

# Towards Predicting Real-Time Properties of a Component Assembly

Egor Bondarev<sup>1,2</sup>, Peter de With<sup>2</sup>, Michel Chaudron<sup>1</sup>

*System Architecture and Networking<sup>1</sup> and Video, Coding and Architectures<sup>2</sup> groups  
Eindhoven University of Technology, P.O. Box 513,  
5600 MB Eindhoven, The Netherlands  
[E.Bondarev@tue.nl](mailto:E.Bondarev@tue.nl)*

## Abstract

*This paper addresses the prediction of timing properties of a component-based application already during the composition phase. At this stage, it is of vital importance to guarantee that the timing requirements (e.g. end-to-end deadlines) of a real-time application that is executed on a target system will be satisfied. This is obtained by predicting the real-time behaviour of a component-based application. In this paper, we extend an already existing scenario-based approach [2] with the possibility to model the behaviour of an application and the behavior of the underlying components. As a result, an application developer can reason accurately about the dynamic resource consumption and real-time properties of a component assembly. The modeling involves the specification of synchronization constraints for tasks and the simulation of application behaviour. A concluding case-study of video encoder development reveals that the approach is not only feasible but also addresses yet unsolved problems of task parallel execution and synchronization.*

## 1. Introduction

Embedded systems are characterized by two closely coupled properties: limited resources and real-time constraints for execution of the applications. The limitation of resources, such as memory size, bus bandwidth and processing power, complicates the satisfaction of the real-time constraints. It is evident that these guarantees are of great importance for e.g. multimedia devices.

For high-volume embedded appliances, such as PDAs and mobile phones, etc, an open, component-based framework for a middleware layer in the

software architecture has been proposed. This framework, known as Robocop [3], was used as a reference for specifying a follow-up ITEA research project, called Space4U. Our research aims at improving the Robocop architecture, by enabling the predictions of application real-time properties.

During the design phase, an application is evaluated to fit on a target system. For an *a-priori* evaluation, this requires prediction of the resource usage of an application. Early prediction of resource usage and timing properties of an application at the design stage increases system robustness and reduces cost and problems in product development.

Component-based technology complicates the prediction of resource usage and timing properties of an application. In component-based systems, the actual behaviour and resource usage are determined by an ensemble of internally and also *externally* developed components. Thus, the prediction task becomes twofold: (1) find and express the *component's* extra-functional properties, and (2) combine these properties to predict the behaviour of the *composition* of the constituent components. In the sequel, we will denote an application also as an *assembly*, because it makes use of the underlying components.

The challenge of predicting real-time properties of a component assembly is of significant interest because of the rapid development of component-based technologies in the embedded systems domain. Some approaches represent an engineering practice [7-9] to the problem. We used [8] as a guideline for our work. A very promising technique that allows design-time estimations of real-time properties of component-based systems is presented in [10]. In this technique, many possible types of software constructions are taken into account, like synchronous and asynchronous communication, as well as synchronization constraints. Recent work on the prediction of performance for

evolving architectures is described in [11]. This approach is based on collecting the component performance data on different platforms and interpolating it for new components or platforms. Real-time frameworks have been introduced in the object-oriented development field. Methods have emerged that enable execution of UML-like specifications, notably Room [12] and Rhapsody [13]. The PRIMA-UML methodology [14] applies queuing networks and extends UML with a real-time performance model for system performance validation. We concentrate on similar methods, but now in the component-based development field. The scenario-based approach proposed in [15] involves estimating static resource usage of a component assembly.

In contrast with this, through our scenario simulation approach, we address a *dynamic* instead of static resource usage, thereby giving more accuracy in the prediction of the assembly behaviour. With respect to *task synchronization*, we adopt the use of synchronization constraints (precedence, critical sections, mutexes) for further adding accuracy in the prediction. The approach still requires little effort from an application developer, because the introduction of application scenarios narrows the state-space and behaviour of an application that the developer should model and simulate. A practical case-study revealed that, besides proving feasibility, the important problem of task parallel execution and synchronization comes to the foreground. This prevents that these usually hidden problems remain unsolved.

The paper is structured as follows. Section 2 refers to the Robocop component model, as a fundament for the new technique. Section 3 addresses various aspects of timing properties of a component assembly. Section 4 discusses the workflow of the approach and gives specifications of the required models. Section 5 clarifies the proposed approach with an encoder application case-study. Section 6 concludes with the *pros* and *cons* of the prediction-enabling approach.

## 2. Robocop component model

The proposed prediction-enabling approach is based on the Robocop component model. The description of the Robocop component model stays out of the scope of this paper. The model concepts are explained in section 2 of [5].

## 3. Timing and Task Synchronization

This section defines basic terms used in the paper, e.g. timing property, task and task synchronization

constraints. There is a clear difference between the component and application timing properties. The *component timing properties* are independent from system run-time execution and scheduling. In most of the cases, these properties are: worst-case, mean-case and best-case execution times per operation. The *application timing properties*, instead, are closely coupled to run-time instances, tasks and scheduling algorithms used in the system. In the real-time application domain, we concentrate on the following timing properties: response time, blocking time of a task, and the number of missed deadlines of a task.

The *response time* of a task is not just a sum of execution times of the operations comprising the task. Usually, the response time is composed of the execution time, blocking time and pre-emption time of the task. Therefore, for the assembly timing property, the task synchronization and scheduling aspects should be considered.

In literature, several definitions of tasks are used. In our context, the *task* is an event-triggered sequence of executed operations. The operations composing a sequence may be implemented by different services. The operations within the sequence may be called synchronous or asynchronous. The tasks may have *synchronization constraints* between them, e.g. precedence, rendezvous and mutual exclusion. Usually, the system resource sharing imposes the synchronization constraints.

All these aspects have been used in this paper in a regular way, confining to the presented definitions.

## 4. Scenario simulation approach

### 4.1. Approach overview

Our approach proposes to combine the *behaviour* and *resource consumption models* of used components with an *application model* constructed for possible critical execution scenarios. The *application scenario model* defines the static structure, internal and external events of the application for critical scenarios. *Critical scenarios* are the scenarios that may introduce CPU or memory overload. The resulting set of models serves as an input for virtual scheduling (simulation). The simulation output data shows execution behaviour of the assembly tasks and timing properties of those tasks, i.e. predicted execution timeline, latency and resource utilization of each task.

Summarizing, the key features of our approach are:

- a) Predictions on timing properties are made by simulation at an early stage of development.


- b) It avoids combinatorial complexity of full state-space analysis by usage of scenarios.
- c) It takes task synchronization and scheduling aspects into account.

The following section describes *how* the approach should be implemented.

## 4.2. Workflow of the approach

The main objective of an assembly developer is, given a set of available components and requirements for an assembly, to embed the components in the assembly satisfying the given requirements.

In the domain of real-time application, a developer needs to focus on satisfying extra-functional requirements like response time, or busload (Figure 1).

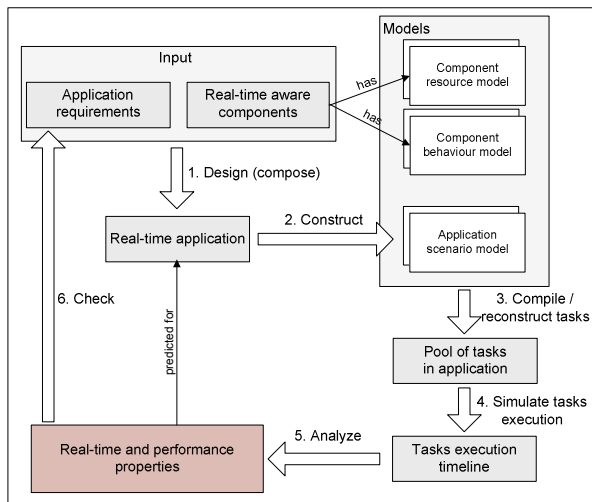


**Figure 1. Conceptual view on the prediction-enabling composition workflow.**

The workflow works is based on two *assumptions*:

- a) Resource usage property and behaviour of the constituent components are specified and available in the used component models.
- b) An application developer is able to find out critical scenarios of the application.

The main steps of the workflow are depicted in Figure 2. The remainder of this section defines the consecutive steps of the workflow.



**Figure 2. Main steps in the predictable RT composition workflow.**

A developer selects and composes a set of available components into an application. According to the above assumption (1), selected components should be *real-time aware*, e.g. have both a *resource model* and a *behaviour model*. These two models are used to accompany an *application scenario model* that is constructed in the next step and, thus, complete the mosaic of the application behaviour. The description of the models is given in Section 4.3.

### Construction of application scenario model

For each critical or commonly used scenario, a developer constructs an *application scenario model* (see Figure 2). The application scenario model consists of two parts: (a) description of service instances and bindings between them, particular for the selected scenario, and (b) description of the application-level events and active threads that trigger execution of operations of the service instances.

### Compilation of models

The *application scenario*, *component resource* and *component behaviour models* are jointly compiled. The goal of the compilation is to reconstruct (generate) the tasks running in the application. Prior to compilation, the task-related data is spread over different models. For instance, the task periodicity may be specified in an *application scenario model*, whereas the operation call sequence comprising the task is specified in relevant *component behaviour models*. The compiler reconstructs all necessary properties of the tasks, like deadline, period, priority and operation call sequence.

### Simulation of tasks execution

An application developer applies a scheduler to the reconstructed task pool, simulating the execution of defined scenario. The scheduling algorithm may vary depending on the algorithm of the operating system, on which the application is supposed to execute. The scheduler should implement prevention of unbounded priority inversion, because the models define various types of synchronization constraints. Resulting data from the scheduler is the *task execution timeline*. This timeline is a subject for schedulability and performance analysis.

### Schedulability analysis

The analysis of the task execution timeline helps to reason about application timing properties like response time, latency of critical tasks, overall schedulability and processor utilization bounds. Many other possible application properties can be derived:

rate of missed deadlines, blocking time, worst and best-case response time per task.

This step results in predicted *real-time and performance properties* of the designed application.

### Checking properties against requirements

The predicted timing properties are checked against the real-time requirements of an application (see Figure 2). For example, worst-case response time of a critical task is verified with its deadline specified in the requirements. If any of the requirements are not met, a developer optimizes the composition and repeats the workflow.

### 4.3. Model description

The purpose of this section is to specify the models introduced in the previous section. It is emphasized here that the models are not a goal by themselves, but are required for obtaining the resource consumption and timing properties.

According to Figure 4, we propose to model application scenarios. This allows decomposing each type of application behavior into a separate simple scenario model. Thus, we can reduce the complexity of the complete behavioral model of the application and partly avoid exploration of all application states.

The following sub-sections specify the above-mentioned models in detail.

#### Component resource model

The *component resource model* (RM) is one of the models of the Robocop component model. RM specifies the predicted resource consumption for all the operations *impl\_opr* implemented by services of an executable component (for certain platform). Resources (*r*) can be memory, CPU, etc. The predicted resource consumption is specified as a (claim, release) tuple for non-processing resources, like memory. For processing resources, like the CPU, the consumption is specified as a single claim.

$m = RM,$   
 where  $m$  is a *Resource Model* and  
 $RM$  is a set of  $rm$  (resource usage of an operation).  
 $rm = (impl\_opr, resource, consumption),$   
 for operation  $impl\_opr$ .  
 $resource = r \in \{memory, cpu, bus, \dots\}.$   
 $consumption = claim,$   
 in case  $resource$  is  $cpu$ .  
 $consumption = (claim, release),$   
 in case  $resource$  is  $memory$ .  
 $consumption = (claim, time),$   
 in case  $resource$  is  $bus$ .

A component developer defines the resource consumption properties of an operation by worst-case analysis. These properties are calculated only for the operation body itself, excluding resource consumption properties of called operations. This approach allows calculating resource consumption of any sequences of operation calls. In this paper, we do not address platform and parametric variations of the operation resource consumption. The resource model should be specified for a particular reference platform.


#### Component behaviour model

The *component behaviour model* (BM) also belongs to the Robocop component model. BM specifies the behaviour of all operations *impl\_opr* implemented by services of an executable component. A semi-formal specification of the model is as follows.

$m = BM,$   
 where  $m$  is a *Behaviour Model* and  
 $BM$  is a set of  $bm$  (behaviour of an operation).  
 $bm = (impl\_opr, mutexed, behaviour, T),$   
 where  $impl\_opr$  is the implemented operation and  
 $behaviour$  is the operation behaviour description,  
 $T$  is a set of  $t$  (task triggers the operation is associated with),  
 $mutexed$  shows if the operation is mutexed.  
 $mutexed = \in \{true, false\}.$   
 $behaviour = (called\_opr1, called\_opr2, \dots called\_oprn, CS),$   
 where  $called\_opr1, \dots called\_oprn$  is a sequence of called operations and  
 $CS$  is a set of  $cs$  (critical sections).  
 $called\_opr = (opr, nmb\_iterations, calling\_type),$   
 where  $opr$  is the called operation and  
 $nmb\_iterations$  - number of times the operation is called,  
 $calling\_type \in \{synchronous, asynchronous\}.$   
 $cs = (called\_opr1, called\_opr2, \dots called\_oprn).$   
 $t = (periodicity, param, PRECED),$   
 where  $periodicity \in \{periodic, sporadic, aperiodic\},$   
 $PRECED$  is a set of  $preced$  (preceding task triggers),  
 $param$  includes various parameters of  $t$ .  
 $param = (period, interarrival\_time, priority, deadline, offset, jitter).$   
 $preced = (t, ratio),$   
 where  $t$  is a task trigger that precedes the specified task trigger.  
 $ratio = \frac{nmb\_jobs\_of\_current\_task}{nmb\_jobs\_of\_preceding\_task}.$

Firstly, for each operation *impl\_opr* implemented by an executable component, a component developer defines its mutual exclusion property. If an operation is *mutexed*, at most one thread can enter the operation at the same time. Secondly, operation *behaviour*

describes a sequence of operation calls to other interfaces made inside the implemented operation. For example in Figure 5, the implemented operation `Decoder.decode()` has a behaviour described by the following call sequence: `IGetElement.getFrame()`, `IStoreElement.storeFrame()`. The `IGetElement` and `IStoreElement` are the interfaces provided by `ReadBuffer` and `WriteBuffer` services correspondingly.



**Figure 3. Sequence of operation calls (behaviour) of `decode()` operation.**

For each called operation *called\_opr* in the sequence, the number of iterations *nmb\_iterations* and calling type *calling\_type* are specified. Additionally, a set of critical sections *CS* can be specified if necessary in *behaviour*. Critical section *cs* points out the operation of which the execution cannot be pre-empted. Please note that each *called\_opr* must belong to one of the required interfaces for the service.

Finally, a component developer must define the operation autonomous behaviour *T*. We consider that an operation has autonomous behaviour if there is at least one task trigger *t* implemented by the operation. One of the examples of the task trigger is an iterative thread, triggered periodically by a timer. In the decoder example, the `decode()` operation can implement an iterative thread, which is triggered by the system timer each 20 ms. Thus, the whole calling sequence repeats each 20 ms. In the model, the task trigger properties can be specified, including *periodicity*, *period*, *deadline*, *offset*, precedence constraints *preced*, etc.

Concluding, these two models describe component resource consumption and behaviour properties independent of the application context where the component is going to be used.

### Application scenario model

The *application scenario model* (SM) specifies application structure and behaviour for a critical or commonly used execution scenario. Several SMs can be built for an application, depending on a number of interesting scenarios. An application developer is in charge of the *scenario models* construction. The semi-formal structure of the model is presented below.

$$SM = (appl, structure, E, T, depend),$$

where *E* is a set of *e* (event coming from outside of the *appl*),

*T* is a set of *t* (task trigger the *appl* implements), *depend* is a set of components used in the *appl*.

$$structure = (SI, B),$$

where *SI* is a set of *si* (service instances) and *B* is a set of *b* (bindings).

$$b = (from, from\_port, to, to\_port).$$

*from, to* = service instance.

*from\\_port, to\\_port* = port name (named interface).

*e* and *t* = (*opr, periodicity, param, PRECED*),

where *opr* is an operation triggered by the *e* or *t*,

*periodicity*  $\in$  {*periodic, sporadic, aperiodic*},

*PRECED* is a set of *preced* (preceding *e* or *t*),

*param* is number of parameters of *e* or *t*.

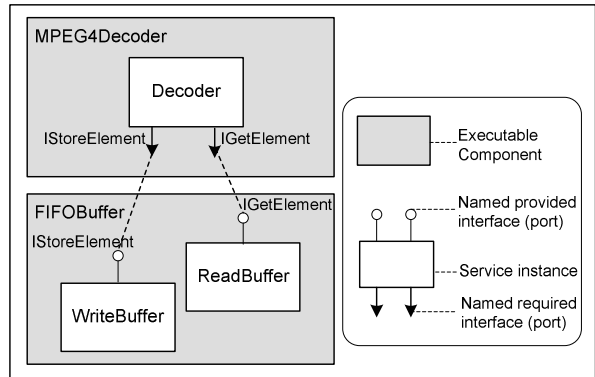
*param* = (*period, interarrival\_time, priority, deadline, offset, jitter*).

*preced* = (*e* or *t, ratio*),

where *e* or *t* is event or trigger which precedes the current one.

*ratio* =  $nmb\_current\_events/nmb\_preceding\_events$ .

Firstly, an application developer specifies an application *structure* for a scenario. The *structure* is represented by a tuple containing *SI* (set of service instances *si*) and *B* (set of *bindings* between the *si*). A *binding* includes information about the bound service instances *from*, and *to*, and in/out ports of the instances *from\\_port, to\\_port*. In Figure 6, dashed lines represent the bindings.



**Figure 4. Example of application structure.**

Secondly, the model defines the components (*depend*) used in the application. This data links the scenario model with the *behaviour* and *resource models* of the corresponding components.

Finally, the application scenario model specifies sets *E* and *T* of events *e* and in-application task triggers *t*, respectively. We define an as event any influence coming from *outside* to an application that changes the current application state. Hardware interrupt, timer or signal from an external sensor can trigger the event.

Normally, this influence is expressed as a call of one of the operations of the application component.

Conceptually, an in-application task trigger is also an event, but it comes from *inside* the application. In other words, this task trigger is implemented by the application. Please recall that we also have a task trigger notion in the *component behavior model*. That task trigger differs by being implemented inside a component. The two types of task triggers are separated into different models, because an in-component task trigger must be specified by a component developer and an in-application task trigger must be specified by an application developer.

The application task trigger calls one of the operations of the application components, thereby starting the task action sequence. Therefore, the  $e$  and  $t$  must be associated with the operation called first (*opr*). In Figure 7, an application periodic task trigger calls `decode()` operation each 40 ms. Thus, in the *scenario model* the trigger shall be associated with the operation.

For each event  $e$  as well the in-application task trigger  $t$ , its *periodicity*, parameters *param* and precedence constraints *preced* are specified.




Figure 5. Task triggered by in-application trigger.

When the scenario models are ready, an application developer proceeds to the simulation phase.

#### 4.4. Model Compilation and Schedulability

In the Space4U project, we have developed a Robocop Integration Environment (RIE) tool that does compilation of the above-mentioned models, simulation of an application scenario and visualization of the simulation data.

In the simulation and schedulability analysis phase, an application developer brings together the *application scenario model* and combined *behaviour-resource models* of the components deployed in the application. At this stage this stack of models can be compiled by RIE. The conceptual goal of the compilation is to identify and reconstruct a set of tasks that the application executes in a particular scenario.

The task-set reconstruction uses only the data from the three above-mentioned models. These models

contain all events; in-application and in-component task triggers, as well as operation call sequences that define a flow of control for the tasks.

For the decoder example, the task reconstruction works as follows: the related *behaviour model* specifies the operation call sequence of the operation `decode()`: `getFrame()`, `storeFrame()` (see Figure 5). Afterwards, the compiler gathers from related *behavior models* the behaviour of these two operations. The operation `getFrame()` calls one operation belonging to other interfaces: `ILogData.logEvent()` (see Figure 8).




Figure 6. `getFrame()` and `storeFrame()` behaviour.

If an operation has an empty operation call sequence (does not call operations belonging to other interfaces), it is considered as a leaf and the task generation proceeds to the next branch. Let us assume that operation `ILogData.logEvent()` is such a leaf. The next operation `storeFrame()` then also calls this leaf operation: `ILogData.logEvent()` (see Figure 8). Thus, the complete reconstructed sequence of the operations executed in the task is as shown in Figure 9.

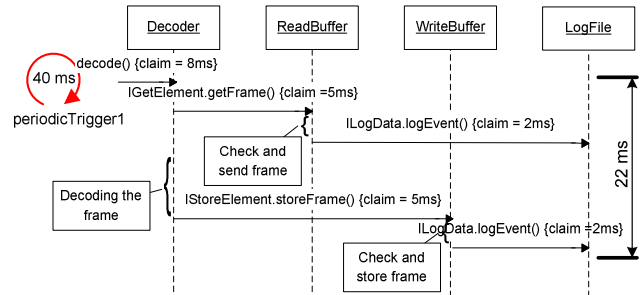


Figure 7. Task generated from the models.

A resource consumption property of each operation in this sequence is specified in the *claim* primitive in the related *component resource model* (see Section 6.1). Knowing this data, we can calculate total resource consumption of the task. For example, the CPU time used by the task (execution time) is the sum of CPU times used by the operations composing the task. In Figure 9, the total execution time of the task amounts to:  $8\text{ms} + 5\text{ms} + 2\text{ms} + 5\text{ms} + 2\text{ms} = 22\text{ms}$ . The other task parameters (period, offset, and deadline) and

precedence are obtained from corresponding task trigger properties specified in the models.

Synchronization constraints for each task are also extracted from the models. The task precedence has been already mentioned. *Mutexed* and critical section *cs*, which are properties of an operation, as well as a task precedence *preced* specified in the *component behavior model*, all define synchronization constraints of tasks. If a mutexed operation of the same service instance is used by two different tasks, then only one of the tasks can execute the operation at the same time.

An execution of the reconstructed tasks of the scenario is simulated by a virtual scheduler. During the simulation, these synchronization constraints are taken into account.

The simulation results are represented as a task execution timeline (see Figure 10).




Figure 8. Task timeline execution of scenario

The schedulability analysis of the simulation data gives us the timing properties of an application. The response time, blocking time, number of missed deadlines can be found for each task. Beside this, the processor utilization bound can be analyzed per application. The predicted properties can be validated with respect to the application requirements.

## 5. Video Encoder Case-study

The objectives of the video encoder case-study are to show practical aspects of the approach utilization and give further clarification. The example starts with requirements, goes through the prediction-enabling composition workflow and ends with predicted timing properties of the application.

### 5.1. Requirements

Taking into account that we do not focus on functional requirements, the required functionality can be expressed in one sentence: *the application shall encode on-the-fly the audio and video signals in MPEG-4 format and subsequently multiplex the compressed signals into one stream* (REQ1).

The extra-functional requirement for the TV-like application: *the number of skipped frames during the encoding on-the-fly should be NULL* (REQ2). This implies that we do not allow missed deadlines for audio and video encoding tasks (real-time application).

### 5.2. Component selection

After the requirements elicitation, the process of the component-based application development continues with component selection. Because our application has a real-time nature, we should select only real-time aware components (resource and behaviour models in their distribution package).

We selected two real-time aware components that bring the required functionality: MPEG4\_Encoder and Data\_Broker having three service each as indicated in Figure 9. Each service has *provides* and *requires* interfaces. For instance, the VideoEnc service provides IVideoEnc interface and requires IBufferAccess and IMux interfaces. The IVideoEnc interface encapsulates the VideoEncode() operation. All public operations are also represented in Figure 9.

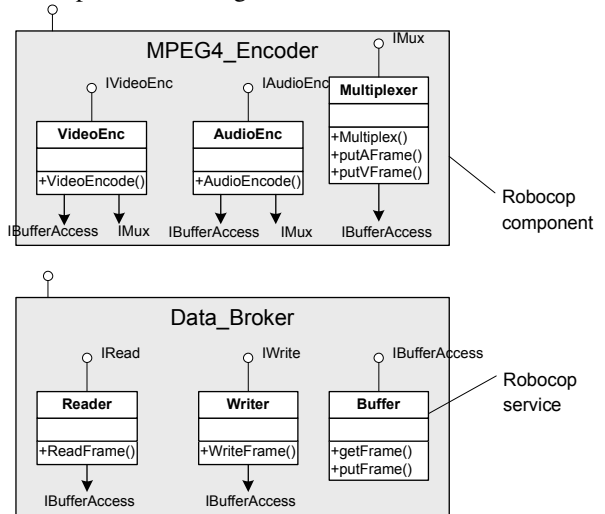


Figure 9. Selected components with services, interfaces and operations.

The corresponding resource and behaviour models (see Figure 10) are constructed according to the rules defined in Section 4. The behaviour (resource) model specifies behavioral (resource usage) aspects of all public operations of the component. Note that there are no task-triggering operations specified in services of both components (fields for task triggers *T* are empty). It means that all operations are passive (have no autonomous behaviour) and should be controlled by application-level events and task triggers.




Figure 10. Behaviour and resource models of the components.

### 5.3. Composing the encoder application

The design (composition) phase, which is the first stage of the workflow, consists of three steps: service instantiation, service instances binding and design of application level events and task triggers.

The *service instantiation* is basically a process of defining a structure of an application depending on required functionality. Our encoder should read, encode AV streams and multiplex them in one MPEG-4 stream. Finally, this stream should be stored. Therefore, the encoder should have at least the following service instances: audio-, video- readers, audio-, video- encoders, multiplexer, and writer. Data communication between the instances can be realized by a set of buffers. This structure (service instantiation) is depicted in Figure 11. As can be noticed, the Reader service is instantiated twice (aReader, vReader) and the Buffer service has three instances (vBuffer, aBuffer, mBuffer).

The second step is *binding the service instances*. *Requires interfaces* are connected to *provides interfaces* of the same type, thus defining data and control flows in the application. Figure 11 depicts the service instance bindings by the connecting arrows.

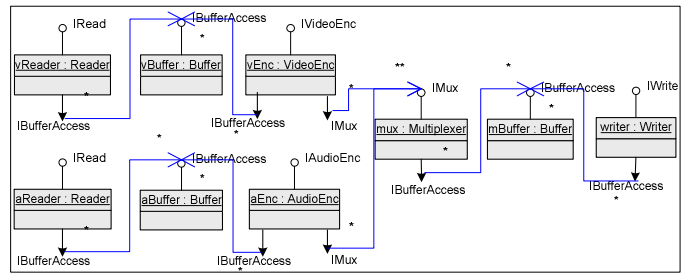


Figure 11. Binding the encoder service instances.

In the third step, a developer identifies necessary application-level *events* and *task triggers*. In component-based systems, an application-level *task trigger* can be implemented in the code of application as a separate thread that wakes up periodically (by timer signals) and invokes one of the component operations. In its turn, an *event* is usually implied by hardware platforms (i.e. interrupts). A developer needs not to implement events, but should take them into account during the design phase.

The services composing the encoder application have no autonomous operations with task triggers inside (all services are passive). In order to make the application alive, we designed six task triggers executing on the application level (see Figure 12). Each of the task triggers periodically invokes one of the operations, thereby creating a separate thread of control. For example, Trigger1 invokes the IRead.readFrame() operation of the vReader service instance. This operation reads one video frame from a file and stores the frame in vBuffer. All triggers are designed to fire with periodicity of 40 ms, since this is common video streaming rate. We defined the deadlines for the triggered tasks to be equal to their periods (40 ms). We specified no precedence constraints for the tasks. Having this information we can construct an *application scenario model*.

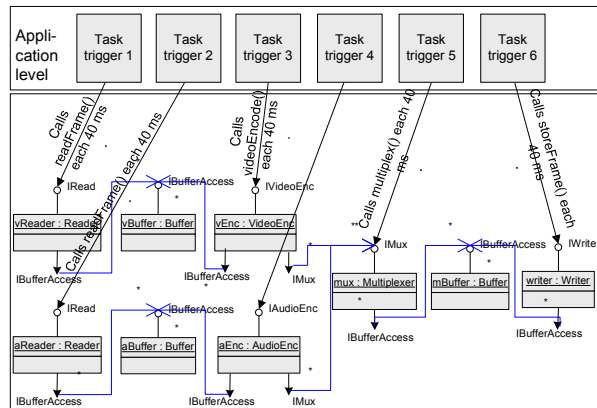


Figure 12. Application level task triggers.

## 5.4. Constructing a scenario model

The construction of a scenario model starts with identification of relevant scenarios. The relevant scenario can be either a common execution scenario or a critical scenario. In the encoder case, the common execution scenario (e.g. encoding mode) is relevant to consider, because it implies high resource usage and correlates with REQ2 (see Section 5.1).

The application scenario model (defined in Section 4.3) requires data about service instances, bindings, events and task triggers for the selected scenario. This data is already known from the above-mentioned design steps, so that we only need to represent this data in the scenario model format. A major part of this scenario model is depicted by Figure 13.

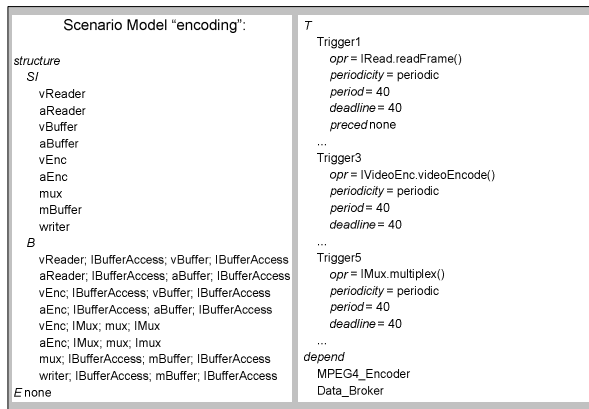


Figure 13. Application scenario model for encoding scenario.

After all related data is inserted in the application scenario model; we give a flow to the RIE for the models compilation and simulation of the compiled tasks.

## 5.5. Model compilation

The RIE compiler reconstructs the tasks in the application scenario (reconstruction process is explained in Section 4.4). Here we graphically represent the result of the task reconstruction (see Figure 14). The tasks are circular lines with arrows showing the control flow directions. For example, video encoding task is triggered by Trigger3 who calls operation VideoEncode(). This operation first calls getFrame() operation of vBuffer, then encodes the received frame and finally calls putVFrame() operation of Mux service instance. This task repeats each 40 ms.

The call sequence diagram for the task is depicted in Figure 15.




Figure 14. Reconstructed tasks after models compilation.

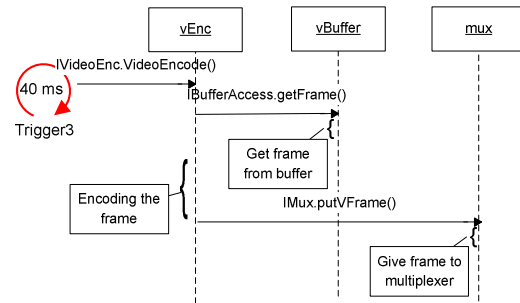


Figure 15. Video encoding task.

## 5.6. Simulation of tasks execution

An execution of the reconstructed tasks can be further simulated by the RIE scheduler. The current algorithm used in the RIE scheduler is rate monotonic with bounded priority inversion. The virtual scheduling of the encoder tasks results in the execution timeline depicted in Figure 16. The three bold vertical lines show: completion, deadlines and triggering moments of each task instance.




Figure 16. Part of the task execution timeline for encoder application.

## 5.7. Schedulability analysis

The schedulability analysis leads to the requirements validation. Our extra-functional REQ2 demands no missing deadlines of the audio and video

encoding tasks (see Section 5.1). According to the generated tasks execution timeline (Figure 16), video encoding (TaskID = 3) and audio encoding (TaskID = 4) tasks meet all deadlines for a simulation period of 10 seconds. Note that this is only true under the condition that the assigned CPU budget is 100 %.

This step ends with the conclusion that the designed application meets its real-time requirements and we can now proceed to the implementation phase.

## 6. Conclusions

We have extended the scenario-based approach for predicting resource usage of component-based systems in [2] with the specifications of task synchronization, component behaviour model and application scenario model. This allows simulation of the real-time task execution per application scenario and handling of synchronization constraints. Based on the simulation results, a developer can derive the behaviour and dynamic resource consumption of an application per scenario. Afterwards, a developer uses this data for prediction of the real-time properties of an application. The method was validated through the Robocop Integration Environment tool that automates complex operations and guides a developer through the composition process.

The proposed prediction approach has a number of benefits. Firstly, it is general and can be applied in different application domains and for various architectural styles. For example, it works for 'blackboard' and 'client-server' architectures. Secondly, the approach allows prediction of dynamically changing resource usage. Thirdly, the approach is more accurate by incorporating task synchronization constraints and distinguishing synchronous and asynchronous communication. Fourthly, the method is compositional, meaning that the resource-usage data of an application can be based on data from its constituent components. Finally, the use of scenarios decreases modeling complexity.

The proposal also has some assumptions and limitations that need further study. Firstly, it assumes that resource usage is constant per operation, whereas it actually may depend on parameter values passed to operations and/or application state. Secondly, the method is restricted to the Robocop component model, which has a notion of 'requires interfaces', whereas other architectures such as COM [4], do not have this notion. Finally, it provides no techniques for specifying the component resource model for different platforms. Extending the relatively simple case in this paper, we

are currently validating the approach on more complex MPEG-4 codec software.

## References

- [1] Ivica Crnkovic and Magnus Larsson. *Building Reliable Component-based Software Systems*, Artech House, 2002, ISBN 1-580-53327-2
- [2] Johan Muskens and Michel Chaudron. Prediction of Run-time Consumption in Multi-task Component-Based Systems. In *Proceedings of 7<sup>th</sup> ICSE Symposium on Component Based Software Engineering*. May, 2004.
- [3] Robocop public homepage. [<http://www.extra.research.philips.com/euprojects/robocop/>]
- [4] D. Box. *Essential COM*. Object Technology Series. Addison-Wesley, 1997.
- [5] Egor Bondarev *et al.* Predicting Real-Time Properties of Component Assemblies: a Scenario Simulation Approach. In *Proceedings of 30<sup>th</sup> EuroMicro Conference on Component Based Software Engineering*. September, 2004.
- [6] R. van Ommering *et al.*, The Koala component model for consumer electronics software. *IEEE Computer*, 33 (3): 78-85, Mar. 2002.
- [7] I. Crnkovic, *et al.*, Anatomy of a research project in predictable assembly. In *5<sup>th</sup> ICSE Workshop on Component Based Software Engineering*. ACM, May, 2002.
- [8] Kurt C. Wallnau. *Volume III: A Technology for Predictable Assembly from Certifiable Components*. April 2003, CMU/ESI-2003-TR-009
- [9] Scott A. Hissam, *et al.*, Packaging Predictable Assembly with Prediction-Enabled Component Technology. November 2001, CMU/ESI-2001-TR-024
- [10] Scott Hissam *et al.*, *Predictable Assembly of Substation Automation Systems: An Experiment Report*. September 2002, CMU/SEI 2002-TR-031
- [11] A. V. Fioukov *et al.*, Estimation of static memory consumption for systems built from source code components. In *Proc. 28th EUROMICRO conference, Component-Based Software Engineering Track*. IEEE Computer Society Press, Sept. 2002.
- [12] B. Selic, *et al.* *Real-Time Object-Oriented Modeling*, Wiley, 1995, ISBN 0471599174.
- [13] B.P. Douglass, *Doing Hard Time. Developing Real-time Systems with UML, Objects, Frameworks and Patterns*, Addison Wesley 1999, ISBN 0-201-49837-5.
- [14] Vittorio Cortellessa, Raffaella Mirandola. *PRIAM-UML: a performance validation incremental methodology on early UML diagrams*. Elsevier Science B.V., 02/2002.
- [15] M. de Jonge, J. Muskens and M. Chaudron. Scenario-based prediction of run-time resource consumption in component-based systems. In *Proceedings of 6<sup>th</sup> ICSE Workshop on CBSE*. ACM. June, 2003.