

Distributed Co-Simulation of Maritime Systems and Operations

Severin Sadjina
SINTEF Ålesund
NO-6009 Ålesund, Norway
Email: severin.sadjina@sintef.no

Lars Tandle Kyllingstad
SINTEF Ocean
NO-7465 Trondheim, Norway
Email: lars.kyllingstad@sintef.no

Martin Rindarøy
SINTEF Ocean
NO-7465 Trondheim, Norway
Email: martin.rindaroy@sintef.no

Stian Skjong
SINTEF Ocean
NO-7465 Trondheim, Norway
Email: stian.skjong@sintef.no

Vilmar Æsøy
Department of Marine Technology
Norwegian University of
Science and Technology
NO-7491 Trondheim, Norway
Email: vilmar.aesoy@ntnu.no

Eilif Pedersen
Department of Marine Technology
Norwegian University of
Science and Technology
NO-7491 Trondheim, Norway
Email: eilif.pedersen@ntnu.no

Here, we present the concept of an open virtual prototyping framework for maritime systems and operations that enables its users to develop re-usable component or subsystem models, and combine them in full-system simulations for prototyping, verification, training, and performance studies. This framework consists of a set of guidelines for model coupling, high-level and low-level coupling interfaces to guarantee interoperability, a full-system simulation software, and example models and demonstrators. We discuss the requirements for such a framework, address the challenges and the possibilities in fulfilling them, and aim to give a list of best practices for modular and efficient virtual prototyping and full-system simulation. The context of our work is within maritime systems and operations, but the issues and solutions we present here are general enough to be of interest to a much broader audience, both industrial and scientific.

1 Introduction

With operations that are becoming increasingly complex and demanding, volatile economic conditions, stricter environmental and safety regulations, and decreasing project lead times, simulation methods have become a key indicator of merit in early design phases within the maritime industry. They allow for a quick exploration of the design space and help to advance concepts into certain directions [1, 2]. Simulation methods also expose concept, interface, and safety flaws, which, in turn, helps reduce risks and enhance opera-

tional performance and efficiency.

Traditional ship design is a sequential process [3]: Development is driven step-by-step and iteratively with each engineering discipline using its own set of tools that are rarely inter-operational. This complicates system level analysis and verification significantly, and obscures errors and issues until late in the design process when they are difficult and expensive to fix. Because of this, there is now an increasing interest in being able to fully integrate simulation techniques into the ship design process for prototyping, verification, training, and performance studies [4]. Through the development of virtual prototypes it is possible to test the characteristics and dynamics of a proposed concept early on, and expose possible issues long before the integration and prototyping stages, during which they usually surface. This is especially important for the maritime industry which—unlike other industries such as automotive, aerospace, and railway—typically has to provide one-of-a-kind solutions. Because of this, it is hard to learn lessons and advance with a traditional sequential process, when designs are unique and different each time around.

Virtual prototyping also holds the promise of substantially simplifying the search for an optimal design and helping to keep costs, risk factors, and environmental impacts low. For a typical offshore supply vessel, for example, the power consumption of the on-board equipment is significant and will, to a large extent, determine the dimensions of the energy systems. At the same time, there is an incentive to

minimize these dimensions from economic and environmental standpoints to keep the entire vessel as small as possible.

If the optimization of the overall system performance is the goal, ship designs can, generally¹, not simply be optimized for the performance of individual components and subsystems, but need to be optimized with respect to total operational performance. Only with the interactions between components, the surroundings, control systems and software, and operators properly accounted for, is it possible to choose a system design with desirable characteristics—such as fuel efficiency, maneuverability, and safety. With more advanced operations requiring more power, interaction, and timing, system performance simulation will become even more important in the future. Full-system simulation, however, remains a challenging and elusive task for at least two reasons:

1. Typical maritime systems and operations are difficult and complex to model and simulate by nature: they are characterized by intricate interactions between a wide range of physical and engineering domains, with dynamics taking place on vastly different time scales, see Fig. 1. Compare, for example, the slow dynamics of large mechanical systems to the fast response of electronic components.
2. This complexity and diversity is reflected in a simulation landscape that is riddled with specialized tools for different physical and engineering domains, with different interfaces and incompatible model representations. Some existing simulation tools for maritime applications are highly advanced in terms of quality, functionality, and usability. But they are mainly developed with research and the optimization of components and subsystems in mind, and lack interconnection capabilities. General software solutions, on the other hand, are too inflexible, and offer model development and configuration that is too time-costly or inaccurate. Moreover, the models themselves span a wide range of complexities and accuracies, including continuous as well as discrete behavior, and different focuses depending on the analyzed phenomena. Consequently, understanding how different subsystems interact with each other and how they influence overall system behavior becomes all the more challenging.

Traditionally, simulations of closely-coupled subsystems are constructed from the ground up, resulting in monolithic simulations for custom interfaces that are too application specific, too customized, and too costly in terms of development time. The ability to assemble re-usable and interchangeable subsystem models into virtual prototypes in a *plug and play* manner—regardless of the environment in which they are developed—should cut down on development times significantly and enable rapid innovation. Undoubtedly, this would be a big step forward for the maritime industry.

¹Overall system performance can only be optimized in terms of the individual subsystems' performance if those subsystems are entirely independent of each other. A simplification which will not hold for all but the most simple systems.



Fig. 1. Maritime systems and operations include a wide range of different engineering domains and physical systems with varying complexity and time scales. This, naturally, makes full-system simulation a challenging endeavor.

The integration of multi-physics simulations, human behavior, and multiple parallel maritime operations has already been successfully demonstrated in operational training simulators [5]. However, to date there are no universally agreed upon methods or standards supporting total systems integration and the analysis of operational performance. The project *Virtual Prototyping of Maritime Systems and Operations* [6, 4] (ViProMa) was initiated within the Norwegian maritime industrial cluster by independent research organizations, universities, and industry partners² with the goal of developing an open, standardized framework and architecture for system simulation and virtual prototyping as a new platform for product development and cooperation: the *Virtual Prototyping Framework* (VPF). This framework includes

- guidelines for model coupling,
- high-level interfaces for coupling models from different engineering and physical domains,
- low-level interfaces for coupling models from different tools,
- a full-system simulation software,
- and example models and demonstrators.

It aims to make the communication between costumers, designers, and product developers more efficient throughout the design process. It also facilitates the consistency and availability of objective *Key Performance Indicators* (KPIs) if they are integrated into the prototyping system.

1.1 Outline

In this paper, we discuss the development of a common technology platform and infrastructure supporting virtual prototyping and simulation-oriented work processes for maritime systems and operations. The ViProMa project serves as an exemplary framework that allows us to have a closer look at the knowledge gaps and challenges with re-

²SINTEF Ocean, NTNU, VARD, Rolls-Royce Marine, and Det Norske Veritas.

spect to the simulation of entire maritime systems and operations. The aforementioned guidelines, which are at the heart of the VPF, are continuously summarized and emphasized throughout the text. While the context of the present work is maritime systems and operations, we would like to emphasize that the challenges discussed and the solutions presented here are broadly applicable. They should, thus, also be of interest and value to people outside of the maritime sector.

First, we give a brief overview of the requirements for a virtual prototyping framework for the maritime industry in section 2. In section 3, we discuss the development of a common architecture for system simulation (the VPF). We review different simulation approaches and justify our choice of using co-simulation (simulator coupling) in ViProMa. We then touch upon the formidable task of defining standardized domain model interfaces in section 4, in order to establish a modular framework with high interoperability and re-usability and a focus on maritime applications. Section 5 brings these concepts together in order to elaborate on the construction of full-system models and the underlying challenges. We then discuss our decision to write our own in-house co-simulation software in section 6, where we also present the challenges faced and the progress made so far. Finally, we provide a conclusion and share our thoughts on future developments in section 7.

2 Virtual Prototyping Framework Requirements

The requirements for a framework for virtual prototyping in the maritime industry are shaped by its most important use cases. In terms of the ViProMa project, these are:

Vessel design. Comparison and optimization of concepts with respect to fuel efficiency, capabilities, operabilities, availability, maintainability, and opportunity for future expansion. This also includes virtual sea trials, and the testing of control systems and station keeping abilities.

Crew training. Extending the possibilities towards higher realism and the inclusion of harsher environments, and increasing the awareness of vessel and equipment limitations.

Decision support. Choice of vessel for a specific voyage, and aligning ship capabilities with weather windows and equipment capabilities with specific operations.

Requirements are also dictated by the desired workflow that such a framework should facilitate: After a model of a subsystem has been developed, it is connected to the full-system simulator. A scenario is then selected, the full-system simulation run, and behavior and capabilities evaluated.

Within the ViProMa project, this led to the following core requirements imposed on the VPF. Note that while these are specific to the project and the maritime industry, most of them are highly relevant to a much broader area of application.

1. There has to be support for distributed and cross-platform operation due to heterogeneous technologies, tools, and platforms, and the possibility for workload

distribution.

2. Human-in-the-Loop and Hardware-in-the-Loop, and, thus, real-time operation, have to be supported for integration with training simulators, dynamic positioning (DP) and other control systems, and various types of machinery and equipment. This also helps to save time in factory acceptance tests.
3. The ability to use components as black boxes to protect intellectual property and sensitive information has to be implemented.
4. The framework must be license-free with no restrictions on commercial use to prevent vendor lock-in, lower the barrier of use, and guarantee a widespread commitment.
5. Increasingly rigid time constraints in the industry demand sufficient performance with regards to the overall prototyping process, as well as the simulations alone. In addition, strategies to achieve reasonable accuracy and stability of full-system simulations need to be established and implemented.
6. The framework has to be sufficiently easy to use and provide well-defined interfaces in order for the industry to actually adopt and use it. This is also crucial in light of future development and maintenance of the framework.
7. It should have a complete component database containing at least generic domain models of varying model complexity. Because final designs are often re-used to save money and time, the database should also be easily searchable.

The following sections are devoted to addressing the challenges and possibilities in realizing these core requirements. Due to limitations in funding and time, however, some of the requirements listed here were only partially or not at all fulfilled within the scope of the ViProMa project. This will be commented on in Sec. 7.

3 Distributed Co-Simulation

In setting out to prototype and simulate a complex maritime system (such as a modern offshore supply vessel), and capture all its relevant dynamics and interactions, one quickly realizes that the traditional monolithic simulation approach is too inflexible, too costly, and too inefficient. A model of an entire vessel usually takes a long time to get ready for simulation, if constructed from the ground up, and re-use is often prohibited when faced with similar problem sets due to the developed model being too customized and too application-specific.

It is clear then, that a modular approach, in which models of the relevant subsystems are interconnected and simulated together, is favorable to cut down on cost and development time. This also means that the process of making changes to a subsystem should be effortless, in the sense that it does not require modifications of any other parts of the full-system model (i.e., the open-closed principle should be adhered to). There are, generally, three methods to combine several models [7]:

1. The use of a common modeling language into which all models are translated for the purpose of the simulation,
2. the exchange of models between tools to run a simulation in one of them (*model exchange*),
3. and *co-simulation* (simulator coupling).

Considering the requirements discussed in Sec. 2, the first two modeling approaches have several drawbacks:

- General software solutions are, typically, too inflexible, and offer model development and configuration that is too time-costly or inaccurate.
- In general, they neglect the availability of matured and specialized domain-specific analysis software for maritime applications.
- Users are reluctant to change modeling languages or tools, as doing so is a major undertaking in practice, and may render investments in tool chains and training worthless.
- The choices for a suitable common modeling language or tool are strongly limited if the development of an *open* framework for systems simulation and virtual prototyping is the goal.

In addition, the use of a common modeling language sometimes means abandoning the black box approach that protects Intellectual Property Rights (IPRs) and sensitive information. Model exchange, however, may play a vital role in future developments for full-system simulation, and we shall touch upon this briefly in Sec. 5.3.

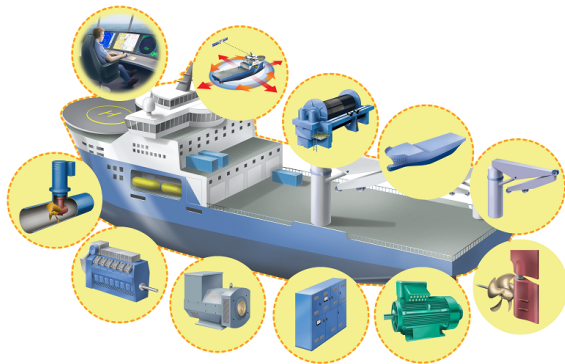


Fig. 2. In a co-simulation setting, different tools and models are interconnected and used independently and in parallel to form a full-system simulation.

As mentioned previously, the fact that typical maritime systems and operations are comprised of a wide range of physical and engineering domains naturally leads to a rather heterogeneous simulation landscape, with specialized tools and proprietary model representations. While this may seem like a rather poor starting point for full-system simulation and virtual prototyping, it is precisely this modular structure of complex engineering systems, in conjunction with

the availability of well-established domain- and application-tailored software, that lends itself quite well to co-simulation, see Fig. 2. The remainder of the present section is devoted to a discussion of the co-simulation approach and a brief review of existing co-simulation software and standards.

3.1 Co-Simulation

The basic idea behind co-simulation is the construction of systems from loosely-coupled stand-alone models and the simulation across different subsystems. Co-simulation facilitates the independent exchange and modification of components, and the use of the most suitable tools and solvers for any given subsystem. This also extends to the possibility of separately taking care of initialization, pre-processing, time integration, and post-processing with different specialized tools. This is very advantageous, because it allows for rapid and accurate model development that makes efficient use of already available modeling software and languages without major investments in new tool chains or training. Co-simulation further has the potential benefit of significantly reducing simulation time by using model-specific solvers and internal step sizes, and by allowing for the distribution of computational loads onto different computers or processor cores. The possibility to conveniently hide internal dynamics and protect sensitive information is another attractive trait of simulator coupling, especially from an industrial perspective.

Guideline 1 (Model coupling). *Use co-simulation to construct full-system models from loosely-coupled stand-alone models and modules.*

Mathematically, co-simulation corresponds to the modular time integration of subsystems that are assumed to be independent in between discrete *communication points* $t_i \in \{t_0, t_1, \dots, t_N\}$. The interactions between the subsystems are only realized at these time points, and are expressed in the form of interface constraint equations,

$$\mathbf{u}(t_i) = \mathbf{L}\mathbf{y}(t_i), \quad (1)$$

where \mathbf{L} is a connection graph matrix relating the inputs \mathbf{u} and the outputs \mathbf{y} . This happens at a rate corresponding to a *macro step size* Δt_i , such that $t_{i+1} = t_i + \Delta t_i$. In general, there is no guarantee that the pieces play together nicely, though. Time synchronization and data exchange are important tasks, consequently, and sound and efficient communication between subsystems implies an adequate understanding of the architecture.

Remark 1.1. Be cautious when selecting coupling method and co-simulation (macro) step size to avoid accuracy and stability issues.

Because input variables are unknown to the subsystems during the time integration $t_i \rightarrow t_{i+1}$ and, thus, need to be approximated (and are often held constant), co-simulation brings its own set of stability and accuracy issues. Most commonly, this is remedied by selecting a sufficiently small

macro step size. At the same time, demands to keep the computational cost within reasonable bounds may, generally, require a lower limit, especially for real-time applications.

Remark 1.2. For linear systems, the macro step size can be chosen from the eigenfrequencies, but it may be very difficult to find a good choice for nonlinear problems.

In addition, there are several other subtleties and challenges to co-simulation accuracy and stability. For example, the presence of algebraic loops can result in instability [8], and the presence of different time integration methods can actually decrease the overall accuracy of the full-system simulation to below the minimum accuracy of the individual subsystems [9]. In general, the development of an efficient and robust co-simulation method that is easy to use and generally applicable is still ongoing research [10]. Additionally, it is not always clear beforehand where to draw subsystem boundaries, and how to choose a set of ‘good’ interface constraint equations. Such factors can play a significant role in determining stability and accuracy, as we shall see in Section 5.

The ViProMa project tried to address many of these issues through original research:

1. A novel *Energy-Conservation-based Co-Simulation* method [11, 12] (ECCO) was developed within ViProMa. It gives readily available feedback on global simulation quality, and significantly improves the accuracy and efficiency of non-iterative co-simulations. We shall discuss it briefly in Sec. 4.1.2.
2. Additionally, an analysis tool for global stability in linear distributed dynamical systems has been proposed by combining dynamic stability and solver stability [13], both of which are intimately linked through local and global time steps. Under certain conditions, an algebraic solution of the total system can be constructed, and probed for global stability. However, this procedure can be very time consuming, and is only applicable for linear dynamical systems. Extensions to nonlinear dynamical systems have been studied as well, but the corresponding work is still ongoing.

The task of establishing an efficient and robust general-purpose co-simulation methodology is far from completed, however. Especially numerical stability is a nontrivial subject to study due to the inherent complexities (different solver methods of various orders, different coupling schemes of various orders, the presence of direct feed-through and algebraic loops, et cetera).

There exist non-iterative (explicit) and iterative (implicit) schemes to couple subsimulators in a co-simulation, and time steps can be performed in parallel (Jacobi scheme) or in serial (Gauss-Seidel scheme). The simplest and most straight-forward of all schemes is the explicit one with constant input approximation. It is easiest to realize, keeps the exchange of coupling data to a minimum, and does not require the repetition of entire macro time steps (rollback). Because of this, it is frequently used in industrial applications, and has also been the main focus for the ViProMa project so

far. It does, however, exacerbate the aforementioned stability and accuracy issues that co-simulation brings about naturally.

3.2 Existing Co-Simulation Software and Standards

Among existing solutions for performing distributed co-simulations, the most prominent one is probably the *High-Level Architecture* (HLA). HLA is not one specific software package; rather, it is a standard which describes a general-purpose co-simulation architecture. It was initially developed by the US Department of Defense for use in wargaming and training simulations, and was eventually made an IEEE standard. The latest version of this is IEEE 1516-2010, commonly called *HLA Evolved* [14]. Several HLA implementations exist today, both commercial and free.

Similar architectures include the *Distributed Interactive Simulation* [15] (DIS), which is the precursor of HLA and is even more geared towards military applications, and the *Common Simulation Interface* [16] (CSI) developed by SINTEF Ocean for the purpose of maritime vessel simulations.

These architectures are designed around the concept of a *federation*, which is a group of independent subsystems (*federates*) that communicate through a common *Run-Time Infrastructure* (RTI). The RTI is responsible for routing signals between the federates and for time synchronization. The federates may be numerical simulations, hardware interfaces, human interfaces, et cetera. Oft-stated advantages of HLA include *interoperability*, in that federates may run on different platforms and use different simulation methods, and *reuse*, in that federates used for one simulation may be easily re-used in another. However, because the wire protocol between the federates and the RTI is not standardised, a federate created for one HLA implementation generally can’t be used with a different implementation. This, along with other reasons which will be explored in Section 6, is why HLA was deemed unsuitable for the VPF.

Another standard we will mention here is the *Functional Mock-up Interface* (FMI), which, unlike HLA, offers a way to make subsystems binary compatible with each other, thus, removing the need for recompilation and facilitating model sharing and co-simulation. FMI has become a key component of the VPF, and we shall therefore describe it in more detail in section 4.2.

4 Simulator Interfaces

In order to ensure interoperability, modularity, and reuse between different models and simulators, well-defined interfaces are needed. Such *simulator interfaces* are a set of conventions that, if adhered to, allow a simulator to be coupled with other simulators using some co-simulation middleware. Here, we distinguish between two levels of interfaces, which we shall refer to as *high-level* and *low-level* interfaces.

A high-level interface is concerned with the concepts which are being modeled and simulated. That is, it deals with what a simulator represents on a physical level and the physical interpretation of data exchanged between simulators. For

example, it is crucial that the value of an output variable which represents a force in units of kN in one simulator is not used for an input variable which represents a force in units of N in another—or one which represents, say, a voltage. A high-level interface could prevent this by either mandating that certain quantities must have specific units, or by defining some mechanism whereby the units can be communicated, so that the co-simulation middleware can make the necessary value adjustments and/or prevent invalid connections. On top of this, the interface can define *groups* of variables which together have some physical significance. An example of these are *power bonds*, which are pairs of variables that represent different means of power transfer between entities. Power bonds are discussed further in section 4.1.

Finally, at the highest level, one can define interfaces that represent categories of components or subsystems in the system being simulated. For example, one could define that an ‘engine’ has a power bond for rotational mechanical power and an output variable which represents fuel consumption. Then, any simulator which adheres to this interface could be used to represent an engine, and could be replaced with any other simulator that has the same interface. This opens great possibilities for ‘plug-and-play’ construction of complex system models and, consequently, rapid evaluation and optimization of different designs.

The high-level interfaces are necessarily underpinned by one or more low-level interfaces, which are concerned with the finer details of how the co-simulation middleware interacts with the simulators. At the lowest level, the physics involved are completely disregarded, and there is nothing preventing one from, for example, coupling a force variable with a voltage variable; the interface deals in bits and bytes, not newton and volt. One example of a low-level interface is what is known as an *application binary interface* (ABI). Among other things, an ABI defines how different types of data, such as integers, real numbers and textual data, are represented in computer memory. Complementary to this is an *application programming interface* (API), which specifies the names of program functions, which data they receive and return, and, to some extent, what the functions do. For the VPF, the choice fell on the *Functional Mock-up Interface*, which defines a simulator API and more. This is discussed further in section 4.2.

4.1 High-Level Interfaces

Power and energy are the universal currencies of physical systems. Energy is conserved and continuous: energy flows out of, or into, a system are always accounted for by appropriate energy storage and dissipation. This is the theoretical foundation of bond graph theory [17, 18], which balances energy flows for each subsystem separately. This way, they can be connected together in a modular fashion, while satisfying energy conservation and continuity for the entire system. The energetic couplings between (sub)systems are realized with so-called *power bonds*, which are defined by a pair of *power variables*: a *flow* and an *effort*. Their product is always a physical power—such as force and veloc-

ity, pressure and flow rate, or voltage and current. The use of power bonds provides a complete and universal, energy-flow-centered connectivity between mathematical models of different engineering and physical domains. As pointed out recently [11], they are thus perfectly suited for high-level interfaces for co-simulation.

Guideline 2 (High-level interfaces). *Properly define and use high-level interfaces to guarantee interoperability of simulation models. Make use of power bonds to model the flow of energy between subsimulators whenever possible. The use of SI units is highly advised. If other units are used, explicitly and clearly document so.*

4.1.1 Power Bonds for Co-Simulation

A power bond between two coupled simulators is realized by connecting two input–output pairs. For example, a model of an electric generator could have a voltage as an output, which a connected electric consumer model would receive as an input. In turn, the consumer would output an electric current, which the generator model accepts as an input. These two exemplary models are then coupled via a power bond.

Demanding that physical couplings between simulators are realized through power bonds has a major advantage: it allows to study the flow of energy between the subsystems directly using nothing but the simulator coupling values. In fact, it makes it possible to directly observe if and where the conservation of energy is violated throughout a co-simulation, which, in turn, helps identify potential simulation issues, and provides valuable and immediate feedback about the quality of the results.

4.1.2 Residual Energies and Energy Conservation

In general, energy is incorrectly transferred between two coupled simulators due to the fact that their states are evolved independently of each other between discrete communication time points, see Fig. 3. In effect, energy is either created or destroyed through the co-simulation coupling during each macro time step [19]. This *residual energy* directly alters the total energy of the overall coupled system [11]. It thus changes its dynamics, deteriorates simulation accuracy, and may pose a threat to numerical stability.

Guideline 3 (Error estimation). *Use an error estimation method to assess and control co-simulation coupling errors, and guarantee the quality and validity of the simulation results.*

Remark 3.1. The use of the *Energy-Conservation-based Co-Simulation* method (ECCO) for reliable co-simulation error estimation is recommended. It is easily implemented, computationally inexpensive, and does not require the repetition of entire co-simulation step sizes.

If power bonds are used, such energy residuals are conveniently calculated from the coupling variable values alone. These concepts are exploited by ECCO to obtain global error estimates. Unlike virtually all other proposed co-simulation

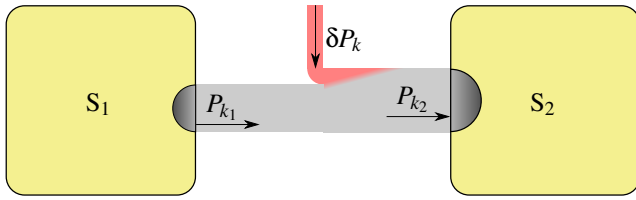


Fig. 3. A residual power $\delta P_k = -(P_{k_1} + P_{k_2})$ emerges and distorts the dynamics of the full system when energy is exchanged between two subsimulators, S_1 and S_2 , in a co-simulation

schemes, ECCO requires neither rollback nor any specific information on model implementations.³ Consequently, it does not prohibit the use of commercial or legacy software (which often makes rollback inefficient or impossible), and it helps protect IPRs. Because of this, it is especially attractive from an industrial perspective. In addition, it accurately tracks coupling errors, even for relatively very large time steps, beyond which stability is already compromised, see Fig. 4.

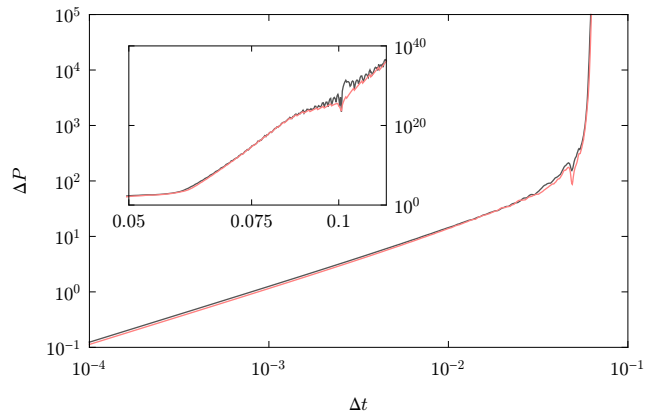


Fig. 4. Energy-conservation-based error estimation (red) compared to the actual error in the power ΔP (gray) as a function of the co-simulation step size Δt for the benchmark model in Ref. [11]. The critical step size is $\Delta t \approx 0.059$ s

Remark 3.2. It is advisable to use methods which adaptively control the co-simulation step size. The use of ECCO for step size control is recommended to ensure approximately accurate energy transfers between subsimulators.

An adaptive control of the macro step size can easily be realized based on ECCO’s error estimate. This can improve the accuracy and efficiency of co-simulations substantially: depending on system reticulation and subsimulator-internal solver accuracies, reductions in the error of 93 % or more can be observed in benchmarks—at no additional computational cost [11, 12]. In practice, models may not support variable macro step sizes, however, generally leaving the potential for more accurate and more efficient co-simulations untouched.

³ To the best of our knowledge, the only other co-simulation method for error estimation and adaptive step size control with these qualities is the one described in Ref. [20].

Remark 3.3. If power bonds are not applicable, other co-simulation methods for error estimation and adaptive step size control can be used, see Ref. [11] and references therein. Almost all of these require the repetition of entire co-simulation step sizes and simulator-internal information, however.³

4.2 Functional Mock-up Interface

The *Functional Mock-up Interface* (FMI) is a tool independent standard for the exchange of dynamic models and for co-simulation [21]. The first version of the standard was published in 2010 as a result of the ITEA2 project MODELISAR. Since 2011, maintenance and development of the standard have been performed by the Modelica Association, and a second version of the standard was released in 2014 [22].

FMI specifies that models should be packaged as *functional mock-up units* (FMUs), which are archive files that contain model code for one or more platforms, along with metadata and documentation. The standard defines the format and structure of files and directories in an FMU, as well as the APIs that must be implemented by the model code. These APIs are defined in terms of the C programming language which, being the *lingua franca* of programming, allows FMUs to be written in, and used from, practically any other language.

The FMI standard consists of two main parts: *FMI for Model Exchange* and *FMI for Co-Simulation*. FMI for Model Exchange specifies an interface for models that represent differential, algebraic, and discrete equations, which are typically coupled with and solved together with other models in some simulation software. Important here is that the solver is supplied by the simulation software and is not part of the FMU code. In contrast, FMI for Co-Simulation, which we shall focus on here, defines an interface for models which are bundled with their own solvers, and which can, therefore, be seen as separate simulators in themselves. Several such ‘subsimsulators’ will typically be coupled together in a co-simulation environment, a piece of software that enables data exchange between the subsimulators, and keeps them synchronized in time. All data exchange takes place at *communication points* (sometimes called *synchronization points*), between which each model is solved independently from the others by its own solver. Note that we use the terms ‘model’ and ‘simulator’ freely here; in practice, there is nothing preventing these entities from being interfaces to hardware such as sensors, actuators, or devices for human input.

FMI for Co-Simulation is based on a master/slave model of communication and control, where subsimulators are slaves that are controlled by a master algorithm. The subsimulators do not have any information about each other, nor about the simulation environment, except for the values they receive for their input variables. Thus, they have no knowledge about or control over which other subsimulators they are coupled to; the data is routed by the master algorithm.

The FMI standard is a fitting choice for maritime applications for three main reasons:

1. It was created in collaboration with the automotive industry for many of the same reasons that we aim to design the VPF for the maritime industry.
2. The standard is already supported by a large number of tools, for example by Dymola, JModelica.org, SIMPACK, SimulationX, and Simulink.
3. FMI is completely open and free to use for any purpose.

Finally, a decisive reason to choose FMI was that the alternative seemed to be to define a new interface and, to a large extent, reinvent the wheel.

Guideline 4 (Low-level interfaces). *Use the Functional Mock-up Interface to ensure compatibility between different simulation tools and languages. Packaging subsimulators as functional mock-up units makes them tool independent and re-usable.*

It should be emphasized at this point that FMI only specifies how the (co-)simulation software interacts with the models; it is *not* in itself a simulation software, nor does it specify or restrict any other parts of the architecture of such a software. More to the point, in a distributed co-simulation setting, FMI does not say how, or in what format, data are transported between the simulation nodes, nor how the nodes are time synchronized. As such, FMI support can be implemented in almost any type of simulation software, and, indeed, the number of tools that support this standard is large and growing quickly [23].

Remark 4.1. Whether or not any given FMU is compliant with the FMI standard can be checked with the free *FMU Compliance Checker* [23].

As an interesting side note, Awais *et al.* have proposed to use HLA as a co-simulation environment for FMI-based components [24]. This is despite the fact that HLA's federation model stands somewhat in contrast to FMI's master/slave structure: federates are not passive slaves that simply wait to be given anonymous input data, but instead actively request, by name, the data they require. The findings by Awais *et al.* appear to reflect this. Their conclusion is that it is technically possible to use HLA as an FMI master, but to be able to make properly generic 'FMU federates', the FMI standard must be extended so that the name mapping between federation data and FMU variables can be specified in the FMU metadata.

One aspect of FMI which should be mentioned here, as it has an impact on the model structures discussed in Sec. 5, is the fact that FMUs are, for most practical purposes, *closed for modification*. That is, once an FMU has been created, there is no simple way to modify its behavior nor its external interface. The model code is typically stored in compiled binary code form, and the numbers, types, and names of input and output variables are fixed. In some contexts, this puts severe limits on scalability. For example, imagine that we have an FMU that models a body which is acted upon by external forces. Then, we have to decide upfront, at model construction time, on a maximum number of force inputs. If more forces are needed, the model source code must be modified and a new FMU compiled, or the force summation (and possibly coordinate transformations) must take place externally.

Also note that while FMI shows enormous potential and is continuously and actively refined, it also has various deficiencies. Some aspects of co-simulation are altogether poorly addressed by the standard, and because limitations are often rather subtle, it is well worth pointing them out. For example, it is possible to design FMUs and master algorithms that are standard compliant but exhibit nondeterministic and unexpected behavior [25]. It is also not directly clear how non-continuous models can be encoded as FMUs [26,27,28]. Furthermore, FMI 1.0 has no general rollback mechanism beyond the previous macro time step, while FMI 2.0 makes rollback only optional, and error control is not addressed by FMI at all [25].

5 Model Construction

Let us now bring together the theoretical concepts and principles from the previous Sections 2–4, and apply them to the construction of full-system models from stand-alone components and submodels. We shall first elaborate on some general considerations and common challenges, before we reap the fruits of our efforts and move on to discuss some model examples at the end of the present section.

5.1 Subsimulators

Subsimulators represent the various components which are connected together to form a full-system simulator. It is important to note that this concept does not only include time-dependent physical systems. The actual functionality of a subsimulator is at the complete discretion of its developer, and includes:

- mathematical models of physical systems evolving over time,
- control system implementations,
- *scenario controllers* (units which control and steer a simulation, e.g., by changing parameters during runtime),
- hardware interfaces (e.g. to allow direct hardware input to the simulation),
- and bridging to other software or co-simulations.

This list is not exhaustive by any means, and several other use cases are possible. That some of these may warrant the use of a different concept will be the subject of Section 5.5. But first, let us discuss the boundaries between subsimulators.

5.2 System Boundaries

When modeling a complex system composed from many different subsystems in a co-simulation environment, one of the many challenges is to determine where to draw the boundaries between the separate modules. To exemplify this, consider simulating a ship which consists of a power system, propulsion units, and a hull. A few examples of how to divide such a system into different subsystems are shown in Fig. 5: One option is to include everything in one single subsimulator, but this makes changes difficult or impossible to perform for users not having access to the subsimulator im-

plementation. By splitting the total system into several parts, the modularity will increase, but so will the complexity.

Where to draw the line is very much up to the designer, but it should be motivated by what is under investigation. Typically, a higher level of modularity makes it easier to include individual models of high fidelity. Therefore, a seemingly good rule-of-thumb is to keep modularity and fidelity high in parts of the system that are of interest, while lowering them for the rest of the system. This is exemplified by Fig. 6 which shows a model with high modularity in the power system. The entire remainder of the ship is captured inside only one subsimulator, by contrast, and its only purpose is to provide ‘good-enough’ dynamics to give studies of phenomena in the power system the necessary embedding. Keep in mind, however, that there needs to be a trade-off between this modeling modularity on one hand, and accuracy and stability on the other: as discussed in Sec. 3.1, higher modularity may, generally, *decrease* the overall accuracy for the part of the system under investigation, and, therefore, for the entire co-simulation.

Guideline 5 (System reticulation). *Try to obtain a good balance between modularity, complexity, accuracy, and numerical stability when splitting up a given system for co-simulation. Beware of time delays between subsimulators, and consider that accuracy may suffer with increasing modularity. At the same time, try to provide a sufficient level of modularity to facilitate interoperability and re-usability of models.*

Remark 5.1. The way a given system is split up for co-simulation generally affects accuracy and stability. An ill-chosen system reticulation can cause a significant deterioration of the quality of the co-simulation results, or even an entirely unstable simulation.

As an example, consider the two different system reticulations for the simple linear quarter car benchmark model studied in Ref. [11]: Mean global errors can easily vary by a factor of ten between different reticulations, and relatively large (constant) macro step sizes, which cause no issues with one reticulation, lead to instability with the other. Another practical example is given by the development of a Dynamic Positioning (DP) controller for ViProMa. The two main parts of such a controller are a high-level motion controller and a *Thrust Allocation Algorithm* (TAA). The initial DP controller design kept the high-level motion controller and the TAA in two separate subsimulators. This seemed like a good idea from a modularity perspective, because it allowed to easily swap out or modify either of the two independently. However, experience showed that the additional macro time step delay introduced by splitting the DP controller this way made the two parts much more difficult to tune, and fragile to changes in the macro time step size. Putting both, the high-level motion controller and the TAA, inside one subsimulator greatly improved the tunability and robustness.

5.3 Tight Coupling

Another important issue, which kept resurfacing when establishing the VPF, is how to sensibly deal with tightly-coupled subsystems. Consider, for example, the rigid mechanical connection between a vessel’s hull and crane. Ideally, such a connection calls for solving the hull–crane system as one, and the straight-forward (explicit) co-simulation of both as separate subsimulators with separate solvers is unfeasible. In other words, it may be best to simply refrain from splitting tightly-coupled systems for co-simulation altogether.

Guideline 6 (Tight coupling). *Avoid exposing tight couplings on the system level whenever possible in order to minimize coupling errors and avoid issues with numerical stability.*

This, however, challenges the general and modular virtual prototyping and full-system simulation approach that is so desirable from a practical point of view. It may still be possible to allow for sufficient flexibility in a large number of applications if generic models are used that offer sufficient parameterization. But in some cases it may simply be impractical to avoid exposing tight couplings between subsimulators, and several strategies have been proposed as solutions:

1. The most common approach is to include dampers and springs in one of the models. This works fine in practice, but has some notable drawbacks: Firstly, suitable parameters need to be found to configure these additional elements. This, secondly, means that the relatively stiff springs and strong dampers required introduce relatively small time constants, which can easily become a challenge for co-simulation.
2. A possibly very fruitful approach is to implement tightly-coupled subsystems using *model exchange*. Then, the models can be shared and modularly integrated, while they are being solved jointly by one solver. While this generally still satisfies the black box criterion, it may mean abandoning well-established simulation tools for some users. It could also prove difficult to find appropriate *open* software solutions enabling the use of model exchange in a co-simulation environment.
3. Yet another way of dealing with rigid coupling is to use advanced iterative co-simulation approaches utilizing subsimulator outputs along with their respective Jacobians. An example of this is the *Interface Jacobian-based Co-Simulation Algorithm* [29], in which coupling conditions are solved iteratively with Newton’s method. Unfortunately, such approaches are presently mainly of academic interest: generally requiring Jacobians from subsimulators—along with their ability to redo macro time steps—is simply an unrealistic condition short-term and medium-term.⁴
4. Finally, one can neglect the effects of one of the subsystem on the other. For example, a hull simply dictates the position of a crane attached to it, while the forces

⁴In the long run, however, such co-simulation techniques are very promising, and will hopefully become the status quo.

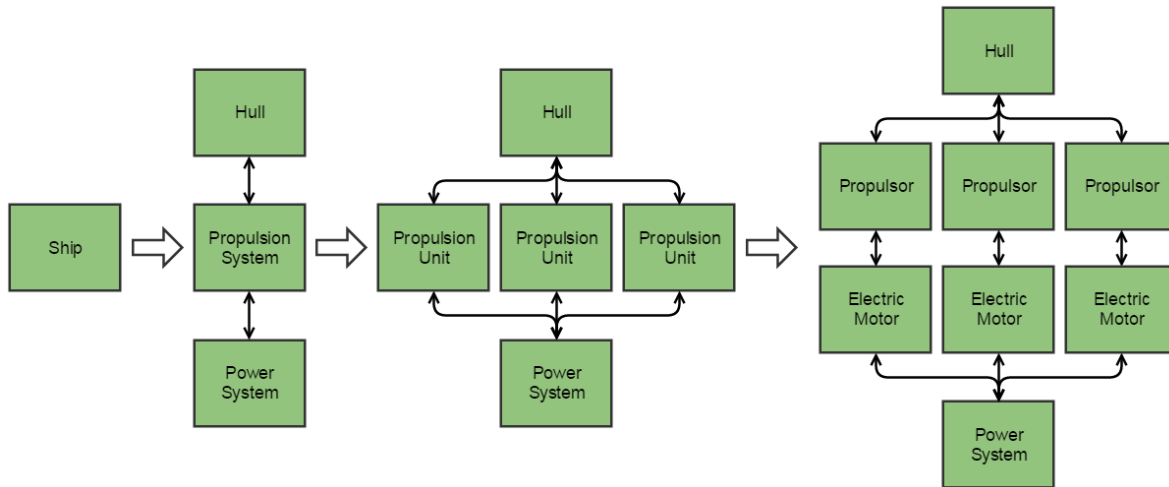


Fig. 5. Different levels of modeled modularity of systems on board a ship

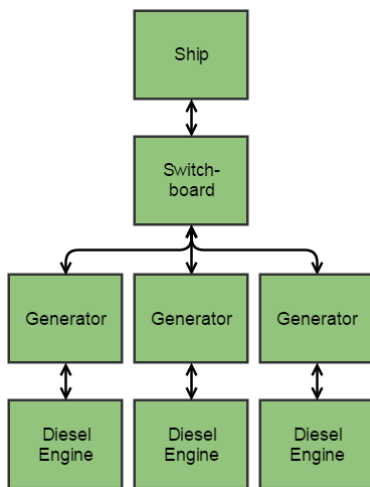


Fig. 6. Example of system modularization for a ship model with a special focus on power-system dynamics

sent back to the hull are only due to the crane's inertia. This approach is only sufficient, however, if an accurate representation of the forces is not desired.

As reflected by these choices, tight coupling is a complex and sensitive issue for virtual prototyping and co-simulation, and requires further research. For the time being, the best course of action will be determined by the case at hand, and can hopefully be found with the examples given here.

5.4 Connections

Connections define the interactions between subsimulators and are intimately interrelated with the system reticulation. Therefore, they need to be treated and implemented carefully. Mathematically, connections are expressed via the connection graph matrix L of Eq. (1), but in a co-simulation they are only enforced at discrete communication time instances. This, naturally, creates challenges for accuracy and stability, as discussed in Sec. 3.1.

As we have just seen, system boundaries need to be chosen sensibly, as to not compromise modularity, simplicity, accuracy, and numerical stability. At the same time, factual interactions between parts of a system and with the environment need to be captured accurately, given their relevance to the use case at hand, available computational capabilities, and modeling abilities and knowledge. Physical connections between subsystems are implemented as high-level interfaces between subsimulators, as discussed previously in Sec. 4.1. In doing so, ensuring the compatibility and correct wiring between the corresponding subsimulator inputs and outputs is of great importance. Let us elaborate on this point in the next two Sections 5.5 and 5.6.

5.5 Function Units

In order to connect the output of one subsimulator to the input of another, it may be necessary to apply some simple transformation in between. For example, in order to match an output's unit to that of an input, a conversion from rad/s to rpm could be needed, or a simple summation, to satisfy Kirchhoff's circuit laws or compute the net force acting on a body (see Fig. 7). Another example includes cases where the number of connections for a given subsimulator is variable or unknown.

In order to maintain modularity and ease-of-use, such transformations should not take place inside the subsimulators, nor should they be taken care of by a master, or hard-coded into the signal routing. One solution is to deploy a subsimulator between the others, dedicated to translating its input into an appropriate output form. While this does work, it has the unfortunate drawback in that it leads to an additional time delay between the original output and input the length of one macro time step.

To avoid this, the theoretical concept of so-called *Function Units* (FUs) was introduced. These act very much like subsimulators, but they are *time independent* and perform their calculations *in between* macro time steps, at communication points. Transformations between inputs and outputs

are, thus, carried out efficiently between time steps, can be easily applied in a modular and flexible fashion, and do not introduce unwanted time delays. In addition, Function Units are also well suited to represent dynamics on timescales that are much shorter than the ones that are relevant for the problem being studied. For example, in a simulation of a ship and its propulsion machinery, many electrical infrastructure components, such as switchboard breakers, would be ideal candidates for Function Units.

Guideline 7 (Function Units). *Use dedicated time independent submodules for unit conversions, coordinate transformations, signal algebra, and other generic operations between subsimulators during a simulation, as well as to include (approximately) instantaneous phenomena. This ensures modularity and ease-of-use, and avoids unnecessary time delays.*

Mathematically, this concept corresponds to a generalization of the connection graph matrix of Eq. (1), such that

$$\mathbf{u}(t_i) = \mathbf{L}(\mathbf{y}(t_i), t_i), \quad (2)$$

where \mathbf{L} is now, in general, a function of the outputs \mathbf{y} and the time. Common linear transformation—such as unit conversions, summations, and coordinate transformations—may still be simply expressed in the form given by Eq. (1), however.

It needs to be emphasized at this point that the concept of Function Units is mainly a theoretical one, and is still awaiting proper practical verification. Yet, their versatility and independence on time makes them essentially indispensable for constructing complex full-system models from reusable stand-alone subsimulators: virtually any type of code can be put inside an FU, and changes can quickly and easily be made on the system level *after* all the subsystem modeling is done.

5.6 Hybrid Causality

The computational causality of a subsystem also plays an important role in defining the connectivity with other subsystems. In some cases, the causality and, thereby, the connectivity of a subsimulator are difficult to determine beforehand without knowledge of the connecting environment.

To illustrate this causality–connectivity dependency, consider a simple mass–damper–spring system. Using integral causality, it may be written as

$$\begin{aligned} \dot{x}_1 &= x_2, \\ \dot{x}_2 &= \frac{1}{m}(\tau - dx_2 - kx_1), \end{aligned} \quad (3)$$

where x_1 and x_2 are the position and the velocity of the system, respectively, m is the mass, d is the damping coefficient, k is the spring stiffness, and τ is an excitation force given to the system as an input. If power bonds are used to connect

to other subsimulators in accordance to Guideline 2, the output of the mass–damper–spring system is given by x_2 . But Eq. (3) can also be written using differential causality, such that

$$\begin{aligned} \dot{x}_1 &= v, \\ \tau &= m \frac{dv}{dt} + dv + kx_1. \end{aligned} \quad (4)$$

Input and output have been switched, with the velocity v now constituting the input, while the output is given by the force τ . Both these causality options, and, thereby, connectivity options, may be equally relevant, depending on the subsimulator environment in the full system. Note that one state was lost from switching from integral to differential causality.

Different strategies exist to resolve conflicts in computationally causality between subsimulators: The model in question can simply be implemented with a specific causality and connectivity. This is a quick and easy fix, but it is dependent on the full system specifics, and not very attractive from a modularity point of view. Alternatively, the interface constraint equations (1) can be solved iteratively. As mentioned in Sec. 3.1, it is often preferred to avoid the repetition of entire co-simulation time steps, however, and to employ explicit co-simulation schemes instead. This leaves us with the final option of implementing the subsimulator as a *hybrid causality model*, that is, with the possibility to switch between different causality options—in some cases even on-line during a simulation.

Guideline 8 (Hybrid causality). *For subsimulators where the preferred causality–connectivity option is difficult or impossible to determine without prior knowledge of the connecting environment, the subsimulators may be implemented as hybrid causality models to ensure compatibility with other subsimulators in the co-simulation environment.*

Remark 8.1. There are a few subtleties when switching causality on-line during a simulation: For one, the state space changes, which either requires a solver that is suited for both systems, or the switching of the solver along with the model. Secondly, initial conditions need to be chosen carefully in order to not violate the conservation of energy, and to avoid discontinuities in the input and output signals.

Hybrid causality models can be implemented in several different ways. For example, two models with different causality options can be implemented in parallel as one subsimulator and switched between. The switching itself can depend on logical choices, and active and passive inputs and outputs. Note that the submodel with differential causality may need to be solved by a more advanced numerical solver. A different approach [13] is the use of a low-pass filter—effectively acting as an integrator—to regain the state that is lost when switching from integral to differential causality, as in the case exemplified by Eqs. (3) and (4). This method keeps the number of states in the system constant and removes any need for iterations.

A typical application of hybrid causality models within the realm of maritime systems is the modeling of weak marine power grids containing more than one energy source (such as generators, batteries, or fuel cells). Because only one source can determine the power grid voltage, the sources need to be implemented as hybrid causality models, in general. This is especially true if they are to be added and removed from the grid during the simulation [30].

5.7 Model Examples

Now that we have set the stage, let us demonstrate the application of the concepts discussed so far by use of examples from the ViProMa project. An extensive presentation and discussion of these use cases is given in Refs. [31, 32]. Here, we will focus on their key points in light of the previous discussions.

5.7.1 Vessel Model

First, consider the system shown in Fig. 7, with a hull model connected to propulsor models and a model of the environment. In addition, a crane model is shown to illustrate the fact that there may be multiple subsystems acting on the hull with a significant force. All of these are realized as subsimulators (green). Additionally, FUs (blue) take care of some simple signal algebra and additional time independent calculations.

The power bonds connecting the crane and trawl models to the hull are realized as *force-velocity* variable pairs, where the velocity is a part of the hull output *states*. The propulsor models here are RPM controlled, and so the variable pairs constituting the corresponding power bonds are of the type *torque-angular velocity*. They are thought to be connected to an actuator which drives them (e.g. el-motor or diesel engine, not shown). All propulsor subsimulators are also connected to an FU labeled *Additional propulsor calculations*. Its function is to sum up all the forces, and to perform additional calculations to determine the thrust loss for each propulsor individually. Another FU, labeled *External hull forces*, sums up all the forces acting on the hull, and applies the resulting net force to its center of gravity. In order to do so, it also requires knowledge about the various points of attack for the forces, which it receives via the additional *position* signals. All of these calculations are time independent and, thus, are conveniently put inside of Function Units. Doing so also brings about all the benefits of FUs discussed in Sec. 5.5: a flexible number of inputs and outputs (to connect any number of subsimulators), the possibility to quickly modify the signal algebra (to alter the thrust loss calculations, for example), and (quasi-)instantaneous execution without any delay in logical time.

Note that the dynamic interaction between the hull and the crane is considered tightly coupled, and has been discussed in detail in Ref. [33] in which the hull-crane system has been modelled within one subsimulator, in accordance with Guideline 6. Alternatively, the hull-crane interaction can be exposed on the co-simulation level [13] (as shown in Fig. 7), keeping in mind the discussions from Sec. 5.3.

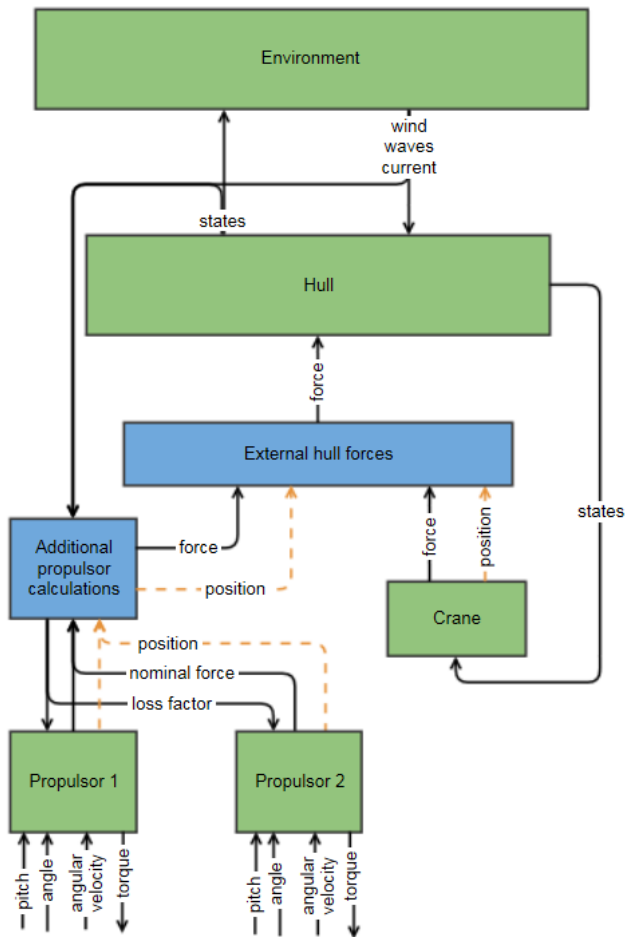


Fig. 7. Illustration showing how propulsors, crane, trawl, and environment subsimulators have been connected for ViProMa. The system includes subsimulators (green) as well as Function Units (blue).

The current state of environment modeling—specifically waves—in the ViProMa project is that principal data, such as wave height, wave period, and wave spread, are specified for each dependent subsimulator individually. [32, 34] In the present example, this is the case for the *Hull* and *Additional propulsor calculations* subsimulators, both of which use the same underlying implementation to realize the wave spectrum from that data. In the future, we would like to centralize wave spectrum construction in one location, and pass it on to subsimulators as required. [31] This is challenging, however, because the amount of data for wave realization can be as large as 16384 wave components.

5.7.2 Power Plant

A central component of an offshore vessel is its power plant which, for example, supplies the vessel's auxiliary and propulsion systems. An example of such a power plant model, which connects *Diesel Electric Generator*, *Switchboard*, and *EL-motor* models, is shown in Fig. 8. The *frequency* signals between the diesel electric generator models and the switchboard are necessary because these specific models are using the direct-quadrature-zero transformation

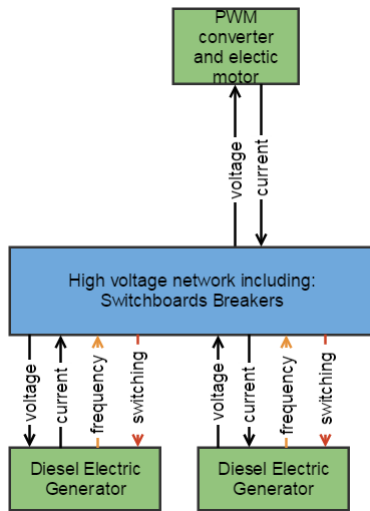


Fig. 8. Illustration showing how a diesel electric generator, a switchboard, and an electric motor subsimulator have been connected in ViProMa. The system includes subsimulators (green) as well as Function Units (blue).

(dq0), and the *switching* signals are used to swap the (computational) causality of the models. As explained in Sec. 5.6, the latter is needed because only one generator model can set the voltage for the switchboard in a weak power grid [35] when they are coupled in the manner depicted. The FU in this example represents a switchboard, which acts as a summer or splitter for the current or voltage, respectively, and also captures simple circuit breaker functionalities. In some cases, a physically more accurate model of a breaker may be needed and a subsimulator deployed instead, but the breaker dynamics are commonly discretized as quasi-instantaneous events and have no effect on the rest of the system (provided they are power conserving).

5.7.3 HI(cs)L: Hardware in the Co-Simulation Loop

An important use case from an industrial perspective is the direct connection between simulation and hardware for design, proofing, and stress testing. In Refs. [32, 30], an Arduino UNO microcontroller [36] was connected via FMI to a co-simulation to act as a DP controller and control the position of a simulated offshore vessel in operation. In addition to the microcontroller, a model-predictive control-based thrust allocation algorithm [37] was included to obtain realistic power consumptions and dynamical responses.

Fig. 9 illustrates the setup of the entire simulator whose aim was to study the effects of control system tuning. It included a vessel model, a position reference system, the thrust allocation algorithm, a power plant, the hardware DP controller with its FMI interface, a wave filter (filtering out most of the oscillatory wave forces), propulsors, and electrical motors controlled by simple local PID controllers. The results from this case study [31, 32] demonstrated that the tuning of the DP controller and the thrust allocation algorithm had a substantial influence on the power consumption.

6 Simulation Software

To support and demonstrate the use of the Virtual Prototyping Framework, the ViProMa project has developed *Coral*, a co-simulation software built from the ground up with FMI support and all the requirements described in section 2 in mind. Being designed for FMI, it has the same master/slave structure that we described briefly in section 4.2.

This software has two primary responsibilities: The first is *communication*, in that it transports data between subsimulators, possibly over a network, ensuring that output values are routed to the correct input variables. The second is *synchronization*, in that it issues commands to all the subsimulators that tell them when to perform a new time step, and how far to simulate before reaching the next communication point.

Since Coral supports distributing simulations over a network, it becomes necessary to actually start the slave programs on each computer when a new simulation is run. This is handled by a small server program called a *slave provider*, which runs on each of the machines that are set up to participate in simulations. This program is responsible for loading the FMUs available on that machine, publishing information about them on the network, and spawning slaves at the request of a master. An FMU is a bit like a *class* in object-oriented programming terminology, in that it represents a ‘blueprint’ for a model, and several instances of that model—slaves, in other words—can usually be created from one FMU. Each such instance then typically has its own state, as well as its own inputs and outputs.

Currently, the Coral slaves simply log their own simulation results directly to file. However, in the near future, a new type of simulation entity called *observer* will be added to the system. Unlike slaves, observers are privy to a lot of information about the structure and state of the simulation: which units are on-line, their inputs and outputs, and so on. Observers have no output values; in fact they have no way of affecting the course of the simulation at all. Typical examples of systems which could be implemented as observers include visualization systems and data loggers.

Coral is implemented as a software library for the C++ programming language, so it can be embedded in programs that need to perform co-simulations. It also comes with a set of command-line tools that allow users to run simple simulations which are configured via text files, and which double as examples that demonstrate how the C++ APIs may be used. See Fig. 10 for an overview of the software structure. The software has been released under a permissive open-source license, and is available for download from the project’s web site [6].

One may rightfully ask why it was deemed advantageous to create a new co-simulation software from scratch, rather than use an existing one, such as, for example, one of the many HLA implementations. The answer is that we were unable to find an existing software that fulfilled all the requirements described in section 2 to our satisfaction. In addition, a goal of the ViProMa project was to research and develop novel simulation methods and technologies, and in this respect it is very useful to start with a blank slate and to

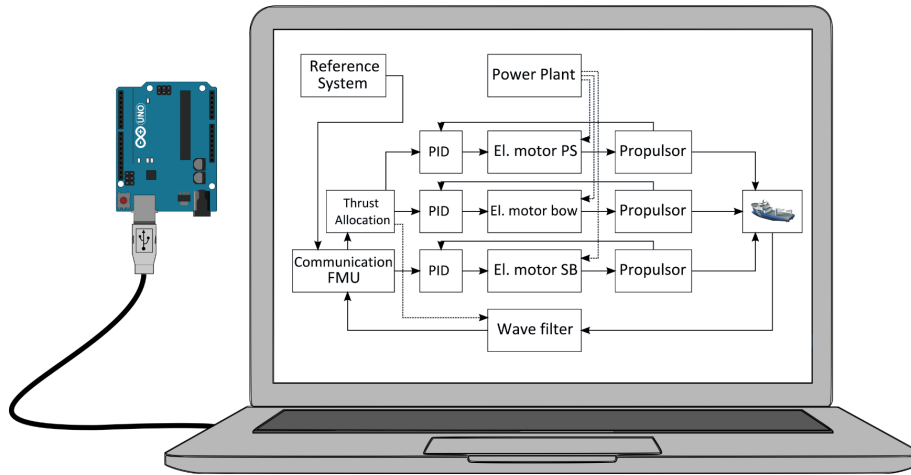


Fig. 9. Hardware-in-the-Loop co-simulation case study with an Arduino UNO microcontroller as a DP controller.

have a code base which is under our full control.

While Coral is still an immature piece of software, it already fulfils the main VPF requirements, and has been used with great success in several case studies [31, 32]. Currently, its most prominent deficiency with respect to the VPF is that the Function Unit concept has yet to be implemented. So far, this has been worked around by carrying out transformations and other time-independent operations within the subsimulators, but, as was pointed out in Sec. 5.5, this solution does not scale. Along with observers, FU support is on the top of the to-do list for future Coral development.

7 Conclusion

While it is widely recognized that new ship designs need to be optimized with respect to the overall operational performance—rather than the performance of components and subsystems—the simulation landscape is riddled with specialized tools for different physical and engineering domains and incompatible model representations. General software solutions, on the other hand, often fall short of providing the flexibility, efficiency, and accuracy needed for model construction in an increasingly competitive environment. To date, there are no universally adopted standards and tools supporting total systems integration and the analysis of operational performance.

In the present paper, we discussed the development of an open and standardized technology platform and infrastructure for virtual prototyping and full-system simulation for maritime systems and operations: the *Virtual Prototyping Framework* (VPF). Its vision is to facilitate the rapid development and sharing of subsystem and component models, and their joint simulation to assess full-system performance, optimize and verify designs, and establish a new arena for collaboration in the maritime industry. The biggest knowledge gaps and challenges were exposed, and solution strategies offered, with a focus on high interoperability, modularity, and re-usability. The application of the VPF, its guidelines, and its interfaces was demonstrated by use of model examples from the *Virtual Prototyping of Maritime Systems*

and Operations [6, 4] (ViProMa) project. Our in-house co-simulation software *Coral* was introduced, and original research towards establishing efficient, accurate, and stable co-simulation methods was discussed.

At the heart of the VPF lies a set of guidelines for model coupling, describing best practices for full-system model construction and simulation. In brief, these are:

Guideline 1. *Couple models and tools using co-simulation.*

Guideline 2. *Use power bonds to model relevant energy transactions between subsimulators whenever possible.*

Guideline 3. *Implement a co-simulation error estimation method.*

Guideline 4. *Use the Functional Mock-up Interface (FMI) to ensure interoperability and re-usability of subsimulators.*

Guideline 5. *Take great care when choosing a system reticulation, and try to balance modularity against complexity and numerical stability.*

Guideline 6. *Try to avoid tight couplings on the system level whenever possible.*

Guideline 7. *Consistently carry out generic signal operations between subsimulators using dedicated time independent modules.*

Guideline 8. *If the computational causality for a given subsimulator can not be determined a priori, implementation as a hybrid causality model is advised.*

These guidelines are by no means meant to be complete or even final. But their introduction into, and adaption by, the maritime industry would be a major step ahead, towards finding tomorrow's design and technology solutions. Note, however, that these guidelines are not specific to the maritime domain and should apply to a wide range of engineering and scientific applications.

The definition of guidelines was motivated by a set of core requirements, which were derived from the most important use cases—vessel design, crew training, and decision support—and the desired work flow for the prototyping

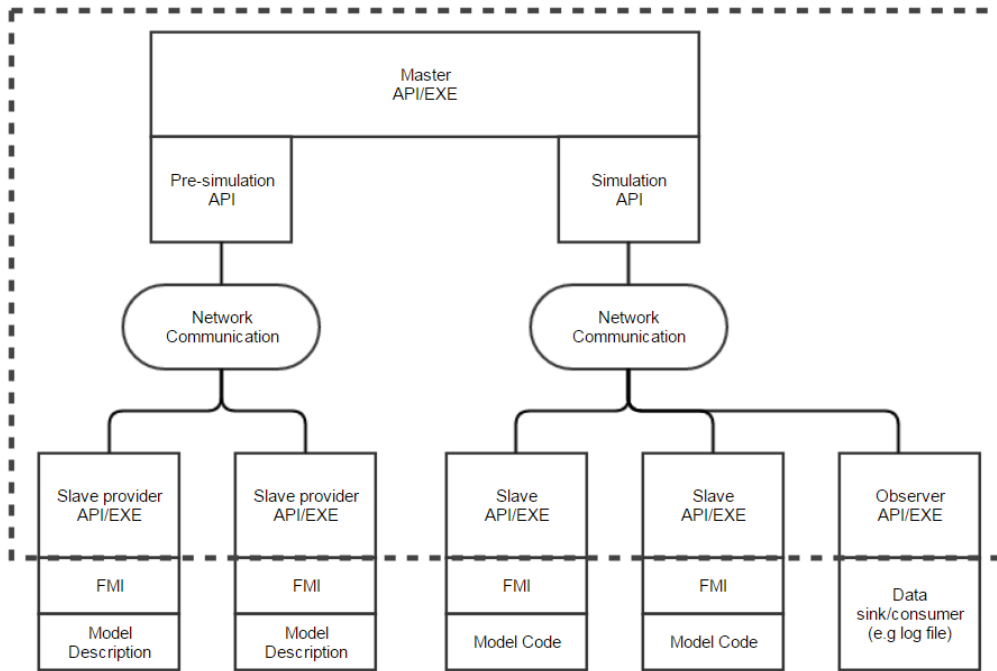


Fig. 10. A diagram that shows the various components in a Coral simulation. Everything inside the dashed rectangle is formally part of Coral. By *API/EXE* we mean that the functionality is offered both in the form of a C++ programming interface and as a ready-made executable application.

framework. As mentioned in Sec. 2, some of these had to be left to future research due to insufficient funding and time. Most notably, we fell short of implementing the following three aspects:

1. A basic open library layer of generic and, to a certain extent, parameterized domain models and components should be available for the VPF to configure ships and operations for simulation and optimization studies. This is also important for demonstration purposes, and to aid in the wide-spread acceptance and adaption of the framework. Creating such a database is an extensive undertaking, however.
2. Only some testing and research [30] was done with respect to real-time capabilities. While, in principle, there is nothing preventing the use of real-time simulations with the VPF as introduced in the present work, more research and verification is surely needed. At any rate, there are a few challenges that need to be considered: For example, computational efficiency can easily become an issue, because real-time subsimulators (such as hardware interfaces) will usually dictate a lower limit for the co-simulation step size. Moreover, network latency effects and noise in measurement signals have to be dealt with. These points can, naturally, also further exacerbate the co-simulation-inherent accuracy and stability issues discussed in Sec. 3.1. A promising project is the *Advanced Co-Simulation Open System Architecture* [38] (ACOSAR), which aims to supplement the FMI standard with a global industry standard for real-time system integration and co-simulation.
3. One important aspect of specifying high-level simu-

lator interfaces is the definition of standardized categories which represent various component and subsystem groups—such as engines or cranes—along with standardized model properties. This would facilitate rapid prototyping in a convenient ‘plug-and-play’ fashion, as mentioned in Sec. 4.

Hopefully, the present work can contribute to the advancement of virtual prototyping and flexible full-system simulation, provide a basis for further research, and have its aforementioned shortcomings amended through future work.

Acknowledgements

This work was funded by the Research Council of Norway (Grant Number 225322), and the industrial partners in the ViProMa project consortium (VARD, Rolls-Royce Marine, and DNV GL). We are grateful for their financial support.

References

- [1] Henry, G., Cox, I., Crossland, P., and Duncan, J., 2009. “Virtual ships: NATO standards development and implementation”. In *Proceedings of the NATO RTO Modelling and Simulation Group Symposium, Systems engineering and assessment LTD, Frome, United Kingdom*, pp. 20–2–20–16.
- [2] Harries, S., Tillig, F., Wilken, M., and Zaraphonitis, G., 2011. “An integrated approach for simulation in the early ship design of a tanker”. In *10th International Conference on Computer and IT Applications*

- in the Maritime Industries, Proceedings of the 10th International Conference on Computer and IT Applications in the Maritime Industries, Technische Universität Hamburg-Harburg, pp. 411–425.
- [3] Evans, J. H., 1959. “Basic design concepts”. *Journal of the American Society for Naval Engineers*, **71**(4), pp. 671–678.
- [4] Hassani, V., Rinderøy, M., Kyllingstad, L. T., Nielsen, J. B., Sadjina, S., Skjong, S., Fathi, D., Johnsen, T., Æsøy, V., and Pedersen, E., 2016. “Virtual prototyping of maritime systems and operations”. In Proceedings of the ASME 2016 35th International Conference on Ocean, Offshore and Arctic Engineering, The American Society of Mechanical Engineers, pp. V007T06A018–V007T06A018.
- [5] Fossen, T. I., 2011. *Handbook of Marine Craft Hydrodynamics and Motion Control*. John Wiley & Sons, Chichester, United Kingdom.
- [6] ViProMa Consortium, 2017. Official web site of ViProMa. <https://viproma.no>. SINTEF Ocean, Trondheim, Norway. Accessed 2018-04-27.
- [7] Clauß, C., Arnold, M., Schierz, T., and Bastian, J., 2012. “Master zur Simulatorkopplung via FMI”. In ASIM-Treffen der Fachgruppen ”Simulation technischer Systeme” und ”Grundlagen und Methoden in Modellbildung und Simulation”, ARGESIM-Verlag, pp. 57–70.
- [8] Kübler, R., and Schiehlen, W., 2000. “Two methods of simulator coupling”. *Mathematical and Computer Modelling of Dynamical Systems*, **6**(2), pp. 93–113.
- [9] Sicklinger, S. A., 2014. “Stabilized co-simulation of coupled problems including fields and signals”. PhD dissertation, Technische Universität München, Munich, Germany.
- [10] Arnold, M., 2009. *Numerical methods for simulation in applied dynamics*, Vol. 507 of *CISM International Centre for Mechanical Sciences*. Springer Vienna, ch. 5, pp. 191–246.
- [11] Sadjina, S., Kyllingstad, L. T., Skjong, S., and Pedersen, E., 2017. “Energy conservation and power bonds in co-simulations: non-iterative adaptive step size control and error estimation”. *Engineering with Computers*, **33**(3), pp. 607–620. arXiv:1602.06434.
- [12] Sadjina, S., and Pedersen, E., 2016. “Energy conservation and coupling error reduction in non-iterative co-simulations”. arXiv:1606.05168.
- [13] Skjong, S., and Pedersen, E., 2016. “The theory of bond graphs in distributed systems and simulations”. In 2016 International Conference on Bond Graph Modeling and Simulation, SIMUL SERIES, Society for Modeling and Simulation International, pp. 147–159.
- [14] IEEE Standards Association, 2010. “IEEE standard for modeling and simulation (M&S) high level architecture (HLA) – framework and rules”. *IEEE Std. 1516-2010*, pp. 1–38. IEEE, New York, New York, USA.
- [15] DIS Steering Committee, 2012. “IEEE standard for distributed interactive simulation–application protocols”. *IEEE Std. 1278.1-2012 (Revision of IEEE Std. 1278.1-1995)*, pp. 1–747. IEEE, New York, New York, USA.
- [16] Husteli, N., 2005. Common Simulation Interface. Documentation and Tutorial. Tech. Rep. 4-2005, MARINTEK, Trondheim, Norway.
- [17] Paynter, H. M., 1961. *Analysis and Design of Engineering Systems: Class Notes for M.I.T. Course 2.751*. M.I.T. Press, Boston.
- [18] Breedveld, P. C., 1984. “Physical systems theory in terms of bond graphs”. PhD dissertation, Twente University, Twente, Netherlands.
- [19] Benedikt, M., Watzenig, D., Zehetner, J., and Hofer, A., 2013. “NEPCE—a nearly energy preserving coupling element for weak-coupled problems and co-simulation”. In V International Conference on Computational Methods for Coupled Problems in Science and Engineering, Coupled Problems, International Center for Numerical Methods in Engineering, pp. 1021–1032.
- [20] Busch, M., 2012. “Zur effizienten Kopplung von Simulationsprogrammen”. PhD dissertation, Universität Kassel, Kassel, Germany.
- [21] Blochwitz, T., Otter, M., Arnold, M., Bausch, C., Elmqvist, H., Junghanns, A., Mauss, J., Monteiro, M., Neidhold, T., Neumerkel, D., Olsson, H., Peetz, J.-V., Wolf, S., and Clauß, C., 2011. “The functional mockup interface for tool independent exchange of simulation models”. In Proceedings of the 8th International Modelica Conference, Linköping Electronic Conference Proceedings, Linköping University Electronic Press, pp. 105–114.
- [22] Blochwitz, T., Otter, M., Åkesson, J., Arnold, M., Clauß, C., Elmqvist, H., Friedrich, M., Junghanns, A., Mauss, J., Neumerkel, D., Olsson, H., and Viel, A., 2012. “Functional Mockup Interface 2.0: The standard for tool independent exchange of simulation models”. In Proceedings of the 9th International Modelica Conference, The Modelica Association, pp. 173–184.
- [23] Modelica Association, 2012. Official web site of the Functional Mock-up Interface. <https://fmi-standard.org>. Modelica Association, Linköping, Sweden. Accessed 2018-04-27.
- [24] Awais, M. U., Palensky, P., Elsheikh, A., Widl, E., and Matthias, S., 2013. “The High Level Architecture RTI as a master to the Functional Mock-up Interface components”. In 2013 International Conference on Computing, Networking and Communications (ICNC), pp. 315–320.
- [25] Broman, D., Brooks, C., Greenberg, L., Lee, E. A., Masin, M., Tripakis, S., and Wetter, M., 2013. “Determinate composition of FMUs for co-simulation”. In 2013 Proceedings of the International Conference on Embedded Software (EMSOFT), pp. 1–12.
- [26] Feldman, Y. A., Greenberg, L., and Palachi, E., 2014. “Simulating Rhapsody SysML blocks in hybrid models with FMI”. In 10th International Modelica Conference, Linköping Electronic Conference Proceedings, Linköping University Electronic Press, pp. 43–52.
- [27] Tripakis, S., and Broman, D., 2014. Bridging the semantic gap between heterogeneous modeling for-

malisms and FMI. Report UCB/EECS-2014-30, EECS Department, University of California, Berkeley.

[28] Broman, D., Greenberg, L., Lee, E. A., Masin, M., Tripakakis, S., and Wetter, M., 2015. “Requirements for hybrid cosimulation standards”. In *Hybrid Systems: Computation and Control (HSCC)*, ACM Press, New York, NY, USA, pp. 179–188.

[29] Sicklinger, S. A., Belsky, V., Engelmann, B., Elmqvist, H., Olsson, H., Wüchner, R., and Bletzinger, K. U., 2014. “Interface jacobian-based co-simulation”. *INT J NUMER METH ENG*, **98**(6), pp. 418–444.

[30] Skjong, S., and Pedersen, E., 2017. “Co-simulation of a marine offshore vessel in DP-operations including hardware-in-the-loop (HIL)”. In *ASME 2017 36th International Conference on Ocean, Offshore and Arctic Engineering (OMAE2017)*, American Society of Mechanical Engineers, ASME, pp. V07AT06A038–V07AT06A038.

[31] Skjong, S., 2017. “Modeling and simulation of maritime systems and operations for virtual prototyping using co-simulations”. PhD dissertation, Norwegian University of Science and Technology, Trondheim, Norway.

[32] Skjong, S., Rindarøy, M., Kyllingstad, L. T., Æsøy, V., and Pedersen, E., 2017. “Virtual prototyping of maritime systems and operations: applications of distributed co-simulations”. *Journal of Marine Science and Technology*, **Online First**, pp. 1–19.

[33] Rokseth, B., Skjong, S., and Pedersen, E., 2017. “Modeling of generic offshore vessel in crane operations with focus on strong rigid body connections”. *IEEE Journal of Oceanic Engineering*, **42**(4), Oct, pp. 846–868.

[34] Yum, K. K., Skjong, S., Tasker, B., Pedersen, E., and Steen, S., 2016. “Simulation of a hybrid marine propulsion system in waves”. In *Proceedings of 28th CIMAC World Congress*, p. 202.

[35] Skjong, S., and Pedersen, E., 2017. “A real-time simulator framework for marine power plants with weak power grids”. *Mechatronics*, **47**, pp. 24–36.

[36] Official web site of Arduino. <https://www.arduino.cc/>. Accessed 2018-04-27.

[37] Skjong, S., and Pedersen, E., 2017. “Nonangular MPC-based thrust allocation algorithm for marine vessels—a study of optimal thruster commands”. *IEEE Transactions on Transportation Electrification*, **3**(3), Sept, pp. 792–807.

[38] ACOSAR Project consortium, 2015. Official web site of the Advanced Co-Simulation Open System Architecture. <http://acosar.eu>. Kompetenzzentrum - Das Virtuelle Fahrzeug Forschungsgesellschaft, Graz, Austria. Accessed 2018-04-27.

List of Figures

1	Maritime systems and operations include a wide range of different engineering domains and physical systems with varying complexity and time scales. This, naturally, makes full-system simulation a challenging endeavor.	2
2	In a co-simulation setting, different tools and models are interconnected and used independently and in parallel to form a full-system simulation.	4
3	A residual power $\delta P_k = -(P_{k_1} + P_{k_2})$ emerges and distorts the dynamics of the full system when energy is exchanged between two subsimulators, S_1 and S_2 , in a co-simulation	7
4	Energy-conservation-based error estimation (red) compared to the actual error in the power ΔP (gray) as a function of the co-simulation step size Δt for the benchmark model in Ref. [11]. The critical step size is $\Delta t \approx 0.059\text{s}$	7
5	Different levels of modeled modularity of systems on board a ship	10
6	Example of system modularization for a ship model with a special focus on power-system dynamics	10
7	Illustration showing how propulsors, crane, trawl, and environment subsimulators have been connected for ViProMa. The system includes subsimulators (green) as well as Function Units (blue).	12
8	Illustration showing how a diesel electric generator, a switchboard, and an electric motor subsimulator have been connected in ViProMa. The system includes subsimulators (green) as well as Function Units (blue).	13
9	Hardware-in-the-Loop co-simulation case study with an Arduino UNO microcontroller as a DP controller.	14
10	A diagram that shows the various components in a Coral simulation. Everything inside the dashed rectangle is formally part of Coral. By <i>API/EXE</i> we mean that the functionality is offered both in the form of a C++ programming interface and as a ready-made executable application.	15