

Reactive Programming for Data-center Management: First Experiences and Prospects

Nicolas BERTHIER

Éric RUTTEN

Noël DE PALMA

November 18-22, 2013

Synchron'13



Outline

- Context & Motivations
- Reactive Programming...
- ... for n -tier Applications Management
- Validation
- Summary

Outline

- Context & Motivations
- Reactive Programming...
- ... for n -tier Applications Management
- Validation
- Summary


Autonomic Management Systems (AMSs)




Management Tasks

- ▶ Facilities: Energy, Cooling...
- ▶ Computation Infrastructure: Physical Machines...
- ▶ Application Management
 - ▶ Resource Allocation (Virtual Machines...), Configuration...
 - ▶ Example 3-tier Application:
→ Apache → Tomcat → Mysql

Autonomic Management Systems (AMSs)



Autonomic Management Systems are Reactive Systems



Noticeable Properties

- ▶ Typical **Feedback Loop**
- ▶ Sensors & Actuators **Spread Over** the Managed System
- ▶ Most Often: **Centralized Execution** of the Software Taking Management Decisions
- ▶ Yet: Management Aspects Addressed Independently from Each Other
 - ▶ Sizing / Repairing /...


Motivating Example: The *Sizing-repair* Case

Sizing

- ▶ Over-load \rightsquigarrow *add*
- ▶ Under-load \rightsquigarrow *remove*

Repair


- ▶ Failure \rightsquigarrow *repair*




Motivating Example: The *Sizing-repair* Case

Tier-level Coordination

- ▶ Concurrency Management
 - ▶ Mutual Exclusions on “Accesses” to Shared (Actuators / Resources)
 - ▶ e.g., *add*, *repair*
 - ▶ Conflicting Behaviors
 - ▶ e.g., *remove*, *repair*




Motivating Example: The *Sizing-repair*⁺ Case



Motivating Example: The *Sizing-repair*⁺ Case

Inter-tier Coordination

- ▶ Remote Impacts of Local Events, e.g.,
 - Failure in $T_1 \rightsquigarrow$ *Transient* Under-load in T_2
 - Failure in $T_1 \rightsquigarrow$ *Transient* (Under-load, then Over-load) in T_2
 - ← Failure in $T_2 \rightsquigarrow$ *Transient* Load Variation in T_1



Coordinating AMSs' Decisions

Observation

- ▶ Systems' Community: Realm of Object/Service/*-oriented Programming
 - ▶ Event-based: Handling One Event at a Time
 - ▶ Missing Global Knowledge

Coordinating AMSs' Decisions

Observation

- ▶ Systems' Community: Realm of Object/Service/*-oriented Programming
 - ▶ Event-based: Handling One Event at a Time
 - ▶ Missing Global Knowledge

Current Solutions...

- ▶ Heavy Use of Synchronization Primitives
 - ▶ Tedious, Error-prone
- ▶ Very *ad hoc*...
 - ▶ Hard to Generalize

Objective & Contributions

Objective

- ▶ Coordination-friendly Programming of AMSs

Proposal: Exploiting the *Reactive Nature* of AMSs

- ▶ New Design Methodology for AMSs
- ▶ Synchronous / Reactive Programming
- ▶ Discrete Controller Synthesis
- ▶ Prototype Implementation & Validation

Outline

- Context & Motivations
- Reactive Programming...
 - Synchronous Programming
 - Discrete Controller Synthesis
- ... for n -tier Applications Management
- Validation
- Summary

Synchronous Programming of Reactive Systems

Benefits

- ▶ High-level Programming Languages
 - ▶ e.g., Automata
- ▶ Efficient Verification Techniques
 - ▶ Design-time Validation
- ▶ Means for Static Property Enforcement
 - ▶ Automatic Generation of *Minimally Restrictive Controllers*
↗ Over-constrained System

Synchronous Programming of Reactive Systems

Benefits


- ▶ High-level Programming Languages
 - ▶ e.g., Automata
- ▶ Efficient Verification Techniques
 - ▶ Design-time Validation
- ▶ Means for Static Property Enforcement
 - ▶ Automatic Generation of *Minimally Restrictive Controllers*
↗ Over-constrained System

Usual Design Technique

- ▶ Model-driven Programming
 - ▶ System Model
 - ▶ Expressing Decisions Based on this Knowledge


Example Synchronous Program

Data-flow Representation





Example Synchronous Program

Data-flow Representation



As a Product of Communicating Boolean Mealy Automata



Executing a Synchronous Program

Typical Execution Pattern

▶ Inputs

- ▶ Impulses (\approx Discrete Events)
- ▶ Measures (\approx Always Relevant Data)


▶ Outputs

- ▶ Actions

Executing a Synchronous Program

Typical Execution Pattern

- ▶ Inputs
 - ▶ Impulses (\approx Discrete Events)
 - ▶ Measures (\approx Always Relevant Data)
- ▶ “Using” the tick
 - ▶ **Decide Relevant Instants (Not Necessarily Periodic!)**
 - ▶ Build an Input Vector *I*
 - ▶ Execute the tick.step() Method
 - ▶ Interpret Output Vector *O*, Send Commands to the Actuators
- ▶ Outputs
 - ▶ Actions



Discrete Controller Synthesis Problem

Principle

- ▶ Behaviors $A_1 \cdots A_n$ (e.g., Automata)
- ▶ Property Φ (e.g., Mutual Exclusions, Avoiding Bad States)

$$A_1 \parallel \cdots \parallel A_n \quad \Phi$$

Discrete Controller Synthesis Problem

Principle

- ▶ Behaviors $A_1 \cdots A_n$ (e.g., Automata)
- ▶ Property Φ (e.g., Mutual Exclusions, Avoiding Bad States)
- ▶ Provided $A_1 \cdots A_n$ are (made) **Controllable** ($\rightsquigarrow A_1' \cdots A_n'$)

$$A_1' \parallel \cdots \parallel A_n' \quad \Phi$$

Discrete Controller Synthesis Problem

Principle

- ▶ Behaviors $A_1 \cdots A_n$ (e.g., Automata)
- ▶ Property Φ (e.g., Mutual Exclusions, Avoiding Bad States)
- ▶ Provided $A_1 \cdots A_n$ are (made) **Controllable** ($\rightsquigarrow A_1' \cdots A_n'$)

\rightsquigarrow Controller C s.t.

$$A_1' \parallel \cdots \parallel A_n' \parallel C \models \Phi$$

Discrete Controller Synthesis Problem

Principle

- ▶ Behaviors $A_1 \cdots A_n$ (e.g., Automata)
- ▶ Property Φ (e.g., Mutual Exclusions, Avoiding Bad States)
- ▶ Provided $A_1 \cdots A_n$ are (made) **Controllable** ($\rightsquigarrow A_1' \cdots A_n'$)

\rightsquigarrow Controller C s.t.

$$A_1' \parallel \cdots \parallel A_n' \parallel C \models \Phi$$

Available Tool

- ▶ In the Synchronous Languages Domain: Sigali (until now...)

Outline

- Context & Motivations
- Reactive Programming...
- ... for n -tier Applications Management
- Validation
- Summary

Principles

Driving the Managed System with *Resource Drivers*

- ▶ **Operating Code** Senses & Acts Upon the Resources
 - ▶ Example of Resources:
 - ▶ Tier of an Application, Sets of Virtual Machines. . .
 - ▶ Pieces of Sequential Code (e.g., *add*)
- ▶ Associated **Behavioral Model** Represents Operating States
 - ▶ e.g., “Stopped”, “Running”, “Adding a node” . . .

Principles

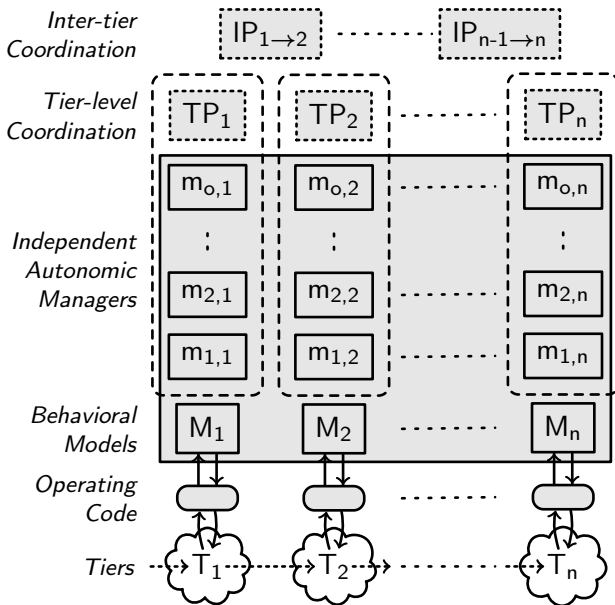
Driving the Managed System with *Resource Drivers*

- ▶ **Operating Code** Senses & Acts Upon the Resources
 - ▶ Example of Resources:
 - ▶ Tier of an Application, Sets of Virtual Machines. . .
 - ▶ Pieces of Sequential Code (e.g., *add*)
- ▶ Associated **Behavioral Model** Represents Operating States
 - ▶ e.g., “Stopped”, “Running”, “Adding a node” . . .

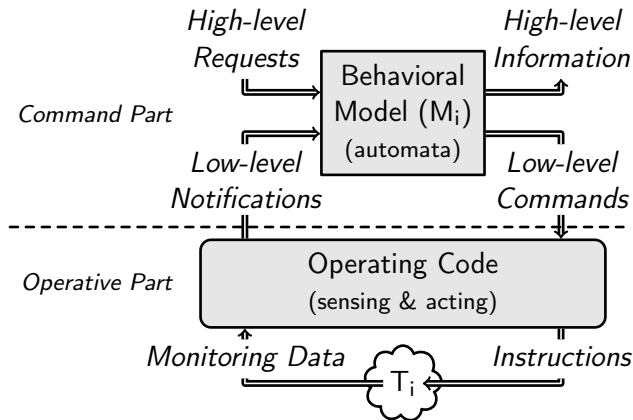
Autonomic Management Strategies & Coordination Policies

- ▶ Expressed in a *Declarative Way*
 - ▶ e.g., Automata, Systems of Equations
 - ▶ Possibly Involving Costs
- ▶ Coordination Policies Enforced by using Discrete Controller Synthesis

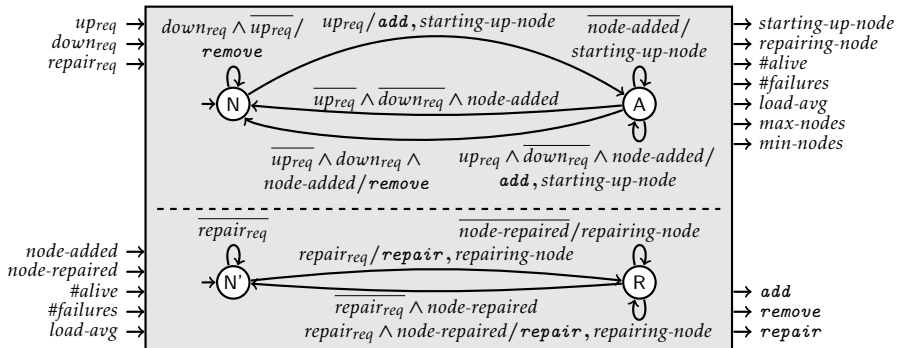
Encoding the AMS Logic for an n -tier Application




Structure of a Resource Driver



Behavioral Model for a Tier



Expressing Autonomic Management Strategies



Example Autonomic Management Strategies

Repair Strategy

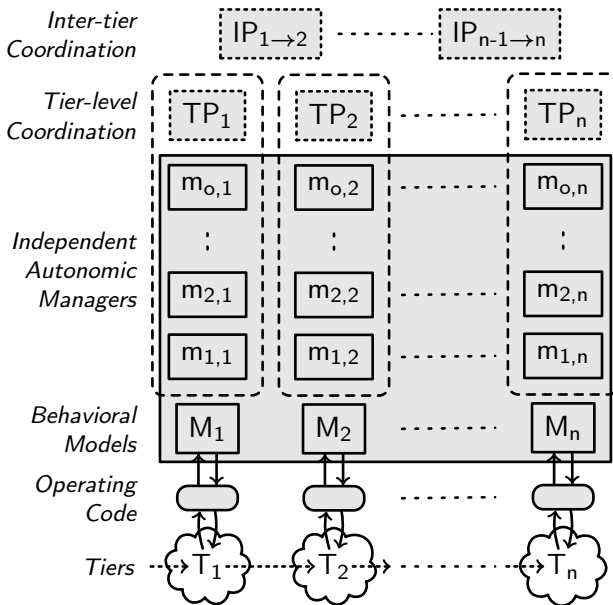
$$\text{repair} = \left(\#failures = 1 \wedge \overline{\text{repairing-node}} \right) \vee \#failures > 1$$

Sizing Strategy

$$\begin{aligned} \text{up} &= \overline{\text{max-nodes}} \wedge \text{upper-threshold}(\text{load-avg}, \#alive) \\ \text{down} &= \overline{\text{min-nodes}} \wedge \text{lower-threshold}(\text{load-avg}, \#alive) \end{aligned}$$

- ▶ Could be Less Trivial Synchronous Programs. . .

Specifying Coordination Policies



Example Tier-level Coordination Property

Allow Repair Operation (*repair*) While Up-sizing (*starting-up-node*),
or Up-sizing Operation (*add*) While Repairing (*repairing-node*)
In Case of Multiple “Recent” Failures Only ($\#failures > 1$)

Implementation

- ▶ Enforce both predicates:

$$(starting-up-node \wedge repair) \Rightarrow \#failures > 1$$

$$(repairing-node \wedge add) \Rightarrow \#failures > 1$$

Example Inter-Tier Coordination Property

For Each Tier T_i ,

Avoid Down-sizing Operations ($remove_i$)

if any of the “Preceding” Tiers T_j both:

- ▶ “Recently” Underwent at Least One Failure ($\#failures_j > 0$)
- ▶ Is Currently Repairing One Node ($repairing-node_j$)

Implementation

Enforce the following predicate, $\forall i \in [2, n]$:

$$\left(\bigvee_{j=1}^{i-1} (repairing-node_j \vee \#failures_j > 0) \right) \Rightarrow \overline{remove_i}$$

Eliciting Control Means

Filtering High-level Requests

$$repair_{req} = ok_{repair} \wedge repair$$

$$up_{req} = ok_{sizing} \wedge up$$

$$down_{req} = ok_{sizing} \wedge down$$

- ▶ High-level Requests: $\{ repair, up, down \}$
- ▶ Controllables: $\{ ok_{repair}, ok_{sizing} \}$
- ▶ Inputs of the Behavioral Model: $\{ repair_{req}, up_{req}, down_{req} \}$

Outline

- Context & Motivations
- Reactive Programming...
- ... for n -tier Applications Management
- **Validation**
 - Prototype Implementation
 - Validation
- Summary

Implementation Choices

Implementing the tick

- ▶ Using BZR
 - ▶ Encapsulating Sigali

Distributed Execution Platform: A^3

- ▶ Message-oriented Middleware (Agent-based)
- ▶ Fault-tolerant
- ▶ Atomic Executions
- ▶ Dynamic Deployments
 - ▶ Allows Adding & Removing Nodes (e.g., Virtual Machines)

Validation

Experimental Setup

Tier	Repair	Sizing	available VMs
T_{apache}	✓	✗	1
T_{tomcat}	✓	✓	3
$T_{\text{mysql-proxy}}$	✓	✗	1
T_{mysql}	✓	✓	3

- ▶ Synthetic Workload (based on Rubis' Auction Site)
- ▶ 9 VMWare Virtual Machines
- ▶ Workload Injection from a Separate Physical Machine
- ▶ Collect Exponentially-Weighted-Moving-Average (EWMA) of CPU Utilization Percentages

Validation


Experimental Setup

Tier	Repair	Sizing	available VMs
T_{apache}	✓	✗	1
T_{tomcat}	✓	✓	3
$T_{\text{mysql-proxy}}$	✓	✗	1
T_{mysql}	✓	✓	3


- ▶ Synthetic Workload (based on Rubis' Auction Site)
- ▶ 9 VMWare Virtual Machines
- ▶ Workload Injection from a Separate Physical Machine
- ▶ Collect Exponentially-Weighted-Moving-Average (EWMA) of CPU Utilization Percentages

Comparing Non-coordinated & Coordinated Executions

Non-coordinated Execution



Coordinated Execution



Outline

- Context & Motivations
- Reactive Programming...
- ... for n -tier Applications Management
- Validation
- Summary

Summary of the 1st Part

Objective

- ▶ Coordination-friendly Programming of AMSs

Proposal: Exploiting the *Reactive Nature* of AMSs

- ▶ New Design Methodology for AMSs
- ▶ Synchronous / Reactive Programming
- ▶ Discrete Controller Synthesis
- ▶ Prototype Implementation & Validation

Future Works

Quantitative Aspects

- ▶ Extending Behavioral Models with Relevant Quantitative Data
 - ▶ Enforce Quantitative Properties
 - ▶ *cf.* Next Part...

Distributed Execution of the AMS / Controller

- ▶ Centralized Design (Scalability Issue)
 - ↳ Synchronous Languages Implementation Problem
- ▶ Decentralized Design (Modeling Issue)
 - ↳ Decentralized Control Problem

Thanks!

Questions?

Outline

- Context & Motivations
- Reactive Programming...
 - Synchronous Programming
 - Discrete Controller Synthesis
- ... for n -tier Applications Management
- Validation
 - Prototype Implementation
 - Validation
- Summary

Outline

- Architecture of the Prototype

Architecture of the Prototype

