

A Co-Simulation Framework for Design of Time-Triggered Automotive Cyber Physical Systems

Zhenkai Zhang, Emeka Eyisi, Xenofon Koutsoukos,
Joseph Porter, Gabor Karsai, Janos Sztipanovits

Institute for Software Integrated Systems (ISIS)

Department of Electrical Engineering and Computer Science

Vanderbilt University

Nashville, TN, USA

Abstract

Designing cyber-physical systems (CPS) is challenging due to the tight interactions between software, network/platform, and physical components. A co-simulation method is valuable to enable early system evaluation. Automotive control system is a typical CPS example and often designed by using time-triggered paradigm. In this paper, a co-simulation framework that considers interacting CPS components for assisting time-triggered automotive CPS design is proposed. Virtual prototyping of automotive vehicle is the core of this framework, which uses SystemC to model the cyber side of the system and integrates CarSim to model the physical side. A network/platform model in SystemC forms the backbone of the virtual prototyping, which bridges automotive control software and physical environment. The network/platform model consists of processing elements abstracted by real-time operating systems, communication systems, sensors, and actuators. The framework is also integrated with a model-based design tool to enable rapid prototyping. The framework is validated by comparing simulation results with the results from a hardware-in-the-loop automotive simulator. According to different design options, design space exploration (DSE) is also demonstrated.

Keywords: Co-Simulation, Virtual prototyping, Model-Based Design, CPS, Automotive Control System, SystemC

1. Introduction

Cyber-physical systems (CPS) are complex systems that are characterized by the tight interactions between the physical dynamics, computational platforms, communication networks, and control software. When designing CPS, a practical approach is to consider three design layers, which include the physical layer, the network/platform layer, and the software layer, as shown in Fig. 1 [26]. The physical layer represents physical components and their interactions, whose behavior is governed by physical laws and is typically described in continuous time using ordinary differential equations. The network/platform layer represents the hardware side of CPS and includes the network architecture and computation platform that interact with the physical components through sensors and actuators.

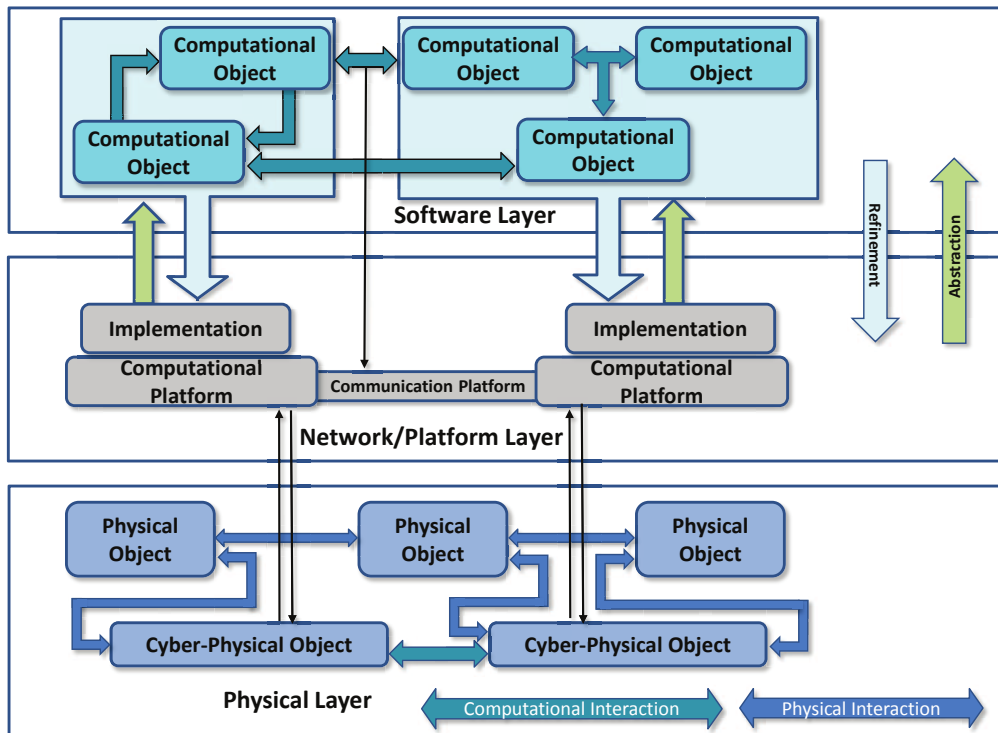


Figure 1: A Simplified View of Designing CPS: Three CPS Design Layers [26]

As a classical CPS domain, automotive systems have been gaining a lot of attention. As automotive system functionalities are increasingly imple-

mented by electronic instead of hydraulic and mechanical systems, up to 70 electronic control units (ECUs) exchanging more than 2500 signals over up to 5 different communication systems can be found in a modern vehicle [9]. The complex cyber-physical interactions make the composability and predictability of these typical safety-critical systems very challenging. Furthermore, the economy factors, such as persistent effort for low production costs and tight time-to-market, further complicate the design of such systems.

Time-triggered architecture (TTA) has been proposed and widely used to address the complexity and composability difficulties posed by automotive control systems by precisely defining the interfaces between components both in the time and value domain in order to provide predictability [12]. In addition, there have been on-going efforts towards the standardization of in-vehicle communication systems based on time-triggered (TT) paradigm (e.g. FlexRay and TTEthernet) with the overall goal of ensuring highly reliable, deterministic, and fault-tolerant system performance [20] [17].

The three-layer CPS design approach can be easily applied to TT automotive CPS design. We often start designing the automotive control system using a high level modeling language such as MATLAB/Simulink. The model serves as an executable specification and the equivalent source code, usually in C, can be generated automatically from the model. At later design stages, the generated source code is deployed on a designed automotive vehicle platform to perform the required functionality. It may not be possible to achieve the required control performance if elements from the three CPS design layers are designed separately and integrated in the end. Interactions between the layers are very tight, so late integration is very likely to result in large design gaps that will be costly to resolve. Moreover, different design options (e.g. processors, communication systems, and software deployments) may need to be explored in order to find trade-offs between performance and economy factors. In order to reduce the efforts and costs and shorten time-to-market, it is important to enable design space exploration (DSE) and get realistic control performance feedback at early design stages. However, the vehicle platform prototype is usually not available at early design stages and even if it is available, testing at the very beginning presents safety and economical challenges.

A cross-layer co-simulation framework that takes into account physical dynamics, control software, computational platforms, and communication networks becomes crucial in the design of automotive CPS. The requirements for such a framework include: (1) it should contain models from three

CPS design layers that can be integrated together; (2) the models should be at appropriate levels of abstraction, so that the simulation is efficient but accurate enough; (3) the scalability of the framework should allow simulation of large distributed automotive CPS; (4) it should allow model-based rapid prototyping to improve the usability.

Co-simulation can be achieved by virtual prototyping. Virtual prototyping can take advantage of different modeling languages/tools and integrate them together to evaluate the whole CPS. Modeling cyber components in SystemC has begun to be dominant in the Electronic System-Level (ESL) design field. SystemC has become a *de facto* system-level design language for hardware/software (HW/SW) co-design and an IEEE standard [10]. SystemC allows modeling at different levels of abstraction. By adding appropriate timing annotations, a SystemC model can reveal timing behavior of the corresponding HW/SW. SystemC also has a standardized library for realization of transaction level modeling (TLM) concepts. TLM focuses on what data is being transferred rather than how it is being transmitted, so a TLM model abstracts away certain communication details to speed up simulation while keeping sufficient accuracy.

The main contributions of the paper include: (1) A co-simulation framework for design of time-triggered automotive CPS that centers on a detailed network/platform layer model in SystemC is proposed. The network/platform layer model, including processing elements (PEs) which are abstracted by real-time operating system (RTOS) models, TTEthernet communication systems, sensors, and actuators, enables TT computation and communication; (2) Rapid prototyping is realized by model transformations from a designed MATLAB/Simulink model to a front-end design environment model to the final virtual prototype. It enables fast generation of executable simulation models with different configurations, including hardware platform, deployment, and timing configurations, so that it makes design space exploration easier; (3) TT communication in terms of its end-to-end latency and jitter is validated against a real implementation of TTEthernet and a TTEthernet model in OMNeT++ INET framework, and the framework is evaluated by using two automotive control system case studies (one is adaptive cruise controller (ACC) and the other one is lane keeping controller with ACC (LKC + ACC), which demonstrates the efficiency and accuracy of the approach. (4) Design space exploration is carried out in terms of using different sampling periods, introducing different additional network traffic, and using different clock synchronization strategies.

In [31], we have introduced a co-simulation framework that is used to facilitate TT CPS design. This paper mainly focuses on TT automotive CPS, which has more detailed virtual prototyping and validation and evaluation results of the TTEthernet model. Moreover, a new automotive case study is also used to show the framework has a good ability to deal with incremental design.

The rest of paper is organized as follows: Section 2 describes the related work; Section 3 introduces the core of the framework, which is the virtual prototyping of TT automotive CPS; Section 4 describes how to achieve rapid prototyping via a model-based design environment; Section 5 uses two automotive control case studies to validate and evaluate the framework; Section 6 illustrates design space exploration; Section 7 concludes this work.

2. Related Work

Co-simulation of CPS requires integrating different models of computation (MoCs). In [7], an operational behavioral semantics integrating discrete event MoC and continuous time MoC is proposed and illustrated by combining SystemC and MATLAB/Simulink. In [28], a similar behavioral semantics is proposed and demonstrated by integrating VDM++ and 20-sim. These papers present formal co-simulation frameworks, but they are not directly applicable to TT CPS design.

In [16], a methodology of virtual prototyping of CPS is proposed which combines SystemC, QEMU, and Open Dynamics Engine to achieve a holistic design view. In [11], a co-simulation environment based on a SH-2A CPU model is demonstrated by combining different design tools including CoMET, Saber, and MATLAB/Simulink. In [19], a co-simulation tool based on SystemC/SCNSL and MATLAB/Simulink is illustrated to facilitate the design of networked control systems. Again, these methods cannot be used for TT CPS design directly, and the approach in [11] does not support simulation of distributed CPS.

The TrueTime toolbox has been proposed and used in MATLAB/Simulink environment to enable CPS simulation [4]. The toolbox considers timing aspects introduced by computation and communication. However, it is difficult to integrate hardware models and support different abstraction levels. Further, preemption can only happen at points between segments which causes timing inaccuracy (e.g. interrupt handling), and the clock synchronization between computation and communication on a node is also implicit.

Utilizing SystemC to help automotive control system design has been investigated in some work. In [13], using SystemC to help simulate and refine automotive software specified by AUTOSAR is presented to deal with the problem which is timing simulation is not supported by AUTOSAR. Other work that introduces virtual prototyping in SystemC to co-simulate the automotive control systems is presented in [25] [30]. However, these approaches only consider the cyber part of the system and do not include a physical dynamics model.

Co-simulation of holistic automotive control systems can date back to late 90's. In [15], a C-VHDL-MATLAB co-simulation approach for automotive control systems is proposed to deal with the joint design of software in C, hardware in VHDL, and mechanical components in MATLAB. However, the framework does not support simulation of distributed automotive CPS, and the timed simulation efficiency is not acceptable. Hardware-in-the-loop (HIL) automotive simulators are also found in [5] and [6]. Compared to software-based simulation frameworks, they are more expensive and usually not available at early design stages. Besides, their network/platform layers are often fixed which limit the initial development. Recently, a co-simulation framework based on a commercial tool called SyAD is mentioned in [1] with the same aim of our work. Their framework uses FlexRay as its communication system, in contrast to which, we use TTEthernet. More importantly, due to using the commercial tool the internal relationship between different components is not described and the technical details have not been unveiled. Our work focuses on TT automotive CPS, and further we integrate the co-simulation framework with a model-based design tool for improving usability.

3. Virtual Prototyping of Automotive CPS

The core of the co-simulation framework is the virtual prototyping of automotive CPS, which is achieved by modeling each CPS design layer and exposing interfaces for integration. A commercial tool called CarSim is used for modeling the physical layer, and SystemC/TLM is applied for modeling the cyber parts including the software layer and the network/platform layer, since it is capable of modeling the cyber system at different abstraction levels and at the same time achieving flexibility. Fig. 2 shows the virtual prototyping architecture of the automotive CPS and the interactions between three design layers. In Fig. 2, there are two nodes (ECUs) connected by a

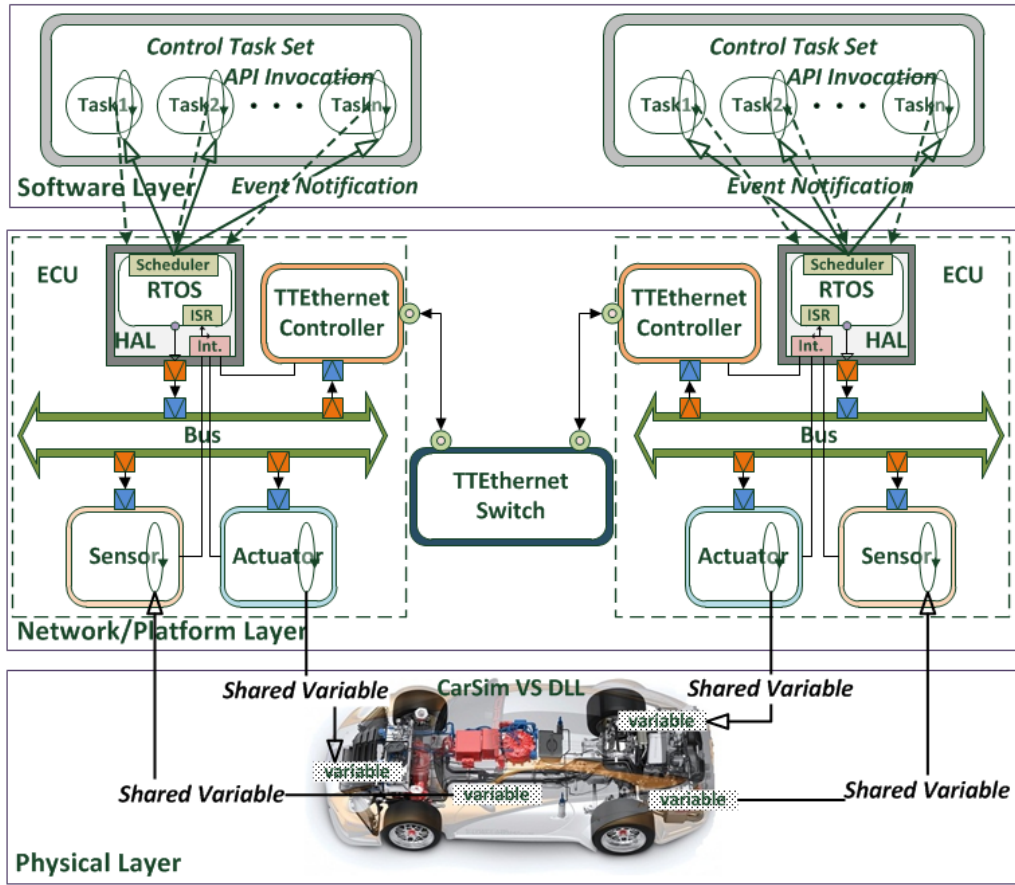


Figure 2: Virtual Prototyping of Automotive Control System by Three CPS Design Layers

TTEthernet switch forming the network/platform layer. In each node, the processing element (PE), TTEthernet controller, sensors, and actuators are connected through a bus.

3.1. Network/Platform Layer

As the backbone of the virtual prototyping of the automotive CPS, the network/platform layer bridges the software layer and the physical layer. The network/platform layer includes the network architecture and hardware platforms that interact with the physical components through sensors and actuators. While executing the software components on processors and transferring data on communication links, their abstract behavior is “translated”

into physical behavior.

The behavior of this layer is captured by several models in SystemC: (1) a PE model for TT computation, (2) a clock model for driving TT operations, (3) a network model compliant with the TTEthernet protocol for TT communication between different nodes, and (4) sensor and actuator models for interaction with the CarSim model. There are various network communication systems that can be used in the TT automotive CPS design, such as TTCAN, FlexRay, and TTEthernet. In this paper, we choose TTEthernet to illustrate the framework, since it has been deployed in many CPS domains, such as automotive, aerospace, and industrial process control.

3.1.1. Processing Element Modeling

Although an instruction set simulator can accurately mimic the behavior of a program running on a specific processor so as to give cycle accurate execution results, many drawbacks impede its use during early CPS design stages, including its low simulation speed for multi-processor simulation and the need to have the final target binary available. In order to accelerate the simulation while preserving accuracy, modeling the PE at higher abstractions levels is needed. An abstract RTOS model with accurate interrupt handling can serve as an efficient and effective model of the PE [29].

The abstract RTOS model in SystemC provides basic services to the software layer, which include task management, scheduling, interrupt handling, and inter-node communication, etc. It also exposes a set of primitive APIs to the software layer to facilitate the use of the model.

TT computation: In this framework, we interpret TT computation as follows: TT tasks are activated by the TT activator of the RTOS at the predefined times and put into the ready queue for scheduling. A TT task can be preempted and put back to the ready queue again, but it should not be blocked on any events (Fig. 3 shows the TT task state transitions). This mechanism can allow a more urgent system service program, such as an interrupt service routine (ISR), to preempt the execution of a TT task, and also allow the design of mixed time-/event-triggered systems. The TT static schedule is computed off-line which also considers the worst case preemption time.

Scheduling: The scheduler is the heart of the abstract RTOS model, whose behavior depends on a specific scheduling algorithm. The scheduling algorithm of the RTOS can use rate monotonic, earliest deadline first, or other real-time scheduling algorithms to schedule the ready queue of the RTOS

model. The ready queue consists of TT tasks and ISRs. As stated in the previous paragraph, there is a TT activator that statically activates the tasks according to an *a priori* schedule table generated by an off-line scheduling tool. The timing properties of the scheduler have two parameters which are scheduling overhead and context switching overhead. These timing properties are RTOS- and hardware platform-specific. Since it is not the focus of this paper, we assume the parameters are already available beforehand.




Figure 3: TT Task State Transitions

Interrupt handling: SystemC has some challenges for RTOS modeling, which can be summarized as non-interruptible *wait-for-delay* time advance and non-preemptive simulation processes. When an interrupt happens, it requires the real-time system to react and handle it in a timely manner. Modeling an accurate preemption mechanism plays an important role in accurate PE modeling. There are several ways to solve this problem: prediction and stepwise method [8]; result oriented modeling method [22]; and *wait-for-event* method [29]. We adopt the method from [29] which makes task use *wait-for-event* other than *wait-for-delay* to advance its execution time (as shown in Fig. 4). A system call of the RTOS model taking execution time as its argument makes the task wait on a *sc_event* object which will be notified after the given execution time elapses if no preemption happens. When an interrupt happens and its corresponding ISR preempts the execution of the task, the notification of the *sc_event* object will be canceled and a new notification time will be calculated according to how much time the preemption took and how much execution time already passed.

Inter-task communication: Within a PE, the communication between TT tasks is through shared memory, since it can be accessed without race-condition. The communication between tasks running on different PEs is

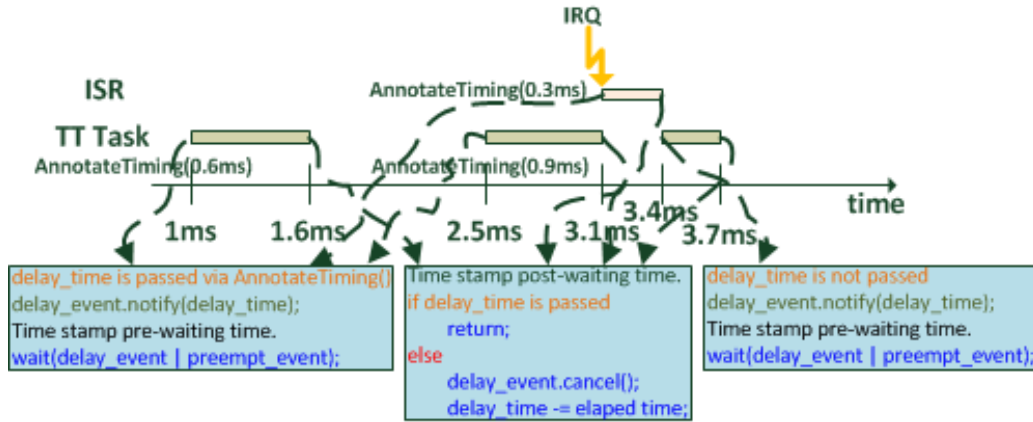


Figure 4: Advancing Execution Time by Using *wait-for-event*

achieved by invoking send/receive APIs of the underlying abstract RTOS model. The corresponding messages are delivered by the underlying TT communication system. State messages are used to prevent the TT tasks from blocking on reading.

PE Integration: In order to integrate the abstract RTOS model as a PE with other models on the network/platform layer through a bus, an additional Hardware Abstraction Layer (HAL) model is added to wrap the abstract RTOS model (as shown in Fig. 5). The HAL model has a multi-port *sc_port* object to collect all the interrupt requests (IRQs) from peripherals in the node, and it is also a hierarchical SystemC channel which implements the pure virtual functions of a HAL interface class. The abstract RTOS model is connected to the HAL model through a *sc_port* object parameterized with the HAL interface class. When the abstract RTOS model communicates with other models, it will send/receive the data via the port by invoking the functions implemented by the HAL model, and the HAL model will initiate a bus transaction.

Clock synchronization: The clock of a PE can be synchronized with the communication controller or be independent according to the configuration. As discussed in [14], if the clock is synchronized with the TTEthernet controller, all the operations are based on a global time base, and the control delay, δ , only depends on the offset and execution time of the actuation in a control period without variation. If the PE model and the network model do not share a global time base, the control delay will have a large variation




Figure 5: Hardware Abstraction Layer for PE Integration

which will be the sum of the periods of the computation and communication.

3.1.2. Clock Modeling

In TTA, time is the driving force for all TT operations. A TT communication system has its own synchronized global time base for correct operation. Computation can be synchronized with the communication or it can be driven by its own independent clock. Since time is the most important notion in TTA, modeling the independent clock and its synchronization service in SystemC becomes necessary. A simple clock can be modeled by using two parameters: a frequency and a drift. However, SystemC uses a discrete event simulation kernel which maintains global simulation time. If we simulate every tick of a clock with a drift, the simulation overhead will be too large, which can seriously slow down the simulation. Instead, we model the clock as follows: A random ppm value is assign to each clock in the interval $[-\text{MAX_PPM}, -\text{MIN_PPM}] \cup [\text{MIN_PPM}, \text{MAX_PPM}]$ (`MAX_PPM` and `MIN_PPM` are set by the user). According to the time-triggered schedule, the duration in clock ticks from the current time to the time when the next time-triggered action needs to take place is calculated. After that, we can get the duration in simulation time by taking into account its clock drift: $\text{duration in simulation time} = \text{duration in clock ticks} \times (\text{tick time} + \text{drift})$, and then we can arrange a clock event with this amount of time by using the notification mechanism of `sc_event` in SystemC.

Because the clock will be adjusted periodically by the synchronization service, the arranged clock event will be affected (its occurrence in simulation time becomes sooner or later). In order to simulate this properly, the arranged clock event and its occurring time in clock ticks is stored in a linked list in order of occurrence. When a clock event occurs or its time has passed due to clock adjustment, it will be deleted from the linked list and processes

pending on it will be resumed. When the clock is corrected, notifications of the arranged clock events are canceled and new simulation times for the notifications of the events are recalculated based on the corrected clock.

A timer model is also built on the clock model, which uses the drift of the clock model to calculate the duration in simulation time and is used for timeout events. In contrast to clock events, timeout events are not affected by clock synchronization and only depend on how many ticks should pass before they occur.

3.1.3. TTEthernet Modeling

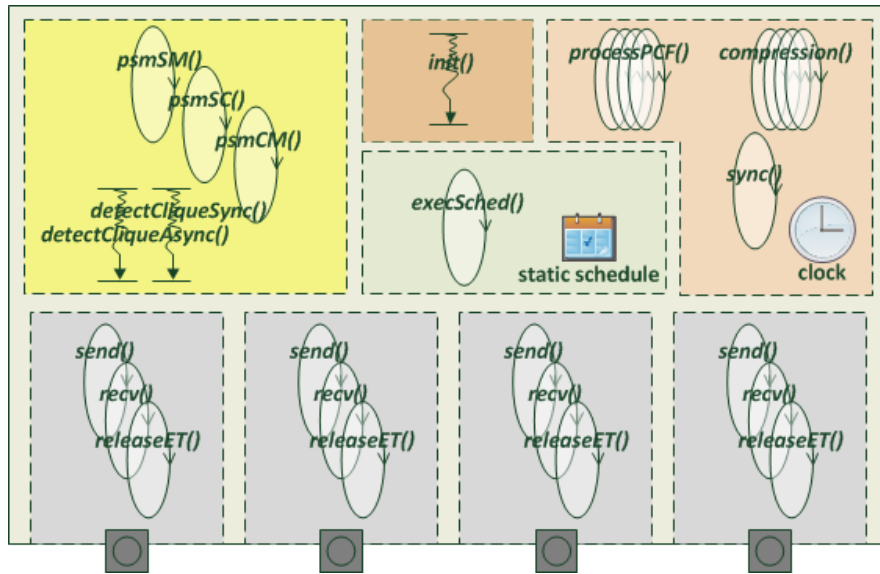


Figure 6: TTEthernet Device Main Structure

We model a concrete network protocol, TTEthernet [21], for inter-node communication. Three traffic classes, which are time-triggered (TT), rate-constrained (RC), and best-effort (BE), are supported in this protocol as well as a transparent traffic called protocol control frame (PCF) that is used for its synchronization service. There are two TTEthernet device types: the TTEthernet controller and the TTEthernet switch. Each node has at least one TTEthernet controller that can be connected by intermediate TTEthernet switch(es). The network topology is star or cascaded star so that the collision domain is segmented and only two TTEthernet devices which are directly connected may contend for the use of the medium.

The model is compliant with the TTEthernet standard [21]. Since the TTEthernet controller and switch have several common functions/services, we extract all the common functions and implement them in a class derived from the *sc_module* class. This class serves as the abstract base class of the TTEthernet controller and switch. The main components of this abstract base class is shown in Fig. 6. It has pure virtual functions that need to be implemented by the controller or switch to define different behaviors of these two different devices. We use *SC_THREAD* processes to model the TT communication behavior and protocol state machines (PSMs) of TTEthernet, and also model its two-step synchronization mechanism that is used to establish the synchronized time base.

The main processes and their functions are listed in Tab. 1. The TT communication is realized by a scheduler process (*execSched()*) which is responsible for signaling the send process (*send()*) to start a TT frame transmission according to a static schedule that relies on synchronized global time. The static schedule guarantees two TT frames never contend for transmission and is used by the TTEthernet device through a configuration file. Each TTEthernet device executes exactly one of the PSMs to maintain its role for synchronization, which are formulated in [21]. All TTEthernet devices can be classified into three different roles: synchronization masters (SMs), compression masters (CMs), and synchronization clients (SCs). Startup service of the PSMs tries to establish an initial synchronized global time to make devices operate in synchronized mode. When a device detects there is a synchronous/asynchronous clique scenario (*detectCliqueSync()/detectCliqueAsync()*), the restart service of PSMs will try to resynchronize itself. When operating in synchronized mode, TTEthernet uses a two-step synchronization mechanism: SMs dispatch PCFs to CMs, and CMs calculate the global time from the PCFs (i.e. “compress”) and dispatch “compressed” PCFs to SMs and SCs. SMs and SCs receive “compressed” PCFs and adjust their clocks to integrate into the synchronized time base. When a PCF arrives, a dynamic PCF handler process (*processPCF()*) will be spawned to cope with this PCF. If the TTEthernet device is a CM, a dynamic compression process (*compression()*) will be spawned if there is no process handling corresponding integration cycle of the PCF. After receiving scheduled PCFs, the synchronization process (*sync()*) will be resumed to calculate the clock correction from the PCFs that are in-schedule, and after a fixed delay the clock will be adjusted by the calculated correction value.

The TTEthernet controller model acts as a TLM-2.0 target which re-

Table 1: Processes in TTEthernet Model

Name	Main Function
send() & recv()	send/receive Ethernet frames
execSched()	signal TT frame transmission
releaseET()	arrange ET frame transmission
sync()	calculate clk. correction & adjust clk.
processPCF()	execute permanence function
compression()	compress PCFs
detectCliqueSync()	detect synchronous cliques
detectCliqueAsync()	detect asynchronous cliques
psmSM()	execute sync master PSM
psmCM()	execute compression master PSM
psmSC()	execute sync client PSM

ceives transactions containing Ethernet frames from the PE model via a target socket. Generic payload extensions are added to show which traffic class the Ethernet frame belongs to. The TTEthernet switch model stores and forwards different traffic class frames using different mechanisms. Since TLM-2.0 of SystemC is mainly for modeling memory-mapped buses, modeling TTEthernet requires some extensions: An Ethernet socket is introduced by deriving from both tagged initiator and target sockets of TLM-2.0 in order to simulate the bidirectional communication link between two ports of TTEthernet devices. Binding and accessing methods of the socket are implemented and new payload type for Ethernet is also added.

3.1.4. Sensor and Actuator Modeling

The cyber components interact with the physical system through sensors and actuators. In our model each sensor/actuator has a *SC_THREAD* process that is responsible for updating the sensing/actuation values. The sensors are modeled as active devices, and the actuators are modeled as passive devices.

Fig. 7 shows the operations of sensors and actuators. As an active device, a sensor will periodically use an IRQ line to inform the PE model to fetch the values through a bus transaction. According to the configuration, the active sensors can use their independent clocks or they can be synchronized with the PE model. On the contrary, the values used by an actuator are fed by the PE model periodically or sporadically.

The interactions between the sensors/actuators and the CarSim model are

simply through shared variables. The pointers to these shared variables are used by the sensors and actuators. When there is an update, the value will be read/written from/to the physical model by dereferencing the corresponding pointer.

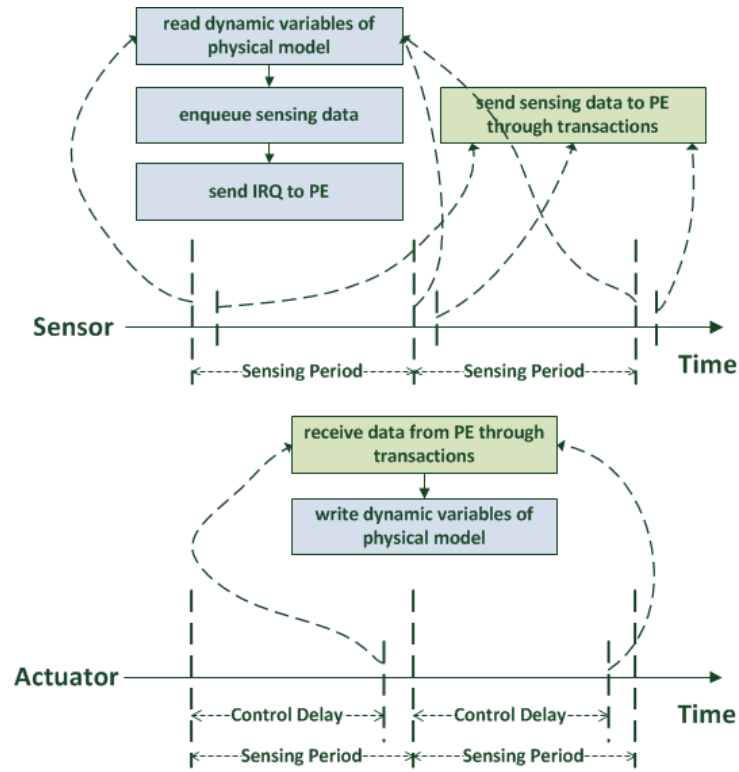


Figure 7: Active Sensor and Passive Actuator

3.2. Software Layer

The software layer comprises the software components with behavior expressed in logical time. Each software component takes the corresponding generated C code from MATLAB/Simulink model to realize its functionality.

All the software components belonging to the same PE are grouped into one task set class which is derived from *sc_module* class. When integrating all the models, the task set will be instantiated and registered to the RTOS model of the corresponding PE, and an off-line defined schedule table for TT activations is also registered to the RTOS model. Each software component

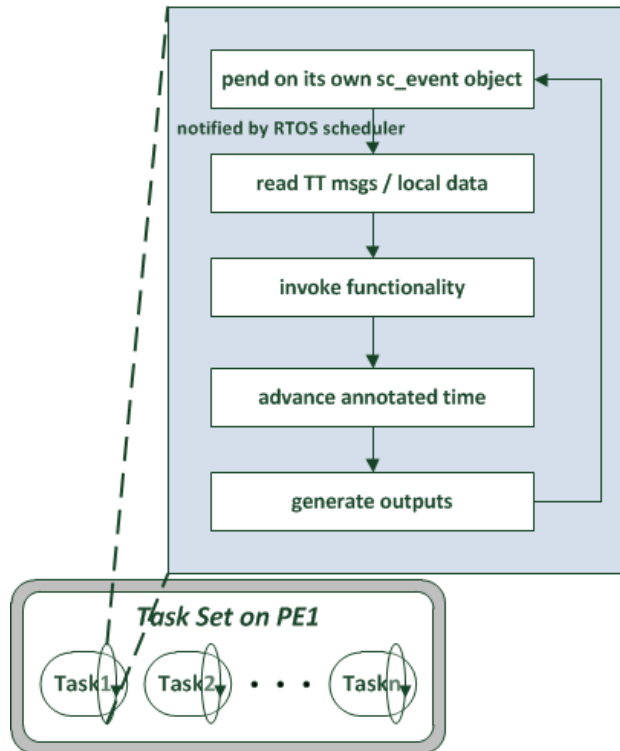


Figure 8: A Software Component

is wrapped into a *SC_THREAD* SystemC process as a task which will be scheduled by the RTOS model. Each task has an *sc_event* object. The execution of a task is pending on its own *sc_event* object which will be notified by the scheduler when the task is scheduled to run. The worst case execution time (WCET) of a software component is needed for the off-line TT paradigm scheduling tool and is annotated to the task. Although the execution of a piece of C code will take zero logical execution time, the task will invoke an RTOS API to delay itself for at least the WCET to generate the outputs. This process is shown in Fig. 8.

As shown in Fig. 2, all the control software tasks constitute the software layer and are distributed over two ECUs. The interactions between the software and network/platform layers are: the scheduler of the RTOS schedules the tasks and informs a task to run by using *sc_event* notification; the tasks acquire RTOS services, such as inter-node communication, via system calls.

3.3. Physical Layer

In order to integrate the physical layer model of an automotive vehicle with the network/platform layer model in SystemC, two requirements for the physical model are: (1) the physical model has to provide input/output interfaces through which we can access the variables representing its dynamics. (2) the physical model also should have an interface for simulation time synchronization. A wrapper module can take advantage of these interfaces and integrate them with the network/platform layer (via sensors and actuators). CarSim is a commercial parameter-based vehicle dynamics modeling software which meets these two requirements [3].

CarSim has a program called VehicleSim (VS) solver used to read and write files, calculate dynamics, and communicate with other software. It has an internal mathematical model that predicts the behavior of vehicles in response to control signals. The solver is in the form of a dynamically linked library (DLL) file with a set of API functions. Integrating CarSim into the co-simulation framework is achieved by a wrapper module that takes charge of synchronization between the SystemC simulation kernel and CarSim VS solver. The solver DLL has a set of time-stepping API functions, and the time of the physical model will be increased with the configured time step by each time-stepping API function call from the wrapper module. The VS solver solves the differential equations according to the current time and update the internal mathematical model. Different dynamics variables reference to the variables in the internal mathematical model through VS APIs. These dynamics variables are revealed by the wrapper module to the sensors/actuators of the network/platform layer through shared variables.




Figure 9: SystemC-CarSim Integration in Time Domain

Due to using the fixed-step solver in CarSim, the interval I between two successive mathematical model updates is fixed. The wrapper module has to

call the CarSim time-stepping function every fixed-step. SystemC uses a discrete event simulator which can process sensing and actuation events of the network/platform layer within an interval. In order to ensure the accuracy of the simulation, we have two restrictions for the SystemC CarSim integration in the time domain: The first one is the sensing period T_S should be greater than the simulation fixed-step, i.e. $T_S > I$; otherwise two successive sensing events within a step will acquire the same dynamics variable evaluations. Similarly for the actuation, whose timing depends on the upper layer computation whose execution time is the control delay δ , the control delay should also be greater than the fixed-step, i.e. $\delta > I$. Since the fixed-step is very small (1 ms in our example), this is not a strong restriction. The second one is after an actuation, next sensing should be at least be separated by an interval boundary. Since we use TTA, every action point in time can be specified at design time; thus, this is not a strong restriction as well.

Fig. 9 shows how the cyber part modeled in SystemC is integrated with the physical layer modeled by CarSim in the time domain (the dashed lines represent the data flow). In this example, the control sampling period is 5ms. S1 acquires the dynamics variable values updated by TS1, and the computation and communication (C1 and N1) take some time before the actuation variables are changed by A1. The changed actuation variables affect dynamics variables updating TS6 and TS7, and then another control period begins.

4. Rapid Prototyping Design Flow

The virtual prototype of a TT automotive CPS can be generated automatically by using the Embedded Systems Modeling Language (ESMoL) environment. ESMoL is a suite of domain-specific modeling languages, providing a single multi-aspect design environment and a set of tools to facilitate the design of embedded real-time control systems [18]. The rapid prototyping design flow is shown in Fig. 10.

The first four steps belong to using ESMoL to facilitate high-confidence control software design. Step 1 specifies the control functionality in the MATLAB/Simulink environment and configure/establish the physical dynamics model. The Simulink model will be imported into the ESMoL automatically to become the functional specification for instances of software components.

Step 2 specifies the non-functional parts of the system in ESMoL which includes: (1) the logical software architecture which captures data depen-

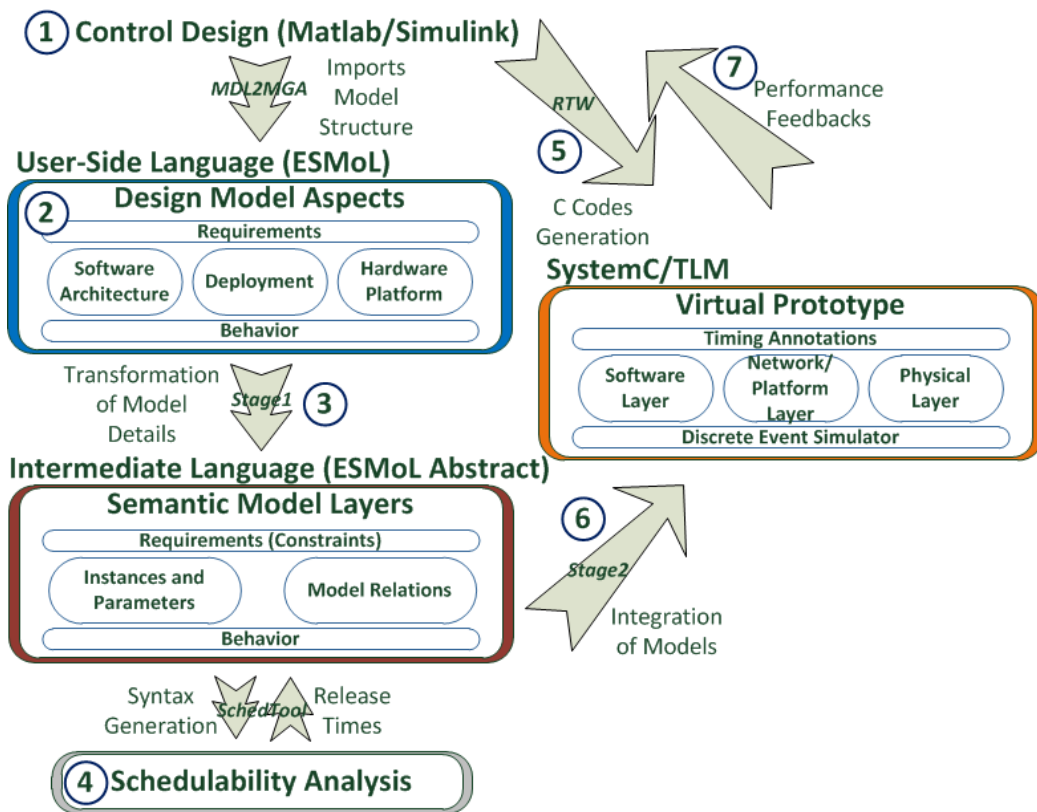


Figure 10: Rapid Prototyping Design Flow Supported by the ESMoL Language and Virtual Prototype

dencies between software component instances independent of their distribution over different nodes; (2) the hardware platforms defined hierarchically as nodes with communication ports interconnected by I/O blocks and networks: Model attributes for a node mainly capture timing resolution, its RTOS scheduling policy, and overhead parameters for data transfers, ISR, and task context switching. The processor speed is indirectly embodied by the WCET of the tasks specified in the timing model. The parameters for an I/O block include its type (input/sensing or output/actuation) and transaction data size. The parameters for a network include its bandwidth and its TDMA slot size. The topology of the network can be obtained from these interconnections; (3) a deployment model set up by mapping software components to nodes, and data messages to communication ports: The deployment model captures the assignment of component instances as periodic tasks running on a particular node. Message ports on component instances are assigned to hardware interface ports in the model to indicate the media through which message are transferred; (4) a timing model established by attaching timing parameter blocks to components and messages: For the time-triggered automotive control software the configuration parameters of the timing model include execution period, desired offset, relative deadline and WCET.

Step 3 translates the ESMoL model into the simpler ESMoL_Abstract model using the *Stage1* interpreter of ESMoL. The model in this intermediate language is flattened and the relationships implied by structures in ESMoL are represented by explicit relation objects in ESMoL_Abstract. Step 4 generates the scheduling problem specification from ESMoL_Abstract model and uses a tool of ESMoL called *SchedTool* to solve the scheduling problem. The results are imported back into ESMoL model and written to the appropriate objects. More details of these four steps can be found in [18].

In order to integrate the co-simulation framework with ESMoL, we extend the ESMoL design flow. Step 5 generates C code from MATLAB/Simulink model using a code generator (e.g. Real Time Workshop (RTW) toolbox). Step 6 uses *Stage2* interpreter of ESMoL to generate the virtual prototype. For each model of the cyber part, there is a corresponding configuration template which can be parameterized by utilizing a template engine (e.g. Google Ctemplate). The interpreter uses the UDM model navigation APIs to traverse the ESMoL_Abstract model to assemble the C code generated by RTW into tasks and parameterize the configuration templates. The template for a task is organized as follows: in an infinite loop it first waits on its own

sc_event object; if the task is a receiver of a remote message, it invokes the read message API with corresponding arguments; then it invokes its generated C function to compute in zero logic execution time; its execution time is enforced by calling the timing annotation API to pass its WCET to the RTOS it belongs to; at last if the task sends a message to the network, it calls the write message API with corresponding arguments; otherwise, it updates the shared memory. All the tasks running on the same node are grouped into one task set class which is derived from *sc_module* class. For each node, the PE, bus, TTEthernet controller, sensors, and actuators are instantiated, connected, and configured in the *sc_main()* function which is the top level of a SystemC program. The configuration files for the model instances are generated according to the specified attributes in the ESMoL model, such as the schedule tables for the RTOSes. The task set class of the node is also instantiated and registered to the RTOS of the PE model. TTEthernet switches are also instantiated and configured. According to the topology defined in the hardware model of ESMoL, all the nodes are connected. The physical model is instantiated and configured. All the pointers to the shared variables of the physical model are passed to the corresponding sensors/actuators. Finally, the co-simulation results of the holistic system provide performance feedback for engineers to revise their designs, which is the step 7.

5. Validation and Evaluation

In this section, we first validate the TTEthernet model by comparing TT traffic delays obtained from our model, a TTEthernet model in OMNeT++ INET framework [24], and a real TTEthernet implementation [27], and also evaluate its scalability and simulation efficiency. Then, we use an adaptive cruise controller (ACC) case study to validate the co-simulation framework by comparing with the results obtained from a hardware-in-the-loop (HIL) automotive simulator. Based on the ACC controller, we incrementally add another automotive system controller which takes charge of lane keeping, and compare the results of the integrated controller under different simulation environments (MATLAB/Simulink, HIL, and co-simulation framework). In terms of measuring control performance, we use the tracking ability of the controller as the performance metric.

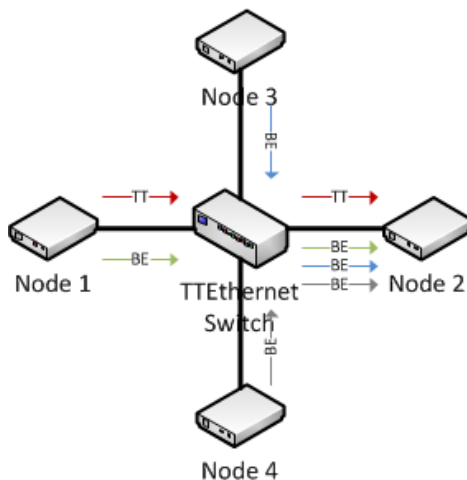


Figure 11: Validation of TT Communication in TTEthernet Model Setup

5.1. TTEthernet Model Validation and Evaluation

We validate the TTEthernet model by comparing the average end-to-end transmission delay and jitter of the TT traffic of different TTEthernet models (including our model in SystemC/TLM, a model in OMNeT++ INET framework [24], and a software-based implementation) under the same experimental scenario. We set up a star topology which has four nodes connected to a central TTEthernet switch with 100Mbit/s links as shown in Fig. 11. Node 1 sends both TT traffic and BE traffic to Node 2, and both Node 3 as well as Node 4 send only BE traffic to Node 2. All the traffic goes through a TTEthernet switch. The communication period is 10ms, and the time slot is $200\mu\text{s}$. The maximum clock drift is set as 200ppm for the models. Node 1 sends a TT frame at 1ms offset of each period. The configuration files including their corresponding XML files for the nodes and switch are generated by the TTTech toolchain [27]. From the generated XML files, we extract parameters such as critical traffic table and schedule table to configure our model and the model in OMNeT++. In this setup, the switch dispatches the TT frame sent by Node 1 at 1.4ms offset of each period. We measure the average end-to-end latency and jitter for different TT frame sizes under full link utilization of BE traffic. Fig. 12 shows the results of our model in SystemC/TLM, the model in OMNeT++ INET framework, and the software stack implementation in Linux from TTTech [27].

From the figure we can see the model in SystemC/TLM and the model




Figure 12: Average End-to-End Transmission Delay and Jitter of Different Frame Sizes

in OMNeT++ INET framework give very similar results. The method of measuring end-to-end transmission delay of software-based TTEthernet implementation is presented in [2], which utilizes two ports on a single box. From the results we can observe there is a latency gap ($90\mu\text{s}$) between frame size of 123 and 124 bytes which is actually caused by the BE-device driver configuration according to [2]. This gap is due to measurement approach limits and will not appear when using the TT communication. The measured jitter of the software-based implementation is bounded by $30\mu\text{s}$. The hardware-based implementations will bound the jitter more tightly [27].

We also evaluate the scalability and simulation efficiency of the TTEthernet model. We set up the evaluation using a central switch, and all the nodes are connected to the switch. The simulation time is 1000s, and increasingly add a pair of nodes into the network. Each pair of the nodes, such as Node 1 and Node 2, communicates with each other using TT, RC, and BE traffic. Each node sends out a TT frame, a RC frame, and a BE frame every 10ms. Thus, there are $300,000 \times \text{number of nodes}$ frames totally. The result is shown in Fig 13.

From the results we can see the model in SystemC/TLM has good simulation efficiency when the number of nodes increases. The simulation speed of the model in OMNeT++ INET framework is also evaluated under the same computation environment. We simulate the same topology and traffic by using the fastest mode in OMNeT++ to get rid of the influence of animation and text outputs.

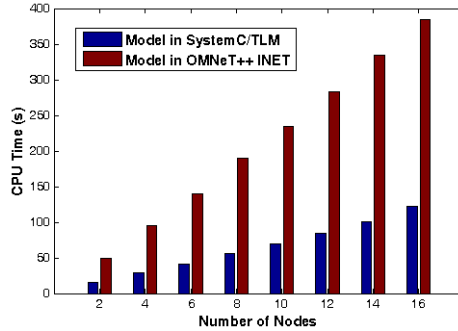


Figure 13: TTEthernet Simulation Efficiency Evaluation: Used CPU Time of Different Number of Nodes.

5.2. Case Studies

We use two automotive case studies to show the co-simulation framework can facilitate the design of automotive CPS efficiently providing more realistic results than MATLAB/Simulink. First, we introduce our hardware-in-the-loop simulator which serves as the basis of the comparisons. Then, we carry out a single adaptive cruise controller (ACC) case study to show the validity and efficiency of the framework. Finally, based on the ACC, we add another automotive controller, lane keeping controller (LKC), to show the co-simulation framework has a good ability to deal with incremental design.

5.2.1. Hardware-in-the-Loop Simulator

In order to test the automotive control system in a more realistic way, we use a HIL automotive simulator. The architecture of the HIL simulator is shown in Fig. 14. The physical dynamics modeled in CarSim is deployed on an RT-Target in a sense that it acts as the real automotive vehicle. The RT-target is also integrated with a TTTech PCIe-XMC card which enables the seamless integration and communication with ECUs on the time-triggered network. The network/platform layer of the HIL simulator is composed of three ECUs which are connected to an 8-port 100Mbps TTEthernet development switch from TTTech. Each ECU is an IBX-530W box with an Intel Atom processor running a RT-Linux operating system. Each ECU is integrated with a TTEthernet Linux driver which is a software-based implementation of TTEthernet protocol to enable communication with other end systems in a TTEthernet network. Automotive control software is distributed over the ECUs and the tasks execute in the kernel space of RT-Linux which

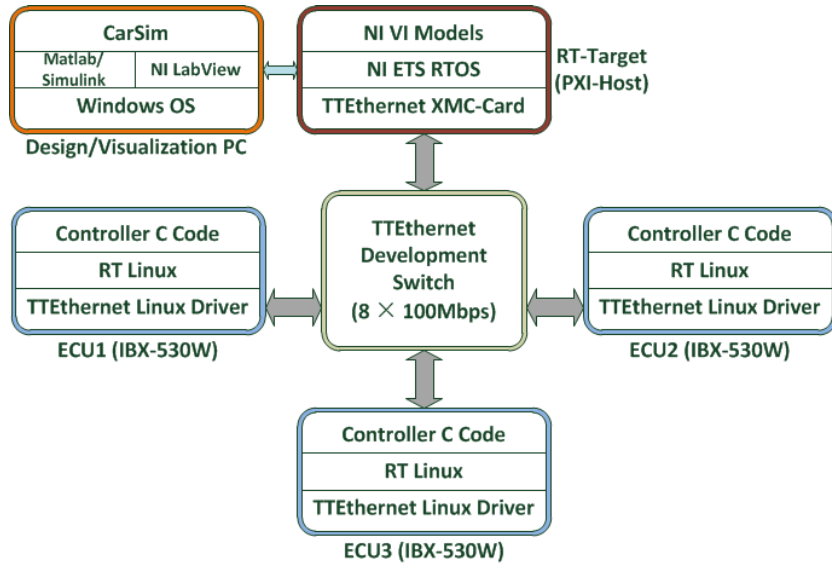


Figure 14: System Architecture of HIL Automotive Simulator

can utilize the synchronized time base of the TTEthernet communication.

5.2.2. ACC Case Study

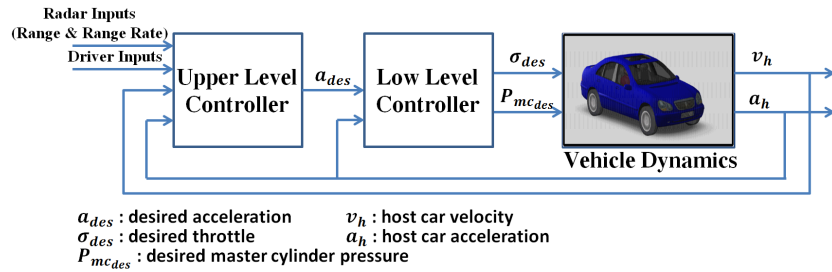


Figure 15: Adaptive Cruise Control System [6]

The control algorithm of ACC is designed in MATLAB/Simulink. Fig. 15 shows a block diagram of the ACC system. The ACC is hierarchically divided into two levels of control: the upper level controller and the low level controller. The main functionality of the upper level controller is to compute the desired acceleration for the ACC-equipped vehicle that achieves the desired spacing or velocity. The main objective of the low controller is two-fold: first, using the desired acceleration command from the upper level

controller, the lower level controller determines whether to apply braking control or throttle control; second, the required control command is applied to the vehicle in order to achieve the desired acceleration. Details about the ACC can be found in [6].

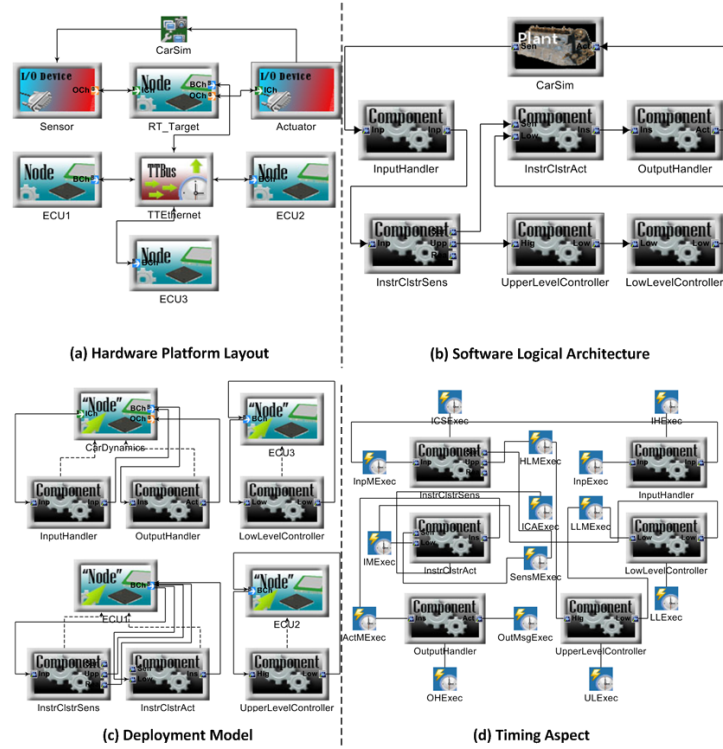


Figure 16: ESMoL Design Models of ACC

The model is imported into ESMoL environment. The four different aspects of the design in ESMoL are shown in Fig. 16. The topology of the network/platform layer is based on the HIL simulator which is shown in Fig. 16 (a). Fig. 16 (b) shows the software logical architecture that depicts the logical interconnections of four ACC tasks, which are *InstrClstrSens*, *UpperLevelController*, *LowLevelController*, and *InstrClstrAct*, and two sensing/actuation tasks, which are *InputHandler* and *OutputHandler*. The deployment of the ACC control software is shown in 16 (c) in which the dashed arrows represent assignment of tasks to their respective ECUs and solid lines represent assignment of message instances to communication channels on the ECU. Finally, the timing and execution model for tasks and message transfers

of the ACC control system are shown in 16 (d).

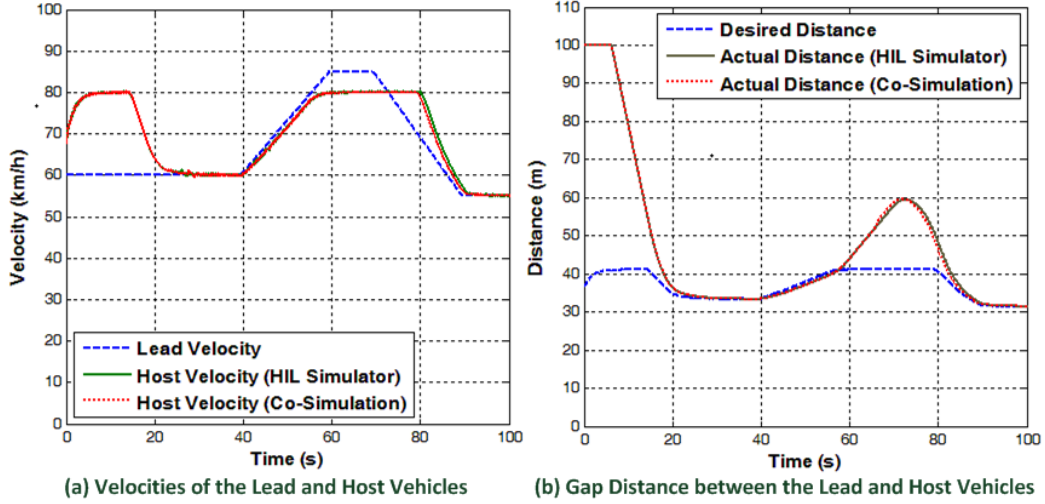


Figure 17: Velocities and Gap Distance from HIL Simulator and Co-Simulation Framework

The sampling period on the HIL simulator is 10ms which is limited by the software-based TTEthernet implementation. In this case study, the velocity of the leading vehicle starts at an initial value of 60km/h. The host vehicle radar range is 100m. The initial global longitudinal positions of the leading vehicle and the host vehicle are 130m and 0m respectively, which means the host vehicle radar is initially out of range. The host vehicle starts at an initial velocity of 65km/h with a driver set target velocity of 80km/h. The expected driving behavior of the host vehicle should be: (1) before the host vehicle detects the leading one, it will speed up to and maintain at most at 80km/h; (2) when the radar detects a slower leading vehicle, the ACC will control the distance between the two vehicles to a driver set time gap, and the desired gap distance is attained when two vehicles travel at the same velocity; (3) when the leading vehicle begins to speed up, the velocity of the host vehicle will also increase in order to achieve a desired velocity (the host maintains at most at 80km/h, even when the leading vehicle exceeds this speed); (4) when the leading vehicle slows down, the host also starts to decrease its velocity in order to maintain the desired space between the vehicles.

The results of the designed ACC running on the HIL simulator and the proposed framework are given in Fig. 17 (since the results are very similar, most of the curves are overlapped). We zoom in on the velocity plots between

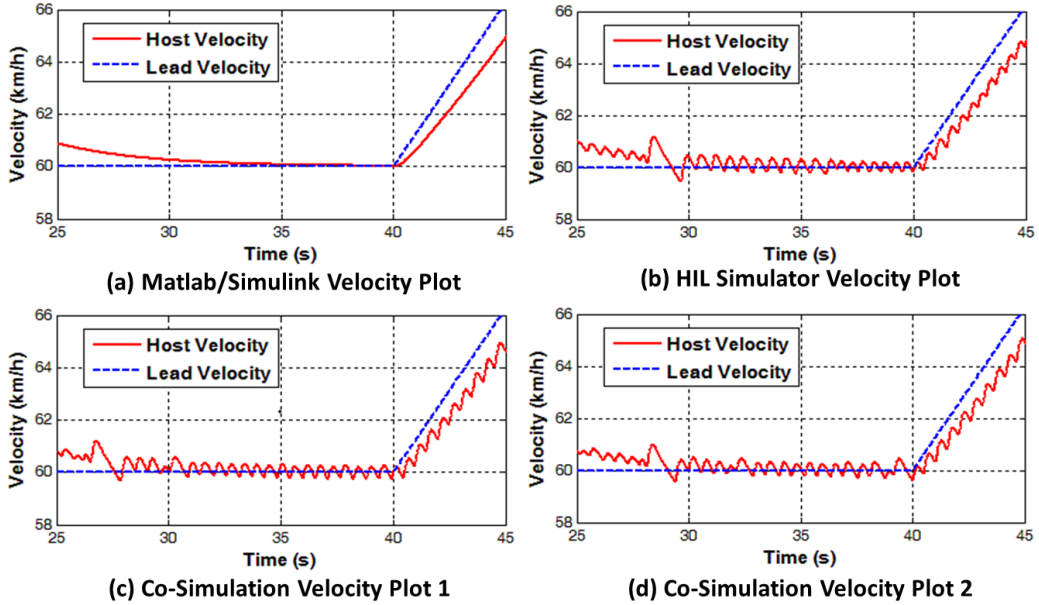


Figure 18: Comparison of Results from MATLAB/Simulink, HIL Simulator, and Co-Simulation Framework

25s and 45s, and also use the result from MATLAB/Simulink as a reference which is shown in Fig. 18 (a). From Fig. 18 (b), we can observe the velocity of the HIL simulator suffers from some oscillations which have a highest value about $0.7\text{km/h} \approx 0.2\text{m/s}$. The co-simulation results also shows these oscillations as shown in Fig. 18 (c) and (d). Due to different randomly assigned clock drifts, the results cannot be exactly the same; yet, the figures show very similar results, especially compared to the result from MATLAB/Simulink. For example, similar but not identical results are given in Fig. 18 (c) and (d) (e.g. the peak of the oscillations is shifted more towards 30s in (d)).

5.2.3. LKC + ACC Case Study

In addition to the ACC case study, we also integrate it with a LKC to carry out an integrated controller case study. Fig. 19 shows the block diagram of the nested PID LKC. The nested PID LKC is composed of two controllers. The outer loop controller, which is the Controller-1, is a PID controller with an additive integral action on the lateral offset to reject the disturbances on the curvature which increase linearly with respect to time. Controller-1 computes a desired reference yaw rate based on the vehicle's




Figure 19: Lane Keeping Control System [23]

lateral displacement. The inner loop controller, which is the Controller-2, is a PI controller and computes the desired steering angle required for achieving zero lateral distance at the look-ahead distance.

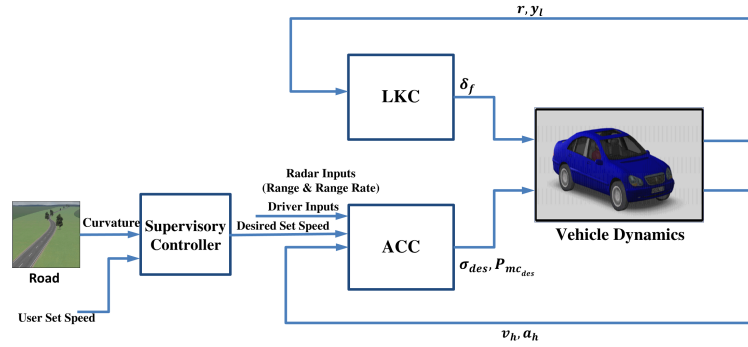


Figure 20: Integrated Control System (LKC + ACC) [23]

The integrated controller’s block diagram is shown in Fig. 20. Although the two controllers affect the behavior of two seemingly different dynamics of the vehicle (the ACC controls the longitudinal dynamics while the LKC controls the lateral dynamics), there exists physical interactions in the both the lateral and longitudinal dynamics of the vehicle. Moreover, changes in the physical environment such as geometry of vehicle path or road curvature highlights certain conflicts in the operation of these two controllers. Thus, we add a supervisory controller whose main objective is to restrict the regions of operations of the integrated system in a safe desirable manner. More details about the LKC and the supervisory controller can be found in [23].

Again, the integrated controller in MATLAB/Simulink is imported into ESMoL environment. The four different aspects of the design in ESMoL are shown in Fig. 21. The topology of the network/platform layer is still the

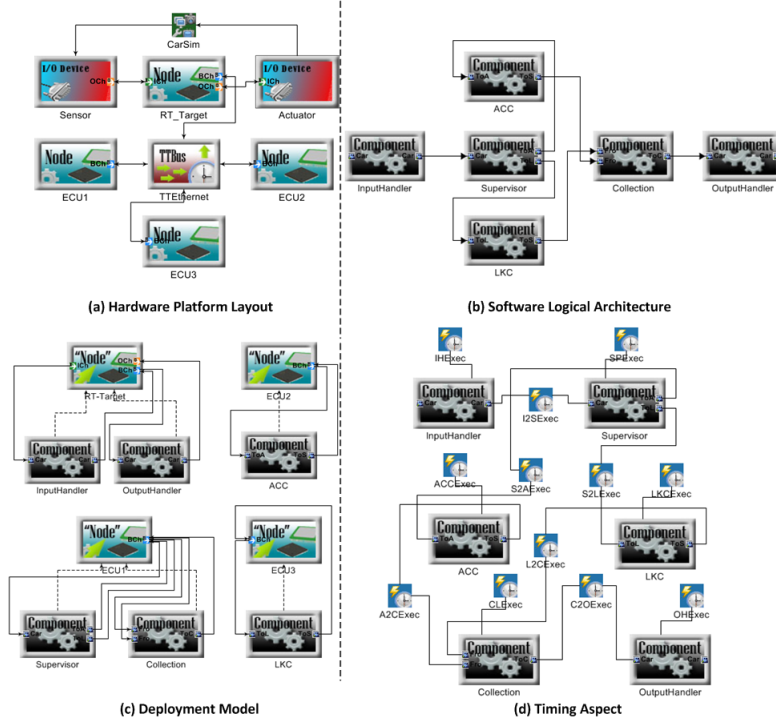


Figure 21: ESMoL Design Models of LKC + ACC

same as the HIL simulator which is shown in Fig. 21 (a). Previously in the ACC case study, we divided the ACC into four tasks and distributed them onto different ECUs. However, in this case study, we keep ACC and LKC as a whole, and distribute ACC and LKC on two different ECUs, which can be seen in Fig. 21 (c). The software logical architecture depicts the logical interconnections six tasks, which are *Supervisor*, *ACC*, *LKC*, and *Collection* and two sensing/actuation tasks, which are *InputHandler* and *OutputHandler*. Finally, the timing and execution model for tasks and message transfers are shown in 21 (d).

In this case study, we establish a test track which has a combination of straight paths and three curved roads with radii of 160m, 200m, and 160m respectively (as shown in Fig. 22). The look-ahead distance of the LKC controller is 5m. The desired time gap of the ACC is set to 1.5s. The leading vehicle starts at an initial position of (0, 0) with an initial speed of 30km/h while the host vehicle, equipped with the integrated control system, starts at an initial position of (-800, 0) with an initial speed of 80km/h.




Figure 22: Test Track with Three Curves

Table 2: Used CPU Time of Different Designs of Automotive Controllers

Name	Sampling Period	Simulation Time	Used CPU Time
ACC	10ms	100s	102s
LKC + ACC	10ms	300s	373s

The mainly concerned results of the integrated controller running on the MATLAB/Simulink, the HIL simulator and the proposed framework are given in Fig. 23 for lateral acceleration and in Fig. 24 for lateral displacement. From the results, we can observe the co-simulation of the controller reveals similar behavior to the controller running on the HIL simulator. Especially in Fig. 24, we can see the trajectories of the co-simulation and HIL simulator have a lot of overlap, whereas the trajectory of MATLAB/Simulink has observable differences from the other trajectories.

5.2.4. Simulation Efficiency

In order to illustrate the simulation efficiency of the whole framework, we provide the CPU time for 100s simulation time of the ACC and 300s simulation time of the integrated controller under a machine with dual cores of 3.40GHz and 8GB memory. The results are shown in Tab. 2.

6. Design Space Exploration

In this section, we explore different design options which include (1) reducing the sampling period to observe its affection on the ACC longitudinal velocity, and (2) introducing additional background network traffic to

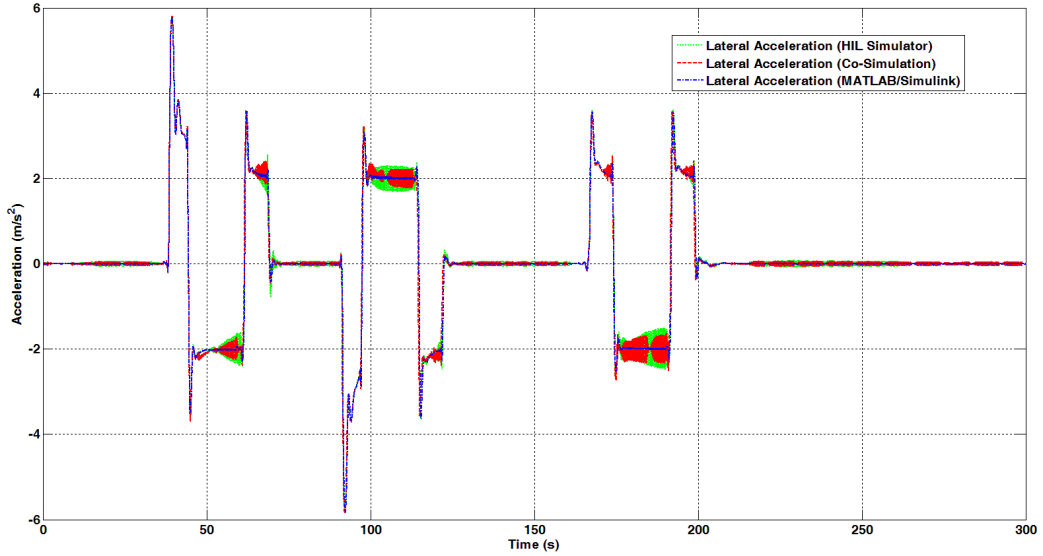


Figure 23: Lateral Acceleration

evaluate its influence on the LKC lateral displacement. In addition, since our HIL simulator has an implementation limitation that does not allow the computation on the RT-Target to be synchronized with the TTEthernet communication, we can use the co-simulation framework to eliminate this implementation limitation to evaluate its impact on the control performance.

6.1. Reducing Sampling Period

In the ACC case study, the ACC software execution on the HIL simulator is not computationally intense. The timing diagram generated by the co-simulation framework (Fig. 25) shows that every task meets its deadline which is represented by the dotted line (due to implementation limitations, on RT-Target the computation is not synchronized with the communication). If the physical layer is not included, like the tools introduced in [25] and [30], the system is perfectly designed. However, when the car dynamics model begins to execute in this simulation, the oscillations of the vehicle velocity can be observed. In order to improve performance, we can increase the sampling rate. The computation of the system is negligible, but the communication system that uses the software-based implementation of TTEthernet becomes an obstacle which limits the fastest reasonable sampling period to 10ms. In order to reduce the sampling period, we need to consider other design

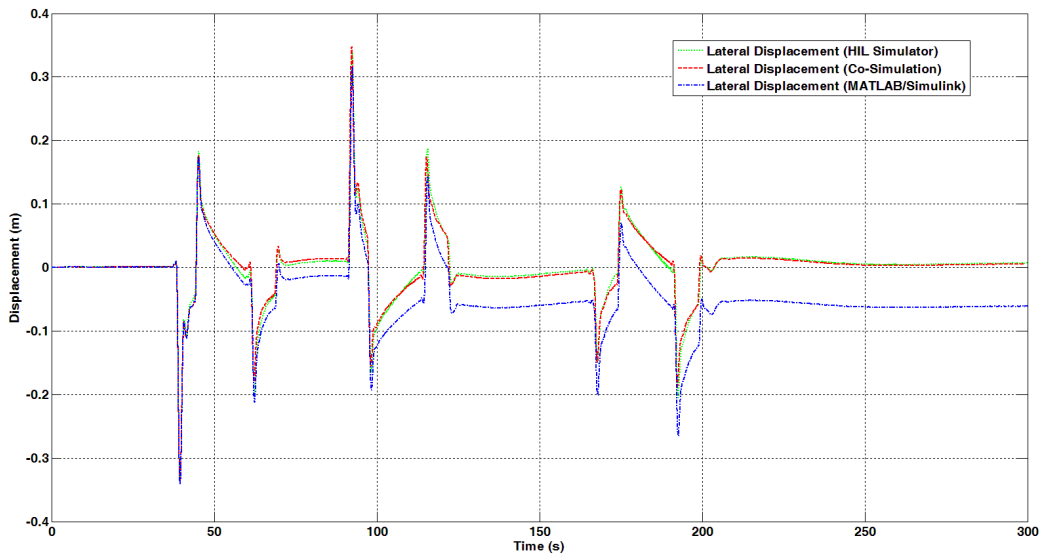


Figure 24: Lateral Displacement




Figure 25: Timing Diagram of ACC Tasks

alternatives in the design space.

By employing the hardware-based implementation of TTEthernet which has a 1Gbits/s bandwidth and more precise clock synchronization, we can achieve the sampling period reduction. A 5ms sampling period is used in the new design and gains in the ACC controller are tuned. The zoomed-in velocity result of the co-simulation is given in Fig. 26, from which we can see the performance of the ACC is much better than the previous one: the oscillations have a highest amplitude about $0.05\text{km/h} \approx 0.015\text{m/s}$.

In this case, the 5ms sampling period causes more computation and communication than in the case of 10ms sampling period, so the used CPU time for 100s simulation time also increases to 194s.




Figure 26: Co-Simulation Velocity Plot of 5ms Sampling Period

Table 3: LKC Control Performance Under Different Network Traffic Scenarios

Scenario	Max. Abs. Err.(m)	Avg. Err.(m)	Max. Osc. Amp.(m)
1	0.3448	0.0267	0.001
2	0.3448	0.0267	0.001
3	0.3768	0.0291	0.012

6.2. Introducing Additional Network Traffic

One advantage of the TTEthernet is that it integrates mixed-criticality traffic classes together which is also supported in our framework. For the LKC + ACC automotive controller, we introduce some background additional traffic to evaluate its influence on the control performance. In this evaluation, instead of using randomly assigned clock drifts, we fix the relative clock drift (50ppm) between the RT-Target computation and the global time base provided by TTEthernet in order to get rid of other affections other than additional network traffic. Since zero lateral displacement is an ideal performance, we use the maximum lateral displacement error, the average error and the maximum oscillation amplitude for 300s running on the three curves test track to compare the control performance. The results are shown in Tab. 3. The first scenario serves as the baseline in which there is no additional traffic but for the network traffic needed for the automotive controller. In the second scenario, we introduce additional BE traffic on the TTEthernet: in each control period, RT-Target sends a BE frame to ECU1, ECU1 sends the frame to ECU2, ECU2 sends the frame to ECU3, and ECU3 sends the frame back to RT-Target. From the result, we can observe that since the automotive controller uses the TT traffic, additional BE traffic does not affect the control performance at all. In the third scenario, we introduce more TT frames so that this can result in generating a new TT schedule

with 15ms sampling period (we assume sampling period should be 5ms's multiple). Since the sampling period is increased, the control performance should be decreased, which can be observed from the increased maximum and average errors and the maximum oscillation amplitude.

In this case, the 15ms sampling period causes less computation and communication than in the case of 10ms sampling period, so the used CPU time for 300s simulation time also decreases to 242s.

6.3. Clock Synchronization Strategies

Our HIL simulator has an implementation limitation that does not allow the computation on the RT-Target to be synchronized with the TTEthernet communication. We can conjecture that the oscillations are mainly due to the delays caused by the non-synchronized computation with the TT communication on the RT-Target. This conjecture can be proved by comparing the co-simulation result of the synchronized setting with the one of the non-synchronized setting of the RT-Target. Fig. 27 shows the co-simulation result of the synchronized setting of the RT-Target, from which we can observe the oscillations are apparently reduced (the highest amplitude is about $0.09\text{km/h} \approx 0.025\text{m/s}$).




Figure 27: Co-Simulation Velocity Plot by Using Synchronized Setting of RT-Target

7. Conclusion

In this paper, we propose a co-simulation framework that can facilitate TT automotive CPS design. A simplified view of designing CPS is to consider three design layers, which include the physical layer, the network/platform layer, and the software layer. The proposed framework contains models from each of the three CPS design layers. SystemC is used to model the cyber

part and CarSim, a commercial automotive simulator, is used to model the physical part of an automotive CPS. Since the network/platform layer is the intermediate layer between the other two layers, it plays an important role in CPS integration and becomes the backbone of the framework models. The models can be configured and integrated to become a virtual prototype of a TT automotive CPS to provide realistic feedback at early design stages. The framework is also integrated with a model-based design tool called ESMoL to enable rapid prototyping. The TTEthernet model of our network/platform model is validated against a real implementation and a TTEthernet model in OMNeT++ INET in term of its end-to-end transmission delay and jitter. Two automotive case studies (ACC and LKC+ACC) are provided to illustrate the framework. The case studies shows that the co-simulation framework provides similar results to a HIL simulator with good efficiency. DSE is also demonstrated taking into account different design options and concerns.

References

- [1] Armengaud, E., Teng, A., Driussi, M., Karner, M., Steger, C., Weiss, R., Jun 2009. Automotive software architecture: Migration challenges from an event-triggered to a time-triggered communication scheme. In: 7th Workshop on Intelligent Solutions in Embedded Systems. pp. 95–103.
- [2] Bartols, F., Steinbach, T., Korf, F., Schmidt, T. C., 2011. Performance analysis of time-triggered ether-networks using off-the-shelf-components. In: Proceedings of the 2011 14th IEEE International Symposium on Object/Component/Service- Oriented Real-Time Distributed Computing Workshops. ISORCW '11. pp. 49–56.
- [3] CarSim, -. [Http://www.carsim.com/](http://www.carsim.com/).
- [4] Cervin, A., Henriksson, D., Lincoln, B., Eker, J., Årzén, K.-E., Jun 2003. How does control timing affect performance? Analysis and simulation of timing using Jitterbug and TrueTime. IEEE Control Systems Magazine 23 (3), 16–30.
- [5] Drolia, U., Wang, Z., Pant, Y., Mangharam, R., oct. 2011. Autoplug: An automotive test-bed for electronic controller unit testing and verification. In: Intelligent Transportation Systems (ITSC), 2011 14th International IEEE Conference on. pp. 1187–1192.

- [6] Eyisi, E., Zhang, Z., Koutsoukos, X., Porter, J., Karsai, G., Sztipanovits, J., 2013. Model-based control design and integration of cyber-physical systems: An adaptive cruise control case study. *Journal of Control Science and Engineering*.
- [7] Gheorghe, L., Bouchhima, F., Nicolescu, G., Boucheneb, H., 2006. Formal definitions of simulation interfaces in a continuous/discrete co-simulation tool. In: *IEEE International Workshop on Rapid System Prototyping*. pp. 186–192.
- [8] He, Z., Mok, A., Peng, C., 2005. Timed rtos modeling for embedded system design. In: *Proceedings of the 11th IEEE Real Time on Embedded Technology and Applications Symposium. RTAS '05*. pp. 448–457.
- [9] Hegde, R., Mishra, G., Gurumurthy, K. S., 2011. An insight into the hardware and software complexity of ecus in vehicles. *Advances in Computing and Information Technology, Communications in Computer and Information Science* 198, 99–106.
- [10] IEEE, 2011. *IEEE Standard SystemC Language Reference Manual*.
- [11] Ishikawa, M., McCune, D. J., Saikalis, G., Oho, S., 2007. Cpu model-based hardware/software co-design, co-simulation and analysis technology for real-time embedded control systems. In: *Proceedings of the 13th IEEE Real Time and Embedded Technology and Applications Symposium. RTAS '07*. pp. 3–11.
- [12] Kopetz, H., Bauer, G., 2003. The time-triggered architecture. *Proceedings of the IEEE* 91 (1), 112–126.
- [13] Krause, M., Bringmann, O., Hergenhan, A., Tabanoglu, G., Rosentiel, W., 2007. Timing simulation of interconnected autosar software-components. In: *Proceedings of the conference on Design, automation and test in Europe. DATE '07*. pp. 474–479.
- [14] Lonn, H., Axelsson, J., 1999. A comparison of fixed-priority and static cyclic scheduling for distributed automotive control applications. In: *ECRTS*. pp. 142–149.
- [15] Marrec, P. L., Valderrama, C. A., Hessel, F., Jerraya, A. A., Attia, M., Cayrol, O., 1998. Hardware, software and mechanical cosimulation for

- automotive applications. In: Proceedings of the Ninth IEEE International Workshop on Rapid System Prototyping. RSP '98. pp. 202–.
- [16] Müller, W., Becker, M., Elfeky, A., DiPasquale, A., 2012. Virtual prototyping of cyber-physical systems. In: ASP-DAC. pp. 219–226.
 - [17] Navet, N., Song, Y., Simonot-Lion, F., Wilwert, C., 2005. Trends in automotive communication systems. Proceedings of the IEEE 93 (6), 1204–1223.
 - [18] Porter, J., Hemingway, G., Nine, H., vanBuskirk, C., Kottenstette, N., Karsai, G., Sztipanovits, J., Sep 2010. The esmol language and tools for high-confidence distributed control systems design. part 1: Language, framework, and analysis. Tech. rep., Vanderbilt University.
URL http://www.isis.vanderbilt.edu/sites/default/files/ESMoL_TR.pdf
 - [19] Quaglia, D., Muradore, R., Bragantini, R., Fiorini, P., 2012. A systemc/matlab co-simulation tool for networked control systems. Simulation Modelling Practice and Theory 23 (0), 71 – 86.
 - [20] Rushby, J. M., 2001. Bus architectures for safety-critical embedded systems. In: Proceedings of the First International Workshop on Embedded Software. EMSOFT '01. pp. 306–323.
 - [21] SAE Standard AS 6802, 2011. Time-Triggered Ethernet.
 - [22] Schirner, G., Dömer, R., 2008. Introducing preemptive scheduling in abstract rtos models using result oriented modeling. In: Proceedings of the conference on Design, automation and test in Europe. DATE '08. pp. 122–127.
 - [23] Shang, D., Eyisi, E., Zhang, Z., Koutsoukos, X., Porter, J., Karsai, G., Sztipanovits, J., 2013. A case study on the model-based design and integration of automotive cyber-physical systems. In: Proceedings of the 21st Mediterranean Conference on Control and Automation. MED '13. pp. 483–492.
 - [24] Steinbach, T., Kenfack, H. D., Korf, F., Schmidt, T. C., 2011. An extension of the omnet++ inet framework for simulating real-time ethernet with high accuracy. In: Proceedings of the 4th International ICST

- Conference on Simulation Tools and Techniques. SIMUTools '11. pp. 375–382.
- [25] Streubühr, M., Jäntschi, M., Haubelt, C., Teich, J., Mar. 2009. From Model-based Design to Virtual Prototypes for Automotive Applications. In: Proceedings of the Embedded World Conference. Nuremberg, Germany, pp. 1–10.
- [26] Sztipanovits, J., Koutsoukos, X., Karsai, G., Kottenstette, N., Antsaklis, P., Gupta, V., Goodwine, B., Baras, J., Wang, S., Jan. 2012. Toward a science of Cyber-Physical system integration. Proceedings of the IEEE 100 (1), 29–44.
- [27] TTEch Computertechnik AG, -. TTEthernet Products. [Http://www.ttech.com/en/products/ttethernet/](http://www.ttech.com/en/products/ttethernet/).
- [28] Verhoef, M., Visser, P., Hooman, J., Broenink, J., 2007. Co-simulation of distributed embedded real-time control systems. In: Proceedings of the 6th international conference on Integrated formal methods. IFM'07. pp. 639–658.
- [29] Zabel, H., Müller, W., Gerstlauer, A., 2009. Accurate RTOS modeling and analysis with SystemC. In: Ecker, W., Müller, W., Dömer, R. (Eds.), Hardware-dependent Software. Springer Netherlands, Ch. 9, pp. 233–260.
- [30] Zeller, M., Weiss, G., Eilers, D., Knorr, R., 2010. Co-simulation of self-adaptive automotive embedded systems. In: Proceedings of the 2010 IEEE/IFIP International Conference on Embedded and Ubiquitous Computing. EUC '10. pp. 73–80.
- [31] Zhang, Z., Porter, J., Eyisi, E., Karsai, G., Koutsoukos, X., Sztipanovits, J., 2013. Co-simulation framework for design of time-triggered cyber physical systems. In: Proceedings of the ACM/IEEE 4th International Conference on Cyber-Physical Systems. ICCPS '13. pp. 119–128.