

**UNIVERSITY OF OSLO**  
**Department of Informatics**

**Formal Modeling  
and Analysis of  
Resource  
Management for  
Cloud Architectures**

Research Report 424

Elvira Albert,  
Frank S. de Boer,  
Reiner Hähnle,  
Einar Broch Johnsen,  
Rudolf Schlatte,  
S. Lizeth Tapia  
Tarifa, and  
Peter Y. H. Wong

ISBN 82-7368-388-5  
ISSN 0806-3036

January 2013



# Formal Modeling and Analysis of Resource Management for Cloud Architectures: An Industrial Case Study using Real-Time ABS\*

Elvira Albert<sup>1</sup>    Frank S. de Boer<sup>2</sup>    Reiner Hähnle<sup>3</sup>    Einar Broch Johnsen<sup>4</sup>  
Rudolf Schlatte<sup>4</sup>    S. Lizeth Tapia Tarifa<sup>4</sup>    Peter Y. H. Wong<sup>5</sup>

<sup>1</sup> DSIC, Complutense University of Madrid, Spain  
elvira@fdi.ucm.es

<sup>2</sup> CWI, Amsterdam, The Netherlands  
f.s.de.boer@cwi.nl

<sup>3</sup> Technical University of Darmstadt, Germany  
haehnle@cs.tu-darmstadt.de

<sup>4</sup>Department of Informatics, University of Oslo, Norway  
{einarj, rudi, sltarifa}@ifi.uio.no

<sup>5</sup> SDL Fredhopper, Amsterdam, The Netherlands  
peter.wong@fredhopper.com

## Abstract

We demonstrate by a case study how performance, resource consumption, and deployment on the cloud for an industrial distributed system can be formally modeled and analyzed using the abstract behavioral specification language Real-Time ABS. These non-functional aspects are integrated with an existing formal model of the functional system behavior, achieving a separation of concerns between the functional and non-functional aspects in the integrated model. The resource costs associated with execution in the system depend on the size of local data structures, which evolve over time; we derive corresponding worst-case cost estimations by static analysis techniques and integrate them into our resource-sensitive model. The model is further parameterized with respect to deployment scenarios which capture different application-level management policies for virtualized resources. The model is validated against the existing system's performance characteristics and used to simulate, analyze, and compare deployment scenarios on the cloud.

## 1 Introduction

Virtualization is a key technology enabler for cloud computing. Although the added value and compelling business drivers of cloud computing are undeniable [16], this new paradigm also poses considerable new challenges that have to be addressed to render its usage effective for industry. Virtualization gives access to elastic amounts of resources to services on the application-level; for example, the processing capacity allocated to a service may be changed according to demand.

Current software development methods, however, do not support the development of virtualized software in a satisfactory way; in particular, they do not help developers model and validate application-level resource management strategies for virtualized resources early in the development process. For example, the concrete and precise cost and resource requirements during deployment of a service in the cloud must be estimated manually in each case based on each

---

\*Partly funded by the EU project FP7-231620 HATS: Highly Adaptable and Trustworthy Software using Formal Models (<http://www.hats-project.eu>).

customer’s requirements and expectations. This seriously restricts the potential for fine-tuning a service to the available virtualized resources during the design phase, and forces the developer to introduce and validate resource management for the service after the development of the overall service logic and control flow. In the worst case, faulty provisioning can make it necessary to redesign an application, which is extremely expensive when discovered at deployment time. This article demonstrates how to overcome these limitations by integrating resource management into a formal, yet realistic, model that permits simulation and analysis during service *design time*.

Our long-term goal is the integration of virtualization into the development process of general purpose software services, by leveraging resources and resource management to the modeling of software. Our starting point in addressing this challenge is the recently developed abstract behavioral specification language ABS [27]. ABS is *object-oriented* to stay close to high-level programming languages and to be easily usable as well as accessible to software developers, it is *executable* to support full code generation and (timed) validation of models, and it has a *formal semantics* which enables the static analysis of worst-case resource consumption in the models.

Real-Time ABS extends ABS with time and with primitives for leveraging resources and their dynamic management to the abstraction level of software models [12, 30]. The extension achieves a separation of concerns between the *application model*, which requires resources, and the *deployment scenario*, which reflects the virtualized computing environment and provides elastic resources. For example, an application model may be analyzed with respect to deployments on virtual machines with varying features: the amount of allocated computing or memory resources, the choice of application-level scheduling policies for client requests, or the distribution over different virtual machines with fixed bandwidth constraints. The simulation and analysis tools developed for ABS may then be used to compare the performance of a service ranging over different deployment scenarios already at the modeling level.

## Summary of Contributions

The overall contribution of this article is a large industrial case study, which demonstrates that it is possible to model aspects of performance, resource consumption, and deployment on the cloud *at design time* based on Real-Time ABS. This article refines and substantially extends a paper published at ESOC 2012 [21], which focused on the modeling and validation of the case study for average costs by means of simulation; in the present article we take worst-case costs into account and apply *static* resource analysis and *hybrid* validation to the case study. The main contributions of the presented work are as follows. The last two points are original contributions of this article.

- *Modeling*. The non-functional aspects are integrated with a model of the functional system behavior, achieving a separation of concerns between the functional and non-functional aspects of the case study. The ABS model is parameterized over deployment scenarios which capture different application-level management policies for virtualized resources.
- *Validation*. The model is validated against the existing system’s performance characteristics and used to simulate and compare deployment scenarios on the cloud. A companion paper [31] details the modeling of the cloud provider and compares our approach to results obtained by simulation tools.
- *Analysis*. For the first time, resource analysis [3] has been applied to a case study of industrial size. In particular, we have been able to analyze statically a large fragment of the model whose Java implementation included an *undetected hot spot* in the program. We succeeded to infer the worst-case resource consumption of this part of the code and the formal analysis explains precisely why its resource consumption is high.

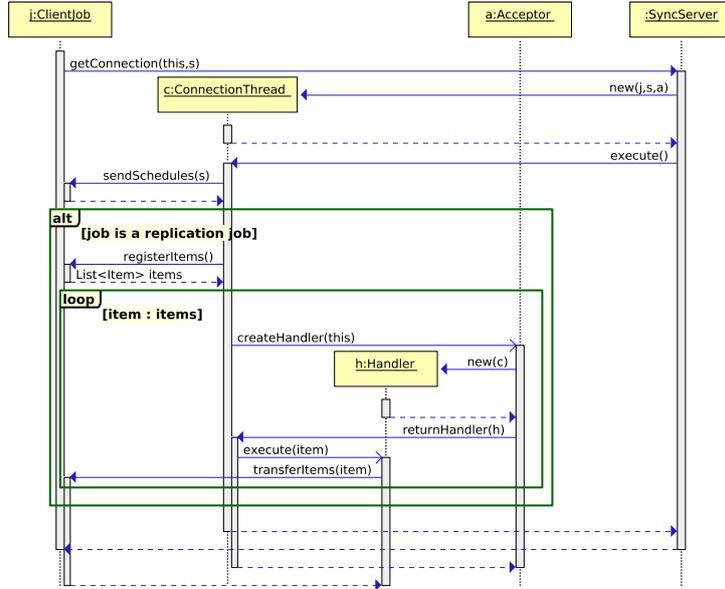


Figure 1: The interaction between the ClientJob, the Acceptor, the Handler, and the Connection-Thread.

- *Hybrid validation.* We have *combined* dynamic and static analysis to validate the model in a novel way. In particular, we have used the resource bound inferred by static resource analysis during the *timed* simulation of deployed, concurrent models. In previous work [5] the results of resource analysis had merely been used combined with dynamic analysis to validate the untimed behavior of sequential fragments of code for small examples.

## Organization of the Article

The article is organized as follows: In Section 2 we motivate and describe the case study. Section 3 presents the modeling layers of Real-Time ABS, starting with the functional and imperative layers before we explain how deployment scenarios are modeled in Real-Time ABS. Section 4 describes the modeling of the case study and proceeds with resource analysis and calibration of the model, before we combine the obtained worst-case bounds with simulation techniques in a hybrid validation approach to the case study. Section 5 discusses related work and Section 6 concludes the article.

## 2 The Case Study: Background & Motivation

The Fredhopper Access Server (FAS) is a distributed, concurrent object-oriented system that provides search and merchandising services to e-Commerce companies. Briefly, FAS provides to its clients structured search capabilities within the client’s data. Each FAS installation is deployed to a customer according to a deployment architecture; see [44] for a detailed presentation of the individual components of FAS and its deployment model.

For the purposes of this case study, FAS consists of a set of live environments and a single staging environment. A *live environment* processes queries from client web applications via web services. FAS aims at providing a constant query capacity to client-side web applications. The *staging environment* is responsible for receiving data updates in XML format, indexing the XML, and distributing the resulting indices across all live environments. The distribution of indexed data is done according to a *Replication Protocol*, which is implemented by the *Replication System*.

The Replication System consists of a `SyncServer` at the staging environment and one `SyncClient` for each live environment. The `SyncClient` initially performs a `Boot` job to identify `Replication` jobs that are relevant for its live environment, and creates listeners for these `Replication` jobs. `Replication` jobs may relate to the *search* index, *business rules* (e.g., presentation and promotions), and *data* (e.g., faceted navigation). The `SyncServer` determines the schedule and content of different `Replication` jobs, while a `SyncClient` receives data and configuration updates.

The `SyncServer` communicates to a `SyncClient` by creating `ConnectionThread` objects. These objects serve as the interface to the server side of the Replication Protocol. A `ConnectionThread` object delegates actual data replication to its `Handler` via the `Acceptor`. A `SyncClient`, on the other hand, schedules and creates `ClientJob` objects to handle communications to the client side of the Replication Protocol; thus, `ClientJob` objects on a client side live environment recursively realize the replication schedules computed by the `ConnectionThread` objects on the server side staging environment. Figure 1 shows a UML sequence diagram which illustrates the replication protocol between a `ClientJob`, a `ConnectionThread`, and an `Acceptor`. In this article, we detail a part of the Replication Protocol that is informally described in Figure 2.

1. `SyncServer` starts by listening for connections from `SyncClient` objects.
2. `SyncClient` creates a `ClientJob` object with a `Boot` job, which connects to the `SyncServer`. This corresponds to the message `getConnection(this,s)` in Figure 1, where `s` is a schedule.
3. `SyncServer` creates a `ConnectionThread` to communicate with this `ClientJob`.
4. The `ClientJob` asks the `ConnectionThread` for *all replication schedules*, denoted by the message `sendSchedule(s)` in Figure 1. Replication schedules dictate when and where the `SyncServer` monitors for changes in the staging environment. These changes are replicated to the live environments through their `SyncClient` objects. Each schedule specifies a *replication type*; i.e., the number of locations and type of data to be replicated. The schedule also specifies the amount of time until the replication must commence and the deadline of each replication.
5. The `ClientJob` receives the replication schedules and creates new `ClientJob` objects representing the different schedules. If the old `ClientJob` object represented a `Boot` job, it releases the `ConnectionThread` and terminates.
6. When a `Replication` job is triggered, its associated `ClientJob` object immediately connects to the `SyncServer`.
7. `SyncServer` creates a `ConnectionThread` to communicate with each `ClientJob`.
8. A `ClientJob` asks the `ConnectionThread` for its replication schedule and recursively creates a new `ClientJob` object to deal with the next schedule. The `ClientJob` then receives a sequence of file updates according to its replication type, after which it terminates.
9. The `ConnectionThread` first sends a replication schedule to the `ClientJob` according to the `ClientJob`'s replication type, the `ConnectionThread` then sends the message `registerItems` to the `ClientJob` to acquire which sets of files to update. For each set of files to be updated, the `ConnectionThread` creates a `Handler` via the `Acceptor` to send the message(s) `transferItem(item)` (See Figure 1). After sending the file updates, the `ConnectionThread` terminates.

Figure 2: Informal description of the interactions in the Replication Protocol.

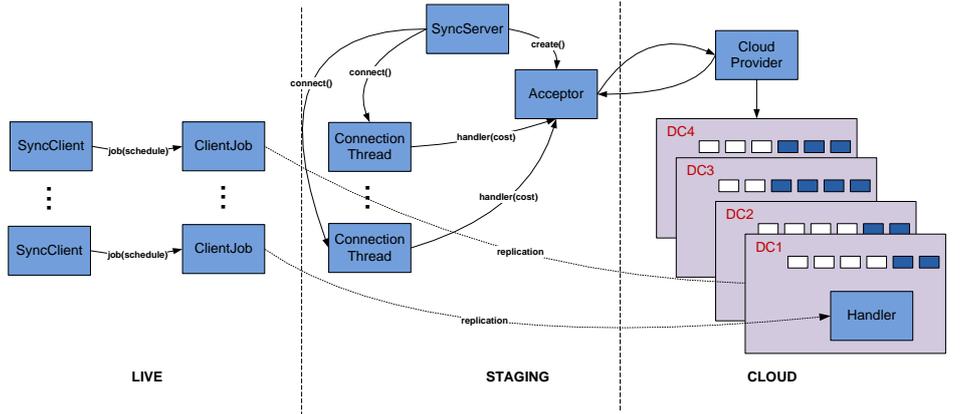


Figure 3: An on-demand deployment architecture for the Replication System using Cloud resources.

**Relevance to Cloud Computing.** FAS provides structured search and navigation capabilities within client data. For the last decade, FAS installations were deployed as server-based products on client premises with fixed hardware resources. However, on-premise deployment does not scale with growing demand on throughput and update frequency. This is visible in particular when customers experience drastic increases on throughput and data updates at certain time periods. For example, retail customers typically expect large throughput during seasonal sales. In these periods, considerably larger amounts of purchases are made and stock units need to be updated very frequently, in order to avoid that customers receive incorrect information. Higher throughput requires a larger number of live environments and, in turn, a larger number of `SyncClient` objects. On-premise deployment cannot cope with varying, on-demand requirements without large up-front investments in hardware which remains unused most of the time.

To cater for these requirements, FAS is today deployed as a service (SaaS) over virtualized resources that provide the necessary elasticity. To this end, virtualization makes elastic amounts of resources available to application-level services; for example, the processing capacity allocated to a service can be changed on demand. Figure 3 shows how an on-demand deployment architecture for the Replication System on virtual environments is implemented using cloud resources. Virtualized resources allow the `SyncServer` (via the `Acceptor`) to elastically allocate resources to each replication job based on the *cost* and the *deadline* of the replication to be conducted by the corresponding `ClientJob` object.

### 3 Modeling Virtualized Resources and Deployment in Real-Time ABS

ABS is an abstract, executable, object-oriented modeling language with a formal semantics [27], targeting distributed systems. ABS is based on concurrent object groups (COGs), akin to concurrent objects [19, 28], Actors [1], and Erlang processes [7]. COGs in ABS support interleaved concurrency based on guarded commands. This allows active and reactive behavior to be easily combined, by means of a cooperative scheduling of processes which stem from method calls. Real-Time ABS extends ABS models with *implicit* time [12]; the execution time is not specified directly in terms of durations (as in, e.g., UPPAAL [34]), but rather *observed* by measurements of the executing model. With implicit time, no assumptions about execution times are hard-coded into the models. Instead, the execution time of a method call depends on how quickly the call is effectuated by the server object. In fact, the execution time of a statement varies

with the *capacity* of the chosen deployment architecture and on *synchronization* with (slower) objects; *similar calls to the same method do not always take the same amount of time*.

### 3.1 Behavioral Modeling in ABS

ABS combines functional and imperative programming styles with a Java-like syntax [27]. COGs execute in parallel and communicate through asynchronous method calls. However, the data manipulation inside methods is modeled using a simple functional language based on user-defined algebraic data types and functions. Thus, the modeler may abstract from the details of low-level imperative implementations of data structures while maintaining an overall object-oriented design close to the target system.

#### 3.1.1 The Functional Layer

The functional layer of ABS consists of algebraic data types such as the empty type `Unit`, booleans `Bool`, integers `Int`; parametric data types such as sets `Set<A>` and maps `Map<A>` (for a type parameter `A`); and functions over values of these data types, with support for pattern matching. To illustrate the definition of data types and functions, *polymorphic sets* can be defined using a type variable `A` and two constructors `EmptySet` and `Insert`. We define a function `contains` which checks whether an element `el` is in a set `set` recursively by pattern matching over `set`:

```
data Set<A> = EmptySet | Insert(A, Set<A>);
def Bool contains<A>(Set<A> set, A el) =
  case set {
    EmptySet => False ;
    Insert(el, _) => True;
    Insert(_, xs) => contains(xs, el);
  };
```

In Real-Time ABS, measurements are additionally obtained by comparing values from a global clock, which can be read by an expression `now()` of type `Time`. Using the built-in function `timeDifference` to compare to values of type `Time`, we can for example define a function to give us the elapsed time since a given starting time `start`, as follows:

```
def Rat elapsed(Time start) = timeDifference(now(),start);
```

#### 3.1.2 The Imperative Layer

The imperative layer of ABS addresses concurrency, communication, and synchronization at the level of objects, and defines interfaces, classes, and methods. ABS objects are *active* in the sense that their `run` method, if defined, gets called upon creation. Communication and synchronization are decoupled in ABS. Communication is based on asynchronous method calls, denoted by assignments `f=o!m(e)` where `f` is a future variable, `o` an object expression, and `e` are (data value or object) expressions. After calling `f=o!m(e)`, the caller may proceed with its execution *without blocking* on the method reply. Two operations on future variables control synchronization in ABS. First, the statement `await f?` *suspends the active process* unless a return value from the call associated with `f` has arrived, allowing other processes in the same COG to execute. Second, the return value is retrieved by the expression `f.get`, which *blocks all execution in the object* until the return value is available. Inside a COG, ABS also supports standard synchronous method calls `o.m(e)`.

A COG can have at most one active process, executing in one of the objects of the COG. This active process can be unconditionally suspended by the statement `suspend`, adding this process to the queue of the COG, from which an enabled process is then selected for execution. The

guards  $g$  in **await**  $g$  control suspension of the active process and consist of Boolean conditions conjoined with return tests  $f?$  on future variables  $f$ . Just like functional expressions, guards  $g$  are side-effect free. The remaining statements of Real-Time ABS are standard; e.g., sequential composition  $s_1; s_2$ , assignment  $x=rhs$ , and **skip**, **if**, **while**, and **return** constructs. Right hand side expressions  $rhs$  include the creation of an object group **new cog**  $C(e)$ , object creation in the group of the creator **new**  $C(e)$ , method calls, and future dereferencing  $f.get$ , in addition to the functional expressions  $e$ .

To illustrate how the two layers of Real-Time ABS can be combined, let us assume that an interface  $l$  implements a method  $m$ . We define a method `timer` which takes as parameter an object of interface  $l$  and returns the elapsed time for the method call to  $o$ , as follows:

```
Rat timer(l o){ Time start, Bool b;
  start = now(); b = o.m(); return elapsed(start);
}
```

## 3.2 Deployment Modeling in ABS

The response time to a request in a distributed system depends not only on the size of the job requested, but also on the amount of available resources and on the usage policy for these resources, which are scattered around the deployment architecture of the distributed system. Deployment architectures express how distributed systems are mapped on physical and/or virtual media with many locations; the planning and validation of a deployment architecture to optimize performance, implies determining the amount of necessary resources at the different locations as well as an optimal usage of these resources, such that the system fulfills its performance requirements.

Real-Time ABS lifts deployment architectures to the abstraction level of the modeling language, where the physical or virtual media are represented as *deployment components* [29]. In a Real-Time ABS model, different deployment components may have different bounds on the locally available resources.

Real-Time ABS introduces a separation of concerns between the resource cost of performing a computation and the resource capacity of a given deployment component. This separation of concerns between resource *cost* and resource *capacity* aids to model and validate different deployment scenarios at an early stage during the software development process. The focus in this article is on CPU resources in virtualized media; we use resource cost annotations to express resource consumption during computation. Deployment components are discussed in Section 3.2.1 and resource consumption in Section 3.2.2.

### 3.2.1 Deployment Components

A *deployment component* in Real-Time ABS captures the execution capacity of a location in the deployment architecture, on which a number of COGs are deployed. The execution capacity is specified as an amount of resources which are available per accounting period; for simplicity, this accounting period is fixed in the semantics of Real-Time ABS and corresponds to the time between integer values in the dense time domain of the language. The main block of a model executes in a root COG located on a default deployment component `environment`, with unrestricted processing capacity. To capture different deployment architectures, a model may be extended with other deployment components with different resource capacities. When COGs are created, they are by default allocated to the same deployment component as their creator, but they may also be allocated to a different deployment component. In contrast, an object which belongs to a COG will always be located on the same deployment components as its

```

data Resource = InfCPU | CPU(Int capacity) ;

interface DC {
  Int total();
  Int load(Int n);
  Int transfer(DC target, Int amount);
}

```

Figure 4: The specification of resources and the interface of deployment components.

group. Thus, in a model without explicit deployment components all objects execute in the default **environment**, which places no restrictions on the processing capacity of the model.

Deployment components are first-class citizens of Real-Time ABS. They may be passed around as arguments to method calls, they support a number of methods. Deployment components may be created dynamically, depending on control flow, or statically in the main block of the model. Syntactically, deployment components in Real-Time ABS are manipulated in a way similar to objects. Variables which refer to deployment components are typed by an interface DC and new deployment components are dynamically created as instances of class `DeploymentComponent`, which implements DC. The definition of the interface DC is given in Figure 4.

The DC interface provides the following methods for resource management: **total()** returns the number of resources currently allocated to the deployment component, **load(n)** returns the deployment component’s average load in percentage during the last *n* accounting periods (i.e., the used compared to the total number of resources for each accounting period in the window), and **transfer(target,r)** reallocates *r* resources from the current deployment component to a **target** deployment component. If the former has less than *r* resources available, the available amount is transferred (and provided as the return value of the call).

The type `Resource`, given in Figure 4, reflects processing capacity. It has the two constructors `InfCPU` for infinite (i.e., unbounded) resource capacity and `CPU(r)` for restricted resource capacity, where *r* represents the amount of processing resources available in an accounting period. Deployment components are created by the expression `new cog DeploymentComponent(d,c)`. Here, the parameter *c* of type `Resource` specifies the initial CPU capacity of the deployment component. The parameter *d* of type `String` is a descriptor mainly used for monitoring purposes; i.e., it provides a user-defined name for the deployment component which facilitates querying the run-time state but that has no semantic effect. Objects are deployed on deployment components when the objects are created. By default an object is deployed on the same deployment component as its creator. However, a different deployment component may be selected by means of an optional *deployment annotation* `[DC: e]` to the object creation statement, where *e* is an expression of type DC. Note that deployment annotations can only occur associated with the creation of COGs.

An object may relocate the COG it belongs to a different deployment component *e* by executing the statement `movecogto(e)`. This is well-defined, because at most one object can be active in a COG at any given time. Since all objects are deployed on some deployment component, we let the expression `thisDC()` refer to the deployment component where the object is currently deployed, similar to the standard self-reference `this` in object-oriented languages.

### 3.2.2 Resource Consumption

The available resource capacity of a deployment component determines the amount of computation which may occur in the objects deployed on that deployment component. Objects allocated to the deployment component compete for the shared resources in order to execute, and they

may execute until the deployment component runs out of resources or they are otherwise blocked. For the case of CPU resources, the resources of the deployment component define its capacity inside an accounting period, after which the resources are renewed.

The resource consumption of executing statements in the Real-Time ABS model is determined by a default cost value which can be set as a compiler option (e.g., `-defaultcost=10`). However, the default cost does not discriminate between the statements, so a more refined cost model will often be desirable. For example, in a realistic model the assignment `x=e` should have a significantly higher cost for a complex expression `e` than for a constant. For this reason, more fine-grained costs can be inserted into Real-Time ABS models by means of *cost annotations* [`Cost: e`]. Note that cost annotations can be associated with any statement, and that a statement may have several annotations.

It is the responsibility of the modeler to specify appropriate resource costs. A behavioral model with default costs may be gradually transformed to provide more realistic resource-sensitive behavior by inserting such cost annotations. The manual estimation of resource cost is time consuming and error-prone. Therefore, it is desirable to have tool support for this activity. COSTABS [3] is an automated static analysis tool that is able to compute a worst-case approximation of the resource consumption of the non-virtualized programs, based on static analysis techniques. In the sequel, we apply COSTABS to obtain such worst-case cost expressions for the parts of the model that is deployed on virtual machines, and use these expressions in our cost annotations (see Section 4.1).

However, the modeler may also want to capture *normative* constraints on resource consumption, such as resource limitations, at an abstract level; these can be made explicit in the model during the very early stages of the system design. To this end, cost annotations may be used by the modeler to abstractly represent the cost of some computation which is not fully specified.

## 4 Case Study: The ABS Model

The Replication System, introduced in Section 2, is part of the Fredhopper Access Server (FAS). The current Java implementation of FAS has over 150,000 lines of code, of which 6,500 are part of the Replication System. The functional aspects of the Replication System have previously been modeled in detail in ABS [44]. This section describes how the model was extended to capture non-functional and resource aspects of the Replication System. The extended model consists of 40 classes, 17 data types, and 80 user-defined functions (in total 5,000 lines of ABS code, 25% of which capture scheduling information as well as file systems and data bases from third party libraries not included in the Java implementation).

Figure 5 shows the main interfaces. The interface `ClientJob` declares two methods. The task of method `executeJob()` is to execute replication schedules, and method `size()` returns the size of the client job's underlying file structure. Scheduling of a given replication schedule `s` is done with method `scheduleJob(Schedule s)` declared in interface `SyncClient`.

Objects of type `ConnectionThread` do not receive method calls, so the interface `ConnectionThread` models its objects as active objects (without methods). The interface `Handler` declares the method `files()` that is responsible for file updating. The `Handler` objects are created by objects of class `ConnectionThread` by invoking the `createHandler(job,cost)` method of the `Acceptor`. On method invocation, the acceptor creates handler objects on virtual machines which are acquired from a `CloudProvider` by the method `createMachine(capacity)`. The methods `acquireMachine(vm)`, `releaseMachine(vm)` are used to start and stop virtual machines (modeled by deployment components) to let replication schedules be conducted by `ConnectThread` objects. For presentation purposes, we focus here on the interface implementations given in the classes `CloudProvider`, `Acceptor`, and `ConnectionThread`.

```

interface ClientJob {
    Unit executeJob();
    Int size();
}
interface SyncClient {
    Unit scheduleJob(Schedule s);
}
interface ConnectionThread { }
interface Handler {
    Unit transfer(Set<File> files);
}
interface Acceptor {
    Handler createHandler(ClientJob job, Int cost);
    Unit finish(Handler h);
    Unit end();
}
interface CloudProvider {
    DC createMachine(Int capacity);
    Unit acquireMachine(DC vm);
    Unit releaseMachine(DC vm);
    Int getAccumulatedCost();
}

```

Figure 5: Interfaces of the Replication System.

The `CloudProvider` interface (shown in Figure 5) is implemented by a class of the same name. Virtual machines are modeled by deployment components in ABS, on which the client application can deploy objects. In addition, the cloud provider keeps track of the *accumulated cost* incurred by the client application. This accumulated cost can be retrieved with the help of the method `getAccumulatedCost()` at any time during execution. The accumulated cost is calculated in terms of the sum of the processing capacities of the *active* virtual machines over time; i.e., a call to the method `acquireMachine(vm)` starts the accounting for machine `vm` and a call to the method `releaseMachine(vm)` stops the accounting for `vm`. Inside the cloud provider, an active `run()` method ensures that the accounting is done for every accounting period. To focus on application-level management of virtualized resources, as implemented by the balancer, and not on a specific strategy for cloud provisioning, we do not detail the cloud provider further (one possible implementation of the `CloudProvider` interface in Real-Time ABS is given in [31]).

We model and compare three potential load balancing strategies offered by different implementations of the `Acceptor` interface, for the application-level management of virtualized resources. When an `Acceptor` receives requests for file updates from `ConnectionThread` objects, it deploys `Handler` objects on cloud instances to conduct file updating with the `ClientJob` objects. The implementations of `Acceptor` reflect different strategies for interacting with the cloud provider to achieve application-level resource management:

**Constant balancing** simply deploys all the objects of type `ConnectionThread` on a single virtual machine sufficient for the expected load, and keeps this machine running;

**As-needed balancing** calculates the needed CPU capacity of the virtual machine for a specific replication schedule with a given deadline, and deploys `Handler` objects to a machine supplying the required resources disregarding the cost; and

**Budget-aware balancing** calculates the CPU capacity of the cloud instance for a given budget. Unused funds can be “saved up” to cope with load spikes, but the cost of running the system is still bounded by the overall budget.

**The Cloud User Account.** Each acceptor encapsulates an `Account` object that realizes the book-keeping for a cloud user account (see Figure 6). The implementation `AccountImpl` maintains

```

interface Account {
    DC getInstance(Int size);
    Unit dropInstance(DC d);
    Int getCostPerTimeUnit();
    Int getLastInstanceStartUpTime();
}

class AccountImpl implements Account {
    Map<Int, Set<DC>> instances = EmptyMap;
    Int costPerTimeUnit = 0;
    Int instanceStartUpTime = 0;

    DC getInstance(Int size) {
        DC d = null;
        Time t = now();
        costPerTimeUnit = costPerTimeUnit + size;
        if (hasSetFor(instances, size)) {
            d = takeOne(lookup(instances, size));
            instances = removeFrom(instances, size, d);
            Fut<Unit> fa = provider!acquireMachine(d); await fa?;
        } else {
            Fut<DC> fdc = provider!createMachine(size);
            await fdc?; d = fdc.get;
        }
        instanceStartTime = timeDifference(t, now());
        return d;
    }

    Unit dropInstance(DC d) {
        Fut<Unit> fr = provider!releaseMachine(d); await fr?;
        Fut<Int> fs = d!total("CPU"); await fs?; Int size = fs.get;
        costPerTimeUnit = costPerTimeUnit - size;
        instances = addToSet(instances, size, d);
    }
}

```

Figure 6: The Cloud User Account.

a data structure in the field `instances`, which sorts the available virtual machines by CPU processing capacity. Furthermore, the class keeps track of the current cost per time unit `costPerTimeUnit` for the cloud user account, and the observed start up time `instanceStartUpTime` of the most recent virtual machine to start up. The method `getInstance(size)` either requests a new virtual machine from the cloud provider or brings online an existing offline machine of the appropriate size. The method `dropInstance(d)` takes the machine `d` offline when it is no longer active.

The load balancing strategies are defined in classes implementing the `Acceptor` interface (two of these implementations are detailed in Figure 7), as follows:

*Constant balancing* over-provisions by processing all replication schedules on a single virtual machine with sufficient capacity. Its implementation is shown in Figure 7 (class `ConstantAcceptor`). The acceptor initially requests a single machine through its cloud user account and deploys all `ConnectionThread` objects to this machine after initialization to process the replication schedules.

*As-needed balancing* receives a `request` for a connection from a `ClientJob` object, calculates the `resources` needed by the virtual machine to fulfill the replication schedule, and requests a machine of appropriate size through the user account. Implementation details are omitted for brevity.

*Budget-aware balancing* is a strategy where the acceptor has a given *budget per accounting period* and may “save resources” for later (class `BudgetAcceptor` in Figure 7). The class parameter `budgetPerTimeUnit` provides this budget and the field `availableBudget` keeps track of the accumulated (saved) unused resources, according to the budget. When the acceptor receives a request from a `ClientJob` object, it calculates the resources needed to fulfill the schedule in `wanted` and the resources it has available on the budget in `maxresources`. If resources are available on the budget, the acceptor calls `getInstance(size)` on the user account to get the best machine according

```

class ConstantAcceptor(Int instanceSize, Account acc)
implements Acceptor {
  DC dc = null;
  Unit run() { dc = acc.getInstance(instanceSize); }
  Handler createHandler(ClientJob job, Int cost) {
    await dc != null;
    [DC: dc] Handler h = new cog HandlerImpl(job, cost);
    return h;
  }
  Unit finish(Handler h) { }
  Unit end() { acc.dropInstance(dc); }
}

class BudgetAcceptor(Account acc,
  Int budgetPerTimeUnit) implements Acceptor {

  Int availableBudget = 1; List<Int> budgetHistory = Nil;

  Unit run() {
    while (True) {
      Int cu = acc.getCostPerTimeUnit();
      availableBudget =
        availableBudget + budgetPerTimeUnit - cu;
      budgetHistory = Cons(availableBudget, budgetHistory);
      await duration(1, 1);}
    }
  Handler createHandler(ClientJob job, Int cost) {
    Int dur = durationValue(deadline());
    Int startUp = acc.getLastInstanceStartUpTime();
    Int wanted = (cost / dur) + 1 + startUp;
    Int maxresource =
      (budgetPerTimeUnit - costPerTimeUnit)
      + (max(availableBudget, 0) / dur);
    Handler handler = null;
    if (maxresource > 0) {
      DC dc = acc.getInstance(min(wanted, maxresource));
      [DC: dc] Handler h = new cog HandlerImpl(job, cost);}
    return h;
  }
  Unit finish(Handler handler) {
    Fut<DC> fdc = handler!.release();
    await fdc?; DC instance = fdc.get;
    acc.dropInstance(instance);
  }
  Unit end() { }
}

```

Figure 7: The classes `ConstantAcceptor` and `BudgetAcceptor`, which implement the *constant* and *budget-aware* strategies.

to the budget. The `run()` method monitors the resource usage and updates the available budget for every accounting period. It also maintains a log `budgetHistory` of the available resources over time.

#### 4.1 Resource Analysis and Calibration

Our initial approach to providing cost annotations for the ABS model of the Replication System, was to average measurements of the execution time of client jobs for the different replication schedules on the Java implementation (reported in [21]). In particular, we were interested in jobs specified by the following two types of replication schedules:

**Search:** A replication job originating from the search schedule replicates changes from the search index, i.e., the underlying data structure providing search capability on a customer’s product items; and

**Data:** A replication job initiated by the data schedule replicates changes concerning the item and navigation indices, i.e., the core index structures and data model for providing faceted

navigation on a customer’s product items.

During the calibration of the simulation constants, we identified a *hot spot* in the Java implementation of the Replication System and located it to be in the method `transferItem(fileset)` of `ConnectionThread`. In ABS this corresponds to the method `transferItem` of the `Handler` interface, implemented by the class `HandlerImpl` shown in Figure 8.

```
class HandlerImpl(ClientJob job, Int cost) implements Handler {
  Unit transfer(Set<File> fileset) {
    [Cost: cost] this.transferItem(fileset);
  }
  Unit transferItem(Set<File> fileset) {...}
}
```

Figure 8: Class `HandlerImpl`.

In our measurements, this method accounts for 99% of the execution time. The hot spot justifies adding a single cost annotation to the corresponding method in the ABS model. As we aim to provide accurate simulation results at the modelling level, we make an in-depth analysis of the hot spot in our model and use the COSTABS tool [3] to *statically analyse* the resource consumption of `transferItem(fileset)`.

COSTABS is a cost analysis tool that, given an input method, is able to automatically infer a sound *upper bound* on its resource consumption. In other words, the analysis guarantees that the execution of the method will never exceed the inferred amount of resources for any input data.

In Sections 4.1.1–4.1.4 we describe in detail how we have used COSTABS to analyse the resource consumption of the `transferItem(fileset)` method and explain the different parameters that occur in the analysis. Using the derived cost expression, in Section 4.1.5 we describe how we synthesize other simulation parameters.

#### 4.1.1 Cost Metrics

The first option to be selected in COSTABS is the *cost metric* (a.k.a. cost model) which specifies the type of resource we want to measure. COSTABS offers a wide range of cost models [2], including traditional cost models for measuring the number of execution steps and memory allocation, and also cost models specific to concurrent and distributed applications like the task-level of the program which estimates the peak of tasks that can be simultaneously spawned in the execution.

In our case study, we are interested in justifying the worst-case execution time of `transferItem(fileset)`. For this purpose, we have selected the cost model that counts the number of execution steps. This includes both, steps performed in the functional and steps in the imperative part of the ABS model. On the one hand, execution time is often directly related to the number of execution steps performed, although even for static deployment some other factors also clearly influence the execution time (see work on WCET [32]). On the other hand, we are also interested in understanding the computational complexity of the method `transferItem`. Thus, we need a cost model which assigns cost to all instructions of the program and does not ignore certain parts. This would not be achieved for instance with a cost model that infers memory consumption, since a loop that does not allocate memory has an associated resource consumption zero, while its computational complexity is not zero.

```

def Int sizeofFiles(Set<File> t) =
  case t {
    EmptySet => 1;
    Insert(f, fs) => 1 + sizeofFile(f) + sizeofFiles(fs);
  };

def Int sizeofFile(File t) =
  case t {
    Pair(fld, fContent) =>
      1 + sizeofFileld(fld) + sizeofFileContent(fContent);
  };

def Int sizeofFileld(Fileld t) = strlen(t);

```

Figure 9: Selected size abstractions for Set of Files.

### 4.1.2 Size Abstraction

The goal of the cost analysis is to infer closed-form upper bounds provided as functions on the data input *sizes*. When a program manipulates terms, its cost usually depends on the *size* of the terms. For instance, if a loop traverses a list, the cost of the loop often depends on its length. COSTABS relies on the notion of *norms* [11] to define the *size of a term*. Norms are functions that map terms to their sizes. Any norm can be used in the analysis, depending on the nature of the data structures used in the program. They can also be synthesized automatically from the program’s type definitions. In what follows, we use the *term-size* norm, which counts the number of type constructors in a given term, defined as:  $size(Co(t_1, \dots, t_n)) = 1 + \sum_{i=1}^n size(t_i)$  and  $size(x) = x$ . Note that the size of a program variable  $x$  is defined as  $x$ . In this way, we account for the size of the term to which  $x$  is bound at runtime.

Our target method `transferItem(fileset)` has a parameter of type `Set<File>`. `Set` is a predefined ABS type that can be `EmptySet` or `Insert(f,fs)`, where `f` is of type `File` and `fs` of type `Set<File>`. The type `File` is defined as `type File = Pair<Fileld,FileContent>` in ABS (we omit the definition of `FileContent` since it is not relevant). The functions provided in Figure 9 define (part of) the size abstraction used by the analyzer for `Set<File>`. These functions will be used later by the simulator to evaluate the cost functions on specific input data.

### 4.1.3 Class Invariants

When the cost depends on the size of data stored in fields, inferring cost often requires *class invariants* that provide guarantees on such fields. In general, such class invariants are a way to incorporate guarantees on the global states when the considered process is resumed. In our case study, the cost of executing the method `transferItem(fileset)` depends on a field `rdir` that has type `Directory` and is declared in class `ClientDataBaselmp`. It represents the directory which keeps all files whose content is to be transferred. Some operations carried out to transfer items traverse the directory `rdir`. To be able to infer an upper bound for the method, we need to specify that the field `rdir` is finite, i.e., bounded from above. This can be specified by means of the following class invariant `[rdir <= max(rdir)]` which states that field `rdir` has a maximum value that is denoted by `max(rdir)` in what follows. This invariant is the only manual input required to infer an upper bound for `transferItem(fileset)`: the analysis is fully automatic otherwise.

### 4.1.4 Upper Bounds

After having selected the cost model and the size abstraction, and having provided the class invariant, we automatically infer the following (asymptotic) upper bound using COSTABS:

$$\text{nat}(\text{fileset})^2 * \text{nat}(\text{max}(\text{rdir})) + \text{nat}(\text{fileset})^3$$

```

def Int cost(Int sizeOfFileSet, Int sizeOfFileStructure) = ...;

class ConnectionThreadImpl(ClientJob job, Acceptor accpt,
  SyncServer server) implements ConnectionThread {
  Maybe<Command> cmd = Nothing;
  Set<Schedule> schedules = EmptySet;
  Unit run() {
    await cmd != Nothing;
    schedules = this.sendSchedule();
    if (cmd != Just(ListSchedule)) {
      ...
      while (hasNext(filesets)) {
        Pair<Set<Set<File>>, Set<File>> nfs = next(filesets);
        filesets = fst(nfs); Set<File> fileset = snd(nfs);
        Fut<Int> sf = job!size(); Int size = sf.get;
        accpt.createHandler(job, cost(sizeOfFileSet, size));
      }
      ...
    }
  }
}

```

Figure 10: Class ConnectionThreadImpl.

This upper bound is a polynomial of degree 3 on the argument `fileset` and the class field `rdir`. Function `nat` is defined as: “`def Int nat(Int a) = if a > 0 then a else 0;`” and used to avoid negative values of the upper bound expression.

To obtain this upper bound, COSTABS has analyzed 37 methods and functions that were transitively invoked from `transferItem(fileset)`. The analysis revealed that the execution time of the method is asymptotically larger than any other bound obtained for any of the remaining functions and methods. This explains the hot spot and gives clear directives for potential optimizations.

#### 4.1.5 Model Calibration

The implementation of the `ConnectionThread` interface is sketched Figure 10 along with the type signature of the `cost` function that is generated by COSTABS and that returns the asymptotic upper bound computed in the previous section of the number of execution steps inside `transferItem(fileset)`.

Schedule	Interval	Deadline
Search	11	3
Data	11	11

Table 1: Derived interval and deadline parameters for simulation.

Using the inferred cost expression, we extend the functional ABS model of the Replication System with annotations for resource and timing information. To determine suitable deadlines for the individual schedules and the intervals between each `ClientJob`, we calibrate the model by considering a reference scenario for which we aim to meet all deadlines. Here, we select as our reference scenario the case of one single typical live environment, and set the capacity of the cloud instance such that it can successfully handle its client. We iteratively simulated a Replication System consisting of one `SyncClient` and a fixed number of replication jobs on a single cloud instance that offers 161600 CPU resources per time unit (corresponding to an accounting period in the model) and the interval to 11 time units between replication jobs. With these parameters, we determine the lowest deadlines of each types of schedules that can be met by all replication jobs, that is, 100% QoS (quality of service). Table 1 also shows the results of this initial simulation.

Schedule	Execution Time	Simulated Cost
Search	34.0s	90682
Data	274.9s	885315

Table 2: Comparison between execution time of the Java implementation of the Replication System against simulated cost.

## 4.2 Simulation Results

Table 2 shows a comparison between the average execution time of replication jobs per type of schedule measured on the Java implementation of the Replication System and the average cost of simulated jobs in the ABS model using the inferred upper bound cost expression from COSTABS. We observe that jobs specified by the **Search** schedule are ca. 8.1 times cheaper than jobs specified by the **Data** schedule in the Java implementation while they are ca. 9.76 times cheaper in the ABS model. This is due to the fact that the inferred cost expression defines the *worst* cost instead of the average cost.

The following figures show the simulation results that we obtained for the different balancing strategies, varying over the number of live environments. Each simulation runs for 100 time units in different scenarios varying from 1 to 20 live environments (in reality, up to 20 environments are typically required to handle large query throughputs over a large number of product items). For each number of environments and for each balancing strategy, Figure 11 shows the quality of service (QoS) as a percentage of replication jobs whose deadlines have been met and Figure 12 shows the total accumulated cost. The results exhibit some of the brittleness of an overloaded system.

As explained in Section 4.1.5, we calibrated the ABS model by establishing a deployment component size for the constant and budget balancing strategies such that the Replication System was able to handle a scenario with one live environment without missing deadlines (100% QoS). We then chose a size and a budget for the balancing strategies that would let us observe behavior both under normal load and overloaded scenarios. The system behavior differed remarkably for the three balancing strategies under higher load. As expected, the *as-needed* balancing strategy exhibits 100% QoS, albeit with simulated cost rising linearly with the number of clients and ending orders of magnitude above the other two strategies.

On the other hand, to model the budget-aware balancing strategy, the **BudgetAcceptor** needs to be able to reject tasks (to stay within the budget). When the system is overloaded, this can be observed through a degraded QoS and, somewhat more surprisingly at first sight, by a decreasing overall cost. We observed that under high load, the **BudgetAcceptor** only managed to fit in jobs from the cheaper **Search** schedule, while rejecting most jobs of the more expensive **Data** schedule (see Table 2). It must be said that this scenario is not conducive for showing the advantages of the budget-aware strategy—in a scenario of permanent high load, this strategy cannot ever accumulate the surplus needed to handle momentary load spikes. On the other hand, below the maximum load the budget-aware strategy performs as well as but much cheaper than the over-provisioning (constant) strategy, while being able to deal with transient higher loads.

To compare the simulation results against the running system, we executed the simulation task loads on the Java implementation of the Replication System and measured the execution time. Figure 13 shows that both simulated cost and measured execution time grow linearly, but the real system exhibits a constant factor resulting from system start-up time, which can be seen most clearly in the case of few live environments. Upon reflection, we decided not to model start-up time, since this one-time cost is amortized among all requests made during the server lifetime and hence is a negligible factor in a long-running system such as the FAS Replication

System.

## 5 Related Work

To reduce complexity, general-purpose modeling languages strive for *abstraction* [33]: descriptions primarily focus on the functional behavior and logical composition of software, largely overlooking how the software’s deployment influences its behavior. However, by using virtualization technology an application can *modify resource parameters of its deployment environment during execution*, e.g., it may dynamically create virtual processors. For cyber-physical and embedded systems it is today accepted that modeling and programming languages need a timed semantics [35]. The Java Real-Time Specification (RTSJ) [15] extends Java with high-resolution time, including absolute time and relative time, and new thread concepts to solve time-critical problems: threads in RTSJ offer more precise scheduling than standard threads, with 28 strictly enforced priority levels. The modeling and analysis of single resources is discussed in, e.g., [3, 20, 42]. *Resource-aware programming* allows users to monitor the resources consumed by their programs, to manage such resources in the programs, and to transfer (i.e., add or remove) resources dynamically between distributed computations [36].

Resource constraints in the embedded systems domain led to a large body of work on performance analysis using formal models based on, e.g., process algebra [10], Petri Nets [40], and priced [14], timed [6], and probabilistic [9] automata and games (an overview of automata-based approaches is [42]). Related approaches are also applied to web services and business processes with resource constraints [24, 37]. These approaches typically abstract away from the data flow and *declare* the cost of transitions in terms of time or in terms of a single resource. The automata-based modeling language MODEST [13] combines functional and non-functional requirements for stochastic systems, using a process algebra with dynamically computed weight expressions in probabilistic choice. Compared to ABS, these approaches do not associate capacities with locations but focus on non-functional aspects of embedded system *without* resource provisioning and management of dynamically created locations as studied in our article.

Work on the modeling of object-oriented systems with resource constraints is scarce. The UML profile for scheduling, performance and time (SPT) describes scheduling policies according to the underlying deployment model [22]. Using SPT, the Core Scenario Model (CSM) [39] is informally defined to generate performance models from UML. However, CSM is not executable as it only identifies a subset of the possible system behaviors [39]. Verhoef’s extension of VDM++ for embedded real-time systems [41] is based on abstract executable specification and models static deployment of fixed resources targeting the embedded domain, namely CPUs and buses.

Related work on simulation tools for cloud computing is mostly reminiscent of network simulators. Testing techniques and tools for cloud-based software systems are surveyed in [8]. In particular, CloudSim [18] and ICanCloud [38] are simulation tools using virtual machines to simulate cloud environments. CloudSim is a mature tool which has been used for a number of papers, but it is restricted to simulations on a single computer. In contrast, ICanCloud supports distribution on a cluster. EMUSIM [17] is an integrated tool that uses AEF (Automated Emulation Framework) to estimate performance and costs for an application by means of emulations to produce improved input parameters for simulations in CloudSim. Compared to these approaches, our work aims to support the developer of client applications for cloud-based environments at an early phase in the software engineering process and is based on a formal semantics.

Resource analysis [4, 25, 26, 43] aims at estimating the resource usage of a given program and providing guarantees that the program will not exceed the inferred amount of resources. Typically, such analyses are either based on *monitoring* the program execution, or on *formal*

*methods* that are able to infer this information *statically*, i.e., without executing the program on concrete data. Static approaches have a clear conceptual advantage over monitoring, since they provide guarantees that are valid for *any input*, not only for specific runs. There exist several powerful static resource analyzers (including COSTA [4], SPEED [25], RAML [26]), however, these have mainly dealt with traditional (i.e., non real-time, non-distributed) applications. A recent extension [3] deals with the theoretical complications that distributed systems pose to static analysis and reports on a prototype implementation. In this article, for the first time, resource analysis has been applied to a case study of industrial size.

In software design, no general, systematic means exists today to model and analyze software in the context of a set of available virtualized resources, nor to analyze resource redistribution in terms of load balancing or reflective operations. None of the cited works directly addresses the challenges raised by virtualization; in particular, they do not model quantitative resources as data inside the system itself, which is a key aspect of virtualized computations.

## 6 Conclusion and Future Work

Industry leaders in cloud computing research recently expressed “the need to describe your infrastructure requirements as executable code” (Thomas Schwindt, IBM Research Labs Zurich in [23]) as well as that “performance modeling is a future key technology” (Bryan Stephenson, HP Labs Palo Alto, also in [23]). These judgments align perfectly with the main result of this article: that it is possible to model and analyze *low-level* software aspects, including performance, resource consumption, and deployment, in an *executable, abstract* language. The Real-Time ABS language used for this purpose shows that it can cope with an *industrial* case study.

As an immediate benefit from executable abstract modeling, it is possible to perform comprehensive simulations that allow to predict and to evaluate the consequences of different scheduling, load balancing, and deployment strategies. We also demonstrated that one can calibrate the model in such a way that meaningful comparisons and predictions relative to deployed code are possible. This led to the identification of a hot spot in the actual production code.

The formal semantics of Real-Time ABS makes it feasible to analyze models *statically*. Using the cost analysis tool COSTABS, it is possible to *automatically* compute symbolic worst-case bounds for resource consumption. To the best of our knowledge, static resource analysis is applied here the first time successfully on an industrial-strength case study.

The inferred bounds were in turn used to *automatically* provide cost annotations for expressions and statements in the Real-Time ABS model, hence making their manual creation unnecessary. This allows for a *hybrid* deployment validation approach where dynamic simulation and static analysis combine their strengths.

In summary, the results of this paper demonstrate that the combination of abstract, executable modeling together with state-of-the-art static analysis tools is well on the way toward *tool-supported software design of virtualized applications*.

**Acknowledgments.** Although the author list of this article is rather long already, we gratefully thank the many more people who have been involved in the development of the ABS and Real-Time ABS languages, their toolset, as well as the COSTABS system. Without their effort, the research reported here would not have been possible.

## References

- [1] G. A. Agha. *ACTORS: A Model of Concurrent Computations in Distributed Systems*. MIT Press, 1986.
- [2] E. Albert, P. Arenas, S. Genaim, M. Gómez-Zamalloa, and G. Puebla. Cost Analysis of Concurrent OO programs. In *The 9th Asian Symposium on Programming Languages and Systems (APLAS'11)*, volume 7078 of *LNCS*, pages 238–254. Springer, Dec. 2011.
- [3] E. Albert, P. Arenas, S. Genaim, M. Gómez-Zamalloa, and G. Puebla. COSTABS: a cost and termination analyzer for ABS. In O. Kiselyov and S. Thompson, editors, *Proc. Workshop on Partial Evaluation and Program Manipulation (PEPM'12)*, pages 151–154. ACM, 2012.
- [4] E. Albert, P. Arenas, S. Genaim, G. Puebla, and D. Zanardini. Cost Analysis of Object-Oriented Bytecode Programs. *TCS*, 413(1):142–159, 2012.
- [5] E. Albert, S. Genaim, M. Gómez-Zamalloa, E. B. Johnsen, R. Schlatte, and S. L. Tapia Tarifa. Simulating concurrent behaviors with worst-case cost bounds. In M. Butler and W. Schulte, editors, *FM 2011*, volume 6664 of *LNCS*, pages 353–368. Springer, June 2011.
- [6] T. Amnell, E. Fersman, L. Mokrushin, P. Pettersson, and W. Yi. TIMES: A tool for schedulability analysis and code generation of real-time systems. In K. G. Larsen and P. Niebert, editors, *Proc. of the First International Workshop on Formal Modeling and Analysis of Timed Systems (FORMATS 2003)*, volume 2791 of *LNCS*, pages 60–72. Springer, 2003.
- [7] J. Armstrong. *Programming Erlang: Software for a Concurrent World*. Pragmatic Bookshelf, 2007.
- [8] X. Bai, M. Li, B. Chen, W.-T. Tsai, and J. Gao. Cloud testing tools. In J. Z. Gao, X. Lu, M. Younas, and H. Zhu, editors, *Proc. 6th Intl. Symposium on Service Oriented System Engineering (SOSE'11)*, pages 1–12. IEEE, 2011.
- [9] C. Baier, B. R. Haverkort, H. Hermanns, and J.-P. Katoen. Performance evaluation and model checking join forces. *Comm. ACM*, 53(9):76–85, 2010.
- [10] F. Barbanera, M. Bugliesi, M. Dezani-Ciancaglini, and V. Sassone. Space-aware ambients and processes. *TCS*, 373(1–2):41–69, 2007.
- [11] F. Benoy and A. King. Inferring argument size relationships with CLP(R). In *Proc. of LOPSTR'97*, volume 1207 of *LNCS*, pages 204–223. Springer, 1997.
- [12] J. Bjørk, F. S. de Boer, E. B. Johnsen, R. Schlatte, and S. L. Tapia Tarifa. User-defined schedulers for real-time concurrent objects. *Innovations in Systems and Software Engineering*, 2012. Available online: <http://dx.doi.org/10.1007/s11334-012-0184-5>. To appear.
- [13] H. C. Bohnenkamp, P. R. D'Argenio, H. Hermanns, and J.-P. Katoen. MODEST: A compositional modeling formalism for hard and softly timed systems. *IEEE Trans. Software Eng.*, 32(10):812–830, 2006.
- [14] P. Bouyer, U. Fahrenberg, K. G. Larsen, and N. Markey. Quantitative analysis of real-time systems using priced timed automata. *Comm. ACM*, 54(9):78–87, 2011.

- [15] E. J. Bruno and G. Bollella. *Real-Time Java Programming: With Java RTS*. Prentice Hall PTR, 2009.
- [16] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic. Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility. *Future Generation Computer Systems*, 25(6):599–616, 2009.
- [17] R. N. Calheiros, M. A. Netto, C. A. D. Rose, and R. Buyya. EMUSIM: an integrated emulation and simulation environment for modeling, evaluation, and validation of performance of cloud computing applications. *Software: Practice and Experience*, 2012. Available online: <http://dx.doi.org/10.1002/spe.2124>. To appear.
- [18] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. F. D. Rose, and R. Buyya. CloudSim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms. *Software, Practice and Experience*, 41(1):23–50, 2011.
- [19] D. Caromel and L. Henrio. *A Theory of Distributed Objects*. Springer, 2005.
- [20] A. Chander, D. Espinosa, N. Islam, P. Lee, and G. C. Necula. Enforcing resource bounds via static verification of dynamic checks. *ACM ToPLaS*, 29(5), 2007.
- [21] F. S. de Boer, R. Hähnle, E. B. Johnsen, R. Schlatte, and P. Y. H. Wong. Formal modeling of resource management for cloud architectures: An industrial case study. In F. D. Paoli, E. Pimentel, and G. Zavattaro, editors, *Proc. European Conference on Service-Oriented and Cloud Computing (ESOCC 2012)*, volume 7592 of *LNCS*, pages 91–106. Springer, Sept. 2012.
- [22] B. P. Douglass. *Real Time UML – Advances in the UML for Real-Time Systems*. Addison-Wesley, 3 edition, 2004.
- [23] ESOCC panel discussion: Global management in service-oriented and cloud computing: Challenges and open issues. European Conference on Service-Oriented and Cloud Computing (ESSOC), Sept. 2012.
- [24] H. Foster, W. Emmerich, J. Kramer, J. Magee, D. S. Rosenblum, and S. Uchitel. Model checking service compositions under resource constraints. In I. Crnkovic and A. Bertolino, editors, *Proc. 6th joint meeting of the European Software Engineering Conf. and the ACM SIGSOFT Intl. Symposium on Foundations of Software Engineering (ESEC/FSE’07)*, pages 225–234. ACM, 2007.
- [25] S. Gulwani, K. K. Mehra, and T. M. Chilimbi. Speed: Precise and Efficient Static Estimation of Program Computational Complexity. In *POPL*, pages 127–139. ACM, 2009.
- [26] J. Hoffmann and M. Hofmann. Amortized Resource Analysis with Polynomial Potential. In *The 19th European Symposium on Programming (ESOP’10)*, volume 6012 of *LNCS*, pages 287–306. Springer, 2010.
- [27] E. B. Johnsen, R. Hähnle, J. Schäfer, R. Schlatte, and M. Steffen. ABS: A core language for abstract behavioral specification. In B. Aichernig, F. S. de Boer, and M. M. Bonsangue, editors, *Proc. 9th International Symposium on Formal Methods for Components and Objects (FMCO 2010)*, volume 6957 of *LNCS*, pages 142–164. Springer, 2011.
- [28] E. B. Johnsen and O. Owe. An asynchronous communication model for distributed concurrent objects. *Software and Systems Modeling*, 6(1):35–58, Mar. 2007.

- [29] E. B. Johnsen, O. Owe, R. Schlatte, and S. L. Tapia Tarifa. Dynamic resource reallocation between deployment components. In J. S. Dong and H. Zhu, editors, *Proc. International Conference on Formal Engineering Methods (ICFEM'10)*, volume 6447 of *LNCS*, pages 646–661. Springer, Nov. 2010.
- [30] E. B. Johnsen, O. Owe, R. Schlatte, and S. L. Tapia Tarifa. Validating timed models of deployment components with parametric concurrency. In B. Beckert and C. Marché, editors, *Proc. International Conference on Formal Verification of Object-Oriented Software (FoVeOOS'10)*, volume 6528 of *LNCS*, pages 46–60. Springer, 2011.
- [31] E. B. Johnsen, R. Schlatte, and S. L. Tapia Tarifa. Modeling resource-aware virtualized applications for the cloud in Real-Time ABS. In T. Aoki and K. Tagushi, editors, *Proc. 14th International Conference on Formal Engineering Methods (ICFEM'12)*, volume 7635 of *LNCS*, pages 71–86. Springer, Nov. 2012.
- [32] R. Kirner and P. P. Puschner. Classification of WCET analysis techniques. In *ISORC*, pages 190–199. IEEE Computer Society, 2005.
- [33] J. Kramer. Is abstraction the key to computing? *Comm. ACM*, 50(4):36–42, 2007.
- [34] K. G. Larsen, P. Pettersson, and W. Yi. UPPAAL in a nutshell. *International Journal on Software Tools for Technology Transfer*, 1(1–2):134–152, 1997.
- [35] E. A. Lee. Computing needs time. *Comm. ACM*, 52(5):70–79, 2009.
- [36] L. Moreau and C. Queinnec. Resource aware programming. *ACM ToPLaS*, 27(3):441–476, 2005.
- [37] M. Netjes, W. M. van der Aalst, and H. A. Reijers. Analysis of resource-constrained processes with Colored Petri Nets. In K. Jensen, editor, *Proceedings of the Sixth Workshop on the Practical Use of Coloured Petri Nets and CPN Tools (CPN 2005)*, volume 576 of *DAIMI*. University of Aarhus, 2005.
- [38] A. Nuñez, J. Vázquez-Poletti, A. Caminero, G. Castañé, J. Carretero, and I. Llorente. iCan-Cloud: A flexible and scalable cloud infrastructure simulator. *Journal of Grid Computing*, 10:185–209, 2012.
- [39] D. B. Petriu and C. M. Woodside. An intermediate metamodel with scenarios and resources for generating performance models from UML designs. *Software and System Modeling*, 6(2):163–184, 2007.
- [40] M. Sgroi, L. Lavagno, Y. Watanabe, and A. Sangiovanni-Vincentelli. Synthesis of embedded software using free-choice Petri nets. In *Proc. 36th ACM/IEEE Design Automation Conference (DAC'99)*, pages 805–810. ACM, 1999.
- [41] M. Verhoef, P. G. Larsen, and J. Hooman. Modeling and validating distributed embedded real-time systems with VDM++. In J. Misra, T. Nipkow, and E. Sekerinski, editors, *Proceedings of the 14th International Symposium on Formal Methods (FM'06)*, volume 4085 of *LNCS*, pages 147–162. Springer, 2006.
- [42] A. Vulgarakis and C. C. Seceleanu. Embedded systems resources: Views on modeling and analysis. In *Proc. 32nd IEEE Intl. Computer Software and Applications Conference (COMPSAC'08)*, pages 1321–1328. IEEE Computer Society, 2008.

- [43] R. Wilhelm, J. Engblom, A. Ermedahl, N. Holsti, S. Thesing, D. B. Whalley, G. Bernat, C. Ferdinand, R. Heckmann, T. Mitra, F. Mueller, I. Puaut, P. P. Puschner, J. Staschulat, and P. Stenström. The worst-case execution-time problem—overview of methods and survey of tools. *ACM Trans. Embedded Comput. Syst.*, 7(3), 2008.
- [44] P. Y. H. Wong, N. Diakov, and I. Schaefer. Modelling Distributed Adaptable Object Oriented Systems using HATS Approach: A Fredhopper Case Study (invited paper). In *2nd International Conference on Formal Verification of Object-Oriented Software*, volume 7421 of *LNCS*. Springer, 2012.

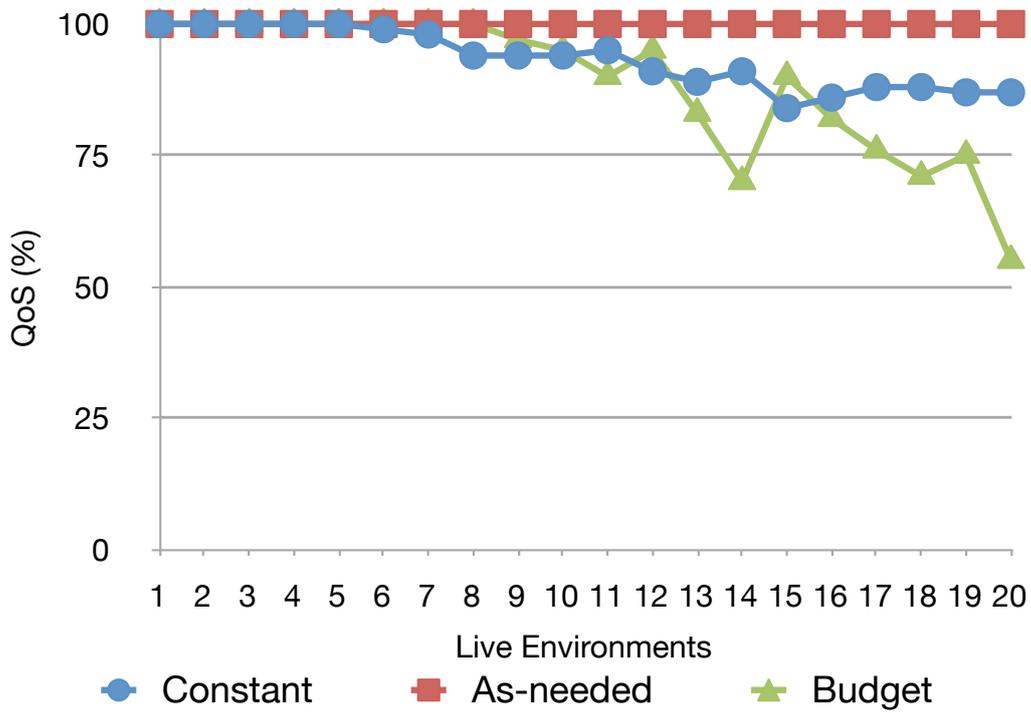


Figure 11: Simulation results: QoS, varying over the number of live environments.

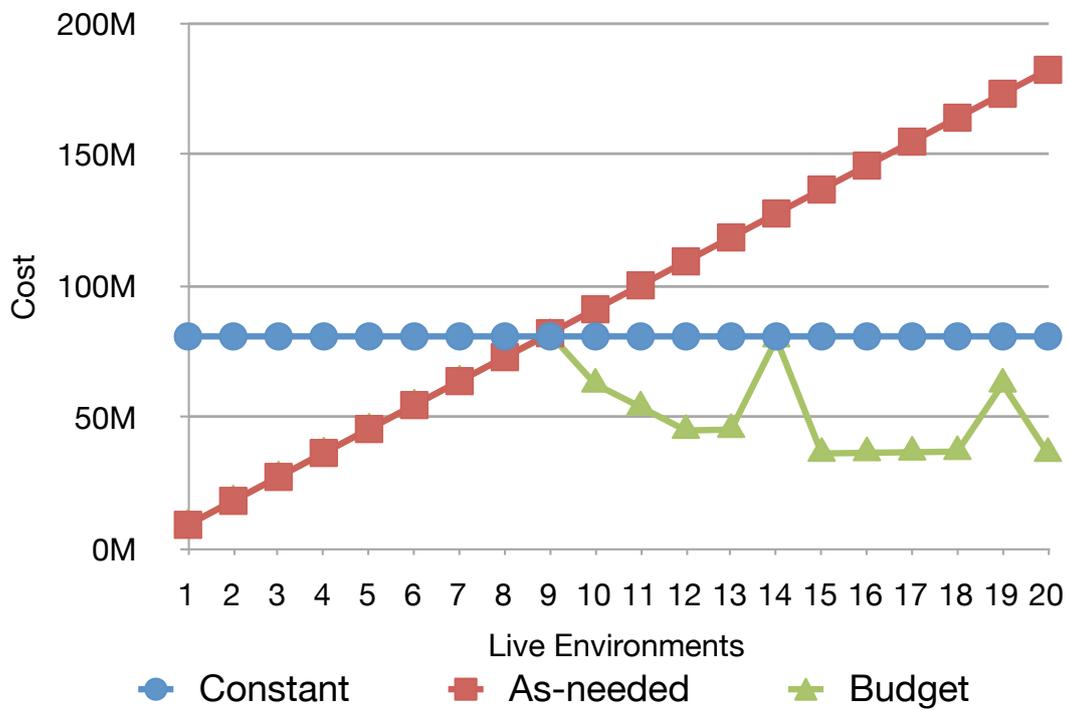


Figure 12: Simulation results: Total cost, varying over the number of live environments.

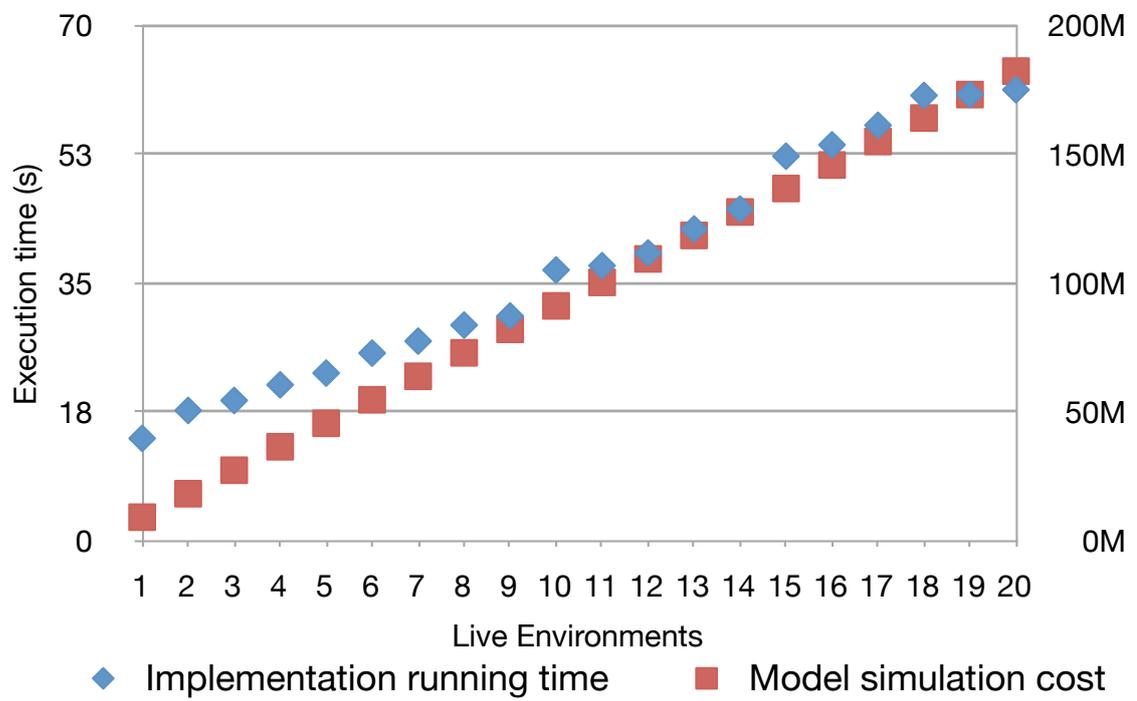


Figure 13: Comparison of the measured execution time of the implementation (left scale) and the accumulated cost of the simulation for the as-needed policy (right scale).