

Co-simulation at Different Levels of Expertise with Maestro2

Simon Thrane Hansen^{a,*}, Casper Thule^a, Cláudio Gomes^a, Kenneth Guldbrandt Lausdahl^a, Frederik Palludan Madsen^a, Giuseppe Abbiati^b, and Peter Gorm Larsen^a

^aDepartment of Electrical and Computer Engineering, Aarhus University, Åbogade 34, Denmark

^bDepartment of Civil and Architectural Engineering, Aarhus University, Inge Lehmanns Gade 10, Denmark

Abstract

When different simulation units are coupled together there are different choices to take, in particular regarding the granularity of such a co-simulation. When prototyping systems, it is typically favourable to get an initial idea of how a collection of simulation units work together without spending too much time setting up the orchestration. However, the granularity of such a simulation may be far away from what is needed in relation to the purpose of the simulation. In order to enable more flexibility and control over the co-simulation it is necessary to be able to steer the orchestration in a more detailed manner. This paper presents an open source co-simulation orchestration engine based on the Functional Mockup Interface standard but with a Domain Specific Language (DSL) enabling detailed control between the individual simulation units. The same tool can thus be used right out of the box for low-granularity co-simulation, and for high-granularity simulation the DSL enable a significant flexibility.

Keywords: co-simulation framework, domain specific language, functional mockup interface standard

1. Introduction

Co-simulation is a technique for simulating complex systems by combining multiple simulation tools into a single simulation [Kübler and Schiehlen \(2000b\)](#); [Gomes et al. \(2018b\)](#)

Interoperability between simulation tools is achieved through the use of Functional Mock-up Units (FMUs) defined by the Functional Mock-up Interface (FMI) standard [Committee \(2014, 2021\)](#). An FMU encapsulates a Simulation Unit (SU) by providing a standardised interface of inputs, outputs, and functions to let a *co-simulation framework* control the simulation of a coupled system of FMUs, referred to as a *scenario*. A scenario is obtained by coupling inputs and outputs of the FMUs in the scenario, as illustrated in [Figure 1](#). A coupling denotes that the output FMU's state influences the input FMU's state. [Figure 1](#) depicts a scenario with two coupled FMUs of a coupled mass-spring-damper, similar to a later example in [Section 4](#).

A co-simulation framework executes the scenario by computing the joint behaviour of the system by coordinating the execution of the FMUs in the scenario ac-




Figure 1: A co-simulation scenario with two SUs MSD1 and MSD2 representing a coupled mass-spring-damper system. The SUs are represented as rectangles, and the arrows f_1 and f_2 denote the connections between the SUs.

ording to an orchestration algorithm (OA). The OA describes how stimuli are exchanged between the FMUs in the scenario and how the state of the FMUs evolves over the course of the co-simulation. Although the OA is not part of the FMI standard, it is a critical component of a co-simulation framework, as studies have shown that the OA can significantly affect the accuracy of co-simulation results [Busch \(2016\)](#); [Kalmar-Nagy and Stanculescu \(2014\)](#); [Schweizer et al. \(2015\)](#); [Arnold \(2010\)](#); [Gomes et al. \(2018e\)](#); [Schweizer et al. \(2016\)](#); [Andersson \(2016\)](#); [Hansen et al. \(2021b\)](#). To obtain accurate co-simulation results, the OA must be tailored to the scenario and the characteristics of the FMUs it contains, as illustrated in [Figure 2](#), which compares the results of two different OAs for the scenario presented in [Figure 1](#) with its analytical solution. Both OAs are compliant with the FMI standard, but the results differ

*Corresponding author

Email address: sth@ece.au.dk (Simon Thrane Hansen)




Figure 2: Comparison of the results of two different OAs for the scenario in Figure 1 with its analytical solution. The step size used in the co-simulation is 0.1 s.

significantly.

Figure 2 shows that the OA significantly affects the accuracy of the co-simulation results and suggests that the OA should be carefully designed to minimise the error in a co-simulation. Similar results have been reported in Busch (2016); Kalmar-Nagy and Stanculescu (2014); Schweizer et al. (2015); Arnold (2010); Gomes et al. (2018e); Schweizer et al. (2016); Andersson (2016); Hansen et al. (2021b), where the authors show empirically that co-simulation results can be highly sensitive to the order in which the FMUs are simulated and how they exchange data.

The sensitivity can be attributed to several factors, but can be summarised as follows: The discrete nature of co-simulation, where data is exchanged between FMUs at discrete points in time, called communication points, can be challenging for FMUs representing continuous processes. Such FMUs typically rely on numerical solvers (variable step or fixed-step) to advance in simulated time and use a variety of approximation techniques, such as extrapolation and interpolation, to reason about the values of the FMU’s inputs between communication points Kübler and Schiehlen (2000a); Gomes et al. (2018b). The approximation techniques used by the FMUs impose constraints on the OA interactions, dictating the order in which the FMUs are simulated and how they exchange data to account for the characteristics of the FMUs Gomes et al. (2019b); Hansen et al. (2022b).

To address these challenges and cater to a wide range of application domains, a co-simulation framework should strike a balance between allowing experts to customise the OA and providing synthesised OAs to new users, while trying to minimise the overhead of or-

chestrating the co-simulation. This balance empowers experts to fine-tune the OA to minimise co-simulation error, while enabling new users to quickly start co-simulating without having to delve into intricate details.

Unfortunately, to the best of our knowledge, there is a lack of open source co-simulation frameworks that offer such flexibility without compromising on usability and performance. As a result, users who wish to customise/fine-tune the OA to accommodate the idiosyncrasies of the SUs, or researchers who wish to experiment with novel co-simulation algorithms, are left with having to develop their own OA from scratch using low-level co-simulation libraries. Developing an OA for a large scenario is a time-consuming task that requires the user to spend many hours laying the groundwork before they even reach the point where they can tune their co-simulation, as the user must code all interactions with the SUs.

In summary, there is a need for a co-simulation framework that provides flexibility while maintaining customisability, expressiveness, verifiability, and speed. Such a framework should strike a balance between providing experts with customisable OAs and providing new users with synthesised OAs, allowing different levels of granularity for co-simulation practitioners at all levels of expertise.

Contribution. To address the above issues, we propose an open source co-simulation framework called Maestro2, which leverages the latest advances in OA synthesis Gomes et al. (2019b); Hansen et al. (2022b) to enable rapid development of customisable OAs and use code generation to enable high-performance co-simulations. Specifically, Maestro2 provides different levels of granularity to describe the co-simulation scenario and the OA, ranging from a high level of granularity suitable for prototyping and new users, to a lower level DSL that describes the OA in detail, suitable for experts fine-tuning the OA to accommodate the idiosyncrasies of the FMUs.

Thus, the main contribution of this manuscript is the co-simulation framework Maestro2, which enables the different levels of granularity to support co-simulation practitioners with different levels of expertise. The manuscript also presents the results of the application of Maestro2 in two case studies.

Prior Work. This manuscript represents the complete implementation of the idea originally proposed in Thule et al. (2020) to maximise reuse among co-simulation frameworks. Since then, Maestro2 has been used in several case studies and research projects (see Section 4).

124 Among other, Maestro2 has been used to experiment 171
 125 with different approaches for handling algebraic loops 172
 126 during the initialisation of a co-simulation Hansen et al. 173
 127 (2021c). While the above works introduce Maestro2, 174
 128 none of them introduces the Maestro2 approach in 175
 129 detail, showcasing how it empowers users of different lev- 176
 130 els of expertise and with different needs to perform co- 177
 131 simulations. The Maestro2 approach is the main contri- 178
 132 bution of this manuscript.

133 *Structure.* The remainder of this manuscript is struc- 181
 134 tured as follows: The next section introduces the main 182
 135 concepts used throughout the manuscript, including the 183
 136 FMI standard, the concept of co-simulation, and the 184
 137 concept of orchestration algorithms. Then, Section 3 185
 138 presents the Maestro2 approach and how it enables dif-
 139 ferent levels of granularity. Section 4 summarises a case
 140 study where Maestro2 has been applied and Section 5
 141 discusses related work. Finally, Section 6 details future
 142 work and concludes the manuscript.

143 2. Background

144 This section serves as an introduction to the funda-
 145 mental concepts of co-simulation, the FMI standard,
 146 and OAs, and provides an overview of the problem
 147 we aim to address. However, due to the breadth
 148 and complexity of co-simulation and the FMI stan-
 149 dard, and the interested reader is referred to Gomes
 150 et al. (2018b,d) for a comprehensive introduction to
 151 co-simulation and Blochwitz et al. (2011); Committee
 152 (2014); Gomes et al. (2021b) for the FMI standard.
 153 The notation and definitions introduced in this section
 154 are adopted from Gomes et al. (2019a); Hansen and
 155 Ölveczky (2022).

156 2.1. Co-simulation

157 Co-simulation is a technique that enables the global
 158 simulation of a system composed of several black-box
 159 SUs, typically developed individually or exported from
 160 different tools (Kübler and Schiehlen (2000b); Gomes
 161 et al. (2018b)). An SU models the behaviour of a *dy-*
 162 *namical system* consisting of inputs and outputs, the state
 163 of the system, and a set of functions to provide stimuli
 164 to the system (by setting the inputs), to retrieve the state
 165 of the system (by getting the outputs), and to evolve the
 166 state of the system in simulated time (by stepping the
 167 model).

168 The evolution of an SU is governed by a set of evolu-
 169 tion rules described by differential and algebraic equa-
 170 tions that define how the state of the system changes

in response to stimuli and the current state of the sys-
 tem. The behaviour trace of an SU is a function that
 maps time to state, and the evolution rules are typically
 obtained using a numerical solver that discretizes the
 continuous-time model of the SU into a discrete-time
 model (a trace that maps time to state), allowing the
 simulation of the SU in discrete time steps.

A formal definition of an SU is given in Definition 1,
 where the evolution rules are captured by the function
 doStep_c , which advances the state of the SU in simu-
 lated time. In addition, the function doStep_c returns a
 step size to accommodate those SUs that use numerical
 solvers with variable step lengths or implement error es-
 timation mechanisms. These SUs may conclude that a
 step size of H will result in an intolerable error.

Definition 1. An SU with identifier c is represented by
 the tuple $\langle S_c, U_c, Y_c, \text{set}_c, \text{get}_c, \text{doStep}_c \rangle$, where:

- S_c represents the state space.
- U_c and Y_c the set of input and output variables,
 respectively.
- $\text{set}_c : S_c \times U_c \times \mathcal{V} \rightarrow S_c$ and $\text{get}_c : S_c \times Y_c \rightarrow \mathcal{V}$
 are functions to set the inputs and get the outputs,
 respectively (we abstract the set of values that each
 input/output variable can take as \mathcal{V}).
- $\text{doStep}_c : S_c \times \mathbb{R}_{>0} \rightarrow S_c \times \mathbb{R}_{>0}$ is a function
 that instructs the SU to compute its state after a
 given time duration. If an SU is in state $s_c^{(t)}$ at time
 t , $(s_c^{(t+h)}, h) = \text{doStep}_c(s_c^{(t)}, H)$ approximates the
 state $s_c^{(t+h)}$ of the corresponding model at time $t+h$,
 with $h \leq H$.

A collection of SUs can be coupled to form a co-
 simulation scenario by connecting the outputs of one
 SU to the inputs of another SU (see Definition 2). A
 coupling means that the state of one SU always depends
 on the state of another SU - this is called a *coupling*
restriction and can be thought of as a system invariant,
 which says that the values of the coupled inputs and out-
 puts must always be equal. Nevertheless, the coupling
 restrictions are only satisfied at specific points in time,
 called *communication points*, when the SUs exchange
 data. The SUs try to compensate for the inconsistency
 between the communication points by making assump-
 tions about the evolution of the values on their inputs.
 Obviously, these assumptions can cause a significant er-
 ror in the co-simulation for large intervals between the
 communication points. In fact, they can be the main
 source of error (Arnold et al. (2014)), so it is essential
 to fine-tune the OA to account for the characteristics of
 the SUs. The characteristics of these approximation func-
 tions are captured by the function R , see Definition 2.

221 The function R links each input to indicate whether the
 222 input SU expects the coupled output to be simulated be-
 223 fore or after the input SU itself.

224 **Definition 2 (Scenario).** A scenario is a structure
 225 $\langle C, L, R \rangle$, where

- 226 • C is a finite set (of SU identifiers).
- 227 • L is a function $L : U \rightarrow Y$, where $U = \bigcup_{c \in C} U_c$
 228 and $Y = \bigcup_{c \in C} Y_c$, and where $L(u) = y$ means that
 229 the output y is coupled to the input u .
- 230 • $R : U \rightarrow \mathbb{B}$ is a predicate, which describes the
 231 SUs' input approximation functions. $R(u) = \text{true}$
 232 means that SU c expect the SU d of the output y
 233 coupled to u to be simulated before c . Similarly,
 234 $R(u) = \text{false}$ means that SU c expect the SU d of
 235 the output y coupled to u to be simulated after c .

236 To illustrate the concepts introduced in [Definitions 1](#)
 237 and [2](#), consider the scenario in [Figure 3](#) and the cor-
 238 responding behavioural trace shown in [Figure 4](#). The
 239 scenario consists of two SUs, a controller SU and
 240 a tank SU, and two couplings, one from the output
 241 valve state of the controller SU to the input valve
 242 state of the tank SU, and one from the output water
 243 level of the tank SU to the input water level of the
 244 controller SU. Each SU has some parameters (e.g. max
 245 level and min level) which are used to configure the
 246 simulation. The function R is omitted from the scenario,
 247 as it is not part of the FMI standard, and is discussed in
 248 [Section 2.2](#).

249 The controller SU is a simple controller that opens
 250 or closes a valve based on the current water level, and
 251 the tank SU is a simple tank that keeps track of the cur-
 252 rent water level based on the flow of water in and out
 253 of the tank through the valve. The behaviour trace of
 254 the scenario is in [Figure 4](#) and is obtained by simulating
 255 the scenario from time 0 to 10 seconds. The trace is the
 256 function σ that maps time to the state of the scenario,
 257 i.e. $\sigma(t) = \langle s_c^{(t)} \mid c \in C \rangle$.

258 The *co-simulation error* is the difference between the
 259 simulated behaviour of the scenario and its ideal be-
 260 haviour, i.e. the behaviour obtained by simulating the
 261 scenario with an infinitely small step size or by solving
 262 the differential equations analytically, which is gener-
 263 ally not possible.

264 The simulation of a scenario is controlled by the OA,
 265 which is the algorithm coordinating the execution of the
 266 SUs in the scenario to obtain the joint behavioural trace
 267 of the system. The OA comprises multiple stages, in-
 268 cluding everything from loading the SUs to terminating
 269 the simulation and performing the co-simulation by in-
 270 voking the set_c , get_c , and doStep_c functions defined

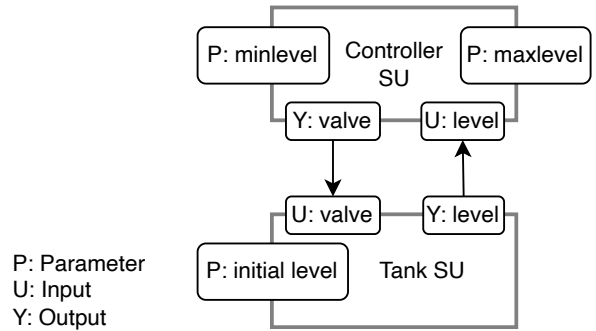


Figure 3: The the co-simulation scenario of the water tank example illustrated as a block diagram. The scenario is adapted from [Mansfield et al. \(2017\)](#).

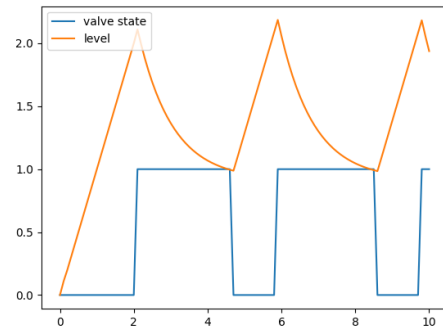


Figure 4: The behavioural trace of the Water Tank example shown as a plot of the water level and valve state over time for a simulation from 0 to 10 seconds with a step size of 0.1 seconds. The scenario is adapted from [Mansfield et al. \(2017\)](#).

271 in [Definition 1](#). The stages of a typical OA are sum-
 272 marised in [Figure 5](#), the colours are used to visually
 273 distinguish the three overarching phases of the OA: *ini-*
 274 *tialisation*, *simulation*, and *termination*. The initial set
 275 of phases *Start*, *Instantiate*, *Setup* and *Initialise* are ex-
 276 ecuted once before the simulation starts and are respon-
 277 sible for loading the SUs, creating instances, setting ini-
 278 tial parameters, and computing the initial state of the
 279 SUs, respectively. The initialisation is followed by the
 280 simulation loop, which consists of the *Step* and *Plotting*
 281 stages, which are responsible for advancing the SUs in
 282 simulated time and reporting results. Finally, the stage
 283 *Free* releases resources and terminates the simulation.

284 Many articles [Broman et al. \(2013a\)](#); [Hansen et al.](#)
 285 [\(2022b\)](#); [Gomes et al. \(2018a\)](#) on OAs do only consider
 286 the *Initialise* and *Step* stages, as these are the most crit-
 287 ical for the correctness of the co-simulation results and
 288 constitutes the phases where the OA interacts with the

289 SUs using the set_c , get_c and $doStep_c$ functions defined in Definition 1. These phases are concerned with
 290 satisfying the coupling restrictions and ensuring that the
 291 SUs move in lockstep, i.e. that the SUs are synchronised
 292 with respect to simulated time.
 293




Figure 5: Generic OA structure. Each round rectangle represents a stage in the execution of the OA. For instance, the plotting stage will query the outputs of the SU and record them in a CSV file. Adapted from Thule et al. (2020).

294 For a better understanding of the OA stages *Initialise* and *Step*, consider the scenario in Figure 3. The
 295 parameters (minimum water level, maximum water level and initial level) of the SUs are set in the
 296 *Initialise* stage, after which the initial state of the SUs is calculated by calling the functions `get` and `set` to
 297 get the valve state output from the controller SU and set it to the valve state input on the water level SU. Then
 298 the water level output from the tank SU is assigned to the water level input on the controller SU. The simulation
 299 then begins. The outputs are retrieved and logged at time 0 in the *Plotting* step. The *Step* stage uses an algorithm
 300 similar to the algorithm shown in Algorithm 1 to compute new states for the SUs by calling the `doStep`
 301 function and to exchange data between the SUs by calling the `get` and `set` functions.
 302
 303
 304
 305
 306
 307
 308
 309

Algorithm 1 Step algorithm for the watertank scenario in Figure 3.

- 1: `doStep(tank,0.1)` \triangleright Advance the tank SU by 0.1s
- 2: $level' \leftarrow getOut(tank,level)$ \triangleright Get water level of tank
- 3: `setIn(ctr,level,level')` \triangleright Set water level on controller
- 4: `doStep(ctr,0.1)` \triangleright Advance the controller SU by 0.1s
- 5: $valve' \leftarrow getOut(ctr,valve)$ \triangleright Get valve state of controller
- 6: `setIn(tank,valve,valve')` \triangleright Set valve state on tank

311 Algorithm 1 shows the order in which the SUs are
 312 simulated and how and when data is exchanged between
 313 the SUs. The algorithm advances the tank SU by 0.1
 314 second, retrieves the water level from the tank SU, sets
 315 the water level on the controller SU, advances the controller SU by 0.1 second, retrieves the valve state from the controller SU, and sets the valve state on the tank SU. We have deliberately omitted the error handling and logging of the outputs for brevity, nevertheless a practical implementation of co-simulation will have to deal

321 with such things. The *Plotting* stage is employed between each iteration of the *Step* stage to log the outputs of the SUs (valve state and water level) at the current communication point. This process is repeated until the simulation is terminated, which in this case is when the simulation time reaches 10 seconds. Once the simulation is finished, the *Free* stage is invoked to release resources and terminate the simulation.

329 Due to the numerous modelling and simulation tools that are capable of producing SUs, and the many *ad-hoc* co-simulation implementations (Gomes et al. (2018c)), the community has proposed a standard for the SU interface: the Functional Mockup Interface (FMI) Blochwitz et al. (2011); Committee (2014); Gomes et al. (2021b) standard to enable interoperability between SUs.

2.2. The FMI Standard

336 The Functional Mock-up Interface (FMI) standard is a tool-independent standard for the exchange of models and co-simulation, originally developed during the European ITEA2 project called MODELISAR (Blochwitz et al. (2011)). The standard provides and describes a compiled C-interfaces, the structure of a static description file, called *ModelDescription*, and a way of packaging these into a zip file according to a predefined structure. Consequently, a component that implements the C-interfaces according to the rules of the FMI standard, compiles its model into a dynamic/shared library, provides a *ModelDescription* and packages these into a zip file according to a predefined structure, is called a Functional Mock-up Unit (FMU). FMUs can be exported from a variety of modelling and simulation tools, such as Dymola Brück et al. (2002), OpenModelica Fritzson (2015), and Simulink. They can be imported into a co-simulation framework, such as INTO-CPS Larsen et al. (2016), to be integrated with other FMUs in a co-simulation scenario. This summarises the main purpose of the FMI standard, which is to enable interoperability between modelling and simulation tools by providing a standardised interface for SUs, which is essential for the simulation of CPSs.

The *ModelDescription* file defines the interface of the FMU and contains, among other things, information about its inputs, outputs and parameters, called *ScalarVariables*. Each *ScalarVariable* has a type, an identifier, a causality, and a variability constraining how the value of the *ScalarVariable* can be obtained and changed during the simulation using a set of functions, defined by the FMI standard, analogous to the set_c , get_c , and $doStep_c$ functions defined in Definition 1. The FMI standard defines a set of functions for getting and setting the values of *ScalarVariables*, e.g. `fmi2GetReal`

for Real and fmi2GetInteger for Integer. The FMI standard also defines a function to advance the state of the FMU in simulated time, e.g. fmi2DoStep.

Algorithm 1 shows how these FMI functions can be used to simulate the scenario in Figure 3. In Algorithm 1 we have deliberately chosen to simulate the tank SU before the controller SU to minimise the co-simulation error. However, the FMI standard does not provide any means to specify such constraints, so the definition of a scenario according to the FMI standard does not include the function R . Consequently, a co-simulation framework compliant with the FMI standard must either simulate the SUs in an arbitrary order or provide a mechanism for the user to specify the order in which the SUs are simulated. Our co-simulation framework provides the latter approach, thus providing the user with more control over the co-simulation, which is essential for fine-tuning the OA to minimise the co-simulation error.

The FMI standard for co-simulation aims to capture the common denominator of co-simulation and therefore strives for simplicity. However, this simplicity comes at a cost, as numerous studies (Gomes et al. (2019b); Oakes et al. (2021); Gomes et al. (2018e); Schweizer et al. (2015); Gomes et al. (2018a); Hansen et al. (2022b)) have shown how the accuracy of co-simulation results can be improved by tailoring the OA to the specific scenario by incorporating domain knowledge/implementation details of the SUs. For example, the OA can be adapted to take into account the implementation of the SUs, e.g. whether an SU interpolates or extrapolates an input, which is not captured by the standard.

By including such details, the OA can be adapted to improve the accuracy of the co-simulation results, as we show in Section 4. By incorporating such details, the OA can be customised to improve the accuracy of the co-simulation results, as we show in Section 4. Another limitation of the standard is in the context of network simulation, where the FMI export tool for the ns-3 network simulator supports simulation of purely discrete behaviour as it allows progression with 0 time (CES et al. (2021)), which is disallowed by the FMI standard¹.

Our goal is to provide a co-simulation framework that leverages both the strength of the FMI infrastructure and community while at the same time providing a framework more flexible than the FMI standard, which has to target a vast audience. We aim to provide a framework that can be customised to support advanced co-simulation based studies by incorporating experimental

¹Section 4.4.2, page 105 (FMI (2020)).

features and research results to support co-simulation practitioners at all levels of expertise. In the long term, we aim to improve the co-simulation support for the following simulation activities, each of which places specific requirements on the co-simulation frameworks.

Optimisation/DSE: Co-simulations are run as part of an optimisation loop, for example, in a Design Space Exploration (DSE) approach. This includes decision support systems, used, for example, in a digital twin (Glaessgen and Stargel (2012)) setting, where a modelled system is updated based on the operating system. Some of the specific requirements include: the ability to define co-simulation stop conditions, the ability to compute sensitivity, high performance, fully automated configuration, faster than real-time computation.

Certification: Co-simulation results can be used as a part of a certification endeavour. Requirements include fully transparent, and formally certified, synchronisation algorithms.

X-in-the-loop: Co-simulations include simulators that are constrained to progress in sync with the wall-clock time, because they represent human operators or physical subsystems.

Fault Injection: Co-simulations provide an additional test environment where all sorts of scenarios can be tested. Fault injection is a specific type of test where faults and other irregularities are injected into the system to investigate the system's behaviour under such conditions.

Last but not least, co-simulation tools have different audiences ranging between researchers, students, and industry from different domains. Each audience has different requirements and expectations from a co-simulation tool. At the same time, students and researchers are interested in transparency and customisation possibilities, while the industry is interested in plug-and-play, stable, scalable, and mature solutions with consumable interfaces. Consequently, we believe that the co-simulation framework, presented in the next section, meets the needs of all audiences by providing a low-level interface for students and researchers and a high-level interface for industry.

3. Co-simulation with Maestro2

This section describes the guiding principles behind Maestro2, namely the separation of concerns between the specification of a co-simulation and the execution of a co-simulation. The separation of concerns is achieved through the use of a Domain Specific Language (DSL) called MaBL, which specifies a co-simulation scenario

472 that can be analysed, verified and optimised prior to execution. The section begins with a brief introduction to the Maestro2 framework before introducing the MaBL DSL and how it is used to describe and analyse a co-simulation scenario. Finally, the section shows how Maestro2 uses code generation to provide a performant and flexible co-simulation engine.

479 3.1. Maestro2

480 Maestro2 is a co-simulation framework based on the FMI standard. It is written in Java and is available as open source software (<https://github.com/INTO-CPS-Association/maestro>). Co-simulation is performed by executing a MaBL specification, a “C-like” DSL for specifying co-simulation scenarios. The specification describes all the steps of the OA (see Figure 5), the FMUs involved and the connections between them. The general idea behind the invention and use of MaBL is to separate the specification of the OA from the execution of the OA. This need arises from the desire to analyse, verify and optimise the OA prior to execution, which was identified based on the experience with Maestro1 (Thule et al. (2019)).

484 The Maestro2 approach, illustrated in Figure 6, can be summarised as follows: The user can either write a MaBL specification manually or generate it using one of the approaches described in Section 3.3. The MaBL specification is then fine-tuned and analysed by a range of expansion plugins, which also can be used to expand the specification with additional functionality. Finally, the MaBL specification is executed using either the MaBL interpreter or a code generator, which generates a high-performance co-simulation engine in C++. Although Figure 6 depicts the Maestro2 approach as a linear process, it is possible to jump back and forth between the different phases to fine-tune the specification.

487 The following sections detail the different phases of the Maestro2 approach, starting with the specification phase. Nevertheless, before diving into these phases, we take a look at the MaBL DSL, the common denominator of the Maestro2 approach.

512 3.2. Maestro Base Language (MaBL)

513 The MaBL DSL is a “C-like” language that is used to specify a co-simulation scenario. A MaBL specification is a collection of modules, functions, and annotations that describe the OA. In the following, we give a brief introduction to the MaBL DSL using a series of small didactic examples. The reader can consult the online documentation for a more detailed description of the MaBL DSL (Association (d)).

521 Each MaBL specification must contain a entity called simulation, which is the entry point of the specification. The simulation block (line 8 in Listing 1) contains a set of imports (lines 9-10), which are used to import so-called runtime modules, and a set of annotations (lines 11-13), which are used to configure the simulation environment available to expansion plugins within the simulation, which are described below. Runtime modules are treated in Section 3.2.1, whereas expansion plugins are described in Section 3.3.1.

531 A module definition can also be part of a MaBL file, as exemplified by module DataWriter in Listing 1 line 1 – 6, which also includes yet another module (DataWriterConfig) in line 2.

Listing 1: MaBL Specification Structure

```

535
536 1 module DataWriter
537 2 import DataWriterConfig;
538 3 {
539 4   DataWriterConfig writeHeader( string headers
540     ↪   [] );
541 5   ...
542 6 }
543 7
544 8 simulation
545 9 import FMI2;
546 10 import Logger;
547 11 @Framework( "FMI2" );
548 12 @FrameworkConfig( "FMI2", "{...}\connections
549     ↪   \":{\\"{ ctrl }. ctrlInstance .valve \":{\\"{
550     ↪   wt}. wtInstance .valvecontrol \"},...}" ) ;
551 13 {
552 14   // Simulation code goes here
553 15 }
554

```

555 MaBL is a statically and strongly typed language that incorporates the type system of the FMI 2.0 standard (Real, Integer, Boolean, String). It extends this type system with the Array type and introduces the FMI2 type to represent an FMU. Additionally, runtime modules can introduce new types.

561 MaBL provides a range of built-in functions, including load and unload, which facilitate the loading and unloading of runtime modules, respectively. In terms of non-module functions, MaBL follows a minimalistic approach, offering basic arithmetic operations such as +, -, *, /, and fundamental Boolean operators such as ==, !=, !, >=, <=, <, >, &&, ||. Furthermore, MaBL offers standard control flow constructs such as if-else, while, try-finally to describe the OA.

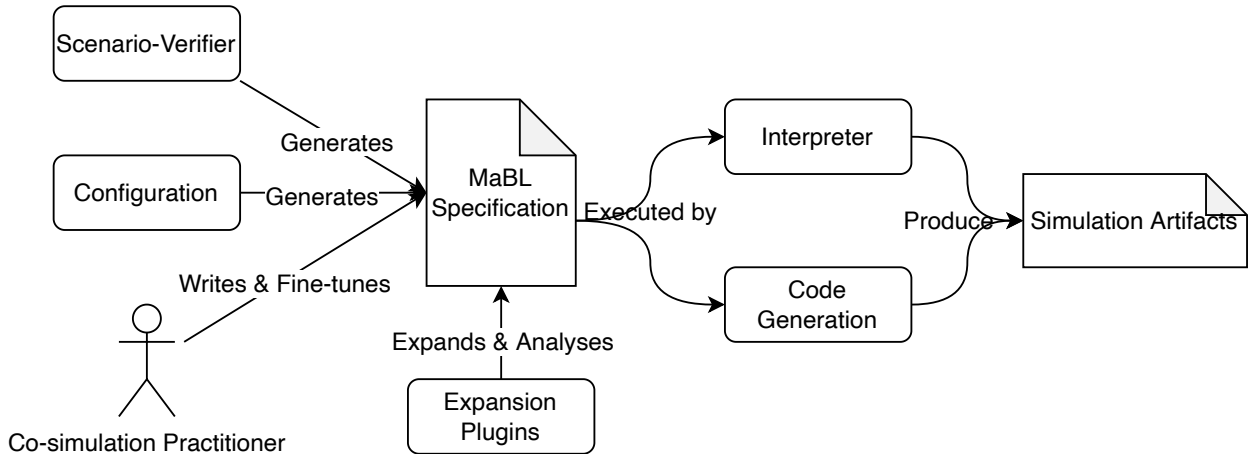


Figure 6: The Maestro2 is centred around the MaBL DSL. Maestro2 provides different approaches to generate a MaBL specification to enable co-simulation practitioners of different levels of expertise. A MaBL specification can be analysed by a range of expansion plugins. Finally, Maestro2 offers two approaches to execute a MaBL specification, namely the MaBL interpreter and a code generator.

3.2.1. Runtime Modules

More advanced features and functionality are typically implemented as runtime modules, which are loaded and executed during the execution of a MaBL specification through function calls.

A *runtime module* is a dynamically linked library that exposes a set of functions that can be called from MaBL to perform specific tasks that are not natively supported. For example, the FMI2 runtime module provides the FMI2 interface and offers various runtime module options. The “regular” runtime module unpacks an FMU, loads its dynamically linked library, and invokes its functions. On the other hand, the JFMI2 module allows loading an entity by specifying the class name instead of the FMU path. This is particularly useful for prototyping, as it enables development in other JVM-based languages such as Java, Kotlin, or Scala, bypassing the need for compilation into shared libraries and packages.

Another example of a runtime module is the Fault Injection module, which allows faults to be injected into the simulation with minimal effort. Specifically, the module wraps around an FMU and intercepts all data to and from the FMU, allowing it to modify the data before it is passed to the FMU based on a given configuration (Frasheri et al. (2021)).

3.3. Generation of Specifications

To cater for users with different levels of expertise, Maestro2 offers a variety of approaches to generating a MaBL specification as illustrated in Figure 6. The most basic approach is to manually write the specification in MaBL, which is a viable option for small specifications




Figure 7: The different approaches to generate a MaBL specification in terms of workload and customizability.

and experienced users. However, writing a large scenario manually is a tedious task, so it is desirable to generate the specification from other sources. Maestro2 offers the following approaches to generating a MaBL specification (based on the amount of work required (from most to least)): (i) Expansion Plugins, (ii) MaBL API, (iii) OA using the Scenario-Verifier, and (iv) Configuration.

These approaches vary in the degree of customisability and effort required to generate the specification, as illustrated in Figure 7. The expansion plugins are deliberately not included in the figure, as they are not a

614 standalone approach, but rather a set of plugins that can
615 be used to extend a MaBL specification generated using
616 one of the other approaches. Nevertheless, all approaches
617 are described below.

618 3.3.1. Expansion Plugins

619 MaBL specifications can be populated using an advanced
620 set of expansion features called *expansion plugins*,
621 which are community developed plugins that generate
622 MaBL based on a given set of parameters provided
623 by the user and the scenario at hand. A MaBL specification
624 can contain `expand` constructs, which are used
625 to invoke expansion plugins. Maestro2 will then invoke
626 the expansion plugin and replace the `expand` construct
627 with the MaBL generated by the corresponding expansion
628 plugin. MaBL expansion plugins serve two purposes:
629 1. to generate MaBL based on a given set of parameters
630 to reduce the amount of manual work required to create
631 a MaBL specification, and 2. to extend the Maestro2
632 framework with new functionality such as fault injection
633 and design space exploration [Ejersbo et al. \(2023\)](#);
634 [Pierce et al. \(2022\)](#); [Frasheri et al. \(2021\)](#).

635 An example of the application of expansion plugins
636 is the `Initializer` plugin ([Hansen et al. \(2021c\)](#))
637 shown in [Listing 2](#). The `Initializer` defines a function
638 called `initialize` which is called in line 3 using
639 the `expand`. The plugin generates the necessary FMI
640 calls to initialise the FMUs in the simulation in MaBL,
641 based on the connections and dependencies between the
642 FMUs. [Listing 2](#) also shows the `@Config` annotation,
643 which is used to provide additional information to the
644 expansion plugin, which it can use to generate the MaBL.
645 The configuration in [Listing 2](#) specifies the values of the
646 parameters `maxLevel` and `minLevel` of the `Controller`
647 FMU.

Listing 2: MaBL Expansion

```
648  
649 1 @Config("{\parameters\":{\crtl }. ctrlInstance  
650     ↪ .maxlevel\":1,\crtl }. ctrlInstance .  
651     ↪ minlevel\":1}") ;  
652 2 Initializer .expand initialize (components,  
653     ↪ START_TIME, END_TIME);  
654
```

655 Expansion plugins can be chained together, so the
656 MaBL generated by the `Initializer` plugin can contain
657 additional `expand` constructs, which are then expanded
658 by other plugins. This feature, while powerful, should
659 be used with care to avoid potential infinite loops.
660 However, when used properly, it allows complex scenarios
661 to be created with minimal effort, and allows expansion
662 plugin authors to leverage existing functionality. To
663 avoid invalid MaBL specifications, the Maestro2 framework
664 performs a type check

665 on the resulting MaBL specification after each extension
666 plugin is applied. However, it is important to note that
667 the expansion plugins are not limited to generating
668 MaBL, but can also perform other tasks such as verifying
669 the modelDescription of an FMU, as done by the
670 `ModelDescriptionVerifier` plugin. Once a fully expanded
671 MaBL specification without `expand` constructs has been
672 generated, it can be fine-tuned manually if necessary
673 before execution.

674 *Commonly Used Expansion Plugins.* Maestro2 is supplied
675 with a number of commonly used extension plugins to
676 minimise the amount of manual work involved in the
677 creation of a MaBL specification.

678 `Initializer` ([Hansen et al. \(2021c\)](#)) generates the
679 initialisation code for a co-simulation scenario in
680 MaBL. The plugin leverages state of the art algorithms
681 for synthesizing the initialisation algorithm of a co-
682 simulation scenario potentially containing algebraic
683 loops ([Hansen et al. \(2021c\)](#); [Gomes et al. \(2019a\)](#);
684 [Hansen et al. \(2022b\)](#)). Concretely, the plugin builds
685 a dependency graph of the FMUs in the scenario and
686 employs graph-based reasoning to determine the order
687 in which the FMUs should be initialised based on the
688 interconnections between the FMUs and their contracts.
689 The plugin contains an optional feature of verifying the
690 calculated initialisation order against the verifier
691 implemented in [Gomes et al. \(2019a\)](#).

692 `JacobianStepBuilder` produces the MaBL necessary
693 to perform *Step* stage of the OA based on the
694 Jacobian iteration method. Jacobian iteration performs
695 a simulation step by prompting all FMUs to progress
696 in time, retrieves the necessary outputs, and sets the
697 necessary inputs ([Gomes et al. \(2018d\)](#)). The plugin
698 can, similarly to Maestro1, be tailored to use different
699 methods for dynamically determining the step size when
700 performing a simulation with a variable step size
701 ([Thule et al. \(2019\)](#)). The current implementation
702 supports the following methods:

703 **Zero Crossing:** Synchronise the FMUs at a point in
704 time where a given signal is zero or two signals
705 intersect.

706 **Bounded Difference:** The bounded difference
707 constraint bounds the difference between two signals
708 or two consecutive observations of the same signal
709 by a pre-defined value.

710 **Sampling Rate:** Ensures that all FMUs synchronise
711 at pre-determined points in time.

712 **FMU Max Step Size:** First proposed in [Broman et al. \(2013a\)](#)
713 this constraint attempts to avoid the need

714 for rollbacks. It requires the FMUs to implement
715 a non-FMI function `getMaxStepSize` that returns
716 the maximum step that the given FMU can perform
717 at the given point in time.

718 The main difference between this plugin and its coun-
719 terpart in *Maestro1* is that the functionality is now re-
720 alised via MaBL and associated runtime modules, while
721 it was previously implemented directly in Scala.

722 Stabilisation is another commonly used feature to en-
723 sure that the co-simulation converges to a steady state.
724 This is accomplished by performing multiple iterations
725 of a given step until convergence is achieved or the
726 maximum attempts are reached. Concretely, the plu-
727 gin performs a simulation step, checks if the simula-
728 tion has converged, and if not, it rolls back the FMUs
729 to their previous state and performs another simulation
730 step with the output values from previous iteration.

731 `ModelDescriptionVerifier` is a verification plu-
732 gin. A verification plugin does not generate MaBL. In-
733 stead, it has access to the AST of the MaBL specifi-
734 cation and can perform various checks on the specifi-
735 cation. The `ModelDescriptionVerifier` plugin ver-
736 ifies the `ModelDescription` files of the FMUs against
737 a formal model of the FMI standard implemented in
738 VDM-SL via the tool `VDMCheck` (Battle et al. (2020)).
739 Moreover, the plugin verifies that all FMUs are properly
740 unloaded after the simulation has finished by analysing
741 the MaBL AST.

742 More information about the expansion plugins avail-
743 able in *Maestro2* can be found in the online documenta-
744 tion (Association (d)). The online documentation also
745 contains a guide for creating new expansion plugins to
746 enable expert users to extend the *Maestro2* framework
747 with new functionality.

748 3.3.2. MaBL API

749 One approach to generating a MaBL specification is
750 to use the MaBL API, a Java API that can be used to
751 generate a MaBL specification programmatically. This
752 approach is typically used by the expansion plugins
753 (e.g. `Initialization` and `JacobianStepBuilder`) to gener-
754 ate MaBL, but users can also use it to generate a MaBL
755 specification. Specifically, the MaBL API hides the
756 complexity of the underlying MaBL syntax and allows
757 the user to generate a MaBL specification using a high-
758 level API, without having to worry about error handling,
759 AST generation, and other complexities of the MaBL
760 language.

761 An example of the MaBL API is shown in [List-
762 ing 3](#), which is used by the `Initialization` plugin
763 to change the mode of an FMU to `Initialisation`

764 Mode. Although the example is simple, it still empow-
765 ers the user as they do not have to worry about the com-
766 plexities of handling all the possible return values of the
767 FMI function `enterInitializationMode`. Instead,
768 the user can concentrate on the task at hand, which is
769 to change the mode of an FMU to `Initialization`
770 Mode.

Listing 3: FMU instance function call using MaBL API

```
771 fmuInstanceVariable . enterInitializationMode ()
```

774 Another valuable feature of the MaBL API is the
775 ability to programmatically link one FMU’s output to
776 another FMU’s input. The MaBL API will, based on
777 these couplings, generate the necessary MaBL code to
778 exchange the data between the linked ports.

779 The MaBL API provides complete flexibility to the
780 user, allowing them to create a MaBL specification tai-
781 lored to their needs, such as the adaptive co-simulation
782 scenario described in [Section 4](#). However, this expres-
783 siveness comes at the cost of increased complexity, as
784 the user has to deal with all the intricacies of the simu-
785 lation.

786 3.3.3. Scenario-Verifier

787 The third approach to generating a MaBL specifi-
788 cation is to use the `Scenario-Verifier` (Hansen et al.
789 (2021b,a, 2022b)), which is a tool for synthesising and
790 verifying OAs for co-simulation scenarios described in
791 a high-level DSL similar to [Definition 2](#).

792 The `Scenario-Verifier` uses the latest advances in OA
793 synthesis (Hansen et al. (2021c); Gomes et al. (2019a);
794 Hansen et al. (2022b)) and constraints declared by the
795 `R` function to synthesise an OA tailored to a given co-
796 simulation scenario, subject to both step rejections and
797 algebraic loops Kübler and Schiehlen (2000b), while
798 ensuring that the OA respects the implementation de-
799 tails of the scenario and correctly implements the FMI
800 standard. The `Scenario-Verifier` synthesises an OA that
801 employs a stabilisation algorithm to handle algebraic
802 loops and a step size adjustment algorithm to handle
803 step rejections. As a result, a co-simulation practitioner
804 does not need to worry about the intricacies of the OA
805 of such complex scenarios, as the `Scenario-Verifier` will
806 synthesise an OA that ensures a co-simulation where
807 all FMUs move in lockstep and algebraic loops are sta-
808 bilised.

809 Furthermore, the tool can verify a given OA against
810 the FMI standard and the scenario description. Con-
811 cretely, the tool uses a symbolic formalisation of the
812 FMI standard and the OA in the model checker UP-
813 PAAL (Behrmann et al. (2006)) to verify that the OA

814 respects the implementation details of the scenario and 863
815 correctly implements the FMI standard. For example, 864
816 the tool verifies that the OA solves the FMUs in an 865
817 optimal order that respects the implementations of the 866
818 FMUs, that all FMUs move in lockstep, and that alge- 867
819 braic loops are stabilised. Errors in the OA are reported 868
820 to the user in the form of a trace, which can be visually 869
821 inspected to debug the OA.

822 However, it is essential to note that the Scenario- 871
823 Verifier tool only synthesises the *Initialise* and *Step* 872
824 stages of the OA in DSL format. Consequently, it is 873
825 used in conjunction with the MaBL API and expansion 874
826 plugins to generate the remaining stages of the OA and 875
827 translate the DSL into MaBL.

828 This approach is recommended for users with little 877
829 experience with MaBL and co-simulation in general, as 878
830 it requires almost no effort to start a co-simulation. Nev- 879
831 ertheless, the Scenario-Verifier also provides a number 880
832 of advanced features for expert users to fine-tune the OA 881
833 by providing additional constraints and expectations to 882
834 the tool.

835 3.3.4. Configuration

836 The last approach to generating a MaBL specifi- 886
837 cation is to use a configuration file called the multi- 887
838 model (Larsen et al. (2016)). The multi-model is a 888
839 JSON file that details the FMUs involved in the co- 889
840 simulation, the connections between them, and the 890
841 parameters of the FMUs. The configuration file is 891
842 then used to generate a MaBL specification using 892
843 the MaBL API, using both the *Initializer* and 893
844 *JacobianStepBuilder* expansion plugins. Neverthe- 894
845 less, the configuration file can also be used to gener- 895
846 ate a MaBL specification using the Scenario-Verifier 896
847 tool, which requires some minor modifications to the 897
848 multi-model by Maestro2 to make it compatible with 898
849 the Scenario-Verifier tool.

850 This approach is recommended for users with little 900
851 experience with MaBL and co-simulation in general, as 901
852 it requires the least effort to start a co-simulation. The 902
853 configuration file can be generated by other tools, such 903
854 as the INTO-CPS Application (Macedo et al. (2020)). 904

855 3.4. Execution

856 The final step in the Maestro2 approach is to execute 908
857 the MaBL specification generated in the previous step. 909
858 Maestro2 provides two execution modes: *interpretation* 910
859 and *code generation* for executing a MaBL specifica- 911
860 tion. 912

861 The interpretation mode is based on an interpreter 913
862 written in Java, which is the default execution mode of 914

Maestro2 as it is the most convenient for development 863
and debugging purposes. Nevertheless, in order to min- 864
imise the overhead of the co-simulation framework, a 865
code generator that translates a MaBL specification into 866
C++ code has been implemented as well. The code gener- 867
ator is based on the Java interpreter, and thus the gener- 868
ated code is functionally equivalent to the interpreted 869
code. The code generator offers a significant perfor- 870
mance improvement over the Java interpreter as shown 871
in Figure 8. The comparison in Figure 8 is based on 872
100 co-simulations of the water tank scenario described 873
in Figure 3. Each simulation has a different end time, 874
starting with an end time of 1s for simulation 1 and in- 875
creasing by one for each simulation so that the last sim- 876
ulation (number 100) has an end time of 100s. All these 877
100 co-simulations were run on Maestro1 (the prede- 878
cessor of Maestro2, see Section 5), Maestro2 with the 879
Java interpreter, and the code-generated C++ version of 880
Maestro2. 881

882 The execution time of the various co-simulations is 882
883 shown in Figure 8. The figure shows two plots, one 883
884 including the time taken to load FMUs and one without 884
885 the time taken to load the FMUs, telling the same story. 885
886 The time taken to generate the specification and compile 886
887 the C++ code is not included in the figure, as it is a one- 887
888 time cost. 888

889 The figure shows that the C++ code generated by 889
890 Maestro2 is significantly faster than the interpreter 890
891 (*maestro2*) and Maestro1. The optimisation of the 891
892 generated code is amplified as the simulation time in- 892
893 creases, as the simulation loop is the primary time ex- 893
894 penditure as the simulation time grows. As the fig- 894
895 ure shows, the interpreter (*maestro2*) is faster than 895
896 Maestro1. This is because Maestro1 performs various 896
897 lookups during runtime, whereas Maestro2 performs 897
898 these lookups during compilation. 898

899 However, the performance gain of the code gener- 899
900 ated does come at a cost in terms of expressivity, as 900
901 it requires used runtime modules to be ported to C++, 901
902 as only the MaBL specification is translated to C++. 902
903 Nevertheless, some runtime modules have already been 903
904 ported to C++, such as the FMI2 runtime module, the 904
905 DataWriter module, and the MEnv module, to test and 905
906 validate the approach. 906

907 Although the execution times in Figure 8 are small, 907
908 and you might think that the performance gains it not 908
909 worth the effort, it is essential to note that the perfor- 909
910 mance gain is amplified when performing DSE. A DSE 910
911 study consists of a series of co-simulations with dif- 911
912 ferent parameters to find the optimal solution. Thus 912
913 the performance gain is amplified as the number of co- 913
914 simulations increases. 914




Figure 8: Performance comparison of Maestro, Maestro2 Interpreter and Maestro2 Code Generated version. In the left figure the time for loading the FMUs is subtracted from the total simulation time. In the right figure it is included.

915 3.5. Utilizing Maestro2

916 Maestro2 offers two interfaces for interacting with
917 the framework: a Command Line Interface (CLI) and a
918 REST interface. Both interfaces offer similar function-
919 ality and strike a balance between new functionality and
920 compatibility with Maestro1 to ensure a smooth transi-
921 tion.

922 The CLI offers a minimalistic interface for interact-
923 ing with Maestro2, which is useful for scripting and au-
924 tomation as it does not require running a web server.
925 The CLI is used by the INTO-CPS DSE functionali-
926 ty (Bogomolov et al. (2020)).

927 The REST interface allows Maestro2 to function as
928 a web server, enabling cloud support and remote ac-
929 cess. It offers the flexibility of accessing Maestro2's
930 functionality through HTTP requests. The REST inter-
931 face is used by applications like the INTO-CPS Appli-
932 cation Macedo et al. (2020). Additionally, the REST in-
933 terface supports live-streaming of data via web sockets,
934 enabling real-time data updates as the simulation pro-
935 gresses. This feature involves adding an extra listener
936 to the DataWriter runtime module, as briefly demon-
937 strated in Listing 1. It showcases the possibilities of combining
938 MaBL with runtime modules.

939 The interfaces are not covered in detail here, but the
940 interested reader can refer to the online documenta-
941 tion (Association (d)) for more information on their us-
942 age and capabilities.

943 4. Case Studies

944 This section presents two case studies that illustrate
945 how Maestro2 can be used to tackle a broad variety
946 of co-simulation scenarios. This section provides a
947 brief overview of the case studies, the essential chal-
948 lenges that cannot be solved by a standard co-simulation
949 framework, and the role of Maestro2 in tackling these
950 challenges. More details about the case studies can be
951 found in the corresponding references.

952 4.1. Adaptive Mass-Spring-Damper Co-simulation

953 The adaptive mass-spring-damper case study, de-
954 tailed in Inci et al. (2021) and illustrated in Figure 9,
955 consists of two linear mass-spring-damper subsystems,
956 connected to rigid walls and coupled with a spring-
957 damper. The system is simulated with two FMUs (MSD1
958 and MSD2) as shown in Figure 11, one for each mass-
959 spring-damper subsystem. Subsystem 1, MSD1, acts as
960 an inert system by inputting the coupling force from
961 Subsystem 2, MSD2, and outputting displacement and
962 velocity to Subsystem 2.




Figure 9: Double mass-spring-damper.

963 Although the system may initially appear simple, it
964 poses challenges when accuracy is critical, as demon-
965 strated in Inci et al. (2021). Their work demonstrates
966 that achieving accurate results for this system requires
967 the use of an adaptive co-simulation algorithm, which
968 dynamically changes the order of actions in the algo-
969 rithm employed in the Step stage of the co-simulation.

970 Changing the order of actions affects the sequence in
971 which FMUs are simulated and how they communicate
972 with each other. Specifically, the system depicted in
973 Figure 9 can, according to the FMI standard, be simu-
974 lated using the two algorithms described in Figure 10.
975 Algorithm 2 simulates MSD1 first, followed by MSD2,
976 while Algorithm 3 simulates MSD2 first, followed by
977 MSD1. The two algorithms can, in fact, be synthesised
978 by Maestro2 using the Scenario-Verifier tool.

979 In [Inci et al. \(2021\)](#) it is suggested that instead of 1029
980 choosing one of the algorithms a priori, it is better to 1030
981 choose the algorithm that minimises the error at each 1031
982 co-simulation step (i.e., adaptive co-simulation). The 1032
983 error is estimated by comparing the results of the two al- 1033
984 gorithms at each co-simulation step against a reference 1034
985 solution obtained by solving the system of equations ana- 1035
986 lytically. To estimate the error and determine the algo- 1036
987 rithm, an additional FMU called `SwitchingDecision` 1037
988 is employed as shown in [Figure 11](#). 1038

989 The adaptive co-simulation algorithm uses the 1039
990 judgement on the output `best_order` of the 1040
991 `SwitchingDecision` FMU to decide which algo- 1041
992 rithm to use at each co-simulation step. Concretely, 1042
993 the adaptive co-simulation algorithm starts by employ- 1043
994 ing [Algorithm 2](#) to simulate the system for a single 1044
995 co-simulation step and then uses the `best_order` 1045
996 output of the `SwitchingDecision` FMU to decide 1046
997 which algorithm to use at the next co-simulation step to 1047
998 minimise the co-simulation error. 1048

999 *Role of Maestro2.* Maestro2 provides the necessary 1050
1000 functionality to implement adaptive co-simulation.
1001 Concretely, MaBL allows the user to implement the 1051
1002 adaptive co-simulation algorithm in [Figure 11](#) by using 1052
1003 a conditional statement to decide between [Algorithm 2](#)
1004 and [Algorithm 3](#), in [Figure 10](#) based on the `best_order` 1053
1005 output of the `SwitchingDecision` FMU. Furthermore, 1054
1006 MaBL support for declaring new variables to store vari- 1055
1007 ables between co-simulation steps, facilitating the im- 1056
1008 plementation of the adaptive co-simulation algorithm.
1009 Finally, Maestro2’s performance made the running time 1057
1010 difference between the adaptive and static algorithms 1058
1011 negligible. 1059

1012 The errors of the adaptive co-simulation algorithm 1060
1013 are compared with those of the static algorithm in [Fig- 1061](#)
1014 [ure 12](#). As can be seen, the adaptive algorithm attempts 1062
1015 to follow the best sequence at any given time, thus free- 1063
1016 ing the user from determining which algorithm to em-
1017 ploy at a specific time. 1064

1018 4.2. Hardware-in-the-loop Co-Simulation

1019 The second case study showcases the use of Maestro2 1066
1020 in a hardware-in-the-loop co-simulation scenario, where 1067
1021 the numerical simulation of a system is coupled with a 1068
1022 physical system. Hardware-in-the-loop co-simulation is 1069
1023 a common practice in seismic testing of civil engineer- 1070
1024 ing structures ([McCrum and Williams \(2016\)](#)). 1071

1025 The reported case study, detailed in [Gomes et al. 1071](#)
1026 ([2021a](#)), consists of a physical cantilever beam coupled 1072
1027 to a linear spring, as illustrated in [Figure 13](#). The can- 1073
1028 tilever beam is excited by a sinusoidal loading applied 1074

via an electric linear actuator to study the dynamic re-
sponse in terms of displacement (u) of the beam to seis-
mic excitations provided by the linear spring. [Figure 13](#)
shows a picture of the experimental setup, along with
a simplified version of the system, as depicted in the
left part of the figure. Finally, the right part of the fig-
ure provides a schematic overview of the experimental
setup.

Hardware-in-the-loop co-simulation is enabled by us-
ing a hybrid testing setup depicted in [Figure 14](#). The hy-
brid testing (HT) setup comprises a three subsystems,
modelled as distinct FMUs, as shown in [Figure 14](#).
The HT setup consists of a physical substructure (PS),
and a numerical substructure (NS) simulated by a finite
element (FE) software, and an FMU that couples the
two substructures to enforce compatibility between the
physical and numerical substructures (Coupling). The
PS structure is equipped with several sensors and actu-
ators, which are connected to a data acquisition system
on an industrial PC via EtherCAT. The sensors and actu-
ators are used to measure the response of the PS and
to provide the excitation to the NS.

Role of Maestro2. There are two challenges in this
case:

1. Mistakes in the co-simulation could lead to physi-
cal consequences on the connected hardware. Here
Maestro2 support for static analysis plays a crucial
role to prevent these.
2. One of the main challenges in this case study is
the need to synchronise the numerical and phys-
ical substructures. This is achieved by using a
custom orchestration algorithm implemented in
MaBL. Concretely, the orchestration algorithm en-
sures that the Coupling FMU is simulated after the
other two FMUs.

A video of the experiment is available online, see [Asso-
ciation \(2021\)](#).

5. Related Work

Co-simulation is a large field and challenging to
cover thoroughly, with co-simulation frameworks being
a moving target. For this reason, we introduce some
of the existing co-simulation frameworks and compare
them to our contribution in [Table 1](#), on the item where
Maestro2 is most novel: its capability for extensive cus-
tomization while maintaining robust verification capa-
bilities. For surveys on the co-simulation topic, we refer

Algorithm 2 MSD1 \rightarrow MSD2	Algorithm 3 MSD2 \rightarrow MSD1	
1: doStep(S_1, H)	1: doStep(S_2, H)	At time t:
2: $x_1' \leftarrow \text{getOut}(S_1, x_1)$	2: $F_k' \leftarrow \text{getOut}(S_2, F_k)$	doStep(S_1, H): Advances the state of FMU S_1 by H
3: $v_1' \leftarrow \text{getOut}(S_1, v_1)$	3: setIn(S_1, F_k, F_k')	getOut(S_1, y): Returns the output y of the FMU S_1
4: setIn(S_2, x_1, x_1')	4: doStep(S_1, H)	setIn(S_1, u, y): Assigns the input u of the FMU S_1 to the value y
5: setIn(S_2, v_1, v_1')	5: $x_1' \leftarrow \text{getOut}(S_1, x_1)$	
6: doStep(S_2, H)	6: $v_1' \leftarrow \text{getOut}(S_1, v_1)$	
7: $F_k' \leftarrow \text{getOut}(S_2, F_k)$	7: setIn(S_2, x_1, x_1')	
8: setIn(S_1, F_k, F_k')	8: setIn(S_2, v_1, v_1')	

Figure 10: Possible algorithms. Adapted from [Inci et al. \(2021\)](#). Both algorithms are valid according to the FMI standard and can be used to simulate the system depicted in [Figure 9](#).

Table 1: Overview of co-simulation frameworks and their customization options.

Tool	FMI	Compiled	Interpreted	License	Customizations
DACCOSIM Evora Gomez et al. (2019)	NG	Yes	No	Yes	Open Source Step size
Dymola Brück et al. (2002)	Yes	Yes	No	No	Proprietary Step size
FIDE Cremona et al. (2016)	No ²	Yes	No	No	Proprietary Step size
VICO Hatledal et al. (2021)	Yes	No	Yes	Yes	Open Source Step size, Runtime behavior through coded extensions.
C2WT Neema et al. (2014)	No ³	No	Yes	Yes	Proprietary Step size
FMPy	Yes	No	Yes	Yes	Open-Source Full customization by coding interaction with FMUs in Python.
OMSimulator Ochel et al. (2019)	Yes	No	Yes	Yes	Open-Source Step size, Algebraic loop solver.
Van Acker et al. (2015)	Yes	Yes	No	No	Open-Source Step-size.
Maestro1	Yes	No	Yes	Yes	Open-Source Step-size, Algebraic loop solver.
Maestro2	Yes	Yes	Yes	Yes	Open-Source Full customization using domain specific language.




Figure 11: Co-simulation scenario. Adapted from [Inci et al. \(2021\)](#).

the reader to [Gomes et al. \(2018b\)](#); [Hafner and Popper \(2017\)](#); [Palensky et al. \(2017\)](#) for related surveys.

Maestro1 [Thule et al. \(2019\)](#), the predecessor of Maestro2, was developed during the INTO-CPS project ([Association \(c\)](#); [Larsen et al. \(2016\)](#)) and is an FMI-based co-simulation orchestration engine. It can perform co-simulations with a fixed algorithm and lacks the customisation and verification abilities of Maestro2. DAC-

COSIM [Galtier et al. \(2015\)](#); [Evora Gomez et al. \(2019\)](#) has been rebuilt as DACCOSIM NG, and extended with additional features, such as the capability of packaging multiple FMUs, including a scenario into a single FMU, referred to as Matryoshka FMU ([Evora Gomez et al. \(2019\)](#)). Another interesting FMI-based co-simulation framework is VICO ([Hatledal et al. \(2021\)](#)). VICO runs on the Java Virtual Machine and is thereby cross-platform. It supports the FMI companion standard System Structure and Parameterisation ([Jochen Köhler et al. \(2016\)](#); [Association \(e\)](#)) for defining the structure of a co-simulation. It strongly focuses on the possibility of composition through a clear separation of objects composed solely of data and systems that act on such objects. This is referred to as Entity-Component-System ([Martin \(2007\)](#)). The promise of this approach is supporting runtime behaviour change and simplicity of use. Furthermore, it features built-in 3D graphics and plotting capabilities. While this shares the goals of Maestro2, the approach is different. One could, for




Figure 12: On the top, error in v_1 for adaptive and static co-simulation sequences. At the bottom, sequence changes of the adaptive co-simulation (corresponds to the `best_order` output in Figure 11). Reproduced from Inci et al. (2021).

example, consider composing a co-simulation through VICO and then use the MaBL API to generate the execution of the composed co-simulation. The compositional features could also be implemented as a runtime plugin. Van Acker et al. (2015) presents a modelling language to model a co-simulation setup and a related transformation to an optimised master algorithm ready for execution. Thus, it is similar in nature to some of the concepts of Maestro2 and the Scenario Verifier. One difference is, for example, that it does not contain an interpreter or a plugin structure. The AVL Model.CONNECT co-simulation tool is a professional tool focusing primarily on the automotive market. It does support FMI, and its strength seems to be its ability to carry out Hardware-In-the-Loop simulations. We have not been able to find indications that it supports the level of customizability we demonstrate with Maestro2. For more related FMI-based tools, the reader is referred to the tools page on the FMI-website (Committee), where several FMI-enabled importing tools are listed.

Crucially, Maestro2 also targets verification efforts, ideally both at the FMU and orchestration level, and as such, contributions within this area are of interest as well. The tool VDMCheck (Battle et al. (2020)), is an example of FMU verification. It verifies the ModelDescription file of an FMU, and has explicit support within Maestro2. This can be applied directly, as it does not require user interaction. Gomes et al. (2018a) considers adapting simulators to correct interaction assumptions based on a different environment. Their approach

is to create a new FMU, referred to as *external FMU* that encloses one or more FMUs, referred to as *internal FMUs*. Via a DSL called *baseSA* it is possible to create rules for mapping actions applied to the external FMU to actions applied to the internal FMUs and interaction between the internal FMUs. This is applied through a sound definition of hierarchical simulators that leaves the internal FMUs unmodified and thus improves modularity and preserves transparency. MaBL is capable of representing a semantically equivalent set of operations, and as such, MaBL and Maestro could function as an execution engine for *baseSA*, if one was inclined to write such an expansion plugin. However, it does not feature the FMU-generation capabilities that generalise the approach in Gomes et al. (2018a) to all FMI-based orchestration engines. The Scenario Verifier (Hansen et al. (2021b)) described in Section 3.3.3 is an example of a tool that calculates and verifies an FMI-based co-simulation OA based on constraints and expectations of the enclosed FMUs. Enriching the environment of such an OA has been proven possible, see Section 3.3.3, to create an executable co-simulation in MaBL. The last example of tooling to be considered in this publication related to verifying the behaviour of a co-simulation and its constituents is within the domain of Test Automation in Ouy et al. (2017). Here, a Test FMU is created in order to stimulate the system under test according to system requirements, whereas other FMUs represent the system under test. An external tool then evaluates the outputs of the test FMU in order to determine whether the system under test expressed correct behaviour.

6. Concluding Remarks

This paper introduced Maestro2, a co-simulation framework designed for running FMI-based co-simulations. The key contribution of this work is the Maestro2 approach, which leverages the MaBL DSL to empower users in customizing the OA and minimizing co-simulation errors.

Maestro2 offers different levels of automation for describing the co-simulation scenario and the OA, catering to both prototyping needs and expert-level fine-tuning. The framework utilizes code generation techniques to ensure minimal overhead and high performance for co-simulation practitioners.

With its plugin architecture, Maestro2 enables users to extend the framework with new capabilities, such as incorporating new FMU types, supporting additional logging formats, and integrating new verification tools.

The framework's effectiveness has been demonstrated through multiple case studies and research

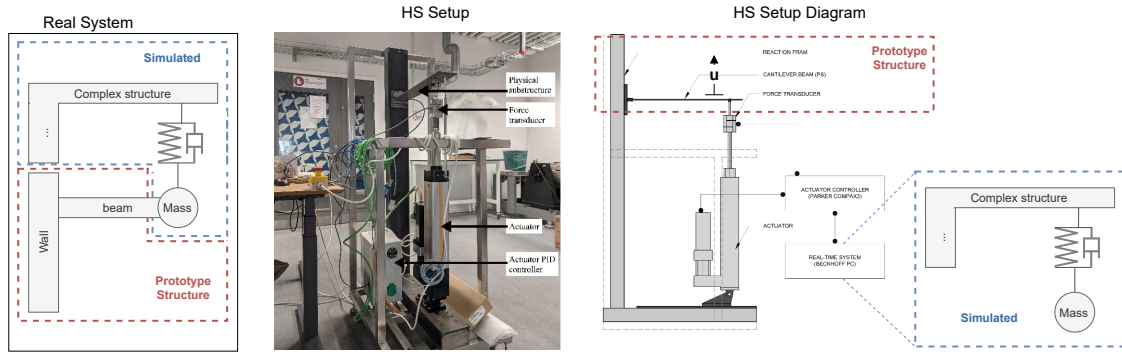


Figure 13: Experimental setup installed at the Dynamisk LAB of Aarhus University. Adapted from Gomes et al. (2021a).




Figure 14: Co-simulation scenarios that implements the setup described in Figure 13. The dashed boxes represent the fact that NSFMU (respectively PSFMU) communicate with a FE Software (resp. the Test Setup), when the `fmi2DoStep` function is invoked.

1184 projects (see Section 4), showcasing its capabilities in 1211
 1185 tackling real-world problems. 1212

1186 To our knowledge, Maestro2 is the only open-source 1213
 1187 co-simulation framework that offers a balance between 1214
 1188 flexibility, usability, and performance. However, we ac- 1215
 1189 knowledge the presence of open challenges highlighted 1216
 1190 by the case studies, including the need for a more ser- 1217
 1191 viced-friendly interface for the MaBL DSL and further en- 1218
 1192 hancements to the verification capabilities of the frame- 1219
 1193 work. Addressing these challenges will be crucial for 1220
 1194 advancing the framework and improving its usability in 1221
 1195 practical applications. 1222

1196 Future work will focus on these challenges and ex- 1223
 1197 plore promising directions to enhance Maestro2. 1224

1198 *Future Work.* Maestro2 provides, due to its plugin- 1225
 1199 based architecture, a solid foundation for future work. 1226
 1200 Some of the most promising directions for future work 1227
 1201 are: Supporting the newest version of the FMI stan- 1228
 1202 dard (Junghanns et al. (2021); Hansen et al. (2022a)) 1229
 1203 to enable co-simulation of a broader range of systems, 1230
 1204 such as hybrid and reactive systems. This work has al- 1231
 1205 ready been initiated, and MaBL is currently being ex- 1232
 1206 tended to support the new features of the FMI stan-
 1207 dard. The work is expected to be completed during 2023, 1233
 1208 making Maestro2 one of the first co-simulation frame-
 1209 works to support the new version of the FMI standard. 1234

1210 Furthermore, we plan to extend Maestro2 to be ap- 1235

1211 plicable in the context of digital twin engineering,
 1212 more specifically, the incubator project (Feng et al.
 1213 (2021a,b)), some initial results of which are presented
 1214 in (Association (b,a)). We expect this will lead to new
 1215 interfaces and functionality, such as the possibility of
 1216 changing the simulation algorithm or replacing FMUs
 1217 at runtime due to external changes while still preserv-
 1218 ing the benefits of calculating a specification before execution.

1219 Finally, we plan to extend the verification capabili-
 1220 ties of Maestro2 to permit verification of the final MaBL
 1221 specification using the Scenario Verifier (Hansen et al.
 1222 (2022b)). This is an appealing challenge, as it be-
 1223 comes available to all other tools using MaBL/Maestro2
 1224 as their execution target. Possible verification efforts
 1225 include verifying that the MaBL specification is well-
 1226 formed, deterministic, and adheres to the FMI stan-
 1227 dard. Nevertheless, the verification efforts must be care-
 1228 fully considered, as the verification capabilities are po-
 1229 tentially at odds with the customizability of MaBL em-
 1230 powering experts to fine-tune the OA to minimise co-
 1231 simulation error.

Acknowledgments

This research was funded by a number of externally funded research projects including DiT4CPS, UPSIM,

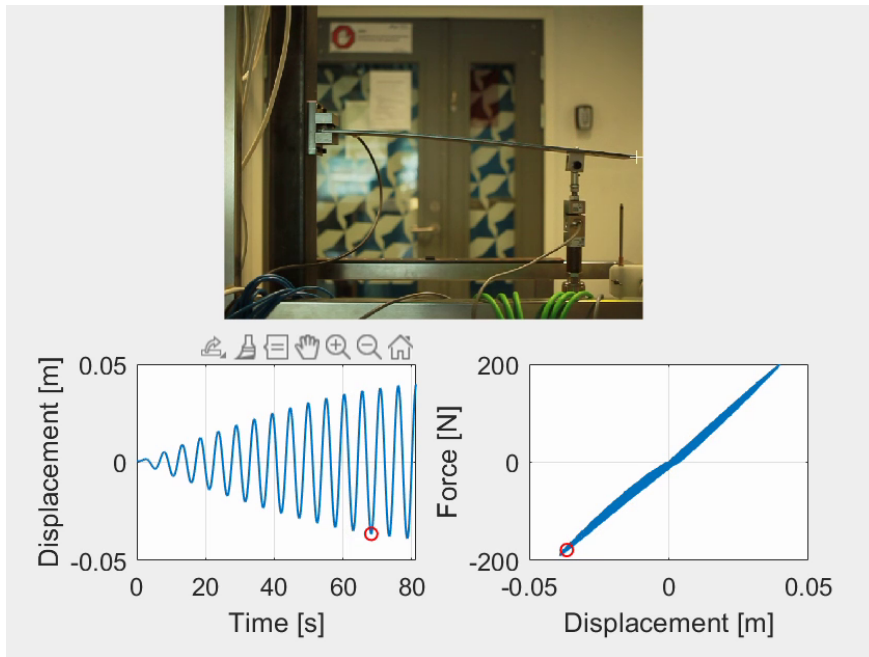


Figure 15: Numerical results as reported in Gomes et al. (2021a). The full video can be seen online Association (2021).

- 1236 HUBCAP and AgroRobottiFleet. Furthermore, we are 1267
 1237 grateful to the Poul Due Jensen Foundation, which has 1268
 1238 supported the establishment of the Center for Digital 1269
 1239 Twin Technology at Aarhus University.
- 1240 **References**
- 1241 Andersson, C., 2016. Methods and Tools for Co-Simulation of Dy- 1267
 1242 namic Systems with the Functional Mock-up Interface. Ph.D. thesis. 1268
 1243 Lund University. 1269
 1244 Arnold, M., 2010. Stability of Sequential Modular Time Integra- 1270
 1245 tion Methods for Coupled Multibody System Models. Journal of 1271
 1246 Computational and Nonlinear Dynamics 5, 9. doi:10.1115/1. 1272
 1247 4001389. 1273
 1248 Arnold, M., Clauß, C., Schierz, T., 2014. Error Analysis and Er- 1274
 1249 ror Estimates for Co-simulation in FMI for Model Exchange and 1275
 1250 Co-Simulation v2.0, in: Progress in Differential-Algebraic Equa- 1276
 1251 tions, Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 107– 1277
 1252 125. doi:10.1007/978-3-662-44926-4_6. 1278
 1253 Association, I.C., a. Digital twin tutorial. URL: [https:// 1279
 1254 sites.google.com/view/fm2021tutorialdt/home](https://sites.google.com/view/fm2021tutorialdt/home). visited 1280
 1255 July 11th, 2023. 1281
 1256 Association, I.C., b. Fm workshops and tutorials. URL: [http:// 1282
 1257 lcs.ios.ac.cn/fm2021/workshops-and-tutorials/](http://lcs.ios.ac.cn/fm2021/workshops-and-tutorials/). vis- 1283
 1258 ited July 11th, 2023. 1284
 1259 Association, I.C., c. Integrated tool chain for model-based design 1285
 1260 of cps. URL: [https://cordis.europa.eu/project/id/ 1286
 1261 644047](https://cordis.europa.eu/project/id/644047). visited July 11th, 2023. 1287
 1262 Association, I.C., d. Maestro2 documentation. URL: 1288
 1263 [https://into-cps-maestro.readthedocs.io/en/ 1289
 1264 latest/user/index.html](https://into-cps-maestro.readthedocs.io/en/latest/user/index.html). visited July 11th, 2023. 1290
 1265 Association, I.C., 2021. Hybrid testing experiment video. URL: 1291
 1266 <https://youtu.be/-VkrQJaUo1o>. visited July 11th, 2023. 1292
 1293
 1294
 1295
 1296
 1297
 1298
 1299
 1300
 1301
- Association, M., e. SSP standard website. URL: [https://ssp- 1267
 standard.org/](https://ssp-standard.org/). visited July 11th, 2023. 1268
 Battle, N., Thule, C., Gomes, C., Macedo, H.D., Larsen, P.G., 2020. 1269
 Towards a Static Check of FMUs in VDM-SL, in: FM 2019 Inter- 1270
 national Workshops, Springer International Publishing, Porto, 1271
 Portugal. pp. 272–288. doi:10.1007/978-3-030-54997-8_18. 1272
 Behrmann, G., David, A., Larsen, K.G., Håkansson, J., Pettersson, P., 1273
 Yi, W., Hendriks, M., 2006. UPPAAL 4.0, in: QEST 2006, IEEE 1274
 Computer Society. pp. 125–126. doi:10.1109/QEST.2006.59. 1275
 Blochwitz, T., Otter, M., Arnold, M., Bausch, C., Clauss, C., 1276
 Elmqvist, H., Junghanns, A., Mauss, J., Monteiro, M., Neidhold, 1277
 T., Neumerkel, D., Olsson, H., Peetz, J.V., Wolf, S., 2011. The 1278
 Functional Mockup Interface for Tool independent Exchange of 1279
 Simulation Models, in: Proc. of the 8th International Modelica 1280
 Conference, Linköping University Electronic Press; Linköpings 1281
 universitet, Dresden, Germany. pp. 105–114. doi:10.3384/ 1282
 ecp11063105. 1283
 Bogomolov, S., Fitzgerald, J., Foldager, F., Larsen, P.G., Pierce, K., 1284
 Stankaitis, P., Wooding, B., 2020. Tuning Robotti: the Machine- 1285
 assisted Exploration of Parameter Spaces in Multi-Models of a 1286
 Cyber-Physical System, in: Fitzgerald, J.S., Oda, T. (Eds.), Proc. 1287
 of the 18th International Overture Workshop, Overture. pp. 50–64. 1288
 Broman, D., Brooks, C., Greenberg, L., Lee, E.A., Masin, M., Tri- 1289
 pakis, S., Wetter, M., 2013a. Determinate composition of FMUs 1290
 for co-simulation, in: 2013 Proc. (EMSOFT), IEEE. pp. 1–12. 1291
 doi:10.1109/EMSOFT.2013.6658580. 1292
 Broman, D., Derler, P., Eidson, J.C., 2013b. Temporal Issues in 1293
 Cyber-Physical Systems. J. Indian Inst. Sci. 93, 389–402. 1294
 Brück, D., Elmqvist, H., Mattsson, S.E., Olsson, H., 2002. Dymola 1295
 for multi-engineering modeling and simulation, in: Proc. of mod- 1296
 elica, Citeseer. 1297
 Busch, M., 2016. Continuous approximation techniques for co- 1298
 simulation methods: Analysis of numerical stability and local er- 1299
 ror. Journal of Applied Mathematics and Mechanics 96, 1061– 1300
 1081. doi:10.1002/zamm.201500196. 1301

- CES, A., Widl, E., Strasser, T.I., 2021. Erigrd/ns3-fmi-export: v1.1. URL: <https://doi.org/10.5281/zenodo.4638103>, doi:10.5281/zenodo.4638103. the time progression of 0 is mentioned in the readme.md file.
- Committee, F.S., . Fmi website. URL: <https://fmi-standard.org/tools/>. visited July 11th, 2023.
- Committee, F.S., 2014. Functional Mock-up Interface for Model Exchange and Co-Simulation. <https://fmi-standard.org/downloads/>.
- Committee, F.S., 2021. Functional Mock-up Interface for Model Exchange, Co-Simulation, and Scheduled Execution. <https://fmi-standard.org/downloads/>.
- Cremona, F., Lohstroh, M., Broman, D., Lee, E.A., Masin, M., Tripakis, S., 2017. Hybrid co-simulation: it's about time. *Software & Systems Modeling* doi:10.1007/s10270-017-0633-6.
- Cremona, F., Lohstroh, M., Tripakis, S., Brooks, C., Lee, E.A., 2016. FIDE, in: Proc. of the 31st Annual ACM Symposium on Applied Computing, ACM. doi:10.1145/2851613.2851677.
- Ejersbo, H., Lausdahl, K., Frasheri, M., Esterle, L., 2023. fmiSwap: Run-time Swapping of Models for Co-simulation and Digital Twins. arXiv preprint arXiv:2304.07328 .
- Evora Gomez, J., Cabrera, J.J.H., Tavella, J.P., Vialle, S., Kremers, E., Frayssinet, L., 2019. Daccosim NG: co-simulation made simpler and faster, in: Linköping electronic conference proceedings, pp. 785–792. doi:10.3384/ecp19157785.
- Feng, H., Gomes, C., Thule, C., Lausdahl, K., Iosifidis, A., Larsen, P.G., 2021a. Introduction to Digital Twin Engineering, in: Martin, C.R., Blas, M.J., Psijas, A.I. (Eds.), 2021 ANNSIM, Virginia, USA. pp. 19–22.
- Feng, H., Gomes, C., Thule, C., Lausdahl, K., Sandberg, M., Larsen, P.G., 2021b. The incubator case study for digital twin engineering. arXiv:2102.10390.
- FMI, 2020. Functional Mock-up Interface for Model Exchange and Co-Simulation. Standard 2.0.2. URL: <https://fmi-standard.org/downloads/>.
- Frasheri, M., Thule, C., Macedo, H.D., Lausdahl, K.G., Larsen, P.G., Esterle, L., 2021. Fault injecting co-simulations for safety, in: Proceedings of the Fifth International Joint Conference on System Reliability and Safety, 2021. ICSRS 2021, pp. 24–26. Accepted for publication in 5th International Conference on System Reliability and Safety.
- Fritzson, P., 2015. Principles of Object-Oriented Modeling and Simulation with Modelica 3.3: A Cyber-Physical Approach. IEEE Press. 2 ed., Wiley. doi:10.1002/9781118989166.
- Galtier, V., Vialle, S., Dad, C., Tavella, J.P., Lam-Yee-Mui, J.P., Plessis, G., 2015. FMI-Based Distributed Multi-Simulation with DACCOSIM, in: Spring Simulation Multi-Conference, Society for Computer Simulation International, Alexandria, Virginia, USA. pp. 804–811.
- Glaessgen, E., Stargel, D., 2012. The Digital Twin Paradigm for Future NASA and U.S. Air Force Vehicles, in: Structures, Structural Dynamics, and Materials Conference: Special Session on the Digital Twin, American Institute of Aeronautics and Astronautics, Reston, Virginia. pp. 1–14. doi:10.2514/6.2012-1818.
- Gomes, C., Abbiati, G., Larsen, P.G., 2021a. Seismic Hybrid Testing using FMI-based Co-Simulation, in: Proc. of the 14th International Modelica Conference, Linköping University Electronic Press, Linköpings Universitet, online. pp. 287–295.
- Gomes, C., Lucio, L., Vangheluwe, H., 2019a. Semantics of Co-simulation Algorithms with Simulator Contracts, in: 2019 ACM/IEEE 22nd International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C), IEEE, Munich, Germany. pp. 784–789. doi:10.1109/MODELS-C.2019.00124.
- Gomes, C., Meyers, B., Denil, J., Thule, C., Lausdahl, K., Vangheluwe, H., De Meulenaere, P., 2018a. Semantic Adaptation for FMI Co-simulation with Hierarchical Simulators. SIMULATION 95, 1–29. doi:10.1177/0037549718759775.
- Gomes, C., Najafi, M., Sommer, T., Blesken, M., Zacharias, I., Kotte, O., Mai, P., Schuch, K., Wernersson, K., Bertsch, C., Blochwitz, T., Junghanns, A., 2021b. The FMI 3.0 Standard Interface for Clocked and Scheduled Simulations, in: Proc. of the 14th International Modelica Conference, Linköping University Electronic Press, Linköpings Universitet, online. pp. 27–36. doi:10.3384/ecp2118127.
- Gomes, C., Oakes, B.J., Moradi, M., Gamiz, A.T., Mendo, J.C., Dutre, S., Denil, J., Vangheluwe, H., 2019b. HintCO - Hint-Based Configuration of Co-Simulations, in: International Conference on Simulation and Modeling Methodologies, Technologies and Applications, Prague, Czech Republic. pp. 57–68. doi:10.5220/0007830000570068.
- Gomes, C., Thule, C., Broman, D., Larsen, P.G., Vangheluwe, H., 2018b. Co-simulation: A Survey. *ACM Computing Surveys* 51, 49:1–49:33. doi:10.1145/3179993.
- Gomes, C., Thule, C., DeAntoni, J., Larsen, P.G., Vangheluwe, H., 2018c. Co-simulation: The Past, Future, and Open Challenges, in: ISOLA 2018, Springer Verlag, Limassol, Cyprus. pp. 504–520. doi:10.1007/978-3-030-03424-5_34.
- Gomes, C., Thule, C., Larsen, P.G., Denil, J., Vangheluwe, H., 2018d. Co-Simulation of Continuous Systems: A Tutorial. Technical Report arXiv:1809.08463. University of Antwerp, Belgium. arXiv:1809.08463.
- Gomes, C., Thule, C., Lausdahl, K., Larsen, P.G., Vangheluwe, H., 2018e. Stabilization Technique in INTO-CPS, in: 2nd Workshop on Formal Co-Simulation of Cyber-Physical Systems, Springer, Cham, Toulouse, France. pp. 45–51. doi:10.1007/978-3-030-04771-9_4.
- Hafner, I., Popper, N., 2017. On the terminology and structuring of co-simulation methods, in: Proceedings of the 8th International Workshop on Equation-Based Object-Oriented Modeling Languages and Tools, ACM, New York, NY, USA. pp. 67–76. URL: <https://doi.org/10.1145/3158191.3158203>, doi:10.1145/3158191.3158203.
- Hansen, S.T., Gomes, C., Larsen, P.G., Van de Pol, J., 2021a. Synthesizing co-simulation algorithms with step negotiation and algebraic loop handling, in: Martin, C.R., Blas, M.J., Psijas, A.I. (Eds.), 2021 ANNSIM, pp. 1–12. doi:10.23919/ANNSIM52504.2021.9552073.
- Hansen, S.T., Gomes, C., Palmieri, M., Thule, C., van de Pol, J., Woodcock, J., 2021b. Verification of co-simulation algorithms subject to algebraic loops and adaptive steps, in: Lluch Lafuente, A., Mavridou, A. (Eds.), FMICS 2021, Springer International Publishing, Cham. pp. 3–20.
- Hansen, S.T., Gomes, C.G., Najafi, M., Sommer, T., Blesken, M., Zacharias, I., Kotte, O., Mai, P.R., Schuch, K., Wernersson, K., Bertsch, C., Blochwitz, T., Junghanns, A., 2022a. The FMI 3.0 Standard Interface for Clocked and Scheduled Simulations. *Electronics* 11, 3635.
- Hansen, S.T., Ölveczky, P.C., 2022. Modeling, algorithm synthesis, and instrumentation for co-simulation in maude, in: Bae, K. (Ed.), *Rewriting Logic and Its Applications*, Springer International Publishing, Cham. pp. 130–150.
- Hansen, S.T., Thule, C., Gomes, C., 2021c. An FMI-Based Initialization Plugin for INTO-CPS Maestro 2, in: Cleophas, L., Massink, M. (Eds.), SEFM 2020 Collocated Workshops, Springer International Publishing, Cham. pp. 295–310.
- Hansen, S.T., Thule, C., Gomes, C., van de Pol, J., Palmieri, M., Inci, E.O., Madsen, F., Alfonso, J., Castellanos, J.Á., Rodriguez, J.M., 2022b. Verification and synthesis of co-simulation algorithms subject to algebraic loops and adaptive steps. *STTT* 24, 999–1024.
- Hatledal, L.L., Chu, Y., Styve, A., Zhang, H., 2021. Vico: An

- entity-component-system based co-simulation framework. *Simulation Modelling Practice and Theory* 108, 102243. doi:<https://doi.org/10.1016/j.simpat.2020.102243>.
- Inci, E.O., Gomes, C., Croes, J., Thule, C., Lausdahl, K., Desmet, W., Larsen, P.G., 2021. The effect and selection of solution sequence in co-simulation, in: Martin, C.R., Blas, M.J., Psijas, A.I. (Eds.), 2021 ANNSIM, Virginia, USA. pp. 1–12.
- Jochen Köhler, Hans-Martin Heinkel, Pierre Mai, Jürgen Krasser, Markus Deppe, Mikio Nagasawa, 2016. Modelica-Association-Project “System Structure and Parameterization” – Early Insights, Tokyo, Japan. pp. 35–42. doi:<http://dx.doi.org/10.3384/ecp1612435>.
- Junghanns, A., Blochwitz, T., Bertsch, C., Sommer, T., Wernersson, K., Pillekeit, A., Zacharias, I., Blesken, M., Mai, P., Schuch, K., Schulze, C., Gomes, C., Najafi, M., 2021. The Functional Mock-up Interface 3.0 - New Features Enabling New Applications, in: Proc. of the 14th International Modelica Conference, Linköping University Electronic Press, Linköpings Universitet, online.
- Kalmar-Nagy, T., Stanculescu, I., 2014. Can complex systems really be simulated? *Applied Mathematics and Computation* 227, 199–211. doi:[10.1016/j.amc.2013.11.037](https://doi.org/10.1016/j.amc.2013.11.037).
- Kübler, R., Schiehlen, W., 2000a. Modular Simulation in Multi-body System Dynamics. *Multibody System Dynamics* 4, 107–127. doi:[10.1023/A:1009810318420](https://doi.org/10.1023/A:1009810318420).
- Kübler, R., Schiehlen, W., 2000b. Two Methods of Simulator Coupling. *Mathematical and Computer Modelling of Dynamical Systems* 6, 93–113. doi:[10.1076/1387-3954\(200006\)6:2;1-M;FT093](https://doi.org/10.1076/1387-3954(200006)6:2;1-M;FT093).
- Larsen, P.G., Fitzgerald, J., Woodcock, J., Fritzson, P., Brauer, J., Kleijn, C., Lecomte, T., Pfeil, M., Green, O., Basagiannis, S., Sadovykh, A., 2016. Integrated tool chain for model-based design of Cyber-Physical Systems: The INTO-CPS project, in: 2nd International Workshop on Modelling, Analysis, and Control of Complex CPS (CPS Data), IEEE, Vienna, Austria. pp. 1–6. doi:[10.1109/CPSData.2016.7496424](https://doi.org/10.1109/CPSData.2016.7496424).
- Macedo, H.D., Rasmussen, M.B., Thule, C., Larsen, P.G., 2020. Migrating the INTO-CPS Application to the Cloud, in: Sekerinski, E., Moreira, N., Oliveira, J.N., Ratiu, D., Guidotti, R., Farrell, M., Luckcuck, M., Marmsoler, D., Campos, J., Astarte, T., Gonnord, L., Cerone, A., Couto, L., Dongol, B., Kutrib, M., Monteiro, P., Delmas, D. (Eds.), FM 2019 International Workshops, Springer International Publishing, Cham. pp. 254–271.
- Mansfield, M., Gamble, C., Pierce, K., Fitzgerald, J., Foster, S., Thule, C., Nilsson, R., 2017. Examples Compendium 3. Technical Report. INTO-CPS Deliverable, D3.6.
- Martin, A., 2007. Entity systems are the future of mmog development - part 1. URL: <http://t-machine.org/index.php/2007/09/03/entity-systems-are-the-future-of-mmog-development-part-1/>. (Accessed on 06/23/2021).
- McCrum, D.P., Williams, M.S., 2016. An overview of seismic hybrid testing of engineering structures. *Engineering Structures* 118, 240–261. doi:[10.1016/j.engstruct.2016.03.039](https://doi.org/10.1016/j.engstruct.2016.03.039).
- Neema, H., Gohl, J., Lattmann, Z., Sztipanovits, J., Karsai, G., Neema, S., Bapty, T., Batteh, J., Tummescheit, H., Sureshkumar, C., 2014. Model-based integration platform for fmi co-simulation and heterogeneous simulations of cyber-physical systems, in: The 10th International Modelica Conference 2014, Modelica Association, Lund, Sweden.
- Oakes, B.J., Gomes, C., Holzinger, F.R., Benedikt, M., Denil, J., Vangheluwe, H., 2021. Hint-Based Configuration of Co-simulations with Algebraic Loops, in: *Simulation and Modeling Methodologies, Technologies and Applications*. Springer International Publishing, Cham. volume 1260, pp. 1–28. doi:[10.1007/978-3-030-55867-3_1](https://doi.org/10.1007/978-3-030-55867-3_1).
- Ochel, L., Braun, R., Thiele, B., Asghar, A., Buffoni, L., Eek, M., Fritzson, P., Fritzson, D., Horkeby, S., Hällquist, R., Kinnander, Å., Palanisamy, A., Pop, A., Sjölund, M., 2019. OMSimulator - integrated FMI and TLM-based co-simulation with composite model editing and SSP, in: Linköping Electronic Conference Proc., Linköping University Electronic Press. doi:[10.3384/ecp1915769](https://doi.org/10.3384/ecp1915769).
- Ouy, J., Lecomte, T., Foldager, F.F., Henriksen, A.V., Green, O., Hallerstedte, S., Larsen, P.G., Couto, L.D., Antonante, P., Basagiannis, S., Falleni, S., Ridouane, H., Saada, H., Zavaglio, E., König, C., Balcu, N., 2017. Case Studies 3, Public Version. Technical Report. INTO-CPS Public Deliverable, D1.3a. URL: https://into-cps.org/fileadmin/into-cps.org/Filer/D1.3a_Case_Studies.pdf.
- Palensky, P., Van Der Meer, A.A., Lopez, C.D., Joseph, A., Pan, K., 2017. Cosimulation of intelligent power systems: Fundamentals, software architecture, numerics, and coupling. *IEEE Industrial Electronics Magazine*, 34–50doi:[10.1109/MIE.2016.2639825](https://doi.org/10.1109/MIE.2016.2639825).
- Pierce, K., Lausdahl, K., Frasheri, M., 2022. Speeding up design space exploration through compiled master algorithms, in: Macedo, H., Pierce, K. (Eds.), Proc. of the 20th International Overture Workshop, pp. 66–81. doi:[10.48550/arXiv.2208.10233](https://doi.org/10.48550/arXiv.2208.10233). 20th Overture Workshop ; Conference date: 05-07-2022 Through 05-07-2022.
- Schweizer, B., Li, P., Lu, D., 2015. Explicit and Implicit Cosimulation Methods: Stability and Convergence Analysis for Different Solver Coupling Approaches. *Journal of Computational and Nonlinear Dynamics* 10, 051007. doi:[10.1115/1.4028503](https://doi.org/10.1115/1.4028503).
- Schweizer, B., Lu, D., Li, P., 2016. Co-simulation method for solver coupling with algebraic constraints incorporating relaxation techniques. *Multibody System Dynamics* 36, 1–36. doi:[10.1007/s11044-015-9464-9](https://doi.org/10.1007/s11044-015-9464-9).
- Thule, C., Lausdahl, K., Gomes, C., Meisl, G., Larsen, P.G., 2019. Maestro: The INTO-CPS co-simulation framework. *Simulation Modelling Practice and Theory* 92, 45–61. doi:[10.1016/j.simpat.2018.12.005](https://doi.org/10.1016/j.simpat.2018.12.005).
- Thule, C., Palmieri, M., Gomes, C., Lausdahl, K., Macedo, H.D., Battle, N., Larsen, P.G., 2020. Towards Reuse of Synchronization Algorithms in Co-simulation Frameworks, in: *Software Engineering and Formal Methods*, Springer International Publishing, Oslo, Norway. pp. 50–66. doi:[10.1007/978-3-030-57506-9_5](https://doi.org/10.1007/978-3-030-57506-9_5).
- Van Acker, B., Denil, J., Meulenaere, P.D., Vangheluwe, H., 2015. Generation of an Optimised Master Algorithm for FMI Co-simulation, in: *Symposium on Theory of Modeling & Simulation-DEVS Integrative*, Society for Computer Simulation International San Diego, CA, USA, Alexandria, Virginia, USA. pp. 946–953.