

Model-Checking of Component-Based Real-Time Embedded Software Based on CORBA Event Service *

Zonghua Gu

Department of Computer Science
University of Virginia
Charlottesville, VA 22903

Kang G. Shin

Department of EECS
University of Michigan
Ann Arbor, MI 48109

Abstract

As the complexity of real-time embedded software grows, it is desirable to apply formal verification techniques to achieve a high level of assurance. We discuss application of model-checking to verification of component-based real-time embedded software based on CORBA Event Service, with the Avionics Mission Computing software as an application example. We first use the process algebra FSP to formalize specification of software components and system architecture, previously only available in the form of natural language and prone to misinterpretation and misunderstanding. We then use model-checking to verify system-level concurrency properties. Finally, we discuss effective techniques for improving scalability. We have applied our modeling and analysis techniques to several application scenarios with satisfactory results.

1. Introduction

The publish/subscribe model of computation, as implemented in CORBA Event Service [11], has been widely adopted in a variety of application domains, including both real-time embedded systems and enterprise distributed systems. One example is the *Avionics Mission Computing* (AMC) [1] software, which is the embedded software onboard a military aircraft for controlling mission-critical functions, such as navigation, target tracking and identification, weapon firing, etc. The software architecture of AMC is also commonly referred to as the *Bold Stroke Framework*. It is modeled with UML, but manually coded with C++. The UML models mainly serve in a documentation role that the engineer refers to while writing code manually. Therefore, the link between model and code is weak and easily broken in the process of system maintenance and evolution, when code is modified or enhanced without making the corresponding changes of the model, or vice versa. Furthermore, UML has little support for analysis that is relevant for embedded systems, such as real-time and concurrency properties, such as schedulability and dead-

lock freedom. As part of the DARPA *Model-Based Integration of Embedded Software*(MoBIES) Program, an end-to-end tool-chain [2] has been developed collaboratively by researchers from Vanderbilt University, Southwest Research Institute, and University of Michigan. The MoBIES tool-chain covers the entire systems development life-cycle, including modeling, code generation and analysis, and provides a more automated and integrated development process than the current industry practice. The central repository of information in the tool-chain is the *Embedded Systems Modeling Language* (ESML) [3], a domain-specific language for modeling the AMC software using the Generic Modeling Environment (GME) [4] from Vanderbilt University. The ESML meta-model [3] defines a comprehensive modeling language that captures essential aspects of the embedded system, including software architecture, timing and resource constraints, execution threads, execution platform (processors and network) information, allocation of components to threads/processors, etc. We have developed a tool AIRES [5] to perform various static analysis tasks on ESML models, such as dependency, timing, schedulability and automated component allocation.

A limitation of ESML is that it mainly focuses on the *static structural* aspects while largely ignoring the *dynamic behavioral* aspects of AMC software. As a result, ESML models are not *executable*. In order to perform deeper semantic analysis, it would be necessary to construct *executable models*, which enables the use of *simulation* or *model-checking* to verify system correctness. One prominent example of executable models is Harel's Statechart [6]. Model-checking can be viewed as exhaustive simulation, which works by exhaustively exploring the system state space to prove certain correctness properties. In this paper, we use the process algebra *Finite State Processes*(FSP) [7], developed by Jeff and Magee Kramer, with formal semantics defined with *Labelled Transition Systems*, to provide formal specification of dynamic behavioral aspects of the AMC software, and use the tool *Labeled Transition System Analyzer*(LTSA) [7] to analyze the resulting FSP specification by simulation or model-checking. If model-checking fails, LTSA gives us an error trace leading to the error state. It is possible to reconstruct the scenario in more user-friendly notations such as the UML Sequence Diagram.

Although our work is initially targeted towards the AMC software, it is applicable to more general component-based real-time embedded software with event-triggered interaction style. The

* The work reported in this paper was supported in part by DARPA under contract F3615-00-1706, by ARO under grant DAAD19-1-1-0473, and by ESCHER.

AMC software architecture is very similar to the *CORBA Component Model (CCM)* [8], which was originally designed for enterprise applications, but has been recently extended to be real-time and QoS-enabled by researchers from Washington University to produce *Component-Integrated ACE ORB (CIAO)* [9]. In fact, there are plans to migrate the next-generation of AMC software to the CIAO platform. Vanderbilt University has developed a model-based tool-chain *Component Synthesis with Model Integrated Computing (CoSMIC)* [10] for applying Model-Integrated Computing to design and configuration of applications based on the CIAO platform. Our approach could be easily adapted to apply to general CCM applications, for example, by generating FSP models from CoSMIC instead of ESML.

This paper is structured as follows. Section 2 provides a brief introduction to AMC. Section 3 provides an introduction to the modeling language FSP. Section 4 describes modeling of AMC with FSP. Section 5 discusses the specification of correctness properties for verification. Section 6 presents techniques for improving model-checking scalability. Section 8 discusses related work, and Section 9 draws conclusions.

2. Introduction to the Avionics Mission Computing Software

The AMC software consists of components interacting with each other using the publish/subscribe paradigm with Real-Time CORBA Event Service [11] as its underlying communications substrate. Event publisher components push events through the event channel to event consumer components, whose execution is triggered by the arrival of events. The system runs at a number of different rates driven by timer event publishers, such as 40Hz, 20Hz, 10Hz, 5Hz, and 1Hz. This corresponds to the *scenario-based multi-threading* [12] approach, where each end-to-end scenario triggered by a timer is assigned its own thread. Thread priorities are assigned rate-monotonically, that is, higher frequency threads are assigned higher priorities. Rate Monotonic Analysis [13] is used to make real-time guarantees.

Components are composite objects with ports, interacting with one another either through event triggers or method invocations. Some terminologies from *CORBA Component Model (CCM)* are adopted. Each component can have the following types of ports:

- *Publish Port* to publish events.
- *Subscribe Port* to subscribe to events.
- *Receptacle* to issue method invocations.
- *Facet* to accept method invocations.

Component interaction typically (but not always) follows the *control-push/data-pull* style, as shown in Figure 1. The data producer component publishes a `DataAvailable` event from its publish port indicating that it has fresh data; when the data consumer component receives the event from its subscribe port, it issues a `GetData()` call from its receptacle to the producer's facet to retrieve the data. Even though control-push/data-pull is the predominant interaction style by convention, there are some less common cases where a *data-push* style is adopted. That is, instead of a `DataAvailable` event followed by a `GetData()` call, the upstream component directly invokes `SetData()` on the




Figure 1. The control-push/data-pull style of interaction.

downstream component. The two styles are functionally equivalent, but control-push/data-pull has been adopted as the convention in order to maintain some level of uniformity in system design.

3. Introduction to Finite State Processes

We only provide a very brief description of FSP, and refer the interested reader to [7] for more details.




Figure 2. A simple FSP model for a light switch.

Primitive components are defined as finite-state processes using event prefix \rightarrow , choice $|$ and recursion. If x is an event and P a process, then $(x \rightarrow P)$ describes a process that initially synchronizes with the event x and then behaves exactly as process P . Figure 2 shows a simple FSP model for a light switch, which toggles between the off and on states:

`SWITCH = (on \rightarrow off \rightarrow SWITCH).`

We can write an equivalent specification using recursion:

`SWITCH = OFF,
OFF = (on \rightarrow ON), ON = (off \rightarrow OFF).`

If x and y are events, then $(x \rightarrow P | y \rightarrow Q)$ describes a process which initially synchronizes with either x or y , and the subsequent behavior is described by P or Q , respectively. For example, a model for a drinks machine is:

`DRINKS = (red \rightarrow coffee \rightarrow DRINKS
| blue \rightarrow tea \rightarrow DRINKS).`

If the red button is pressed, it dispenses a cup of coffee; if the blue button is pressed, it dispenses a cup of tea.

The alphabet of a process is the set of events it can synchronize with. The interface operator $@$ is used to specify the set of event labels that are visible at the interface of the component and thus may be shared with other components. The alphabetic extension operator $+$ is used to specify extension of process alphabet to include a set of event labels.

Primitive processes can be composed with the parallel composition operator `||` to form a *composite process*. Processes interact via synchronization on common event labels in the traditional style of process algebra. That is, if processes in a composition have a common shared event, all processes must synchronize on the shared event at the same step. For example:

```
MAKER = (make->ready->MAKER) .
USER = (ready->use->USER) .
||MAKER_USER = (MAKER||USER) .
```

The MAKER process and USER process share a common event `ready`, so they must synchronize on that event while the other events can be interleaved. We can use *relabeling* to model synchronization between events with different names. For example, an equivalent specification is:

```
MAKER = (make->done->MAKER) .
USER = (ready->use->USER) .
||MAKER_USER =
(aMaker:MAKER|aUser:USER)
/{aMaker.done/aUser.ready} .
```

4. Modeling AMC with FSP

4.1. Modeling of Component Types

The AMC software is component-based with many different types of components, each with its unique functionality and interfaces, acting as basic building blocks of a complete system. The documentation provided by Boeing [14] provides detailed descriptions for the various component types in natural language. However, these descriptions are not formal and prone to misinterpretation or misunderstanding. We use FSP to provide an unambiguous, formal description for each component *type* based on the natural language descriptions, and instantiate each component *instance* to form a system architecture. In what follows, we describe each component type by excerpting its description from the Boeing documentation, and then presenting its corresponding FSP specification. Note that this is not an exhaustive list of all component types, but only includes those needed for understanding the rest of this paper, plus a few other interesting ones from a modeling viewpoint.

The following naming conventions are used.

- `inEvt` denotes action of the subscriber component to receive an input event.
- `outEvt` denotes action of the publisher component to issue an output event.
- `issueGDCall` denotes action of the caller component to issue a `GetData()` call.
- `receiveGDCall` denotes action of the callee component to receive a `GetData()` call.
- `issueGDReply` denotes action of the callee component to issue a `GetData()` reply.
- `receiveGDReply` denotes action of the caller component to receive a `GetData()` reply.

Similar naming conventions hold for “`SetData()`” calls such as `issueSDCall`, `receiveSDCall`, `issueSDReply` and `receiveSDReply`.

“*DeviceComponent* is used to simulate a device that generates its own data (as in sensor reports). Upon receiving a `Push()`, this component does a `Push()` if it is specified as an event supplier. In a typical scenario, this component is specified as an event consumer of interval timeouts, so as to simulate the device generating information on a periodic basis.”

```
DeviceComp = (inEvt->outEvt->DeviceComp
|receiveGDCall->issueGDReply->DeviceComp) .
```

“*DisplayComponent* is used to display information to the console window. It is used to simulate any output device in a system. Upon receiving a `Push()`, this component does a `Get()` on each component specified in its receptacles. This component then displays the results on the console.”

```
DisplayComp =
(inEvt->issueGDCall->receiveGDReply->display
->DisplayComp) .
```

“*ClosedEDComponent* is closed in the sense that other components cannot alter its data via `Set()` operations. *ED* stands for *event driven*. Upon receiving a `Push()`, this component does a `Get()` on each component specified in its receptacles. This component then generates a `Push()` if it is specified as an event supplier.”

```
ClosedEDComp =
(inEvt->issueGDCall->receiveGDReply->outEvt
->ClosedEDComp
|receiveGDCall->issueGDReply->ClosedEDComp) .
```

“*OpenEDComponent* is open in the sense that other components can set its data. Upon receiving a `Set()`, this component does a `Get()` on each component specified in its receptacles. This component then generates a `Push()` if it is specified as an event supplier.”

```
OpenEDComp =
(inEvt->issueGDCall->receiveGDReply->outEvt
->OpenEDComp
|receiveGDCall->issueGDReply->OpenEDComp
|receiveSDCall->issueGDCall->receiveGDReply
->issueSDReply->outEvt->OpenEDComp) .
```

“*LazyActiveComponent* is used to simulate delayed response to acquiring data. As an optimization strategy, if a component is updated more than it is read, the Lazy Active pattern may be used to only update the data is a request is made. Upon receiving a `Push()`, this component flags its data as invalid. When this component’s `Get()` is called this triggers a the *LazyActiveComponent* to call `Get()` on the components attached to the receptacles of the *LazyActiveComponent*.”

```
LazyActiveComp = (inEvt->outEvt->DataStale
|receiveGDCall->issueGDReply->LazyActiveComp) ,
```

```
DataStale=
(receiveGDCall->issueGDCall->receiveGDReply
->issueGDReply->LazyActiveComp) .
```

“*ModalComponent* is used to alter the flow of events. The component can be enabled and disabled via the facet method

ChangeMode(). When it is enabled, it will update and generate an event when it receives an event. When it is disabled, it will not update or generate an event.”

```

ModalComp = Enabled,

Disabled = (enable->Enabled|disable->Disabled
|inEvt->Disabled),

Enabled = (enable->Enabled|disable->Disabled
|inEvt->issueGDCall->receiveGDReply->outEvt
->Enabled
|receiveGDCall->issueGDReply->Enabled).

```

4.2. Modeling of Component Interactions

4.2.1. Control-Push/Data-Pull Below is the FSP model for the control-push/data-pull interaction style discussed in Figure 1, Section 2.

```

Publisher = (outEvt->Publisher |
receiveCall->issueReply->Publisher).

Subscriber = (inEvt->issueCall->receiveReply
->Subscriber).

||ControlPushDataPull =
(pub:Publisher||sub:Subscriber)
/{pub.outEvt/sub.inEvt,
sub.issueCall/pub.receiveCall,
sub.receiveReply/pub.issueReply}.

```

This modeling approach treats the interaction between an event publisher and an event subscriber as *synchronous*, that is, the outEvt of the publisher synchronizes with the inEvt of the subscriber directly. This is not entirely accurate, since in the physical system the published events go through the CORBA event service and are buffered at the input port of the subscribe component. We can obtain a more accurate model by using separate processes to model the queues/buffers in the middleware infrastructure, and decouple the interactions to make them asynchronous, but that will have a significant impact on scalability in terms of the maximum size of the system that we can check. Our focus is on verification of *application-level* concurrency properties when the application is operating under normal conditions, assuming that the middleware behaves correctly, with no buffer overflows or deadline misses (see Section 4.2.3 for more details on this assumption). This synchronous modeling style has turned out to be at an adequate level of abstraction for the types of concurrency properties we are interested in, including deadlock freedom, event reachability, sequencing constraints, and progress property. That is, adopting a more detailed modeling approach would not change the verdict of the model-checker on these properties. This may not be generally true if we expand the range of properties to include other properties involving the middleware, for example, buffer overflow detection.

The words *synchronous* and *asynchronous* are overloaded terms with different meanings to different people. Here we use the word *synchronous* to mean that pairwise interactions, such as event delivery and method invocation, between

components happen instantaneously without the delays introduced by the middleware infrastructure. This is very different from its meaning in *synchronous formalisms* such as Esterel, which describes a time-triggered system with a global clock tick, commonly found in hardware and safety critical software systems.

4.2.2. Input Event Correlation When a component subscribes to multiple events, there may be two synchronization patterns: AND synchronization means that the component must receive all input events to be triggered; OR synchronization means that the component only needs to receive one of the input events to be triggered. In order to model AND synchronization, we add a new process type called InputANDCorrelator, as shown below and in Figure 3:

```

Event(ID=1) = (inEvt[ID]->GotOne),
GotOne = (inEvt[ID]->GotOne |
matched->Event).

||InputANDCorrelator(NumInputs=2) =
(if(NumInputs>0) then
(forall [i:1..NumInputs] Event(i))).

```




Figure 3. The correlator for two input events with AND synchronization.




Figure 4. The correlator for two input events with OR synchronization.

It models parallel composition of NumInputs number of processes Event, which all synchronize on the same event matched. This ensures that the matched event is emitted only when all input events inEvt[i..NumInputs] occur. The event matched event is in turn used to trigger the downstream subscriber component.

Input events may arrive at different rates. For example, a component may subscribe to inEvt[1] arriving at 20Hz rate, and inEvt[2] arriving at 1Hz rate. Only 1 out of every 20 inEvt[1] is paired up with 1 inEvt[2] to generate a matched event; the other 19 are silently discarded, as modeled by the self loop in state 1.

OR synchronization is simpler, shown below and in Figure 4.

```

Event(ID=1) = (inEvt[ID]->matched[ID]->Event).

||InputORCorrelator(NumInputs=2) =
(if(NumInputs>0) then
(forall[i:1..NumInputs] Event(i))).

```

4.2.3. Real-Time Issues The typical way to model real-time in FSP is to discretize time into uniform segments by using a global event `tick` shared among all the processes in the system to provide a system-wide heartbeat. The typical component execution time in an AMC system is fairly small, in the range of microseconds, while the typical period of execution is fairly large, in the range of milliseconds or even seconds. If we model quantitative time accurately by using a fine-grained partitioning of time on the microsecond scale, the system state space will quickly explode even for trivial examples. Instead, we only ensure that the relative execution frequencies of different rate groups are correct, e.g., the 20Hz thread should execute 20 times more frequently than the 1Hz thread. This can be achieved by using a shared event `tick`. If the event `timeout20hz` is emitted at every tick, then the event `timeout1hz` is emitted every 20 ticks.

```

Timer20hz = (timeout20hz->tick->Timer20hz).
Timer1hz = (timeout1hz->Delay20[1]),
Delay20[t:1..20] = (when(t==20)tick->Timer1hz
|when (t < 20) tick->Delay20[t+1]).

```

Note that the global event `tick` is only shared among all the timers, not the application components. Therefore, even though the periodic timer triggers synchronize to a system-wide heartbeat, application components interact with each other in an asynchronous fashion.

Even though we cannot model quantitative time, we make the assumption that the system is *schedulable*, i.e., all threads meet their deadlines. Without this assumption, we would have a much larger state space due to deadline misses, an error condition that should never arise in a production system, without gaining any additional insight into the system's normal operation. We can achieve separation of concerns by using real-time scheduling theory to verify the schedulability assumption, using for example AIREs [5], and model-checking to verify concurrency properties. We encode this assumption in the model by adding an explicit synchronization between the timer and the *terminal events*, the leaf events of the event dependency graph rooted at the timeout event, in order to ensure that the next `timeout` event will not occur until all events belonging to the current execution frame have been processed. For example, timer-triggered sensor data may go through some processing stages and eventually trigger both the Flight Plan Display and the Navigation Display. We insert an AND correlator to make sure that both Display components have been triggered before the next 20Hz timeout.

```

Timer20hz = (timeout20hz->timer20hzDone->tick
->Timer20hz).

Thread20hz = (...
||fltPlanDisplay:displayComp
||navDisplay:displayComp
||correlator:InputANDCorrelator)
/{fltPlanDisplay.display/correlator.inEvt[1],
navDisplay.display/correlator.inEvt[2],

```

```

correlator.matched/timer20hzDone
}.

```

There may be *aperiodic* and *sporadic* external interrupts that act as thread triggers in addition to periodic timers. For sporadic interrupts, there is a bound of *minimum inter-arrival time* (MIAT) between interrupts, so we make a pessimistic assumption and model the interrupt source as a periodic timer with period equal to MIAT, as is commonly done in real-time scheduling analysis. For aperiodic interrupts, we make the interrupt source *not* synchronize with the global event `tick`, meaning that the interrupts can happen at arbitrary points in time without any timing constraints, which is exactly the definition of aperiodic interrupts.

4.3. An Example Application Scenario

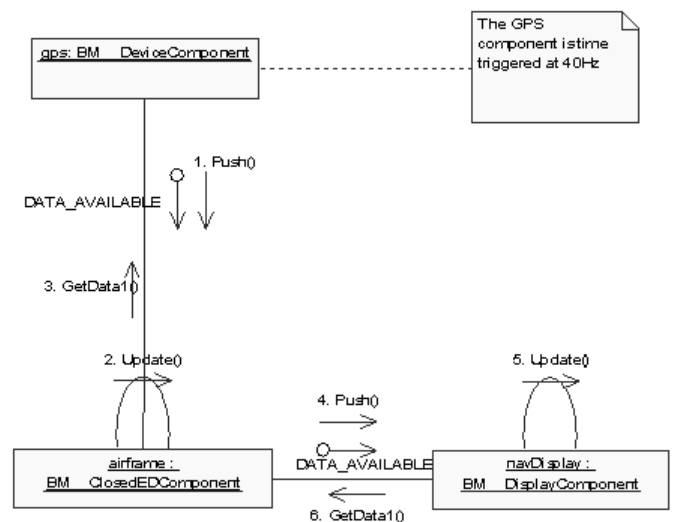


Figure 5. The basic single-processor scenario.

As an illustrative example, we consider the *Basic Single-Processor* application scenario in Figure 5. At a 40Hz rate, the system must update navigation displays with timely airframe position information using inputs from navigation sensors. Triggered by the 40Hz timer, the `gps` component pushes a `DataAvailable` event to the `airframe` component, which updates its state by getting data from `gps`. The `airframe` component then pushes a `DataAvailable` event to the `navDisplay` component, which then updates the display by getting data from `airframe`. The prefix `BM_` for the component types is a naming convention, meaning that these component types are basic *Building Block* models for any application, as opposed more application specific components such as `OM_`, for *Operator Interface* models. We omit these prefixes in FSP specifications. Below is the complete FSP specification for this scenario:

```

Timer40hz = (timeout40hz->timer40hzDone->tick
->Timer40hz).

DeviceComp = (inEvt->outEvt->DeviceComp |

```

```

receiveGDCall->issueGDReply->DeviceComp).

ClosedEDComp =
(inEvt->issueGDCall->receiveGDReply->outEvt
->ClosedEDComp
|receiveGDCall->issueGDReply->ClosedEDComp).

DisplayComp =
(inEvt->issueGDCall->receiveGDReply->display
->DisplayComp).

||Thread40hz = (Timer40hz
||gps:DeviceComp
||airframe:ClosedEDComp
||navDisplay:DisplayComp)
/{timeout40hz/gps.inEvt,
gps.outEvt/airframe.inEvt,
airframe.issueGDCall/gps.receiveGDCall,
airframe.receiveGDReply/gps.issueGDReply,
airframe.outEvt/navDisplay.inEvt,
navDisplay.issueGDCall/airframe.receiveGDCall,
navDisplay.receiveGDReply/airframe.issueGDReply,
navDisplay.display/timer40hzDone }.

```

LTSA has a built-in functionality to perform simulation as user-controller animation. Once the models are developed, we can use interactive simulation to gain deeper understanding of the system dynamics, or model-checking to verify concurrency properties. We will not elaborate on simulation due to space limitations. Instead, we will focus on model-checking in the following sections.

In order to integrate model-checking into the MoBIES toolchain [2], we have developed a model interpreter with C++ using a set of APIs provided by GME to translate an ESML model into its equivalent FSP model. The resulting FSP model consists of two parts. The first part is the fixed *preamble* with FSP specifications for all possible component types, as discussed in Section 4.1. This information is not contained in the ESML model, and exists as a form of library that is inserted into every FSP model. The second part is generated from the ESML model, containing composition of components instantiated from component types contained in the preamble.

5. Specification of System Properties

There are two types of properties that can be checked: *safety* and *liveness*. A safety property asserts that nothing bad happens, and a liveness property asserts that something good eventually happens. We consider the following safety properties: *deadlock freedom*, *event reachability* and *sequencing constraints*. We consider one liveness property related to *progress*. Besides these generic properties, it is possible to specify and verify other application-specific properties, for example, a certain component method is only invoked after a number of other component methods are invoked for a specific number of times and in a specific order.

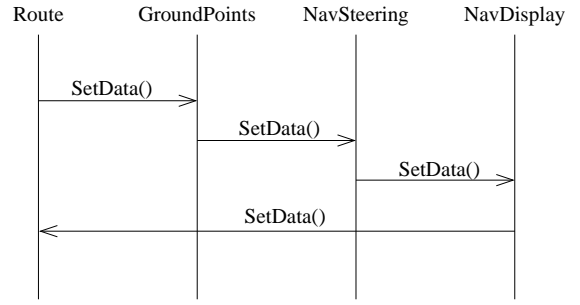


Figure 6. A deadlock situation caused by a dependency cycle.

5.1. Deadlock Freedom

The most important safety property is deadlock freedom. Let's take an scenario of four components forming a chain of *data-push* interactions, as shown in Figure 6. For illustration purposes, we artificially introduce a deadlock situation by adding an extra method call from `navDisplay` to `route`, as shown in Figure 6. When the `route` component's `SetData()` call is invoked, it is still blocked waiting for its method invocation to `groundPoints` to return. This is the classic deadlock situation caused by a circular dependency.




Figure 7. A fragment of the 20Hz thread in the MediumSP application Scenario.

Another possible cause of deadlock is when a component subscribes to multiple input events with AND synchronization, but for certain reasons not all of the input events are available. Figure 7 shows an example application scenario. Component `trackSensor` is triggered periodically by the 20Hz timeout, and issues an output event that is subscribed to by both `track1` and `track2`. Component `tacSteering` subscribes to output events of both `track1` and `track2` with AND synchronization, and issues an event to trigger `tacDisplay`. LTSA reveals no deadlocks in this scenario. Here is the FSP specification:

```

Timer1hz = (timeout1hz->timer1hzDone->Timer1hz).

DeviceComp = (inEvt->outEvt->DeviceComp
|receiveGDCall->DeviceComp).

ClosedEDComp = (inEvt->issueGDCall->outEvt
->ClosedEDComp|receiveGDCall->ClosedEDComp).

```

```
DisplayComp = (inEvt->issueGDCall->display
->DisplayComp).
```

```
Event(ID=1) = (inEvt[ID]->matched->Event).
||InputCorrelator(NumInputs=1)=
(if(NumInputs>0) then (forall[i:1..NumInputs]
Event(i))).
```

```
||System = (Timer1hz||trackSensor1:DeviceComp
||track1:ClosedEDComp||track2:ClosedEDComp
||correlatorTS:InputANDCorrelator(2)
||tacSteering:ClosedEDComp
||tacDisplay:DisplayComp)
/{timeout1hz/trackSensor1.inEvt,
trackSensor1.outEvt/track1.inEvt,
trackSensor1.outEvt/track2.inEvt,
track1.issueGDCall/trackSensor1.receiveGDCall,
track2.issueGDCall/trackSensor1.receiveGDCall,
track1.outEvt/correlatorTS.inEvt[1],
track2.outEvt/correlatorTS.inEvt[2],
correlatorTS.matched/tacSteering.inEvt,
tacSteering.issueGDCall/track1.receiveGDCall,
tacSteering.issueGDCall/track2.receiveGDCall,
tacSteering.outEvt/tacDisplay.inEvt,
tacDisplay.issueGDCall/tacSteering.receiveGDCall,
tacDisplay.display/timer1hzDone}.
```

For illustration purposes, let's make a change to the system by letting trackSensor1 publish two types of events outEvt1 and outEvt2, choosing non-deterministically which event to output at runtime. This could be used to model a modal component which outputs different events depending on its active mode:

```
DeviceComp = (inEvt->outEvt1->DeviceComp|
inEvt->outEvt2->DeviceComp
|receiveGDCall->DeviceComp).
...
...
||System = (Timer20hz||trackSensor1:DeviceComp
||track1:ClosedEDComp||track2:ClosedEDComp
||correlatorTS:InputANDCorrelator(2)
||tacSteering:ClosedEDComp
||tacDisplay:DisplayComp)
/{timeout20hz/trackSensor1.inEvt,
trackSensor1.outEvt1/track1.inEvt,
trackSensor1.outEvt2/track2.inEvt,
...}.
```

This change results in a deadlock situation since tacSteering can only receive one of its two input events. LTSA outputs this trace to deadlock:

```
Trace to DEADLOCK:
  timeout20hz
  trackSensor1.outEvt1
  track1.issueGDCall
  track1.outEvt
```

5.2. Event Reachability

Each component should be triggered/invoked at least once during each execution cycle. Otherwise, the component is redundant, which could signal a design error or inefficiency that wastes system resources. In order to prove that a component's method

C.m is reachable, we introduce a property NotReachable stating that the event C.m never occurs. If this property holds, then C.m is indeed not reachable; otherwise, LTSA returns a counter example showing the path of execution leading to the event C.m. For example, if we would like to check that the action navDisplay.display is executed/reachable, we add the following:

```
property NotReachable = STOP+{reachable}.

||CheckReachability = (System || NotReachable)
/{navDisplay.display/reachable}.
```

Checking this property for the MultirateSP scenario yields this chain of events that lead to the triggering of navDisplay.display:

```
Trace to property violation in NotReachable:
  timeout40hz
  gps.outEvt
  airFrame.issueGDCall
  airFrame.receiveGDReply
  airFrame.outEvt
  navDisplay.issueGDCall
  navDisplay.receiveGDReply
  navDisplay.display
```

This means that navDisplay.display is indeed reachable, which is the correct behavior.

5.3. Sequencing Constraints

Certain events should happen in sequence. For example, the events in a linear chain of event triggers should happen in the order of precedence relation from the head to the tail of the chain. Below is the property specification used to check the correct ordering of events in the 40Hz thread:

```
property SeqConstraint =
(evt1->evt2->evt3->evt4->SeqConstraint).
||CheckSeqConstraint = (SYSTEM||SeqConstraint)
/{timeout40hz/evt1, gps.outEvt/evt2,
airframe.outEvt/evt3, navDisplay.display/evt4
}.
```

LTSA reports no violations for this property. Suppose we change the sequencing order of gps.outEvt and airframe.outEvt:

```
property SeqConstraint =
(evt1->evt2->evt3->evt4->SeqConstraint).
||CheckSeqConstraint = (SYSTEM||SeqConstraint)
/{timeout40hz/evt1, airframe.outEvt/evt2,
gps.outEvt/evt3, navDisplay.display/evt4}.
```

Then, LTSA produces an error trace:

```
Trace to property violation in SeqConstraint:
  timeout40hz
  gps.outEvt
```

If some events may happen in parallel, that is, the events form a general graph instead of a linear chain, then we can only specify

those events that do form a linear chain, since LTSA does not allow non-determinism in property specifications. For demonstration purposes, assume that there is no precedence relation between `airframe.outEvt` and `gps.outEvt`, but both must follow `timeout40hz` and precede `navDisplay.display`, then we should write the sequencing constraint as follows:

```
property SeqConstraint =
  (evt1->evt2->evt4->SeqConstraint).
  ||CheckSeqConstraint = (SYSTEM||SeqConstraint)
  /{timeout40hz/evt1, gps.outEvt/evt2,
  navDisplay.display/evt4}.

property SeqConstraint2 =
  (evt1->evt3->evt4->SeqConstraint2).
  ||CheckSeqConstraint2 = (SYSTEM||SeqConstraint2)
  /{timeout40hz/evt1, airframe.outEvt/evt3,
  navDisplay.display/evt4}.
```

LTSA reveals no violations of these sequencing constraints.

5.4. Progress Property

The properties discussed so far are all *safety properties*, that is, they can be verified by detecting if a bad state is reached given a *finite* execution sequence. On the other hand, *liveness properties* can only be verified for an infinite execution sequence. A general treatment of liveness involves using a temporal logic to specify liveness properties. A restricted class of liveness properties is the *progress* property in the form of `progress P = a1, a2, ..., an`, which asserts that in an infinite execution of a system, at least one of the actions `a1, a2, ..., an` will be executed infinitely often. It is useful for verifying that a system does not contain starvation of certain actions. It is a stronger assertion than reachability, which only requires that certain actions are executed *at least once* during the system's lifetime.

For example, in order to check that the `+display+` methods of both Flight Plan Display and Navigation Display are executed infinitely often in any infinite execution of the MultirateSP scenario, we can add this to its FSP model:

```
progress P1 = {fltPlanDisplay.display}
progress P2 = {navDisplay.display}
```

Note that this is different from:

```
progress P1 = {fltPlanDisplay.display,
navDisplay.display}
```

which states that at least one of `fltPlanDisplay.display` and `navDisplay.display` are executed infinitely often.

6. Scalability Improvements

Perhaps the single biggest impediment to industry adoption of model-checking is lack of scalability due to state-space explosion. We have constructed the FSP model for the *Medium Single-Processor* (MediumSP) scenario, which consists of two threads running at 20Hz and 1Hz. Since there is no synchronization between `Thread20hz` and `Thread1hz`, we can compose and check each one separately. Let's focus on `Thread1hz`, shown in Figure 8. This scenario causes *out-of-memory* error on a state-of-the-art PC workstation. We have applied some techniques to

improve scalability of model-checking, as explained in the following sections. After applying these techniques, we were able to compose and check this scenario.

6.1. Exploiting Domain-Specific Constraints

We can take advantage of certain domain-specific constraints to simplify the model. Normally method calls are modeled with a two-way synchronization between the *caller* component and the *callee* component, as shown below in the FSP model:

```
Publisher = (outEvt->Publisher|receiveGDCall
->issueGDReply->Publisher).

Subscriber = (inEvt->issueGDCall->receiveGDReply
->Subscriber).

||CtrlPushDataPull =
(pub:Publisher||sub:Subscriber)
/{pub.outEvt/sub.inEvt,
sub.issueGDCall/pub.receiveGDCall,
sub.receiveGDReply/pub.issueGDReply}.
```

However, for all practical purposes we can treat the `GetData()` call and reply as an atomic operation, and omit the synchronization action on the method call reply:

```
Publisher = (outEvt->Publisher
|receiveGDCall->Publisher).

Subscriber = (inEvt->issueGDCall->Subscriber).

||CtrlPushDataPull =
(pub:Publisher||sub:Subscriber)
/{pub.outEvt/sub.inEvt,
sub.issueGDCall/pub.receiveGDCall}.
```

This optimization may not be generally applicable to all method calls, only to the `GetData()` call in the *control-push/data-pull* interaction style, where there is no action in between the `GetData()` call and reply. This involves modifying definition of each component type. After applying this optimization, the state space of `Thread1hz` has been reduced considerably. However, this is still too large for LTSA to handle on our PC workstation.

6.2. Compositional Analysis

Construction of the global state space of an application usually causes state-space explosion. We can take advantage of inherent modularity within the application to compose and check the system hierarchically, instead of composing the entire system state-space all at once. This is the typical divide-and-conquer approach. The compositional analysis technique allowed us to successfully compose and check the MediumSP scenario consisting of `Thread1hz` and `Thread20hz`.

LTSA has a built-in capability for *Compositional Reachability Analysis* (CRA) [7]. After a set of components have been checked to be correct, we can abstract and reuse them in other contexts by hiding irrelevant events and only exposing those events that may be of interest to other surrounding components, resulting in a simplified and minimized automaton. We can then reuse this automaton as a module in other contexts.

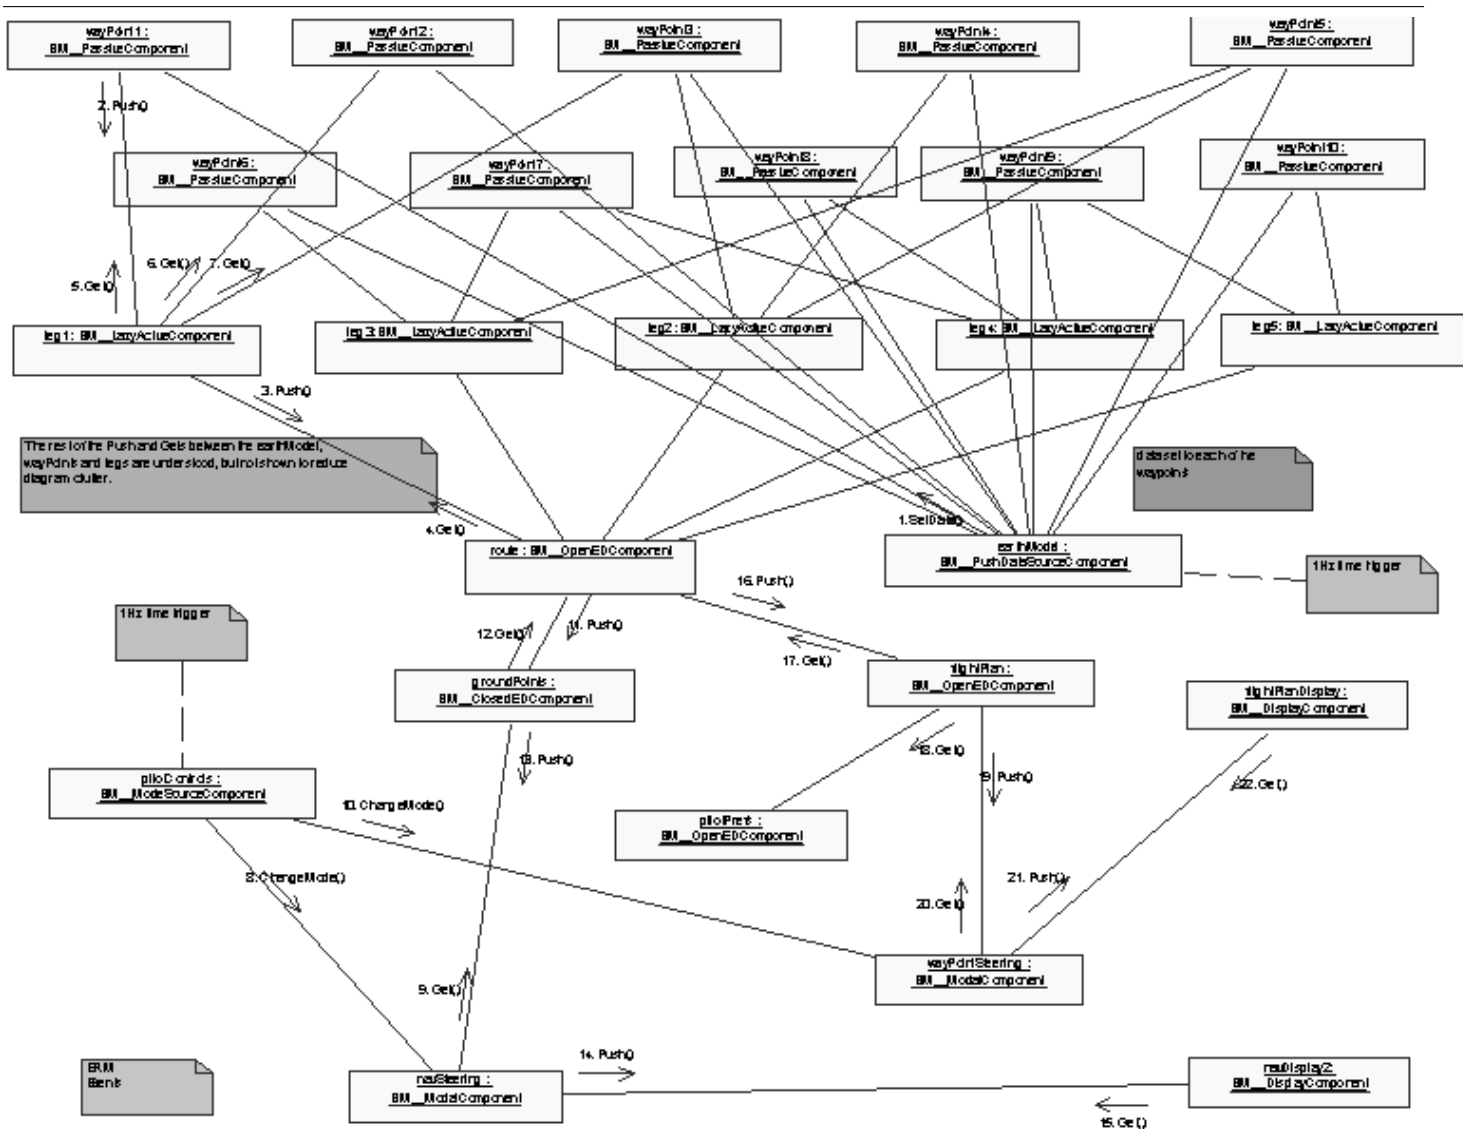


Figure 8. UML model for the 1Hz thread of medium single-processor scenario, taken from Boeing documentation [14].

Looking at the 1Hz thread in Figure 8, we can see that there is a natural separation into three groups of components:

1. earthModel, route, wayPoint1, wayPoint2, wayPoint3, wayPoint4, wayPoint5, wayPoint6, wayPoint7, wayPoint8, wayPoint9, wayPoint10, leg1, leg2, leg3, leg4, leg5.
2. pilotControls, groundPoints, fltPlan, navSteering, waypointSteering, pilotPref, fltPlanDsply, navDisplay2.
3. timer1hz.

We can compose Group1 and Group2 individually while minimizing each group by hiding events that are not relevant to the surrounding components, and finally composing them together with the 1Hz timer. For example, for Group1, we have:

minimal || Group1 =

```

(Timer1hz || earthModel:PushDataSourceComp
|| wayPoint1:PassiveComp || wayPoint2:PassiveComp
|| wayPoint3:PassiveComp || wayPoint4:PassiveComp
|| wayPoint5:PassiveComp || wayPoint6:PassiveComp
|| wayPoint7:PassiveComp || wayPoint8:PassiveComp
|| wayPoint9:PassiveComp || wayPoint10:PassiveComp
|| leg1:LazyActiveComp || leg2:LazyActiveComp
|| leg3:LazyActiveComp || leg4:LazyActiveComp
|| leg5:LazyActiveComp
|| correlatorLeg1:InputCorrelator(3)
|| correlatorLeg2:InputCorrelator(3)
|| correlatorLeg3:InputCorrelator(3)
|| correlatorLeg4:InputCorrelator(3)
|| correlatorLeg5:InputCorrelator(2)
|| route:OpenEDComp
|| correlatorRoute:InputCorrelator(5))
/{...event equivalence specifications omitted}

```

```
//Only three interface events are exposed
@{earthModel.inEvt, route.outEvt,
route.receiveGDCall}.
```

The minimized state machine is quite simple, consisting of only three states. Intuitively, the whole group of components behaves like a single component that accepts a timeout trigger that synchronizes with `earthModel.inEvt`, does some internal processing that is hidden from outside view, issues `route.outEvt`, and finally, receives a `GetData()` call from its downstream component. The FSP specification for `Group2` is similar and is omitted here. The overall system specification is:

```
||Thread1hz = (Timer1hz||Group1||Group2
|correlator:InputANDCorrelator)
/{timeout1hz/earthModel.inEvt,
timeout1hz/pilotControl.inEvt,
route.outEvt/groundPoint.inEvt,
groundPoints.issueGDCall/route.receiveGDCall,
route.outEvt/fltPlan.inEvt,
fltPlan.issueGDCall/route.receiveGDCall,
navDisplay2.display/correlator.inEvt[1],
fltPlanDisplay.display/correlator.inEvt[2],
correlator.matched/timer1hzdone}.
```

However, there is one drawback to the compositional analysis approach. Since internal events are hidden inside of each group of components, we cannot check for end-to-end sequencing constraints that span multiple groups and involves internal events from these groups. We can only check constraints that involve interface events that are exposed by the component group, or those that involve internal events of a single group. All the other properties are not impacted.

7. Performance Evaluation

We have applied model-checking successfully to a number of application scenarios. The experiments were performed on a PC workstation with 512MB of memory and Pentium IV processor running Windows XP. Obviously, using a more powerful computer with lots of memory would help with the scalability problem. The scenarios range from the Basic Single-Processor scenario (BasicSP) with 3 components, to the Medium Single-Processor scenario (MediumSP) with more than 50 components. However, scenarios larger than MediumSP are still beyond the reach of the model-checker despite our state-space reduction techniques. Model-checking generally finishes within seconds or at most a few minutes when the main memory is large enough, otherwise the computer goes into virtual memory thrashing mode and eventually gives out the *memory exhausted* message. We believe 50+ components is already a reasonable size to make this approach useful. A realistic Avionics system has up to hundreds of thousands of components, and no model-checker is expected to be able to scale up to that size. We will have to rely on the designer's manual work to separate out fragments of scenarios that are relatively isolated from the rest of the system and model-check them individually. This is a very reasonable thing to do, as the current application scenarios are just fragments taken from a production system.

Not surprisingly, we have found no errors in these application scenarios, which are fragments taken from a mature, tried-and-true production system. However, we believe the model-checking approach can act as a valuable debugging tool for uncovering subtle concurrency bugs during early design stage of a new system, or maintenance stage of a legacy system.

8. Related Work

Cadena [15] is an an integrated development, analysis, and verification environment for CORBA Component Model (CCM) systems, also using AMC as the application example. The model-checker Bogor [16] is integrated with Cadena for verification of system functional properties. Garlan [17] described a model-checking framework for publish/subscribe systems. The key feature of this framework is a reusable, parameterized state machine model that captures pub-sub runtime event management and dispatch policy. Generation of models for specific systems is then handled by a translation tool that accepts as input a set of component descriptions together with a set of properties, and maps them into the framework where they can be checked using the model-checker SMV [18]. Compared to [15] and [17], our modeling approach adopts a higher level of abstraction and ignores details related to the internals of middleware such as queuing and dispatch policies. This significantly reduces system state-space, and has turned out to be adequate for our purpose of verifying application-level concurrency properties, assuming that the middleware behaves correctly. We also take advantage of LTSA's compositional analysis functionality to help improve scalability, which is not present in [15] and [17]. Some of the property specifications, such as sequencing constraints and progress property, are also unique to our approach.

There are a number of interesting approaches on using model-checking for real-time scheduling analysis. ACSR [19] is a resource-aware real-time process algebra for specification and formal verification of distributed real-time systems. Main features include the ability to specify resources and their usage by system components, and prioritized execution that allows to express different preemptive and non-preemptive scheduling policies. Madl [20] developed automated translation from ESMML to Timed Automata, in order to use the model-checker UPPAAL [21] for verification of real-time properties. One limitation of their approach is that they can only model *non-preemptive* scheduling within a single rate group/thread, but not *preemptive* scheduling between threads, due to lack of expressive power of the Timed Automata formalism. The preemption effects of higher priority threads obviously have an impact on the timing behavior of the lower priority threads since they share the same CPU. As discussed in Section 4.2.3, we believe a more practical approach is to use real-time scheduling theory to prove timing and schedulability properties [22] [5], and use model-checking to prove concurrency properties. Li [23] developed *meta scheduler*, a middleware framework for implementing custom real-time scheduling algorithm on top of a POSIX-compliant operating system, and used UPPAAL to formally verify its correctness. According to our understanding, this modeling approach also seems to be limited to non-preemptive

scheduling.

Karamanolis [24] used LTSA to model and verify workflow schemas by mapping the Workflow Definition Language into FSP models. There are some similarities between the computational models of workflow schemas and AMC software. Both consist of components interacting with events sent and received from output and input ports. However, there are also important differences due to the different application domains. For example, AMC software typically contains several threads executing periodically, while the life-cycle of a workflow schema only consists of one execution from start to finish. The different execution frequencies of multiple threads cause the state space of an AMC application to be much larger than a workflow schema specification with similar complexity. From a modeling perspective, a self-loop has to be added to each legal terminating state in [24] in order to avoid false alarms when checking for deadlocks, while this is not necessary for AMC since it is a reactive system that should never terminate.

9. Conclusions and Future Work

In this paper we have used the model-checking for formal verification of component-based real-time embedded software based on CORBA Event Service. The documentation provided by our industrial partner describes the application components and scenarios with natural language, which is prone to misunderstanding and misinterpretation. Using the FSP modeling language, we were able to formally model the application, and use the LTSA model-checker to verify concurrency properties such as deadlock freedom, event reachability, sequencing constraints and progress property. We also applied effective techniques to cope with the state-space explosion problem. First, we exploit domain-specific properties to reduce the call-return two-way synchronization into a one-way synchronization, thus reducing the number of states of each component. Second, we take advantage of inherent modularity within the application scenario, and use the divide-and-conquer approach to compose the system hierarchically. These techniques showed significant benefits in reducing system state-space, and allowed us to check relatively complex application scenarios on a PC workstation with a relatively modest memory size of 512MB. When the system size is too large for the model-checker to handle, the designer can at least use simulation to gain some insight into system behavior.

As mentioned in the introduction, the UML models for AMC mainly serve in an informal documentation role that the engineer refers to while writing code manually. This is one of the major motivations for developing domain-specific modeling languages and tools to replace UML, such as the MoBIES tool-chain [2], Cadena [15], VEST [25], Time Weaver [26] and CoSMIC [10], in order to have a more formal, automated and integrated software development process. However, there has been recent progress on making UML more formal and suitable for modeling real-time embedded and component-based software. Some examples include the UML Profile for Scheduability, Performance and Time [27] and the UML Profile for CORBA Components [28]. UML 2.0 has adopted the concept of *capsules* communicating with message passing through *ports*, originally de-

veloped in *Real-Time Object-Oriented Modeling (ROOM)* [29], which bears some similarity to the model of computation in AMC and CCM. An interesting question to ask is, can we not give up on UML and develop custom, proprietary modeling notations, but leverage the body of work from the UML community to develop tools based on standards? One argument for preferring a custom modeling approach to UML is that we can achieve better domain-specificity by customizing the meta-model, which is more powerful and flexible than the UML profiling mechanism. Another argument is that embedded systems are so diverse that it is next to impossible to have one standard notation that is suitable for all application domains. We plan to investigate these interesting issues in our future work.

References

- [1] D. Sharp, "Object-oriented real-time computing for reusable avionics software," in *Proc. IEEE International Symposium on Object-Oriented Real-Time Distributed Computing*, 2001, pp. 185–192.
- [2] Z. Gu, S. Wang, S. Kodase, and K. G. Shin, "An end-to-end tool chain for multi-view modeling and analysis of avionics mission computing software," in *Proc. IEEE Real-Time Systems Symposium*, 2003, pp. 78–81.
- [3] G. Karsai, S. Neema, A. Bakay, A. Ledeczki, F. Shi, and A. Gokhale, "A model-based front-end to tao/ace," in *Proc. Second Workshop on the ACE ORB (TAO)*, 2002. [Online]. Available: <http://www.cs.wustl.edu/schmidt/TAOWS02/>
- [4] A. Ledeczki, M. Maroti, G. Karsai, and G. Nordstrom, "Metaprogrammable toolkit for model-integrated computing," in *Proc. IEEE International Conference on Engineering of Computer-Based Systems*, March 1999, pp. 311–317.
- [5] Z. Gu, S. Kodase, S. Wang, and K. G. Shin, "A model-based approach to system-level dependency and real-time analysis of embedded software," in *Proc. IEEE Real-Time Technology and Applications Symposium*, 2003, pp. 78–85.
- [6] D. Harel, "Statecharts: A visual formalism for complex systems," *Science of Computer Programming*, vol. 8, no. 3, pp. 231–274, June 1987. [Online]. Available: citeseer.nj.nec.com/harel87statecharts.html
- [7] J. Magee and J. Kramer, *Concurrency: State Models and Java Programs*. Wiley, 2000.
- [8] OMG, "CORBA Component Model, v3.0," Object Management Group, Tech. Rep., 2002.
- [9] N. Wang and C. Gill, "Improving real-time system configuration via a qos-aware corba component model," in *Proc. IEEE Hawaii International Conference on Systems Sciences*, 2004, pp. 273–282.
- [10] K. Balasubramanian, J. Balasubramanian, J. Parsons, A. Gokhale, and D. C. Schmidt, "A platform-independent component modeling language for distributed real-time and embedded systems," in *Proc. IEEE Real-Time and Embedded Technology and Applications Symposium*, 2004.
- [11] D. Schmidt, D. Levine, and T. Harrison, "The design and performance of a real-time CORBA object event service," in *Proc. ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications*, 1997, pp. 434–445.
- [12] J. Masse, S. Kim, and S. Hong, "Tool set implementation for scenario-based multithreading of uml-rt models and experimental validation," in *Proc. IEEE Real-Time and Embedded Technology and Applications Symposium*, 2003, pp. 70–77.

- [13] C. Liu and J. W. Layland, "Scheduling algorithms for multi-programming in a hard real-time environment," *Journal of the ACM*, vol. 20, no. 1, pp. 46–61, 1973.
- [14] W. Roll and G. Heineman, "Product scenario description document for the model-based integration of embedded software weapon system open experimental platform," Boeing, Tech. Rep., 2002.
- [15] J. Hatcliff, W. Deng, M. Dwyer, G. Jung, and V. Prasad, "Cadena: An integrated development, analysis, and verification environment for component-based systems," in *Proc. IEEE International Conference on Software Engineering*, 2003.
- [16] W. Deng, M. Dwyer, J. Hatcliff, G. Jung, Robby, and G. Singh, "Model-checking middleware-based event-driven real-time embedded software," Kansas State University, Tech. Rep. SAnToS-TR2003-2, April 2003.
- [17] D. Garlan, S. Khersonsky, and J. S. Kim, "Model checking publish-subscribe systems," in *Proc. International SPIN Workshop on Model Checking of Software*, 2003, pp. 166–180.
- [18] K. McMillan, *Symbolic Model Checking*. Kluwer Academic Publishing, 1993.
- [19] H. Ben-Abdallah, D. Clarke, I. Lee, and O. Sokolsky, "Paragon: A paradigm for the specification, verification, and testing of real-time systems," in *Proc. IEEE Aerospace Conference*, Feb 1997, pp. 469–488.
- [20] G. Madl, S. Abdelwahed, and G. Karsai, "Automatic verification of component-based real-time corba applications," in *Proc. IEEE International Real-Time Systems Symposium*, 2004.
- [21] J. Bengtsson, K. G. Larsen, F. Larsson, P. Pettersson, and W. Yi, "UPPAAL - a tool suite for automatic verification of real-time systems," in *Proc. Workshop on Verification and Control of Hybrid Systems*, October 1995, pp. 232–243.
- [22] K. Bryan, M. Murphy, P. Gupta, Y. Liu, L. DiPippo, and V. Fay-Wolfe, "The design, implementation, and performance of the uri static scheduling and real-time binding services for tao," in *Proc. 3rd Workshop on The ACE ORB (TAO)*, 2003. [Online]. Available: <http://www.cs.wustl.edu/cdgill/TAO03/>
- [23] P. Li, B. Ravindran, S. Suhaib, and S. Feizabadi, "A formally verified application-level framework for real-time scheduling on posix real-time operating systems," *IEEE Trans. Software Eng.*, vol. 30, pp. 613–629, 2004.
- [24] C. Karamanolis, D. Giannakopoulou, J. Magee, and S. M. Wheeler, "Model checking of workflow schemas," in *Proc. IEEE International Enterprise Distributed Object Computing Conference*, 2000, pp. 170–179.
- [25] J. A. Stankovic, R. Zhu, R. Poornalingam, C. Lu, Z. Yu, M. Humphrey, and B. Ellis, "Vest: an aspect-based composition tool for real-time systems," in *Proc. IEEE Real-Time and Embedded Technology and Applications Symposium*, 2003, pp. 58–69.
- [26] D. de Niz and R. Rajkumar, "Time weaver: A software-through-models framework for embedded real-time systems," in *Proc. ACM Conference on Languages, Compilers and Tools For Embedded Systems*, 2003, pp. 133–143.
- [27] OMG, "Uml profile for schedulability, performance, and time specification," Object Management Group, Tech. Rep., 2003. [Online]. Available: <http://www.omg.org/technology/documents/formal/schedulability.htm>
- [28] —, "Uml profile for corba components specification," Object Management Group, Tech. Rep., 2004. [Online]. Available: <http://www.omg.org/cgi-bin/doc?ptc/2004-03-04>
- [29] B. Selic, G. Gullekson, and P. T. Ward, *Real-Time Object Oriented Modeling*. Addison Wesley, 1994.