [68] due to its lack of attention to the semantics of heterogeneous subsystem composition.

Consider, for example, UML 2 [5], [6], widely used for modeling software systems. Its derivative SysML [62] has many features that are potentially useful for CPS. The internal block diagram notation of SysML, which is based on the UML 2 composite structure diagrams, particularly with the use of flow ports, has great potential for modeling complex systems, such as aircraft fuel systems. But it has severe weaknesses that limit its ability to address system design problems. The SysML standard defines the syntax of these diagrams, not their semantics. Although the SysML standard asserts that “flow ports are intended to be used for asynchronous, broadcast, or send-and-forget interactions” [62], the standard fails to define the semantics of such interactions. Implementers of tools are free to interpret this intent, resulting in a modeling language whose semantics is defined by the tools rather than by the language itself. There are many semantic alternatives [41], consequently the same SysML diagram may be interpreted very differently by different observers.

Modeling and Analysis of Real-Time and Embedded (MARTE) system [61] also specifically avoids “constraining” (or even defining) the execution semantics of models. Instead, it focuses on providing alternative ways of describing today’s *ad hoc*, noncomposable design practices such as concurrency based on threads [39]. Standardizing notation is not sufficient to achieve effective analysis methods and unambiguous communication among designers. More importantly, without semantics, the modeling framework fails to provide a *platform* for design. Unfortunately, the flexibility of these modeling notations may account for some of their success, because designers can become “standards compliant” without changing their existing practice by just adapting their notation.

Further weakening their semantics, UML notations can be freely reinterpreted by defining a “profile,” greatly re-

ducing the value of the notation as an effective communication vehicle and design tool. We believe that *constraints that lead to well-defined and interoperable models have potentially far greater value*. More importantly, such constraints are essential for these modeling frameworks to become a central part of a platform-based engineering practice [69]. The challenge is to identify which constraints provide the most value while still admitting useful designs.

The inclusion by OMG of Statecharts [27] in the UML standard has helped to narrow the variability for some modeling problems, but in many cases, the exact semantics are determined by the implementation details of the supporting tools rather than by an agreed-upon semantics [20]. In fact, Statecharts also suffers from inadequate standardization. Despite their common origin, variants have proliferated [75]. Even the most widely used implementations of Statecharts that claim to be standards compliant have subtle semantic differences big enough “that a model written in one formalism could be ill-formed in another formalism” [13]. In many implementations, including the widely used RHAPSODY tool from IBM, the semantics is (probably inadvertently) nondeterminate [71].

## V. CONCLUSION

The intrinsic heterogeneity, concurrency, and sensitivity to timing of CPSs poses many modeling challenges. Much of the current work in modeling has insufficiently strong semantics to adequately address these problems. We have described some promising technologies that can help, including hybrid system modeling and simulation, concurrent and heterogeneous MoCs, the use of domain-specific ontologies to enhance modularity, and the joint modeling of functionality and implementation architectures. ■

## REFERENCES

- [1] H. Abelson and G. J. Sussman, *Structure and Interpretation of Computer Programs*, 2nd ed. Cambridge, MA: MIT Press, 1996.
- [2] F. Balarin, M. Chiodo, P. Giusto, H. Hsieh, A. Jurecska, L. Lavagno, C. Passerone, A. Sangiovanni-Vincentelli, E. Sentovich, K. Suzuki, and B. Tabbara, *Hardware-Software Co-Design of Embedded Systems—The POLIS Approach*. Norwell, MA: Kluwer, 1997.
- [3] F. Balarin, H. Hsieh, L. Lavagno, C. Passerone, A. L. Sangiovanni-Vincentelli, and Y. Watanabe, “Metropolis: An integrated electronic system design environment,” *Computer*, vol. 36, no. 4, pp. 45–52, 2003.
- [4] A. Benveniste, B. Caillaud, L. P. Carloni, and A. Sangiovanni-Vincentelli, “Tag machines,” in *Proc. Int. Conf. Embedded Softw.*, Jersey City, NJ, pp. 255–263, 2005.
- [5] C. Bock, “SysML and UML 2 support for activity modeling,” *Syst. Eng.*, vol. 9, no. 2, pp. 160–185, 2006.
- [6] G. Booch, I. Jacobson, and J. Rumbaugh, *The Unified Modeling Language User Guide*. Reading, MA: Addison-Wesley, 1998.
- [7] C. Brooks, A. Cataldo, C. Hylands, E. A. Lee, J. Liu, X. Liu, S. Neuendorffer, and H. Zheng, (2005, Jul. 15). HyVisual: A hybrid system visual modeler. Univ. California, Berkeley, Tech. Rep. UCB/ERL M05/24. [Online]. Available: <http://ptolemy.eecs.berkeley.edu/publications/papers/05/hyvisual/index.htm>
- [8] C. Brooks, C. Cheng, T. H. Feng, E. A. Lee, and R. von Hanxleden, “Model engineering using multimodeling,” in *Proc. Int. Workshop Model Co-Evolution Consistency Manage.*, Toulouse, France, pp. 114–129, 2008.
- [9] J. T. Burch, S. Ha, E. A. Lee, and D. G. Messerschmitt. (1994). Ptolemy: A framework for simulating and prototyping heterogeneous systems. *Int. J. Comput. Simul.* [Online]. 4, *Special Issue on Simulation Software Development*, pp. 155–182. Available: <http://ptolemy.eecs.berkeley.edu/publications/papers/94/JEurSim/>
- [10] J. R. Burch, R. Passerone, and A. L. Sangiovanni-Vincentelli, “Overcoming heterophobia: Modeling concurrency in heterogeneous systems,” in *Proc. Int. Conf. Appl. Concurrency Syst. Design*, 2001, p. 13.
- [11] L. P. Carloni, R. Passerone, A. Pinto, and A. Sangiovanni-Vincentelli, “Languages and tools for hybrid systems design,” *Found. Trends Electron. Design Autom.*, vol. 1, no. 1/2, pp. 1–204, 2006.
- [12] L. P. Carloni, K. L. McMillan, and A. L. Sangiovanni-Vincentelli, “The theory of latency insensitive design,” *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 20, no. 9, pp. 1059–1076, Sep. 2001.
- [13] M. L. Crane and J. Dingel, “Uml vs. classical vs. rhapsody statecharts: Not all models are created equal,” *Int. Conf. Model Driven Eng. Lang. Syst.*, vol. 3713. Berlin, Germany: Springer-Verlag, 2005, pp. 97–112, ser. Lecture Notes in Computer Science.
- [14] B. Crossette. (2001). Jet pilot who saved 304 finds heroism tainted. *New York Times*. [Online]. Available: <http://query.nytimes.com/gst/fullpage.html?res=9504EFD1738F933A2575AC0A9679C8B63>
- [15] A. Davare, D. Densmore, T. Meyerowitz, A. Pinto, A. Sangiovanni-Vincentelli, G. Yang, and Q. Zhu, “A next-generation design framework for platform-based design,” in *Proc. Design Verif. Conf.*, San Jose, CA, pp. 239–245, 2007.
- [16] S. Edwards, L. Lavagno, E. A. Lee, and A. Sangiovanni-Vincentelli, “Design of embedded systems: Formal models, validation, and synthesis,” *Proc. IEEE*, vol. 85, no. 3, pp. 366–390, Mar. 1997.

- [17] S. A. Edwards and E. A. Lee. (2003). The semantics and execution of a synchronous block-diagram language. *Sci. Comput. Programm.* [Online]. 48(1), pp. 21–42. Available: [http://dx.doi.org/10.1016/S0167-6423\(02\)00096-5](http://dx.doi.org/10.1016/S0167-6423(02)00096-5)
- [18] J. Eidson, E. A. Lee, S. Matic, S. A. Seshia, and J. Zou, “A time-centric model for cyber-physical applications,” in *Proc. Workshop Model Based Architecting Construction of Embedded Syst.*, 2010. [Online]. Available: <http://chess.eecs.berkeley.edu/pubs/791.html>
- [19] J. Eker, J. W. Janneck, E. A. Lee, J. Liu, X. Liu, J. Ludvig, S. Neuendorffer, S. Sachs, and Y. Xiong. (2003, Feb.). Taming heterogeneity—The Ptolemy approach. *Proc. IEEE* [Online]. 91(2), pp. 127–144. Available: <http://www.ptolemy.eecs.berkeley.edu/publications/papers/03/TamingHeterogeneity/>
- [20] H. Fecher, J. Schönborn, M. Kyas, and W. P. de Roever, “29 new unclaritys in the semantics of UML 2.0 state machines *International Conference on Formal Engineering Methods (ICFEM)*, vol. 3785. Berlin, Germany: Springer-Verlag, 2005, ser. Lecture Notes in Computer Science.
- [21] T. H. Feng and E. A. Lee, “Scalable models using model transformation,” in *Proc. Workshop Model Based Architect. Construction of Embedded Syst.*, 2008. [Online]. Available: <http://chess.eecs.berkeley.edu/pubs/487.html>
- [22] P. A. Fishwick and B. P. Zeigler, “A multimodel methodology for qualitative model engineering,” *ACM Trans. Model. Comput. Simul.*, vol. 2, no. 1, pp. 52–81, 1992.
- [23] P. Fritzson, *Principles of Object-Oriented Modeling and Simulation With Modelica 2.1*. New York: Wiley, 2003.
- [24] G. Goessler and A. Sangiovanni-Vincentelli, “Compositional modeling in Metropolis,” in *Proc. 2nd Int. Workshop Embedded Softw.*, Grenoble, France, pp. 7–9, 2002.
- [25] Z. Gu, S. Wang, S. Kodase, and K. G. Shin, “An end-to-end tool chain for multi-view modeling and analysis of avionics mission computing software,” in *Proc. Real-Time Syst. Symp.*, 2003, pp. 78–81.
- [26] C. Hardebolle and F. Boulanger, “ModHelX: A component-oriented approach to multi-formalism modeling,” in *Proc. Workshop Multi-Paradigm Model.*, Nashville, TN, 2007.
- [27] D. Harel, “Statecharts: A visual formalism for complex systems,” *Sci. Comput. Programm.*, vol. 8, no. 3, pp. 231–274, 1987.
- [28] F. Herrera and E. Villar, “A framework for embedded system specification under different models of computation in SystemC,” in *Proc. Design Autom. Conf.*, San Francisco, CA, pp. 911–914, 2006.
- [29] 1588: *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, Standard specification, IEEE Instrumentation and Measurement Society, Nov. 8, 2002.
- [30] A. Jantsch, *Modeling Embedded Systems and SoC—Concurrency and Time in Models of Computation*. San Mateo, CA: Morgan Kaufmann, 2003.
- [31] J. F. Jimenez, J. M. Giron-Sierra, C. Insaurralde, and M. Seminario, “A simulation of aircraft fuel management system,” *Simul. Model. Practice Theory*, vol. 15, pp. 544–564, 2007.
- [32] G. Karsai, A. Lang, and S. Neema. (2005). Design patterns for open tool integration. *Softw. Syst. Model.* [Online]. 4(2), pp. 157–170. Available: <http://dx.doi.org/10.1007/s10270-004-0073-y>
- [33] G. Karsai, J. Sztipanovits, A. Ledeczi, and T. Bapty, “Model-integrated development of embedded software,” *Proc. IEEE*, vol. 91, no. 1, pp. 145–164, Jan. 2003.
- [34] G. Kiczales, J. Lamping, A. Mendhekar, C. Maeda, C. V. Lopes, J. M. Loingtier, and J. Irwin, “Aspect-oriented programming *European Conference in Object-Oriented Programming*, vol. 1241. Berlin, Germany: Springer-Verlag, 1997, ser. Lecture Notes in Computer Science.
- [35] E. A. Lee. (1999). Modeling concurrent real-time processes using discrete events. *Ann. Softw. Eng.* [Online]. 7, pp. 25–45. Available: <http://dx.doi.org/10.1023/A:1018998524196>
- [36] E. A. Lee, “Model-driven development—From object-oriented design to actor-oriented design,” in *Proc. Workshop Softw. Eng. Embedded Syst., From Requirements to Implementation*, Chicago, IL, 2003. [Online]. Available: <http://ptolemy.eecs.berkeley.edu/publications/papers/03/MontereyWorkshopLee/>
- [37] E. A. Lee. (2003, Jul. 3). Overview of the Ptolemy project. Univ. California, Berkeley, CA, Tech. Rep. UCB/ERL M03/25. [Online]. Available: <http://ptolemy.eecs.berkeley.edu/publications/papers/03/overview/>
- [38] E. A. Lee, “Cyber-physical systems—Are computing foundations adequate?” in *Proc. Workshop Cyber-Physical Syst., Res. Motiv. Tech. Roadmap*, Austin, TX, 2006. [Online]. Available: <http://ptolemy.eecs.berkeley.edu/publications/papers/06/CPSPositionPaper/>
- [39] E. A. Lee. (2006). The problem with threads. *Computer* [Online]. 39(5), pp. 33–42. Available: <http://dx.doi.org/10.1109/MC.2006.180>
- [40] E. A. Lee, “Cyber physical systems: Design challenges,” in *Proc. Int. Symp. Object/Component/Service-Oriented Real-Time Distrib. Comput.*, Orlando, FL, 2008, pp. 363–369. [Online]. Available: <http://dx.doi.org/10.1109/ISORC.2008.25>
- [41] E. A. Lee, “Disciplined heterogeneous modeling,” in *Proc. Model Driven Eng. Lang. Syst.*, D. C. Petriu, N. Rouquette, and O. Haugen, Eds., 2010, pp. 273–287. [Online]. Available: <http://chess.eecs.berkeley.edu/pubs/679.html>
- [42] E. A. Lee, X. Liu, and S. Neuendorffer. (2009). Classes and inheritance in actor-oriented design. *ACM Trans. Embedded Comput. Syst.* [Online]. 8(4), pp. 29:1–29:26. Available: <http://ptolemy.eecs.berkeley.edu/publications/papers/07/classesandInheritance/index.htm>
- [43] E. A. Lee and E. Matsikoudis, “The semantics of dataflow with firing,” in *From Semantics to Computer Science: Essays in Memory of Gilles Kahn*, G. Huet, G. Plotkin, J.-J. Lévy, and Y. Bertot, Eds. Cambridge, CA: Cambridge Univ. Press, 2009. [Online]. Available: <http://ptolemy.eecs.berkeley.edu/publications/papers/08/DataflowWithFiring/>
- [44] E. A. Lee, S. Neuendorffer, and M. J. Wirthlin, “Actor-oriented design of embedded hardware and software systems,” *J. Circuits Syst. Comput.*, vol. 12, no. 3, pp. 231–260, 2003.
- [45] E. A. Lee and T. M. Parks, “Dataflow process networks,” *Proc. IEEE*, vol. 83, no. 5, pp. 773–801, May 1995, DOI: 10.1109/5.381846.
- [46] E. A. Lee and A. Sangiovanni-Vincentelli. (1998, Dec.). A framework for comparing models of computation. *IEEE Trans. Comput.-Aided Design Circuits Syst.* [Online]. 17(12), pp. 1217–1229. Available: <http://ptolemy.eecs.berkeley.edu/publications/papers/98/framework/>
- [47] E. A. Lee and S. Tripakis, “Modal models in Ptolemy,” in *Proc. 3rd Int. Workshop Equation-Based Object-Oriented Model. Lang. Tools*, Oslo, Norway, 2010, vol. 47, pp. 11–21. [Online]. Available: <http://chess.eecs.berkeley.edu/pubs/700.html>
- [48] E. A. Lee and Y. Xiong, “A behavioral type system and its application in Ptolemy II,” *Formal Aspects Comput. J.*, vol. 16, no. 3, pp. 210–237, 2004.
- [49] E. A. Lee and H. Zheng, “Operational semantics of hybrid systems,” in *Hybrid Systems: Computation and Control (HSCC)*, vol. 3414, M. Morari and L. Thiele, Eds. Zurich, Switzerland: Springer-Verlag, 2005, pp. 25–53. [Online]. Available: [http://dx.doi.org/10.1007/978-3-540-31954-2\\_2](http://dx.doi.org/10.1007/978-3-540-31954-2_2)
- [50] E. A. Lee and H. Zheng, “Leveraging synchronous language principles for heterogeneous modeling and design of embedded systems,” in *Proc. Int. Conf. Embedded Softw.*, Salzburg, Austria, 2007, pp. 114–123. [Online]. Available: <http://dx.doi.org/10.1145/1289927.1289949>
- [51] M. K. Leung, T. Mandl, E. A. Lee, E. Latronico, C. Shelton, S. Tripakis, and B. Lickly, “Scalable semantic annotation using lattice-based ontologies,” in *Proc. Int. Conf. Model Driven Eng. Lang. Syst.*, Denver, CO, pp. 393–407, 2009.
- [52] J. Liu, “Responsible frameworks for heterogeneous modeling and design of embedded systems,” Ph.D. dissertation, Electr. Eng. Comput. Sci. Dept., Univ. California, Berkeley, CA, Dec. 20, 2001. [Online]. Available: <http://ptolemy.eecs.berkeley.edu/publications/papers/01/responsibleFrameworks/>
- [53] J. Liu, B. Wu, X. Liu, and E. A. Lee, “Interoperation of heterogeneous CAD tools in Ptolemy II,” in *Proc. Symp. Design Test Microfabricat. MEMS/MOEMS*, Paris, France, pp. 249–258, 1999.
- [54] X. Liu and E. A. Lee. (2008). CPO semantics of timed interactive actor networks. *Theor. Comput. Sci.* [Online]. 409(1), pp. 110–125. Available: <http://dx.doi.org/10.1016/j.tcs.2008.08.044>
- [55] X. Liu, J. Liu, J. Eker, and E. A. Lee, “Heterogeneous modeling and design of control systems,” in *Software-Enabled Control: Information Technology for Dynamical Systems*, T. Samad and G. Balas, Eds. New York: Wiley-IEEE Press, 2003.
- [56] O. Maler, Z. Manna, and A. Pnueli, “From timed to hybrid systems,” in *Proc. Real-Time, Theory Practice REX Workshop*, 1992, pp. 447–484.
- [57] I. Moir and A. Seabridge, *Aircraft Systems: Mechanical, Electrical, and Avionics Subsystems Integration*, 3rd ed. New York: Wiley, 2008.
- [58] P. J. Mosterman and H. Vangheluwe, “Computer automated multi-paradigm modeling: An introduction,” *Simulation, Trans. Soc. Model. Simul./Int. J. High Performance Comput. Appl.*, vol. 80, no. 9, pp. 433–450, 2004.
- [59] G. Nordstrom, J. Sztipanovits, G. Karsai, and A. Ledeczi, “Metamodeling—Rapid design and evolution of domain-specific modeling environments,” in *Proc. Conf. Eng. Comput.-Based Syst.*, Nashville, TN, 1999, pp. 68–74.
- [60] Object Management Group (OMG). (2008). System modeling language specification v1.1, OMG, Tech. Rep. [Online]. Available: <http://www.sysmlforum.com>

- [61] Object Management Group (OMG). (2008, Aug.). A UML profile for MARTE, beta 2, OMG adopted specification ptc/08-06-09, OMG, Tech. Rep. [Online]. Available: <http://www.omg.org/omgmarte/>
- [62] Object Management Group (OMG). (2010, Jun.). System modeling language specification v1.2, standard specification, OMG, Tech. Rep. [Online]. Available: <http://www.sysmlforum.com>
- [63] T. M. Parks and D. Roberts, "Distributed process networks in Java," in *Proc. Int. Parallel Distrib. Process. Symp.*, Nice, France, pp. 138–146, 2003.
- [64] R. Passerone, L. de Alfaro, T. A. Henzinger, and A. Sangiovanni-Vincentelli, "Convertibility verification and converter synthesis: Two faces of the same coin," in *Proc. Int. Conf. Comput. Aided Design*, San Jose, CA, pp. 132–139, 2002.
- [65] H. D. Patel and S. K. Shukla, *SystemC Kernel Extensions for Heterogeneous System Modelling*. Norwell, MA: Kluwer, 2004.
- [66] W. Pree and J. Templ, "Modeling with the timing definition language (TDL)," *Automotive Software Workshop San Diego (ASWSD) on Model-Driven Development of Reliable Automotive Services*. Berlin, Germany: Springer-Verlag, 2006, ser. Lecture Notes in Computer Science.
- [67] I. Sander and A. Jantsch, "System modeling and transformational design refinement in ForSyDe," *IEEE Trans. Comput.-Aided Design Circuits Syst.*, vol. 23, no. 1, pp. 17–32, Jan. 2004.
- [68] A. Sangiovanni-Vincentelli, "Defining platform-based design," *EEDesign of EETimes*, 2002. [Online]. Available: <http://www.eetimes.com/electronics-news/4141729/Defining-platform-based-design>
- [69] A. Sangiovanni-Vincentelli, "Quo vadis, SLD? Reasoning about the trends and challenges of system level design," *Proc. IEEE*, vol. 95, no. 3, pp. 467–506, Mar. 2007.
- [70] A. Sangiovanni-Vincentelli, G. Yang, S. K. Shukla, D. A. Mathaikutty, and J. Sztipanovits, "Metamodeling: An emerging representation paradigm for system-level design," *IEEE Design Test Comput.*, vol. 26, no. 3, pp. 54–69, May-Jun. 2009.
- [71] W. Schamai, U. Pohlmann, P. Fritzon, C. J. J. Paredis, P. Helle, and C. Strobel, "Execution of umlstate machines using modelica," in *Proc. 3rd Int. Workshop Equation-Based Object-Oriented Model. Lang. Tools*, Oslo, Norway, 2010, vol. 47, pp. 1–10. [Online]. Available: <http://www.ep.liu.se/ecp/047/>
- [72] B. Selic, "The pragmatics of model-driven development," *IEEE Softw.*, vol. 20, no. 5, pp. 19–25, Sep./Oct. 2003.
- [73] J. Sztipanovits and G. Karsai, "Model-integrated computing," *IEEE Computer*, vol. 34, no. 4, pp. 110–111, Apr. 1997.
- [74] S. Tripakis, D. Bui, M. Geilen, B. Rodiers, and E. A. Lee, "Compositionality in synchronous data flow: Modular code generation from hierarchical SDF graphs," *Electr. Eng. Comput. Sci. Dept., Univ. California, Berkeley, CA, Tech. Rep. UCB/EECS-2010-52*, May 7, 2010.
- [75] M. von der Beeck, "A comparison of statecharts variants *Third International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems*, vol. 863, H. Langmaack, W. P. de Roever, and J. Vytopil, Eds. Berlin, Germany: Springer-Verlag, 1994, pp. 128–148, ser. Lecture Notes in Computer Science.
- [76] Wikipedia, *Air Transat Flight 236*. [Online]. Available: [http://en.wikipedia.org/wiki/Air\\_Transat\\_Flight\\_236](http://en.wikipedia.org/wiki/Air_Transat_Flight_236)
- [77] Y. Xiong, E. A. Lee, X. Liu, Y. Zhao, and L. C. Zhong, "The design and application of structured types in Ptolemy II," in *Proc. IEEE Int. Conf. Granular Comput.*, Beijing, China, 2005, vol. 2, pp. 683–688.
- [78] Y. Zhao, E. A. Lee, and J. Liu, "A programming model for time-synchronized distributed real-time systems," in *Proc. Real-Time Embedded Technol. Appl. Symp.*, Bellevue, WA, 2007, pp. 259–268. [Online]. Available: <http://ptolemy.eecs.berkeley.edu/publications/papers/07/RTAS/>

## ABOUT THE AUTHORS

**Patricia Derler** (Member, IEEE) did her undergraduate studies at the University of Hagenberg, Austria and received the Ph.D. degree in computer science from the University of Salzburg, Salzburg, Austria, in 2010.

She is a Postdoctoral Researcher at the University of California Berkeley, Berkeley. Her research interests are in design and simulation of cyber-physical systems, deterministic models of computation, and the use of predictability in software, hardware, and the environment towards efficient simulations and executions.



**Edward A. Lee** (Fellow, IEEE) received the B.S. degree in computer science from Yale University, New Haven, CT, in 1979, the S.M. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 1981, and the Ph.D. degree in electrical engineering and computer science from the University of California Berkeley, Berkeley, in 1986.

He is the Robert S. Pepper Distinguished Professor and former chair of the Electrical Engineering and Computer Sciences (EECS) Department, University of California Berkeley. His research interests center on design, modeling, and analysis of embedded, real-time computational systems. He is a Director of Chess, the Berkeley Center for Hybrid and Embedded Software Systems, and is the Director of the Berkeley Ptolemy project. He is coauthor of six books and numerous papers. He has led the development of several influential open-source software packages, notably Ptolemy and its various spinoffs. From 1979 to 1982, he was a member of technical staff at Bell Telephone Laboratories, Holmdel, NJ, in the Advanced Data Communications Laboratory. He is a cofounder of BDTI, Inc., where he is currently a Senior Technical Advisor, and has consulted for a number of other companies.



Dr. Lee was a National Science Foundation (NSF) Presidential Young Investigator, and won the 1997 Frederick Emmons Terman Award for Engineering Education.

**Alberto Sangiovanni Vincentelli** (Fellow, IEEE) received an electrical engineering and computer science degree ("Dottore in Ingegneria") *summa cum laude* from the Politecnico di Milano, Milano, Italy, in 1971.

He holds the Edgar L. and Harold H. Buttner Chair of Electrical Engineering and Computer Sciences at the University of California Berkeley, Berkeley. He was a cofounder of Cadence and Synopsys, the two leading companies in the area of electronic design automation. He is a member of the Board of Directors of Cadence, Sonics, and Accent. He was a member of the HP Strategic Technology Advisory Board, and is a member of the Science and Technology Advisory Board of General Motors and of the Technology Advisory Council of UTC. He is a member of the High-Level Group, of the Steering Committee, of the Governing Board and of the Public Authorities Board of the EU Artemis Joint Technology Undertaking. He is an author of 880 papers, 15 books, and 3 patents.



Dr. Sangiovanni Vincentelli is member of the Scientific Council of the Italian National Science Foundation (CNR) and a member of the Executive Committee of Italian Institute of Technology. He received, among others, the Kaufman Award of the Electronic Design Automation Council for pioneering contributions to Electronic Design Automation and the IEEE/Royal Society of Edinburgh James Clerk Maxwell Medal for groundbreaking contributions that have had an exceptional impact on the development of electronics and electrical engineering or related fields. He holds an Honorary Doctorate from University of Aalborg, Denmark. He is a Member of the National Academy of Engineering.