

Article

# Verification of Cyberphysical Systems

Marjan Sirjani <sup>1\*</sup>, Edward A. Lee <sup>2</sup> and Ehsan Khamespanah <sup>3</sup>

<sup>1</sup> School of IDT, Mälardalen University, Sweden; marjan.sirjani@mdh.se

<sup>2</sup> Department of EECS, University of California at Berkeley, USA; eal@berkeley.edu

<sup>3</sup> Department of ECE, University of Tehran, Iran; e.khamespanah@ut.ac.ir

\* Correspondence: marjan.sirjani@mdh.se; Tel.: +46-73-662-0517

Version June 6, 2020 submitted to Mathematics

**Abstract:** The value of verification of cyberphysical systems depends on the relationship between the state of the software and the state of the physical system. This relationship can be complex because of the real-time nature and different timelines of the physical plant, the sensors and actuators, and the software that is almost always concurrent and distributed. In this paper, we study different ways to construct a transition system model for the distributed and concurrent software components of a CPS. We describe a logical-time based transition system model, which is commonly used for verifying programs written in synchronous languages, and derive the conditions under which such a model faithfully reflects physical states. When these conditions are not met (a common situation), a finer-grained event-based transition system model may be required. We propose an approach for formal verification of cyberphysical systems using Lingua Franca, a language designed for programming cyberphysical systems, and Rebeca, an actor-based language designed for model checking distributed event-driven systems. We focus on the cyber part and model a faithful interface to the physical part. Our method relies on the assumption that the alignment of different timelines during the execution of the system is the responsibility of the underlying platforms. We make those assumptions explicit and clear.

**Keywords:** Cyberphysical systems, Verification, Lingua Franca, Model checking, Rebeca.

## 1. Introduction

Cyberphysical systems (CPSs) are all around us, as in industrial control systems, robotics, smart grids, autonomous cars, and medical devices. Cyberphysical systems are integrations of computation, networking, and physical processes where physical and software components are deeply intertwined. Cyberphysical systems include networked embedded computers monitoring and controlling the physical processes. They also include mechanical, electrical, chemical, or biological components that are controlled or monitored by computer-based algorithms. A study of CPS may emphasize one or the other perspective. Here, we focus on verification of the software controlling the physical processes not the physical processes being controlled by the software.

Formal verification is about assuring properties of models. A holistic approach to verifying CPSs requires models of both the distributed software and physical processes. However, commonly used models for software are incompatible with commonly used models for physical processes [1]. An alternative is to clearly define the interfaces between the cyber and the physical parts of the system and separate the verification problem, from each side relying on the other side to faithfully carry out the semantics of the interfaces. When verifying software, we rely on the hardware to faithfully carry out the operations specified by the software. Hence, when we prove that the software has some property, such as never reaching some undesired state, we can assume that, with high probability, the physical system that executes the software will reflect a corresponding property. The nature of these

35 interfaces, however, and the underlying assumptions they entail become extremely important. For  
36 CPS, verification is ultimately about assuring properties of the physical world. This means that it is  
37 not sufficient to study the software alone. We need to also study its interactions with its environment.

38 We use a simple example of a train door controller from Sirjani et al in [2] as our running example.  
39 Consider a train door that needs to be locked before the train starts moving. The software controlling  
40 train systems is able to lock the door and then send a command to the train to start moving. We can  
41 build a model of the software, or write a simple program, and formally verify its correctness. But if we  
42 do not know how and when the door gets locked and the train starts moving in response to a software  
43 command, then it will do little good to prove that the software never enters a state where it thinks the  
44 door is unlocked while the train is moving. The necessity to include the physical aspects of the system,  
45 not just its logical ones, is what makes this a CPS.

46 We can verify that the door's software component is never in the unlocked state while the train's  
47 software component is in the moving state. Depending on how the physical interfaces are realized,  
48 however, this may or may not align with the physical world. The state of the software system and  
49 the state of the physical world are not assured of aligning. What if the door component and the train  
50 component are executing on two different microprocessors separated by a network? What does it  
51 mean, in this case, for the two to simultaneously be in some state? To have a useful solution we need  
52 to address the problem of different timelines in distributed systems and different timelines between  
53 the software and the physical world.

54 A cyberphysical system can be viewed as an interacting pair of reactive systems, one defined in the  
55 world of software, and the other in the world of physics. The semantic worlds of physics and software  
56 are radically different and often mutually incompatible. So, to prove properties of cyberphysical  
57 systems, we may not want to combine models from physics with those of software. Our approach is  
58 instead to build a model with a focus on the software side and an abstract (but faithful) model of the  
59 physical side. We model the distributed software system that monitors and receives the data from the  
60 physical processes and sends the control commands to the physical processes. Using this model we  
61 can verify whether, upon receiving certain data, the software system is producing correct output based  
62 on the specified requirements. Modeling the input from the physical world "faithfully," and producing  
63 "correct" output needs more elaboration. The model of the physical world is reduced to its interface to  
64 the software. One major issue in properly modeling this interface is timing.

65 Time is a critical feature in cyberphysical systems. There is the issue of time in distributed software,  
66 and also in the interface of software and the physical world. In order to effectively couple models of  
67 software with models of the physical world, we need modeling frameworks that support more than  
68 one timeline. We will explain later how we rely on certain assumptions to consider one logical timeline  
69 in the model of the software, and how certain guarantees from the selected programming language  
70 and the underlying platform allow us to assume that the logical time and the physical time are aligned  
71 in such a way that the model of the physical interface stays faithful to the physical world. In other  
72 words, our model of the physical inputs is faithful to the physical world and the physical outputs are  
73 created correctly. Note that faithfulness and correctness here depends not only on the *values* of the  
74 inputs and outputs but also on their timing.

75 There are alternative approaches in analyzing CPS that are not the topic of this paper (see  
76 comprehensive overviews in [3,4]). The focus of the approach can be on modeling the physical  
77 processes, the dynamics of the physical quantities. The theory of dynamical control systems is a  
78 well-developed discipline rooted in continuous-time models. In a cyberphysical system, the controller  
79 consists of discrete software with concurrent components operating in multiple possible modes,  
80 interacting with the continuously evolving physical environment. Such systems are often modeled  
81 with a mix of finite automata and continuous dynamics, where mode transitions are modeled by  
82 discrete, instantaneous state transitions in an automaton, and each state of the automaton is associated  
83 with a distinct model of the continuous dynamics. Such models are called hybrid systems [5,6]. We

84 will not consider hybrid system models here. We will instead assume a particular style of software,  
85 embodied in the Lingua Franca language, that yields useful and realistic models.

86 Model checking is a method for formal verification of reactive systems. A model checking tool  
87 receives two inputs: a model of the behavior of the system, and a set of properties represented as  
88 temporal logic formula showing the desired specification of the same system. A state transition  
89 diagram is generated based on the interleaved semantics of the input model and the properties are  
90 checked against this state transition diagram. Here we use the Reactive Object Language, Rebeca [7,8],  
91 and its model checking tool Afra [9] for formal verification of cyberphysical systems. For doing so,  
92 we map Lingua Franca [10–12] programs to an extended version of Timed Rebeca. Lingua Franca is a  
93 programming language based on the Reactor model of computation [11] for building cyberphysical  
94 systems. Both Rebeca and Lingua Franca are actor-based languages.

95 The Hewitt Actor model [13,14] is a reference model for concurrent distributed systems with an  
96 asynchronous event-driven model of computation. Its event-driven concurrent semantics makes it a  
97 natural choice for modeling cyberphysical systems, but it needs to be extended with timing properties.  
98 To model the unknown factors in a system, like the possible inputs from the environment, we can  
99 use nondeterminism in our Timed Rebeca model. Timed Rebeca [15,16] extends Rebeca to model the  
100 timing features. In Timed Rebeca the events are triggered based on their timetag order, and when there  
101 is more than one enabled event with the same logical time tag, they are triggered in nondeterministic  
102 order. Lingua Franca ensures determinism in a similar way, by first ensuring that messages are handled  
103 in timetag order and then also prioritizing reactions within each reactor. To handle simultaneous  
104 messages to distinct reactors, LF uses the precedence graph relation between reactors to constrain the  
105 order of execution. To model the deterministic behavior of LF, we have extended Timed Rebeca with  
106 priorities on message handlers and priorities on actors so that simultaneous messages (those with the  
107 same timetag) are handled in a deterministic order. For ordering the execution of actors, priorities are  
108 too strong, but they work for the purpose of this paper, which is verification.

109 There are multiple timelines involved in cyberphysical systems. To have a faithful model of  
110 time for cyberphysical systems we need to address both (1) the asynchrony in distributed systems,  
111 and (2) the mismatch between physical and logical time. To make analysis possible we need to  
112 build layers of abstraction and use assumptions by relying on the other layers. In Timed Rebeca we  
113 assume synchronized local clocks for actors that gives us a notion of global time across the model.  
114 We use logical timetags, and logical timetags are comparable across all actors in the model. But in  
115 distributed systems we cannot assume synchronized clocks for distributed software components, at  
116 least not perfect ones. We need certain mechanisms to be able to have such assumption. Ptimes[17]  
117 and Spanner[18] are two examples that assume synchronized clocks (up to an error bound) and use  
118 logical timetags. For distributed actors (as faithful representatives of distributed software components)  
119 to be able to have synchronized clocks and comparable timetags we rely on the lower-level network  
120 protocols to provide that for us. The second issue is the two timelines of logical world and physical  
121 world.

122 Lingua Franca includes a notion of “logical time” and binds that notion to “physical time” only  
123 where the software interacts with the physical world. In Lingua Franca the logical timetag of the input  
124 events are assigned based on the physical time of the physical processes. We also need to make sure  
125 that the logical timetag of the output events and the physical time of the actuated physical processes  
126 have the desired relation. We assume that the actuated outputs affect the physical world within a  
127 certain deadline.

128 Our model stays faithful to the system itself only based on the set of assumptions mentioned  
129 above. These assumptions allow us to reason about the system based on the logical timetags in our  
130 Rebeca model. Our model may not be a model with the least semantic gap with cyberphysical systems,  
131 but we will show in this paper that using model checking we are able to catch many subtle design  
132 problems. We show how these problems may exist in very simple examples that exhibit how building  
133 such systems can be extremely error-prone. Many of these problems may be related to the timing

134 configurations. We believe that no simple approach exists for verification of cyberphysical systems.  
135 Several complimentary methods need to be used to cover the analysis of different aspects of such  
136 systems, and in each method we rely on certain assumptions that may be guaranteed by other methods.  
137 In a shorter conference paper version of this paper [19], we raise the interesting questions involved  
138 in verification of cyberphysical systems and we used a couple of examples to show how we move  
139 towards solving some of the problems. Here we explain the problem and the solutions in a more  
140 extensive and structured way using mostly the same examples. The paper is organized as follows.  
141 Section 2 introduces the programming model we assume (reactors) and the language in which programs  
142 are written (Lingua Franca, LF). It sketches the source code in LF for a running example, a train door  
143 controller. Section 3 introduces the Rebeca language and its timed extension, Timed Rebeca, which  
144 we further extend here to express temporal properties of Lingua Franca programs. Section 4 explains  
145 a translation of the train door running example into this extended Timed Rebeca. Section 5 studies  
146 the problem of model checking concurrent LF programs and explains two approaches based on two  
147 different semantics with different levels of granularity. Section 6 refines the train door example with  
148 programming constructs to control timing and increased interactivity and shows how the Rebeca  
149 model checking tool Afra can help identify subtle defects in the design. Section 7 concludes with a  
150 discussion of problems that remain open.

## 151 2. Lingua Franca and Reactors: Building Cyberphysical Systems

152 Lingua Franca (LF) [10–12], is a coordination language designed for embedded real-time systems.  
153 Software components are called “reactors.” The messages exchanged between reactors have logical  
154 timetags drawn from a discrete, totally ordered model of time. Any two messages with the same  
155 timetag are logically simultaneous, which means that, for any reactor with these two messages as  
156 inputs, if it sees that one message has occurred, then it will also see that the other has occurred.  
157 Moreover, every reactor will react to incoming messages in timetag order. If the reactor has reacted to  
158 a message with timetag  $t$ , no future reaction in the same reactor will see any message with a lesser  
159 timetag.

160 If a reactor produces output messages in reaction to an input, then, by default, the logical time  
161 of the output will be identical to the logical time of the input. This principle is borrowed from  
162 synchronous languages [20]. The Lingua Franca compiler ensures that all logically simultaneous  
163 messages are processed in precedence order, so the computation is deterministic. At a logical instant,  
164 the semantics of the program can be given as a unique least fixed point of a monotonic function on a  
165 lattice [21], so the computation is deterministic, even if it is distributed across a network. We call this  
166 semantics, based on the semantics of synchronous languages, the “logical-time-based semantics.” Here,  
167 we also consider an event-based semantics, which becomes useful when an interleaved execution of  
168 events with the same logical timetag becomes observable. Event-based semantics has finer granularity  
169 compared to logical-time-based semantics.

170 The syntax of a subset of Lingua Franca is given in Figure 1. The model consists of a set of reactors  
171 and a main reactor. Reactors contain state variables, input and output ports, physical actions and  
172 reactions. The body of reactions can be written in the target language. As of this writing, LF supports  
173 C, C++, and TypeScript. In each case, the LF compiler generates a standalone executable in the target  
174 language. A reactor may also react to a “physical action,” which is typically triggered by some external  
175 event such as a sensor [11]. The physical action will be assigned a timetag based on the current physical  
176 clock on the machine hosting the reactor.

177 A key semantic property of Lingua Franca is that every reactor reacts to events in timetag order.  
178 Preserving this order in a distributed execution is a key challenge. One technique that has proven  
179 effective is Ptides [22], a decentralized and fault-tolerant coordination mechanism that relies on  
180 synchronized physical clocks with bounded error. The Ptides technique has been applied on a global  
181 scale in Google Spanner [23].

182       Lingua Franca includes a notion of a deadline, which is a relation between logical time and  
 183 physical time, as measured on a particular platform. Specifically, a program may specify that the  
 184 invocation of a reaction must occur within some physical-time interval of the logical timestamp of  
 185 the message. This, together with physical actions, can be used to ensure some measure of alignment  
 186 between logical time and some measurement of physical time.

```

    Model ::= Target Reactor* MainReactor
    Target ::= target targetLanguageName;
    Reactor ::= reactor reactorName { StateVar* Input* Output* Action* Reaction* }
    StateVar ::= state varId : typeId (initialValue);
    Input ::= input inputId : typeId;
    Output ::= output outputId : typeId;
    Action ::= physical action actionId : typeId;
    Reaction ::= reaction (Trigger*) [-> outputId(, outputId)*] {= Code*=}
    Trigger ::= inputId | actionId
    Code ::= Target – Language – Statement
    MainReactor ::= main reactor mainReactorName { Instantiation* Connection* }
    Instantiation ::= id = new reactorName();
    Connection ::= id.inputId -> id.outputId [after delayValue];
  
```

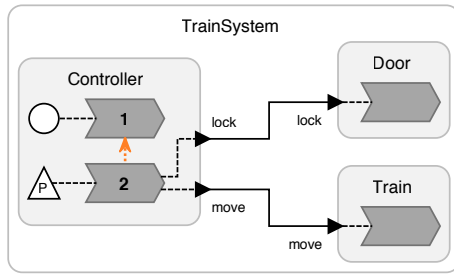
**Figure 1.** Syntax of a subset of Lingua Franca that we use in our examples in this paper (adapted from Lingua Franca Github [24]). The syntax is written in a slightly revised version of Extended BNF where instead of putting terminals in quotations we use words in “**bold**” format. Angled brackets ⟨...⟩ are used as meta parenthesis, superscript + for repetition at least once, superscript \* for repetition zero or more times, whereas using ⟨...⟩ with repetition denotes a comma separated list. Brackets [...] indicates that the text within the brackets is optional. In the syntax, *targetLanguageName*, *reactorName*, *mainReactorName* stand for the target language for LF, the name of the reactor, and the name of the main reactor, respectively. The *varId*, *typeId*, *inputId*, *outputId*, *actionId* stand for the names of a variable, a type, an input and an output, respectively; and *id* stands for the name of an instance of a reactor. *Target-Language-Statement* stands for the statements of the target language. In the *Reactor* rule the components do not need to come in the presented order.

## 187 2.1. The Simple Train Door Controller in Lingua Franca

188       Consider a train door that needs to be locked before the train starts moving. The software  
 189 controlling train systems is able to lock the door and then send a command to the train to start moving.  
 190 Consider in Figure 2b the sketch of an implementation of a highly simplified version of such train  
 191 controller software in Lingua Franca. In this use, the code shown in the figure gets translated into C  
 192 code that can run on a train’s microcontrollers. Similar realizations could be built in any of a number  
 193 of model-based design languages, including any of the synchronous languages [20] (SCADE, Esterel,  
 194 Lustre, SIGNAL, etc.), Simulink, LabVIEW, ModHel’X [26], Ptolemy II [27], or ForSyDe [28], to name a  
 195 few. All will raise similar issues to those we address in this paper.

196       The structure of the code is illustrated in Figure 2a. It consists of three components called  
 197 “reactors,” instances of the reactor classes Controller, Door, and Train. The main reactor (starting on  
 198 line 30) instantiates and connects these components so that the controller sends a messages to both the  
 199 door and the train. These components could be implemented on a single core, on multiple cores, or on  
 200 separate processors connected via a network.

201       Let’s focus first on the interaction between these components and the physical world. The  
 202 Controller reactor class defines a `physical action` named “`external_move`” (line 5), which in Lingua



(a) Structure of the simple door controller example. This image is rendered automatically by the Lingua Franca IDE using the KIELER Lightweight Diagrams framework [25].

```

1 target C;
2 reactor Controller {
3   output lock:bool;
4   output move:bool;
5   physical action external:bool;
6   reaction(startup) {=
7     ... Set up sensing.
8   =}
9   reaction(external)->lock, move {=
10    set(lock, external_value);
11    set(move, external_value);
12  =}
13 }
14 reactor Train {
15   input move:bool;
16   state moving:bool(false);
17   reaction(move) {=
18     ... actuate to move or stop
19     self->moving = move;
20  =}
21 }
22 reactor Door {
23   input lock:bool;
24   state locked:bool(false);
25   reaction(lock) {=
26     ... Actuate to lock or unlock door.
27     self->locked = lock;
28  =}
29 }
30 main reactor System {
31   controller = new Controller();
32   door = new Door();
33   train = new Train();
34   controller.lock -> door.lock;
35   controller.move -> train.move;
36 }

```

(b) Lingua Franca code for the simple door controller example in Figure 2a with a potential defect.

```

1 reactiveclass Controller(5) {
2   knownrebecs {
3     Door door;
4     Train train;
5   }
6   statevars { boolean moveP; }
7   Controller() {
8     self.external();
9   }
10  msgsrv external() {
11    boolean oldMoveP = moveP;
12    moveP = ?(true,false);
13    if(moveP != oldMoveP) {
14      door.lock(moveP);
15      train.move(moveP);
16    }
17    self.external() after(1);
18  }
19 }
20 reactiveclass Train(5) {
21   statevars { boolean moving; }
22   Train() {
23     moving = false;
24   }
25   msgsrv move(boolean tmove) {
26     if (tmove) {
27       moving = true;
28     } else {
29       moving = false;
30     }
31   }
32 }
33 reactiveclass Door(5) {
34   statevars { boolean is_locked; }
35   Door() {
36     is_locked = false;
37   }
38   msgsrv lock (boolean lockPar) {
39     is_locked = lockPar;
40   }
41 }
42 main {
43   @priority(1) Controller controller(door,
44     train):();
45   @priority(2) Train train():();
46   @priority(2) Door door():();
47 }

```

(c) Timed Rebeca model (extended with priorities) for the simple door controller example in Figure 2a.

**Figure 2.** The structure, Lingua Franca program, and Timed Rebeca model for the simple door controller example.

203 Franca is an event that is triggered by something outside the software system and is then assigned a  
 204 logical timetag that approximates the physical time at which that something occurred in the physical  
 205 world [11]. In practice, in the `reaction(startup)` block of code (starting on line 6), which executes  
 206 upon startup of the system, the reactor could set up an interrupt service routine (ISR) to be invoked  
 207 whenever the driver pushes a button to make the door lock and train move. The ISR would call an LF  
 208 function `schedule` to trigger the action and assign it a timetag. The reaction to the `external_move`

209 action (starting on line 9) will be invoked when logical time reaches the assigned timetag. This reaction  
210 sets the outputs named “lock” and “move” to the Boolean value true. Since that outputs are connected  
211 to the input named “lock” of the door component (line 34) and the input named “move” of the train  
212 component (line 34), respectively, this results in a message to the door component and a message to  
213 the train component at the logical time of the timetag.

214 The train component has a state variable named “moving” (line 16) that changes value when it  
215 receives a message on its “move” input port (line 19). The variable has value true when the train is  
216 moving and false when the train is stopped. The door component has a state variable named “locked”  
217 (line 24) that changes value when it receives a message on its “lock” input port (lines 23 and 27).

### 218 3. Timed Rebeca: Model Checking Cyberphysical Systems

219 The Reactive Object Language, Rebeca [7,8], is an actor-based [13,14] modeling language  
220 supported by a model checking tool Afra [9]. Rebeca is used for modeling and formal verification  
221 of concurrent and distributed systems. The model of computation in Rebeca is event-driven and the  
222 communication is asynchronous. The grammar is shown in Figure 3. Actors have message queues;  
223 each actor takes the message on the top of the queue, executes the method related to that message  
224 (called message server) in an atomic and non-preemptive way. While executing a method, messages  
225 can be sent to other actors (or itself), and the values of the state variables can change. Sending messages  
226 is non-blocking, and there is no explicit receive statement.

227 In Timed Rebeca [15,29,30] three keywords are added to model logical time: `delay`, `after` and  
228 `deadline`. Timetags are attached to messages and states of each actor. Here we have a buffer of  
229 timetagged messages instead of a message queue. Using the keyword `delay`, one can model progress  
230 of time while executing a method. If a `send` statement is augmented by `after(t)`, the timetag of the  
231 message when it is put in the queue of the receiver is  $t$  units more than the timetag of the message  
232 when it is sent. The timetag of the message when it is sent is the current logical time of the sender. By  
233 using `after`, one can model the network delay; periodic events can be modeled using `send` messages  
234 to itself augmented by `after`. The `deadline` keyword models the timeout; if the current time of the  
235 receiver actor at the time of triggering the event (taking the message to handle it) is more than the  
236 expressed deadline then the model checking tool will complain and raise the deadline-miss warning.  
237 While mapping Lingua Franca programs to Timed Rebeca we only use the `after` construct and it is  
238 used to increase the value of the logical timetag of the message, like in LF.

239 The original Rebeca language does not have a model of time and handles incoming messages in  
240 nondeterministic order. Timed Rebeca adds a model of time, but still handles incoming messages at  
241 each logical time in nondeterministic order. Our extension supports annotating Rebeca actors, and also  
242 their message servers, with priorities. These priorities can enforce the ordering constraints on message  
243 handlers that are defined by the Lingua Franca language.

244 The external physical inputs in Lingua Franca are modeled as sending those messages to self.  
245 These messages are sent to self augmented with the `after` construct. We can assign nondeterministic  
246 values to `after` and hence the messages are received some time later nondeterministically. Of course  
247 we can also model a periodic physical input by assigning the period of arrival as the value of the `after`  
248 construct.

#### 249 3.1. A Simple Train Door Controller in Timed Rebeca

250 A (slightly simplified) Timed Rebeca model of the program in Figure 2b is shown in Figure 2c.  
251 Given this model, we can use the Afra model checking tool to get the transition system model and to  
252 check safety properties. An interesting point in this Rebeca code is modeling the production of the  
253 stimulus that triggers reactions. We need to model the environment or the interface to the physical  
254 world. On line 8, the constructor for the Controller sends itself the message `external`. On line 12  
255 in the `external` method the value of `moveP` is set to `true` or `false` nondeterministically to show the  
256 possibility of presence or absence of the external message. This is how we model possible external

```

Model ::= Class* Main
Class ::= reactiveclass className (queueLength) { KnownRebecs Vars Constructor MsgSrv* }
KnownRebecs ::= knownrebecs { VarDcl* }
Vars ::= statevars { VarDcl* }
VarDcl ::= type ⟨v⟩+;
Constructor ::= className ((type v)*) { Stmt* }
MsgSrv ::= msgsrv methodName((type v)*) { Stmt* }
Stmt ::= v=e; | v=?(e⟨,e⟩+); | Call; | if (e) { Stmt* } [else { Stmt* }]; | delay(t);
Call ::= rebecName.methodName((e)*) [after(t)][deadline(t)]
Main ::= main { InstanceDcl* }
InstanceDcl ::= className rebecName ((rebecName)*):(⟨literal⟩*);

```

**Figure 3.** Syntax of Timed Rebeca (adapted from [30]). The notation is the same as that in Figure 1. Identifiers *className*, *rebecName*, *methodName*, *queueLength*, *v*, *literal*, and *type* denote class name, rebec name, method name, queue length, variable, literal, and type, respectively; and *e* denotes an (arithmetic, boolean or nondeterministic choice) expression. In the instance declaration (rule *InstanceDcl*), the list of rebec names ((*rebecName*)\*) passed as parameters denotes the known rebecs of that instance, and the list of literals ((*literal*)\*) denotes the parameters of its constructor.

257 stimulus at different times. If this value is changed from the previous period (comparing moveP and  
 258 oldMoveP on line 13) then the two message servers lock and move are called to lock (or unlock) the  
 259 door and move (or stop) the train (lines 14 and 15). This external message is sent to itself every one  
 260 time unit by the controller (line 17).

#### 261 4. Mapping of Reactors to Timed Rebeca with Priorities

262 Table 1 shows the mapping between Reactors and Timed Rebeca (extended with priorities). Each  
 263 reactor in Lingua Franca is mapped to a reactive class, and each reaction is mapped to a message server  
 264 in Rebeca. The trigger in a reaction is the name of the message server, and states in LF are mapped to  
 265 state variables in Rebeca. Rebeca is an object-based language, not a component-based one. Actors call  
 266 each other instead of writing on a port. In Lingua Franca we build the bindings between inputs and  
 267 outputs explicitly in the connection part of the program. In LF a reaction reacts to a trigger, and the  
 268 trigger is one of the inputs to the reactor. A reaction has outputs and those outputs are set by assigning  
 269 values to them. Then in the connection part of the main reactor, all the bindings are set by defining  
 270 which input of which reactor is connected to which output of which reactor. This way the flow of data  
 271 is realised. You can change the topology by changing the connections. In Rebeca, a message server of  
 272 other rebecs (or self) is called, and that is how the binding and the flow is realised. There is also a list  
 273 of known rebecs in a reactive class that shows the rebecs to whom you may send messages to.

274 For the timing issues, there is an after keyword in Lingua Franca that has the same semantics as  
 275 in Timed Rebeca. The timetag of the sent message is increased by the value of the after. Rebeca has a  
 276 delay construct which is not used in LF. Delay in Rebeca increases the timetag within a message server.  
 277 This has no use in synchronous languages.

278 In Lingua Franca the messages are handled in timetag order, for the messages with the same  
 279 timetag the reactions are prioritized within each reactor. To handle simultaneous messages to distinct  
 280 reactors, LF uses the precedence graph relation between reactor to constrain the order of execution. To  
 281 faithfully model the LF programs, Timed Rebeca is extended with priorities. The priorities are added by  
 282 annotations to both message servers and rebecs. The precedence graphs in LF cannot necessarily be  
 283 mapped into priorities, but priorities are enough for the purpose of this paper. Adding the information



284 of precedence graphs to a Rebeca model in the main can be done with no difficulty and is considered  
285 as a future work.

286 We can also describe part of the mapping using the structure diagram in Figure 2a. The triangle  
287 with the “P” is the physical action in LF and external message server in Rebeca, the circle is the  
288 “startup” event in LF and the message sent in the constructor message server in Rebeca, the V-shape  
289 arrows are reactions in LF and message servers in Rebeca, and the red arrows between the reactions  
290 (message servers) are dependencies in LF, and priorities in Rebeca.

291 The mapping between Reactors and Timed Rebeca is natural and can easily be done. In Lingua  
292 Franca we can write the body of reactions in any target language that LF supports. In this work we  
293 write the body of reactions in Rebeca. After the code is model checked and debugged, then the Rebeca  
294 code needs to be translated to one of the languages supported by LF to be able to execute the LF  
295 program. Many design problems can be revealed by model checking the abstract model when the  
296 complicated target code is not yet in place. We can also consider mapping the target codes to Rebeca,  
297 that is left for the future work.

Lingua Franca Construct/Features	Timed Rebeca Construct/Features
<i>reactor</i>	<i>reactiveclass</i>
<i>reaction</i>	<i>msgsrv</i>
<i>trigger</i>	<i>msgsrv name</i>
<i>state</i>	<i>statevars</i>
<i>input</i>	<i>msgsrv</i>
<i>output</i>	<i>known rebecs</i>
<i>physical action</i>	<i>msgsrv</i>
implicit in the topology	<i>Priority</i>
<i>main</i>	<i>main</i>
<i>instantiation (new)</i>	<i>instantiation of rebecs</i>
<i>connection</i>	<i>implicit in calling message servers</i>
<i>after</i>	<i>after</i>
–	<i>delay</i>

Table 1. The mapping between Lingua Franca and Timed Rebeca

## 298 5. Logical-time-based and Event-based Semantics

299 A transition system model, which is needed for model checking, requires a concept of the “state”  
300 of a system at a particular “instant in time.” It does not require that “time” be Newtonian time,  
301 measured in seconds, minutes, and hours and aligned to the Earth’s orbit around the sun. Instead,  
302 it only requires a concept of simultaneity, where the “state” of the system is the composition of the  
303 states of its components at a “simultaneous instant,” whatever that means in the model. In Lingua  
304 Franca, we can define a “simultaneous instant” to be the endpoint when all reactions at a logical time  
305 have completed. The “state” at that “instant” can be defined to be the combination of the state variable  
306 valuations of all the reactors at that “instant.” This is the approach commonly used in synchronous  
307 languages, where transient states during the computation at a logical time are ignored. We call this  
308 interpretation a **logical-time-based semantics**.

309 To perform verification formally, we need to build a state-transition model of the program. Figure  
310 4b gives the logical-time-based semantics of the program in Figure 2b. In the initial state, the door  
311 is unlocked and the train is not moving. This state transition system shows that at each logical time,  
312 the program will nondeterministically either remain in the same state (indicated by the self-loop  
313 transitions) or change to the other state. Once the program is in the new state, at subsequent logical  
314 times, it will similarly nondeterministically remain in the same state or transition back to the initial  
315 state. This transformation relies on the semantics of Lingua Franca being rooted in the fixed-point  
316 semantics of synchronous languages [21].

317 Looking at Figure 4b, it is obvious that the model never enters a state where the train is moving  
318 and the door is unlocked. The transition system model is so simple in this case that there is no need for  
319 a model checker to verify this property.

320 This approach to verification is sound because it accurately and correctly models the semantics  
321 of the program. But the astute reader should be nervous. What if the door component and the train  
322 component are executing on two different microprocessors separated by a network? In this case, there  
323 will be a physical time delay between when the train begins moving and the door gets locked, even  
324 if there is no logical time delay. In this case, the verification exercise is simply misleading unless we  
325 consider this delay in our model.

326 In the Lingua Franca software, the offending physical state of the system, where the train is  
327 moving and door is unlocked, is a transitory state occupied briefly during the computation at a logical  
328 time instant. Its duration in logical time is exactly zero. If the physical system is designed in such a  
329 way that the physical environment can only observe states with non-zero logical time duration, then  
330 we can have confidence in the safety conclusion.

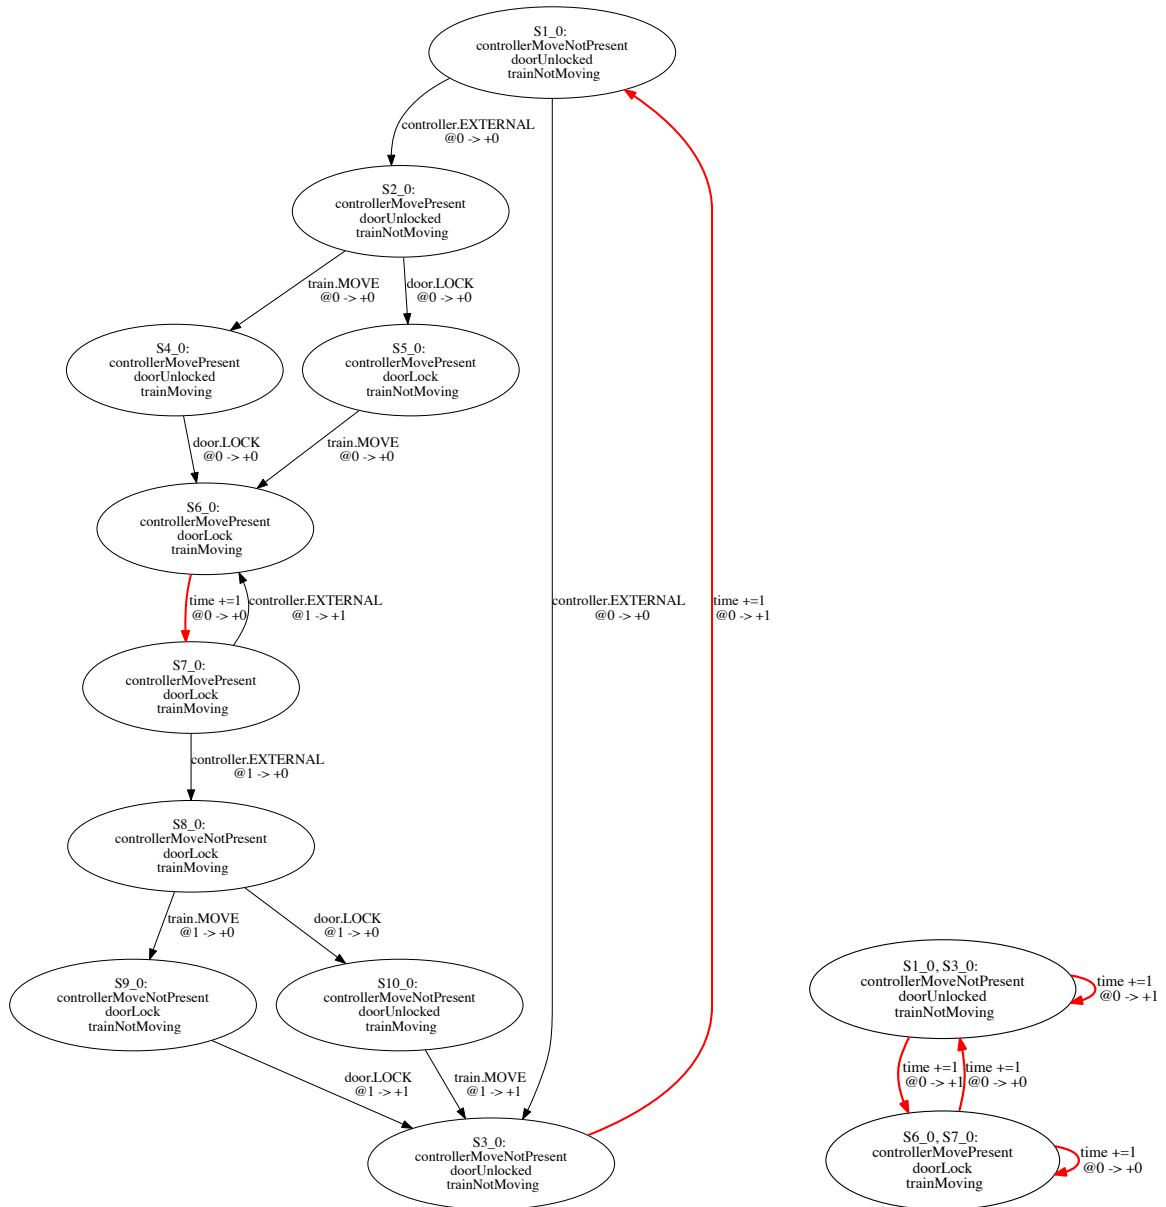
331 It is not uncommon to design control system hardware precisely to make such guarantees.  
332 Programmable Logic Controllers (PLCs), which are widely used to control machinery in industrial  
333 automation, have mechanisms that provide such guarantees [31,32]. In particular, PLC software  
334 does not directly interact with physical actuators. Instead, during a cycle of execution, the software  
335 components write commands to a buffer in memory, and only after the cycle is complete does the  
336 hardware read from that memory and drive the physical actuators. If the memory goes through  
337 transitory unsafe states during the execution of a cycle, those unsafe states are guaranteed to have no  
338 effect on the physical world. If Lingua Franca were to be deployed on hardware with such an I/O  
339 system, where a “cycle” is defined by the completion of all reactions at a logical time, then no safety  
340 violation would occur. However, this conclusion is not based on the program alone, but rather on a  
341 deep and tricky analysis of the program and the hardware on which it is executing. Moreover, the  
342 PLC-style semantics is difficult to realize on a distributed system. If the Door component and the Train  
343 component are executing on distinct microprocessors, then ensuring that their actuations occur only  
344 after a logical-time cycles has been completed requires fairly sophisticated distributed control over the  
345 program execution. Perhaps a better approach is to model the steps in the execution in more detail and  
346 attempt to design the program to be safe even without such a sophisticated I/O system. We will do  
347 that next.

348 A Lingua Franca execution can be modeled as a sequence of reaction invocations, where each  
349 reaction is atomic. We call such a model an **event-based semantics**. It is more fine grained than the  
350 logical-time-based semantics and it includes a sequence of steps performed during a logical time  
351 instant. Each step is one invocation of a reaction in the Lingua Franca program. Each reaction is  
352 triggered by one or more “events,” where an “event” is either a message sent between components  
353 or an action that has been scheduled by a call to the **schedule** function in Lingua Franca. Every such  
354 event occurs at a logical time instant.

### 355 5.1. The State Transition Diagram and the Safety Property of the Example

356 For this simple system, the safety property of interest is that the door be locked while the train is  
357 moving. This can be posed as a formal verification problem, where the goal is to prove this property.  
358 In the program shown in Figure 2b, the door and train components have state variables, and we can  
359 attempt to verify that the door is never in the unlocked state while the train is in the moving state.  
360 Depending on how the physical interfaces are realized, however, this may or may not align with the  
361 physical world. We can use the features of Lingua Franca to assure that the state of the software system  
362 and the state of the physical world are aligned.

363 We can use Afra model checking tool to get the event-based state transition system of the Rebeca  
364 model in Figure 2c and to check safety properties. The event-based transition system is shown in Figure  
365 4a. The transitions shown in black in Figure 4a are transitions that all occur at the same logical time.



**Figure 4.** Transition system model of the Timed Rebeca model in Figure 2c

366 The transitions shown in red coincide with the advancement of logical time. Thus, Figure 4b can be  
367 understood to be an abstraction of this transition diagram that aggregates all the intermediate states at  
368 each logical time into one single state. The self-loops in Figure 4b are represented as the transitions  
369 from S6\_0 to S7\_0 and back, and S1\_0 to S3\_0 and back in Figure 4a.

370 The transition system of Figure 4a is a slightly revised version of the transition system generated  
371 automatically by Afra. In this transition system, the state labeled “S4\_0” violates our safety requirement.  
372 The train is moving and the door is unlocked. There is a safe trace, going through S5\_0 instead of  
373 S4\_0, but the interleaving semantics allows either trace. Similarly, the state labeled “S10\_0” is also not  
374 safe. Here we see the so-called diamond effect that is well-known in the model checking domain and  
375 may be created when two transitions are enabled in the same state (like in states “S2\_0” and “S8\_0” )  
376 and are chosen nondeterministically. If the I/O system makes these transitory states invisible to the  
377 environment, as could be done using the PLC style of I/O, then we do not need this finer grained  
378 transition system model and could instead have verified the safety property using the much simpler  
379 logical-time-based model of Figure 4b. Without such an I/O system, however, we have more work to  
380 do before we can have confidence in this system.

## 381 6. Extending the Simple Train Door Controller

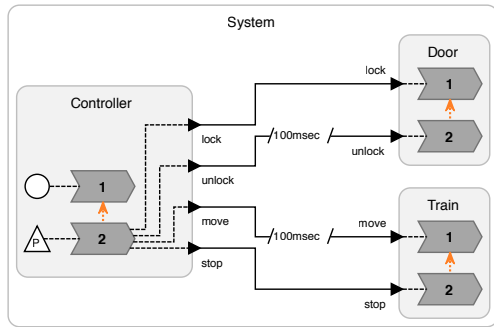
382 We use variations of a simple train door example from Sirjani et al. [2] to show how we address  
383 different questions raised in verification of CPS. In Figure 2, we show a model with three components,  
384 a train, an external door in the train, and a controller that commands the door to lock, and the train to  
385 move. In Section 6.1, we add timing features to the example and show how we can fix the problem of  
386 program in Figure 2 by the proper timing features. We also show the subtleties with the timing and  
387 how we can easily make design mistakes. We also show how the external physical triggers can put  
388 the system at risk and jeopardize the safety property. In Section 6.2, we show an example where an  
389 external physical action can block the progress of system. For example a passenger can keep pressing  
390 the open door button and hence stopping the door from being closed and locked, and as a consequence  
391 prevent the departure of the train.

### 392 6.1. The Train Door Controller with Timing Features

393 The flaw identified by the Afra tool in the program in Figure 2b can be corrected with a  
394 slightly more sophisticated Lingua Franca program. Note that the flaw only exists if we consider the  
395 event-based semantics of the program. A simple way is to define two reactions of move and stop in  
396 the train reactor (instead of just one reaction of move that decides to actuate move or stop based on  
397 the input parameter), and lock and unlock in the door reactor (instead of just one reaction of lock  
398 that decides to actuate lock or unlock based on the input parameter), and increment the timetag of an  
399 unlock or move message so that it has a logical timetag that is strictly larger than the corresponding  
400 stop or lock message. Such a Lingua Franca program is shown in Figure 5b. It has the structure shown  
401 in Figure 5a.

402 Here, we use the `after` keyword on lines 44 and 45 to increment the timetag of the messages by  
403 a specified amount (100 msec). This keyword has exactly the same semantics in Lingua Franca and  
404 Timed Rebeca, so it creates no complications in translation. With these changes, when the Controller  
405 requests that the train move, it issues a lock message with the timetag of the original request and  
406 a move message with a timetag incremented by 100 msec. When it requests that the train stop, the  
407 unlock message is similarly delayed. This change required separating the lock from the unlock signal  
408 and the move from the stop signal because the logical time properties of these pairs of signals differ. In  
409 Figure 2a, by contrast, lock and unlock are carried by a single Boolean, as are move and stop.

410 We can adjust the Timed Rebeca model to match this new design (see Figure 5c) and re-run the  
411 model checker. This time, Afra reveals a more subtle problem that can occur if the system has no  
412 constraints on the spacing between timetags of successive external events. Suppose that the train is  
413 stopped and the door is unlocked and we received `external = true` at logical time 0. This will result



(a) Structure of the door controller example

```

1 target C;
2 reactor Controller {
3   output lock:bool; output unlock:bool;
4   output move:bool; output stop:bool;
5   physical action external:bool;
6   reaction(startup) {=
7     ... Set up external sensing.
8   =}
9   reaction(external)
10    ->lock, unlock, move, stop {=
11    if (external_value) {
12      set(lock, true); set(move, true);
13    } else {
14      set(unlock, true); set(stop, true);
15    }
16  =}
17 }
18 reactor Train {
19   input move:bool; input stop:bool;
20   state moving:bool(false);
21   reaction(move) {=
22     self->moving = true;
23   =}
24   reaction(stop) {=
25     self->moving = false;
26   =}
27 }
28 reactor Door {
29   input lock:bool; input unlock:bool;
30   state locked:bool(false);
31   reaction(lock) {=
32     ... Actuate to lock door.
33     self->locked = true;
34   =}
35   reaction(unlock) {=
36     ... Actuate to unlock door.
37     self->locked = false;
38   =}
39 }
40 main reactor System {
41   c = new Controller(); d = new Door();
42   t = new Train();
43   c.lock -> d.lock;
44   c.unlock -> d.unlock after 100 msec;
45   c.move -> t.move after 100 msec;
46   c.stop -> t.stop;
47 }

```

(b) Variant of Figure 2b that manipulates timetags. The values for **after** are set to 100.

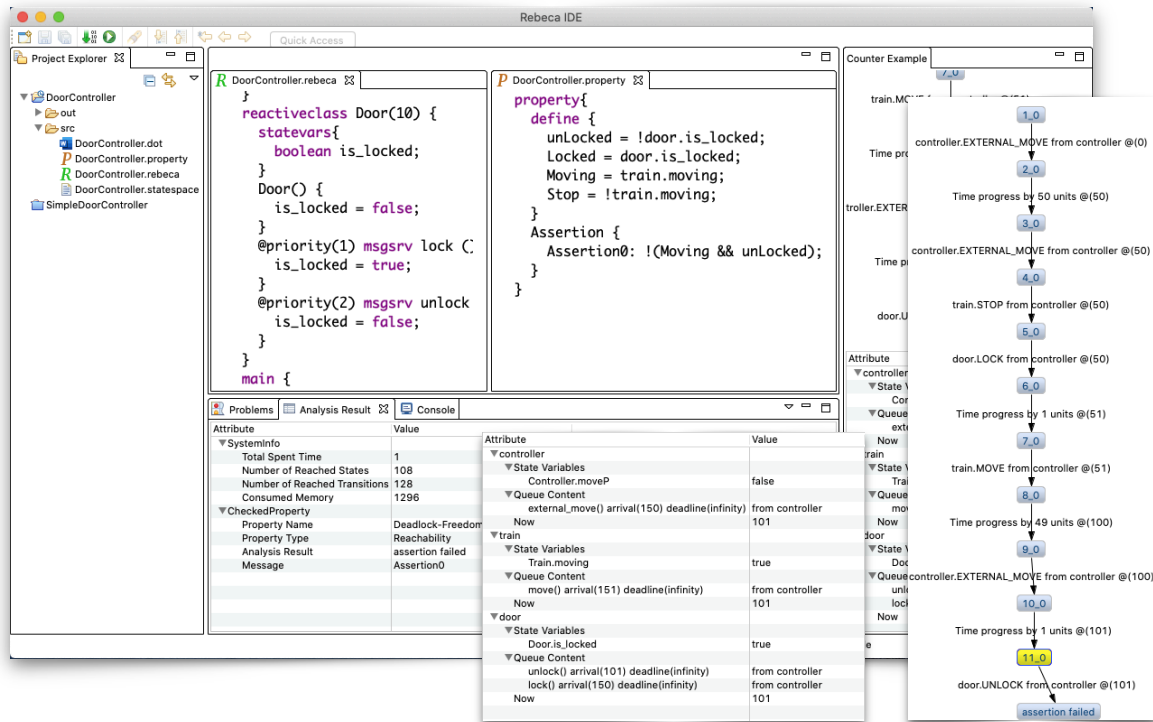
```

1 reactiveclass Controller(5) {
2   knownrebecs{
3     Door door; Train train;
4   }
5   statevars { boolean moveP;}
6   Controller() {
7     moveP = true;
8     self.external_move();
9   }
10  msgsrv external_move() {
11    int d = ?(0, 50); // lock, stop
12    int x = ?(51, 99); // move, unlock
13    int extd = 100; // external_move
14    if (moveP) {
15      door.lock() after(d);
16      train.move() after(x);
17    } else {
18      door.unlock() after(x);
19      train.stop() after(d);
20    }
21    moveP = !moveP;
22    self.external_move() after(extd);
23  } }
24 reactiveclass Train(10) {
25   statevars{
26     boolean moving;
27   }
28   Train() {
29     moving = false;
30   }
31   @priority(1) msgsrv stop() {
32     moving = false;
33   }
34   @priority(2) msgsrv move() {
35     moving = true;
36   } }
37 reactiveclass Door(10) {
38   statevars{
39     boolean is_locked;
40   }
41   Door() {
42     is_locked = false;
43   }
44   @priority(1) msgsrv lock () {
45     is_locked = true;
46   }
47   @priority(2) msgsrv unlock () {
48     is_locked = false;
49   }
50 }
51 main {
52   @priority(1) Controller controller(door,
53     train):();
54   @priority(2) Train train():();
55   @priority(2) Door door():();
56 }

```

(c) The Rebeca model for the code in Figure 5b. Note that here we put different values for **after** constructs comparing to the LF code in Figure 5b.

Figure 5. The Train Door Example with delays in lock/unlock and stop/move



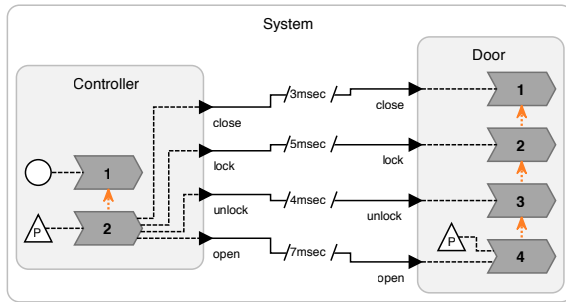
**Figure 6.** An snapshot of Afra finding the counterexample for the model in Figure 5c where the physical external event occurs every 50 units of time (variable `extd` in the model is set to 50). The value of the variables and the contents of the message buffer in state `11_0` is shown in the snapshot.

414 in a lock message to the Door with timetag 0 and a move message to the Train with timetag 100 msec.  
 415 Suppose that we then receive `external = false` at logical time 50 msec. This will result in a stop  
 416 message to the Train with timetag 50 msec, overtaking the move message! But worse, it will send an  
 417 `unlock` message with timetag 150 msec, and the door will unlock while the train is moving! This new  
 418 flaw is revealed by a counterexample generated by Afra shown in Figure 6. In Figure 6, you may see  
 419 the Afra interface with the Rebeca model and the property file including the assertions. The Boolean  
 420 variables defined in the “define” part of the property file are those that are shown in the states in the  
 421 transition system. In the right corner of the figure, the trace of the counterexample is shown and the  
 422 values of variables in state `11_0` (the state right before the assertion is failed) are shown in another  
 423 window.

424 This new flaw is not correctable by simply manipulating logical timetags. The flaw pertains  
 425 to the relationship between physical time and logical time (having no constraints on the spacing  
 426 between timetags of successive `external` events that represent physical actions), and our verification  
 427 strategy here stays entirely in the world of logical time. A similarly cross-cutting flaw could occur if  
 428 the later timetag of the move event does not result in a later occurrence of the train moving physically.  
 429 Again, this flaw pertains to the relationship between physical and logical times, a relationship that  
 430 is ultimately established not only by the software in the systems, but rather by the combination of  
 431 software and hardware.

## 432 6.2. The Train Door Controller and a Passenger

433 Another example is shown in Figure 7. In this example we only show the controller and the  
 434 door. Here the door accepts four commands of `unlock`, `open`, `close` and `lock`. When the train stops at a  
 435 platform the controller `unlock` and then `open` the door. When the train is ready to move the controller  
 436 first `close` and then `lock` the door. The train can only start moving if the doors are locked. But the



(a) Structure of the door controller example

```

1 target C;
2 reactor Controller {
3   output lock:bool; output unlock:bool;
4   output open:bool; output close:bool;
5   physical action external:bool;
6   reaction(startup) {=
7     // ... Set up external sensing.
8   =}
9   reaction(external)->close, lock, open,
10    unlock {=
11     if (external_value) {
12       set(close, true); set(lock, true);
13     } else {
14       set(open, true); set(unlock, true);
15     }
16   =}
17 }
18 reactor Door {
19   input lock:bool; input unlock:bool;
20   input open:bool; input close:bool;
21   physical action ext_open:bool;
22   state locked:bool(false);
23   state is_open:bool(false);
24   reaction(close) {=
25     ... Actuate to close door.
26     self->is_open = false;
27   =}
28   reaction(lock) {=
29     ... Actuate to lock door.
30     if(!self->is_open)
31       self->locked = true;
32   =}
33   reaction(unlock) {=
34     ... Actuate to unlock door.
35     self->locked = false;
36   =}
37   reaction(open, ext_open) {=
38     ... Actuate to open door.
39     if(!self->locked)
40       self->is_open = true;
41   =}
42 }
43 main reactor System {
44   c = new Controller();
45   d = new Door();
46   c.lock -> d.lock after 5 msec;
47   c.unlock -> d.unlock after 4 msec;
48   c.open -> d.open after 7 msec;
49   c.close -> d.close after 3 msec;
50 }

```

(b) Variant of Figure 5b with open door.

```

1 reactiveclass Controller(5) {
2   knownrebecs{
3     Door door;
4   }
5   statevars {
6     boolean moveP;
7   }
8   Controller() {
9     moveP = true;
10    self.external_move();
11  }
12  msgsrv external_move() {
13    int closingDelay = 3;
14    int lockingDelay = 5;
15    int unlockingDelay = 4;
16    int openingDelay = 7;
17    if (moveP) {
18      door.close() after(closingDelay);
19      door.lock() after(lockingDelay);
20    } else {
21      door.unlock() after(unlockingDelay);
22      door.open() after(openingDelay);
23    }
24    moveP = !moveP;
25    self.external_move() after(10);
26  }
27 }
28 reactiveclass Door(10) {
29   statevars{ boolean is_locked, is_open; }
30   Door() {
31     is_locked = false; is_open = true;
32     self.external_open() after(1);
33   }
34   @priority(1) msgsrv close() {
35     is_open = false;
36   }
37   @priority(2) msgsrv lock (){
38     if(!is_open)
39       is_locked = true;
40   }
41   @priority(3) msgsrv unlock() {
42     is_locked = false;
43   }
44   @priority(4) msgsrv open() {
45     if(!is_locked)
46       is_open = true;
47   }
48   @priority(1) msgsrv external_open() {
49     int retryDelay = 2;
50     self.open();
51     self.external_open() after(retryDelay);
52   }
53 }
54 main {
55   @priority(1) Controller controller(door):();
56   @priority(2) Door door():();
57 }

```

(c) The Rebeca model for the code in Figure 7b.

**Figure 7.** The Door Control Example with an external passenger pressing the **open** button. Here we added **open** and **close** commands for the door and do not show the train.

437 door can only be locked if the door is closed. Here we assume that there is an open button for the  
438 passengers also. So, a passenger can press the open button before the door is locked. The controller  
439 sends the close and then the lock command, but there is a scenario in which the external open button is  
440 pressed quickly enough that the door can never be locked, and consequently the train can never move.

441 We can change the values of the after construct for different commands and see various behaviors  
442 of the system for different configurations. With the configuration shown in Figure 7 the door will never  
443 be locked. We checked the assertion of door being unlocked and it is satisfied in this scenario showing  
444 that the door stays unlock all the time. If we increase the value of after in sending `external_open` it  
445 will eventually be large enough to allow the door to close and then lock.

## 446 7. Discussion and Future Work

447 The combination of a language like Lingua Franca with an explicit model of time, and a model  
448 checking tool like Timed Rebeca with Afra can prove quite effective for finding a number of bugs.  
449 Although Rebeca language, which is similar to Java, is expressive enough, it is not clear whether it  
450 would be accepted by designers as a target language, and the toolchains do not currently exist to  
451 compile it down to code that could execute in microcontrollers as would be needed to deploy the train  
452 controller. If these toolchains are created, however, the result could be a very effective package for  
453 designing and deploying formally verifiable CPS software. However, there are some serious limitations  
454 that warrant further research.

455 Based on our (limited number of) experiments and our insights, the mapping between Lingua  
456 Franca and Timed Rebeca can be simple as long as we stay in the logical time domain of Lingua Franca  
457 (and as long as the reaction code in Lingua Franca can be translated to message server code in Timed  
458 Rebeca). By mapping LF to Rebeca and through our examples we demonstrated a set of problems that  
459 can be found using model checking. In the first example we show the model and how logical-time-base  
460 semantics and event-based semantics are different, and how we rely on the underlying platform to  
461 guarantee that the observable behavior is base on the for example logical-time-base semantics. When  
462 we have distributed CPS this may become a nonrealistic assumption. The second example shows how  
463 timing comes in, and how we can rely on different timing configurations to build correct cyberphysical  
464 systems. It also shows how subtle problems may raise that are not easy for a designer to notice,  
465 and how model checking helps in revealing such problems. In the last two examples we show the  
466 connection between the logical and the physical timelines, and the effect of physical events on the  
467 logical behavior of the software and how by using model checking we can move towards finding such  
468 problems.

469 Because Rebeca is designed for model checking, Rebeca models are closed, meaning that there  
470 are no external inputs. The reactions that can be triggered from outside of the Lingua Franca code  
471 (like the **physical actions** named `external` in Figure 2b and `ext_open` in Figure 7b) can be modeled as  
472 message servers that are invoked nondeterministically. This nondeterministic call can be modeled as  
473 a self-call from within the same message server, and there is no need to introduce an extra actor to  
474 model the environment. This message server is first called in the constructor of the rebec, as shown for  
475 `external` on line 8 of Figure 2c, and for `ext_open` on line 32 of Figure 7c.

476 Because the Timed Rebeca code will be used for model checking, we need to be careful regarding  
477 the state space explosion. The external method calls can be problematic here, and the Timed Rebeca  
478 models may have to be carefully crafted in some places. The logical time intervals over which these  
479 methods can be called has a great effect on the state space size. If the state space gets too large, model  
480 checking becomes intractable.

481 Although we performed the mapping from Lingua Franca to Timed Rebeca by hand, it should be  
482 possible to create a Rebeca target for Lingua Franca and then automate the translation. When using  
483 this target, the body of each reaction will need to be written in Rebeca's own language for writing  
484 message servers. This is necessary because Afra analyzes this code to build the transition system



485 model, and as for now Afra is not capable of analyzing arbitrary C, C++, or TypeScript code, the target  
486 languages currently supported by Lingua Franca.

487 One subtle point in model checking of CPS that we presented is that we only checked how the  
488 state of the program evolves in logical time, not how it evolves in physical time. Every model checking  
489 tool that we know of assumes a single timeline, but our systems always have at least three. There is the  
490 logical timeline of timestamps, and programs can be verified on this timeline, proving for example that  
491 a safety condition is satisfied by a state trajectory evolving on this logical timeline. But in a concurrent  
492 and distributed CPS, the state trajectory is also evolving along a physical Newtonian timeline, and our  
493 proof says nothing about its safety on that timeline. Moreover, every clock that measures Newtonian  
494 time will differ from every other clock that measures Newtonian time, so any constraints we impose  
495 on execution based on such clocks may again lead to proofs of safety even though the physical system  
496 is capable of entering unsafe states. Our approach in this paper is relying on a set of assumptions  
497 mainly based on alignment of logical and physical time at the execution time.

498 When we assert that a design has been “verified” against a set of formal requirements, we need to  
499 make every effort to make as clear as possible what are the assumptions about the physical system that  
500 make our conclusions valid. There will *always* be assumptions, and in any real system deployment,  
501 *any* assumption may be violated. There is no such thing as a provably correct system.

502

503

504 **Funding:** The work of the first author is supported in part by KKS SACSys Synergy project (Safe and Secure  
505 Adaptive Collaborative Systems), and KKS DPAC Project (Dependable Platforms for Autonomous Systems and  
506 Control) at Mälardalen University, and MACMa Project (Modeling and Analyzing Event-based Autonomous  
507 Systems) at Software Center, Sweden. The second author receives support from the National Science Foundation  
508 (NSF), award #CNS-1836601 (Reconciling Safety with the Internet) and the iCyPhy Research Center (Industrial  
509 Cyber-Physical Systems, supported by Camozzi Industries, Denso, Siemens, and Toyota).

## 510 References

- 511 1. Lee, E.A. Cyber Physical Systems: Design Challenges. Int. Symp. on Object/Component/Service-Oriented  
512 Real-Time Distributed Computing (ISORC). IEEE, 2008, pp. 363 – 369. doi:10.1109/ISORC.2008.25.
- 513 2. Sirjani, M.; Provenzano, L.; Asadollah, S.A.; Moghadam, M.H. From Requirements to Verifiable Executable  
514 Models using Rebeca. International Workshop on Automated and verifiable Software sYstem DEvelopment,  
515 2019.
- 516 3. Lee, E.A.; Seshia, S.A. *Introduction to Embedded Systems - A Cyber-Physical Systems Approach*; MIT Press,  
517 2017.
- 518 4. Alur, R. *Principles of Cyber-Physical Systems*; MIT Press, 2019.
- 519 5. Henzinger, T.A. The Theory of Hybrid Automata. Proceedings, 11th Annual IEEE Symposium on Logic in  
520 Computer Science, 1996. IEEE Computer Society, 1996, pp. 278–292.
- 521 6. Carloni, L.P.; Passerone, R.; Pinto, A.; Sangiovanni-Vincentelli, A. Languages and Tools for Hybrid Systems  
522 Design. *Foundations and Trends in Electronic Design Automation* **2006**, *1*, 1–204. doi:10.1561/1000000001.
- 523 7. Sirjani, M.; Movaghar, A.; Shali, A.; de Boer, F.S. Modeling and Verification of Reactive Systems using  
524 Rebeca. *Fundam. Inform.* **2004**, *63*, 385–410.
- 525 8. Sirjani, M.; Jaghoori, M.M. Ten Years of Analyzing Actors: Rebeca Experience. Formal Modeling: Actors,  
526 Open Systems, Biological Systems, 2011, pp. 20–56.
- 527 9. Rebeca. Afra Tool, 2019. Available at <http://rebeca-lang.org/alltools/Afra>, Retrieved July, 2019.
- 528 10. Lohstroh, M.; Schoeberl, M.; Goens, A.; Wasicek, A.; Gill, C.; Sirjani, M.; Lee, E.A. Invited: Actors Revisited  
529 for Time-Critical Systems. Design Automation Conference (DAC), 2019.
- 530 11. Lohstroh, M.; Romeo, I.n.I.; Goens, A.; Derler, P.; Castrillon, J.; Lee, E.A.; Sangiovanni-Vincentelli, A.  
531 Reactors: A Deterministic Model for Composable Reactive Systems. Model-Based Design of Cyber  
532 Physical Systems (CyPhy’19), 2019. Held in conjunction with ESWEEK 2019.
- 533 12. Lohstroh, M.; Lee, E.A. Deterministic Actors. Forum on Specification and Design Languages (FDL),, 2019.
- 534 13. Hewitt, C. Viewing control structures as patterns of passing messages. *Journal of Artificial Intelligence* **1977**,  
535 *8*, 323–363.

- 536 14. Agha, G.A. *ACTORS - a model of concurrent computation in distributed systems*; MIT Press series in artificial  
537 intelligence, MIT Press: Cambridge, MA, 1990.
- 538 15. Reynisson, A.H.; Sirjani, M.; Aceto, L.; Cimini, M.; Jafari, A.; Ingólfssdóttir, A.; Sigurdarson, S.H. Modelling  
539 and simulation of asynchronous real-time systems using Timed Rebeca. *Sci. Comput. Program.* **2014**,  
540 *89*, 41–68.
- 541 16. Sirjani, M. Rebeca: Theory, Applications, and Tools. Formal Methods for Components and Objects,  
542 International Symposium, FMCO 2006, 2006, pp. 102–126.
- 543 17. Derler, P.; Lee, E.A.; Matic, S. Simulation and Implementation of the PTIDES Programming Model.  
544 International Symposium on Distributed Simulation and Real-Time Applications, 2008, pp. 330–333.
- 545 18. Corbett, J.C.; et.al. Spanner: Google Globally Distributed Database. *ACM Trans. Comput. Syst.* **2013**,  
546 *31*, 8:1–8:22.
- 547 19. Sirjani, M.; Khamespanah, E.; Lee, E.A. Model Checking Software in Cyberphysical Systems. COMPSAC  
548 2020, 2020.
- 549 20. Benveniste, A.; Berry, G. The Synchronous Approach to Reactive and Real-Time Systems. *Proceedings of the*  
550 *IEEE* **1991**, *79*, 1270–1282.
- 551 21. Edwards, S.A.; Lee, E.A. The Semantics and Execution of a Synchronous Block-Diagram Language. *Science*  
552 *of Computer Programming* **2003**, *48*, 21–42. doi:10.1016/S0167-6423(02)00096-5.
- 553 22. Zhao, Y.; Lee, E.A.; Liu, J. A Programming Model for Time-Synchronized Distributed Real-Time Systems.  
554 Real-Time and Embedded Technology and Applications Symposium (RTAS). IEEE, 2007, pp. 259 – 268.  
555 doi:10.1109/RTAS.2007.5.
- 556 23. Corbett, J.C.; Dean, J.; Epstein, M.; Fikes, A.; Frost, C.; Furman, J.; Ghemawat, S.; Gubarev, A.; Heiser,  
557 C.; Hochschild, P.; Hsieh, W.; Kanthak, S.; Kogan, E.; Li, H.; Lloyd, A.; Melnik, S.; Mwaura, D.; Nagle,  
558 D.; Quinlan, S.; Rao, R.; Rolig, L.; Saito, Y.; Szymaniak, M.; Taylor, C.; Wang, R.; Woodford, D. Spanner:  
559 Google’s Globally-Distributed Database. OSDI, 2012. doi:10.1145/2491245.
- 560 24. Lingua Franca Grammar. Lingua Franca Github . Available at [https://github.com/icyphy/lingua-franca/  
561 blob/master/xtext/org.icyphy.linguafranca/src/org/icyphy/LinguaFranca.xtext](https://github.com/icyphy/lingua-franca/blob/master/xtext/org.icyphy.linguafranca/src/org/icyphy/LinguaFranca.xtext), Retrieved May, 2020.
- 562 25. Schneider, C.; Spönemann, M.; von Hanxleden, R. Just Model! – Putting Automatic Synthesis of  
563 Node-Link-Diagrams into Practice. Proceedings of the IEEE Symposium on Visual Languages and  
564 Human-Centric Computing (VL/HCC ’13), , 2013; pp. 75–82. doi:10.1109/VLHCC.2013.6645246.
- 565 26. Hardebolle, C.; Boulanger, F. ModHel’X: A Component-Oriented Approach to Multi- Formalism Modeling.  
566 MODELS 2007 Workshop on Multi- Paradigm Modeling. Elsevier Science B.V., 2007.
- 567 27. Ptolemaeus, C. *System Design, Modeling, and Simulation using Ptolemy II*; Ptolemy.org: Berkeley, CA, 2014.
- 568 28. Jantsch, A. *Modeling Embedded Systems and SoCs - Concurrency and Time in Models of Computation*; Morgan  
569 Kaufmann, 2003.
- 570 29. Sirjani, M.; Khamespanah, E. On Time Actors. Theory and Practice of Formal Methods - Essays Dedicated  
571 to Frank de Boer on the Occasion of His 60th Birthday, 2016, pp. 373–392.
- 572 30. Khamespanah, E.; Sirjani, M.; Sabahi-Kaviani, Z.; Khosravi, R.; Izadi, M. Timed Rebeca schedulability and  
573 deadlock freedom analysis using bounded floating time transition system. *Sci. Comput. Program.* **2015**,  
574 *98*, 184–204.
- 575 31. International Electrotechnical Commission. *International Standard IEC 61131: Programmable Controllers*, 4.0  
576 ed.; IEC, 2017.
- 577 32. Berger, H. *Automating with SIMATIC S7-1500: Configuring, Programming and Testing with STEP 7 Professional*,  
578 1st ed.; Publicis MCD Werbeagentur GmbH, 2014.