Karlsruhe Institute of Technology

Introduction to the Descartes Meta-Model (DMM)

Samuel Kounev

Charles University Prague, June 4, 2012

DESCARTES RESEARCH GROUP
INSTITUTE FOR PROGRAM STRUCTURES AND DATA ORGANIZATION, FACULTY OF INFORMATICS

References

- Papers can be downloaded from http://www.descartes-research.net
- Vision of Self-Aware IT Systems, Infrastructures and Services
  - S. Kounev. *Self-Aware Software and Systems Engineering: A Vision and Research Roadmap*. In Softwaretechnik-Trends 31(4), November 2011, ISSN 0720-8928.
  - S. Kounev, F. Brosig, N. Huber, and R. Reussner. *Towards self-aware performance and resource management in modern service-oriented systems*. In Proceedings of the 7th IEEE International Conference on Services Computing (SCC 2010), Miami, Florida, USA, July 5-10, 2010.
  - S. Kounev, F. Brosig, and N. Huber. *Self-Aware QoS Management in Virtualized Infrastructures* (Poster Paper). In 8th International Conference on Autonomic Computing (ICAC 2011), Karlsruhe, Germany, June 14-18, 2011.
- Descartes Meta-Model (DMM) / Online Models for Self-Awareness
  - N. Huber, F. Brosig and S. Kounev. *Modeling Dynamic Virtualized Resource Landscapes*. In Proceedings of the 8th ACM SIGSOFT International Conference on the Quality of Software Architectures (QoSA 2012), Bertinoro, Italy, June 25-28, 2012.
  - F. Brosig, N. Huber, and S. Kounev. *Modeling Parameter and Context Dependencies in Online Architecture-Level Performance Models*. In Proceedings of the 15th ACM SIGSOFT International Symposium on Component Based Software Engineering (CBSE 2012), Bertinoro, Italy, June 26-28, 2012.
  - N. Huber, A. van Hoorn, A. Koziolek, F. Brosig, and S. Kounev. *S/T/A: Meta-modeling Run-time Adaptation in Component-based System Architectures*. Under review.
  - N. Huber, F. Brosig, and S. Kounev. *Model-based Self-Adaptive Resource Allocation in Virtualized Environments*. In Proceedings of the 6th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS 2011), Honolulu, HI, USA, May 23-24, 2011.

References (2)

- Automatic Model Extraction (Model Inference) based on Online Monitoring

  - F. Brosig, N. Huber and S. Kounev. *Automated Extraction of Architecture-Level Performance Models of Distributed Component-Based Systems*. 26th IEEE/ACM International Conference On Automated Software Engineering (ASE 2011), Oread, Lawrence, Kansas, November 2011.
  - S. Kounev, K. Bender, F. Brosig, N. Huber, and R. Okamoto. *Automated Simulation-Based Capacity Planning for Enterprise Data Fabrics*. In 4th International ICST Conference on Simulation Tools and Techniques (SIMUTools 2011), Barcelona, Spain, 2011. **Best Paper Award**.
  - F. Brosig, S. Kounev, and K. Krogmann. *Automated Extraction of Palladio Component Models from Running Enterprise Java Applications*. In Proceedings of the 1st International Workshop on Run-time mOdels for Self-managing Systems and Applications (ROSSA 2009). In conjunction with 4th Intl. Conference on Performance Evaluation Methodologies and Tools (VALUETOOLS 2009), Pisa, Italy, October 19, 2009. ACM, New York, NY, USA, October 2009.

- Modeling Virtualization Platforms

  - N. Huber, M. Quast, M. Hauck, and S. Kounev. *Evaluating and Modeling Virtualization Performance Overhead for Cloud Environments*. In Proceedings of the 1st International Conference on Cloud Computing and Services Science (CLOSER 2011), Noordwijkerhout, The Netherlands, May 7-9 2011. **Best Paper Award**.
  - N. Huber, M. von Quast, F. Brosig and S. Kounev. *Analysis of the Performance-Influencing Factors of Virtualization Platforms*. In 12th International Symposium on Distributed Objects, Middleware, and Applications (DOA 2010), Crete, Greece, October 2010. Springer Verlag.

- Miscellaneous

  - S. Kounev, P. Reinecke, K. Joshi, J. Bradley, F. Brosig, V. Babka, S. Gilmore, and A. Stefanek. Resilience Assessment and Evaluation of Computing Systems, *Chapter Providing Dependability and Resilience in the Cloud: Challenges and Opportunities*. Dagstuhl Seminar 10292, Springer Verlag, 2012. To appear.
  - P. Meier, S. Kounev and H. Koziolek. *Automated Transformation of Palladio Component Models to Queueing Petri Nets*. In 19th IEEE/ACM International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2011), Singapore, July 25-27, 2011.
  - R. Krebs, C. Momm, and S. Kounev. *Metrics and Techniques for Quantifying Performance Isolation in Cloud Environments*. In Proceedings of the 8th ACM SIGSOFT International Conference on the Quality of Software Architectures (QoSA 2012), Bertinoro, Italy, June 25-28 2012.

Agenda

- Motivation
- Approach & Methodology
- Exemplary Results
- Vision
- Conclusion

Descartes Research Group @ KIT

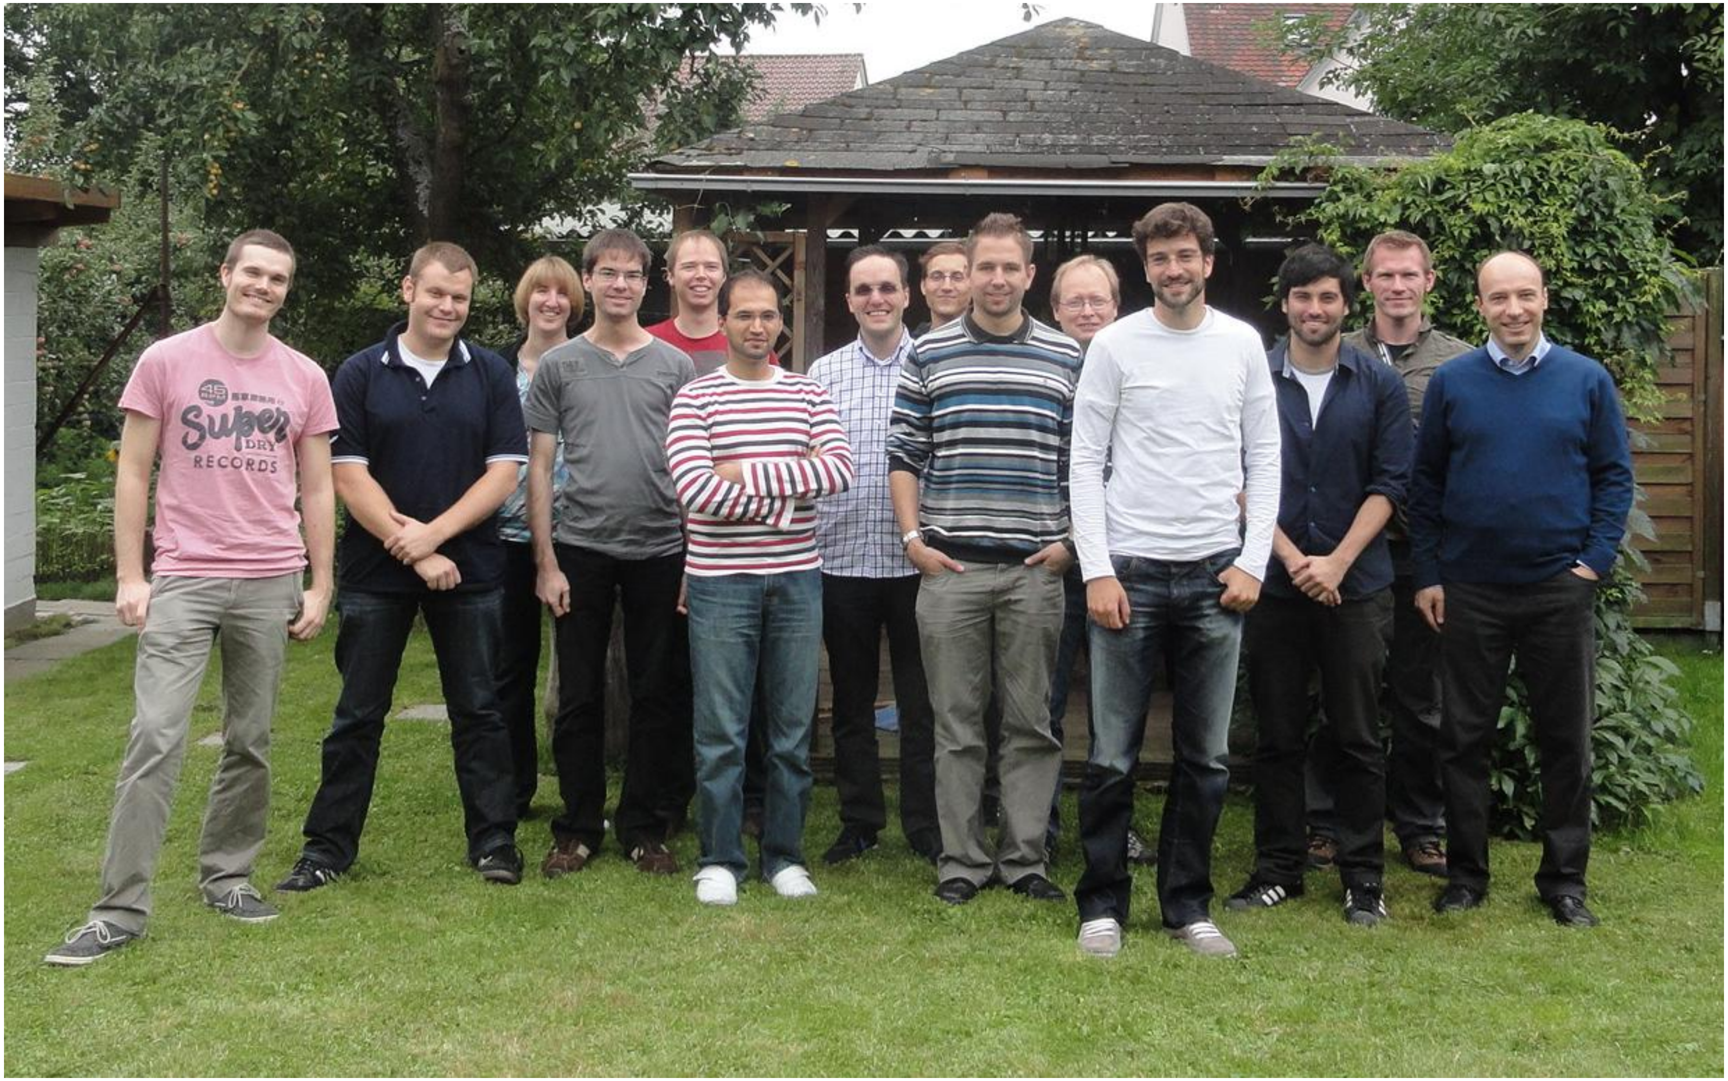

Research Areas

Dependability | Quality-of-Service | Elasticity | Efficiency

Research Areas

Systems Design, Measurement, Monitoring and Analysis
- Benchmarking
- Workload characterization
- Instrumentation & profiling
- Experimental analysis
- Online monitoring

Systems Modeling for Predictability at Design- and Run-Time
- Meta-models for dynamic software systems
- Analytical and simulation-based prediction models
- Model extraction, calibration & maintenance
- Predictability at design-time
- Predictability at run-time

Autonomic and Self-Adaptive Systems Management
- Dynamic resource provisioning & capacity mgmt
- Application quality-of-service management
- Elastic scalability
- Cost and efficiency management
- Power/energy management

Technology Domains

Cloud Computing, Virtualization & Green IT
- SaaS, PaaS, IaaS

Service-oriented Computing
- Web Services, SOA, ESB

Distributed Component-based Systems
- Java EE, MS .NET

Event-based Systems
- EDA, MOM, distributed pub/sub

Grid/Cluster Computing
- Service-oriented Grids

Motivation

- **Increasing data center operating costs**
  - System management costs
  - Power consumption costs
  - Cooling costs

- **Gartner (2009)**
  - Power consumption doubled from 2000 to 2005
  - By 2025, an increase by 1600% expected!
- **Increasing carbon footprint of ICT**
  - 2% to 4% of global CO2 emissions
  - **Projected to rise to 10% in 10 years**

Conventional Data Centers

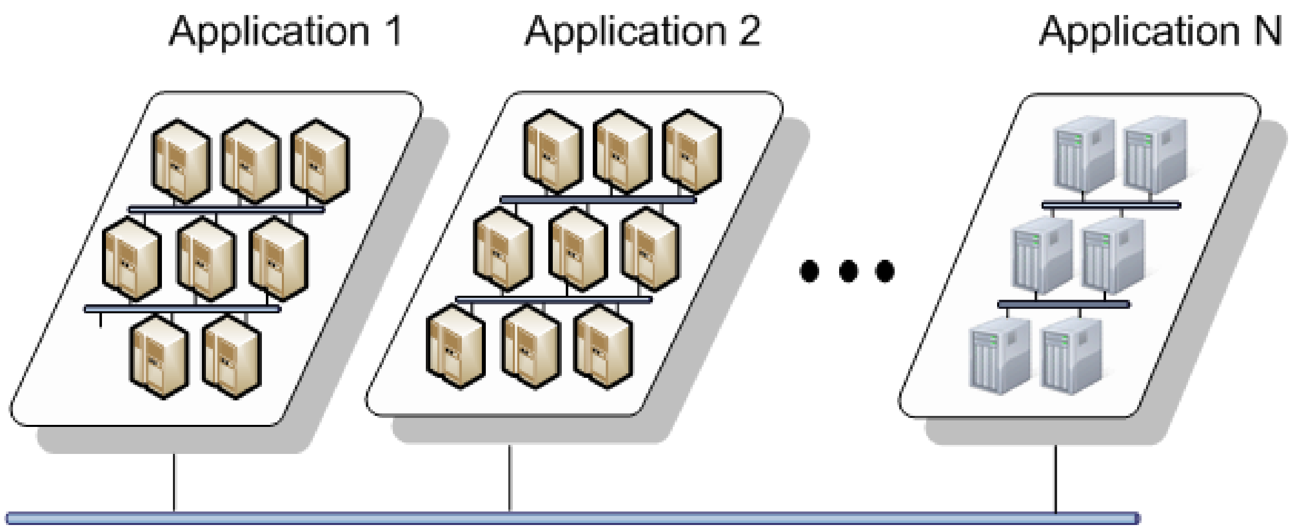

- Applications running on dedicated hardware
- Over-provisioned system resources
- Poor resource utilization and energy efficiency
- Increasing number of servers → rising operating costs

Dynamic Virtualized Infrastructures

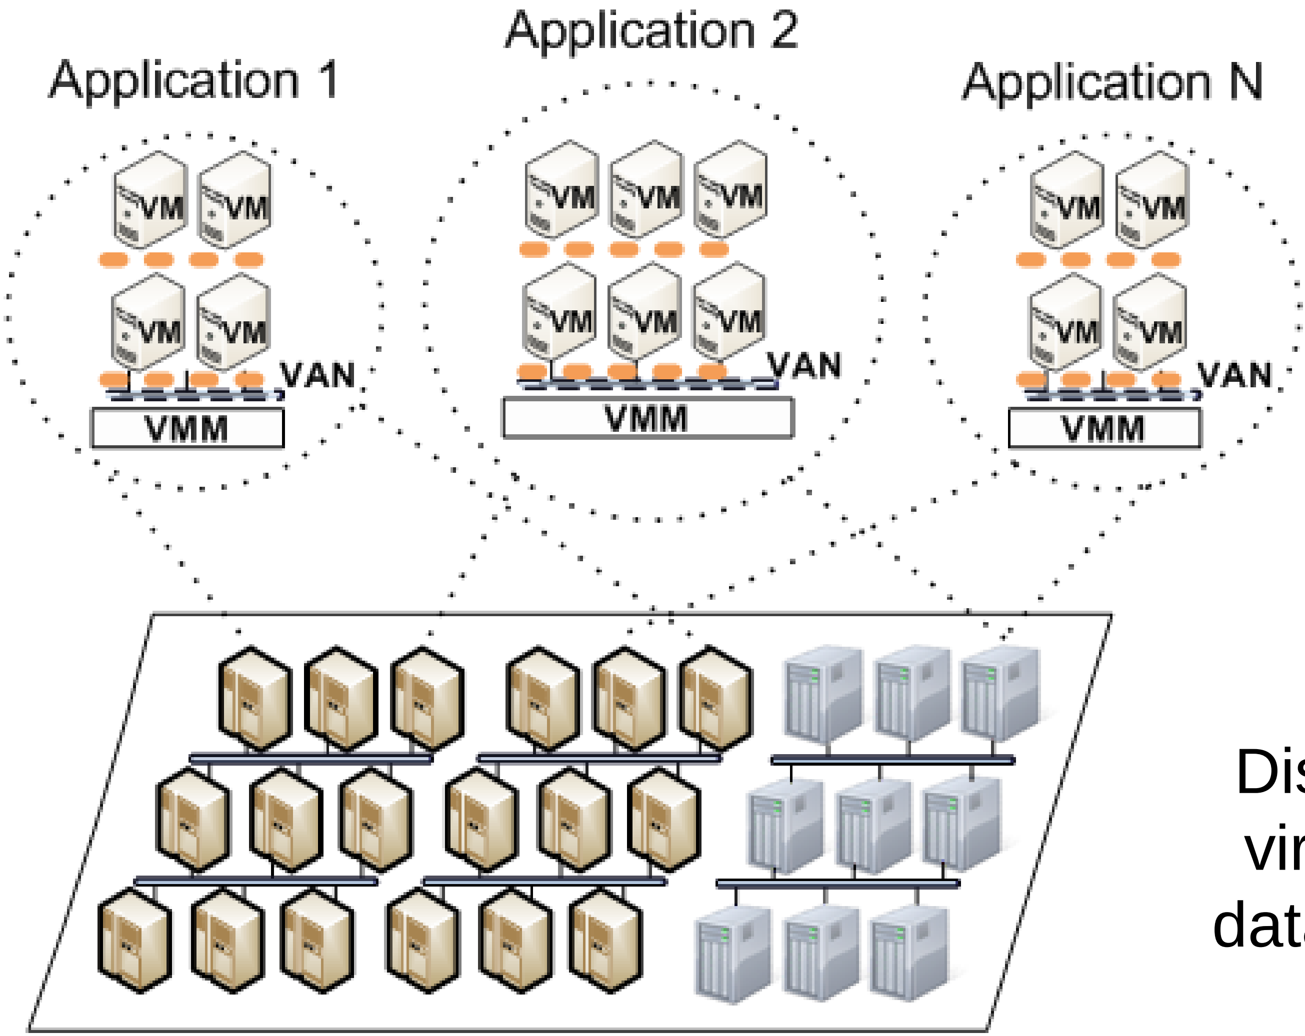

Distributed virtualized data centers

Challenges

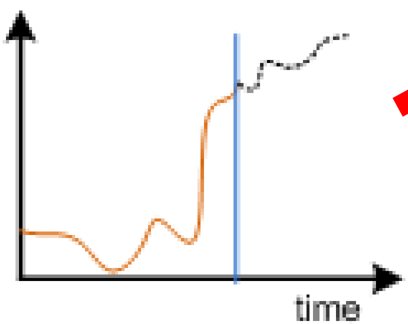

Load Spike

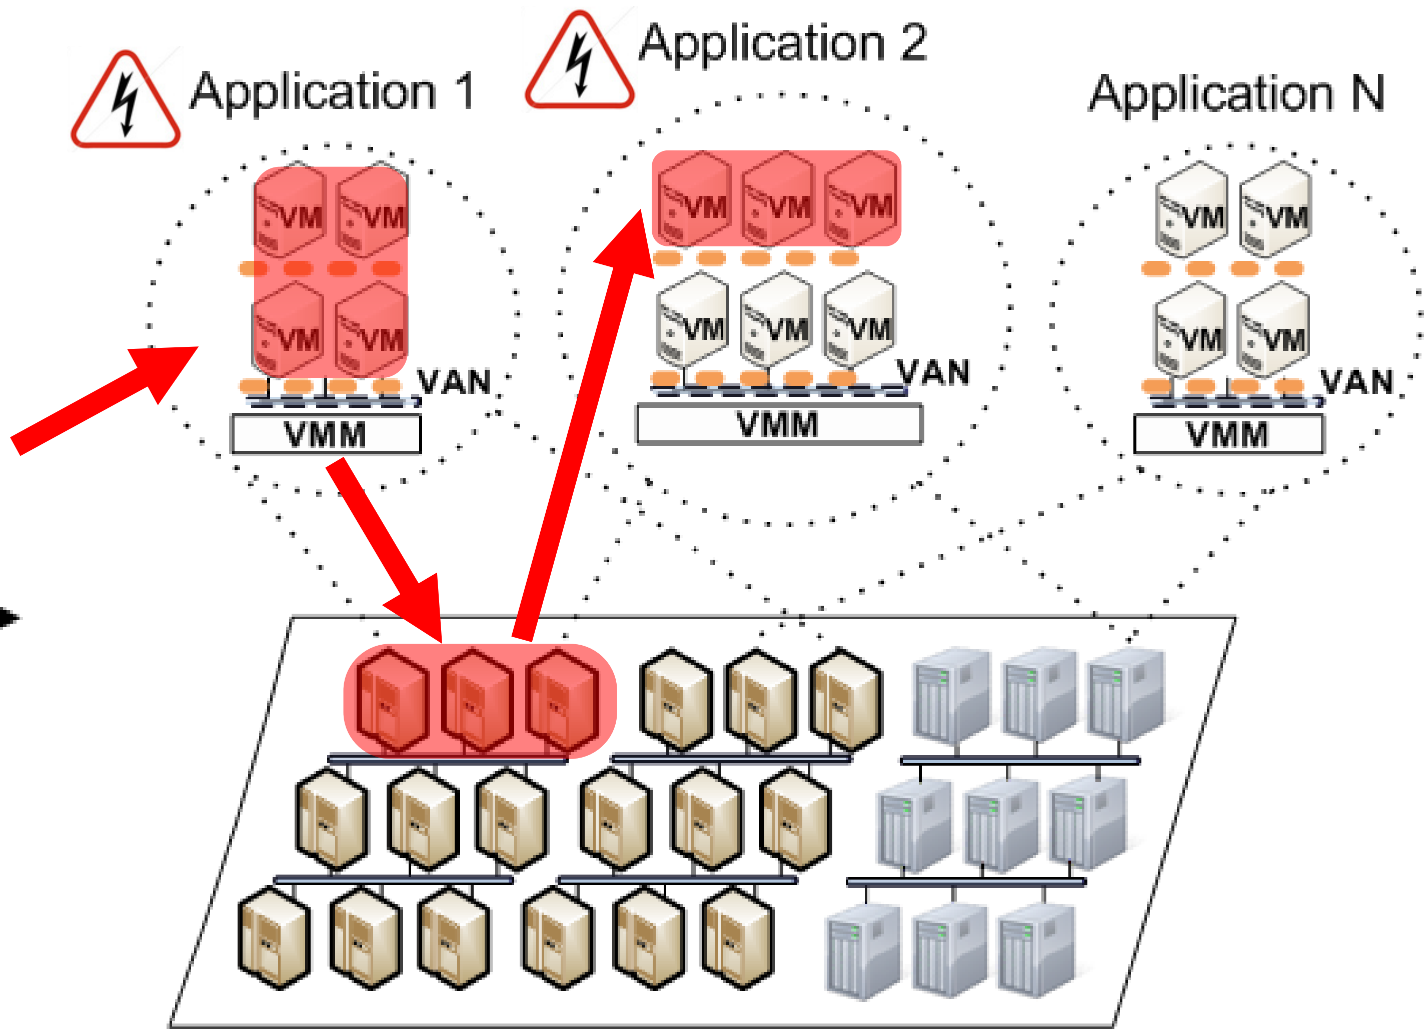

Challenges (2)

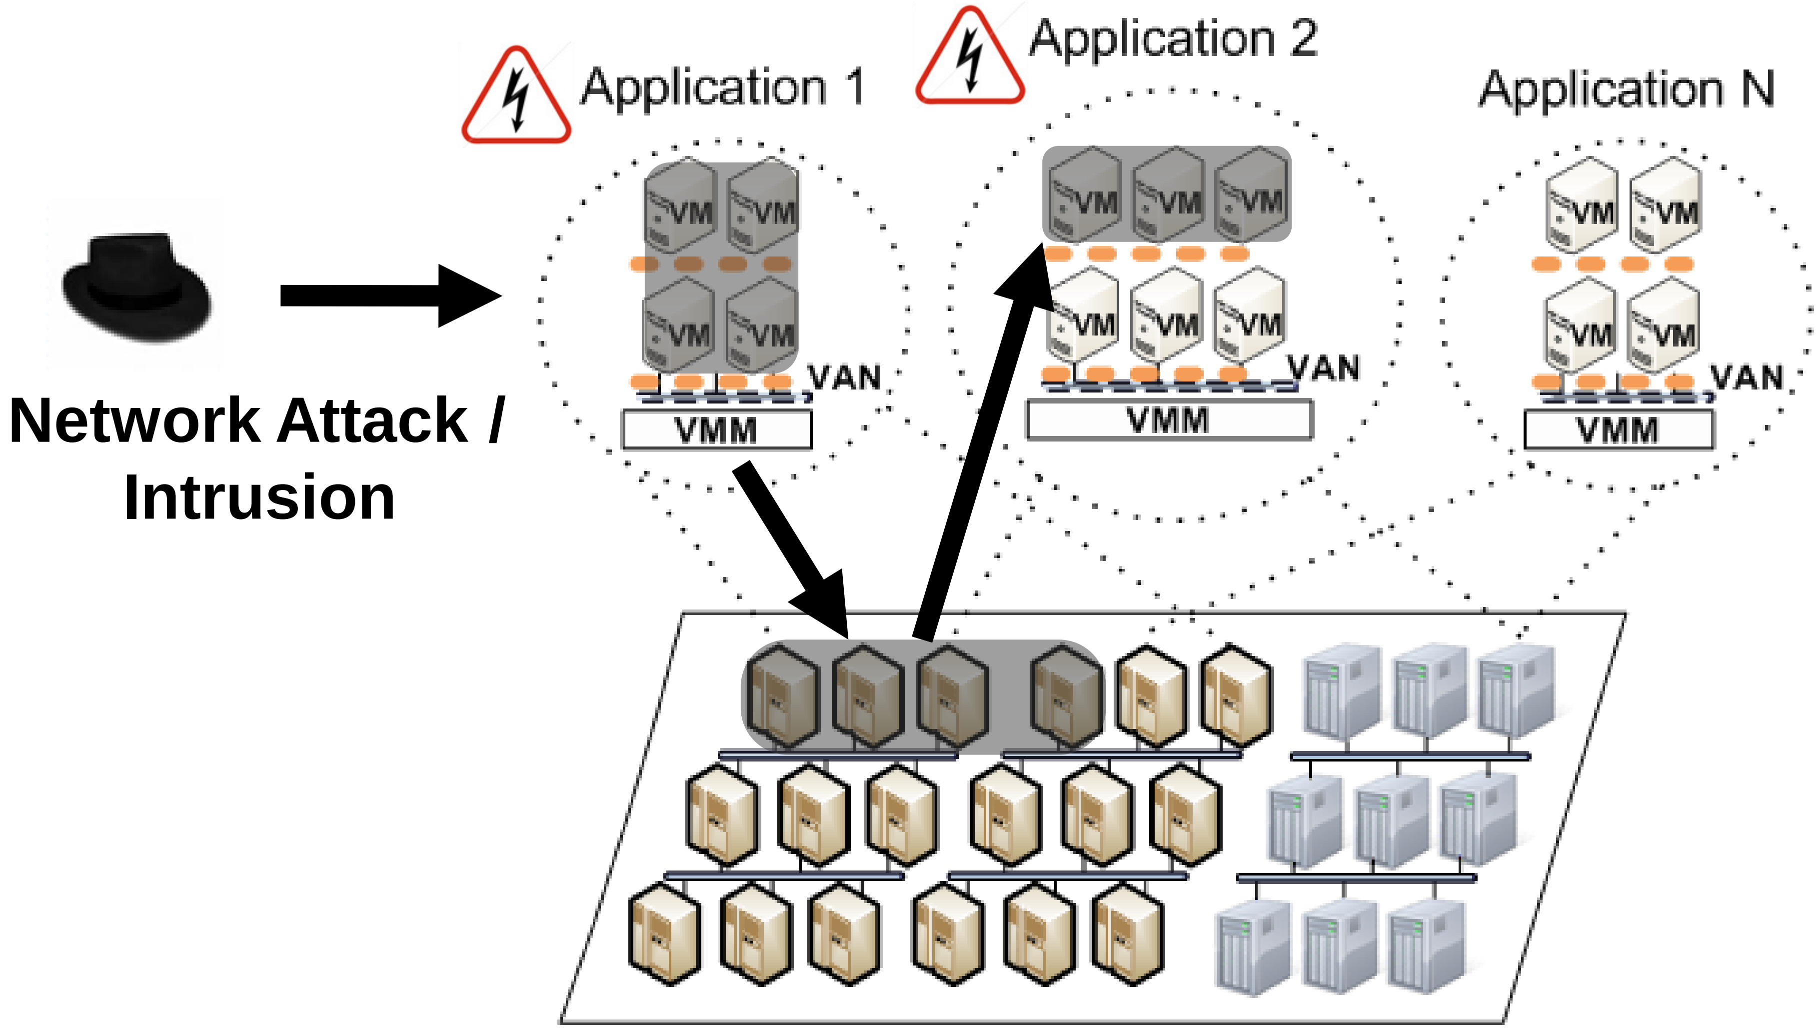

Challenges (3)

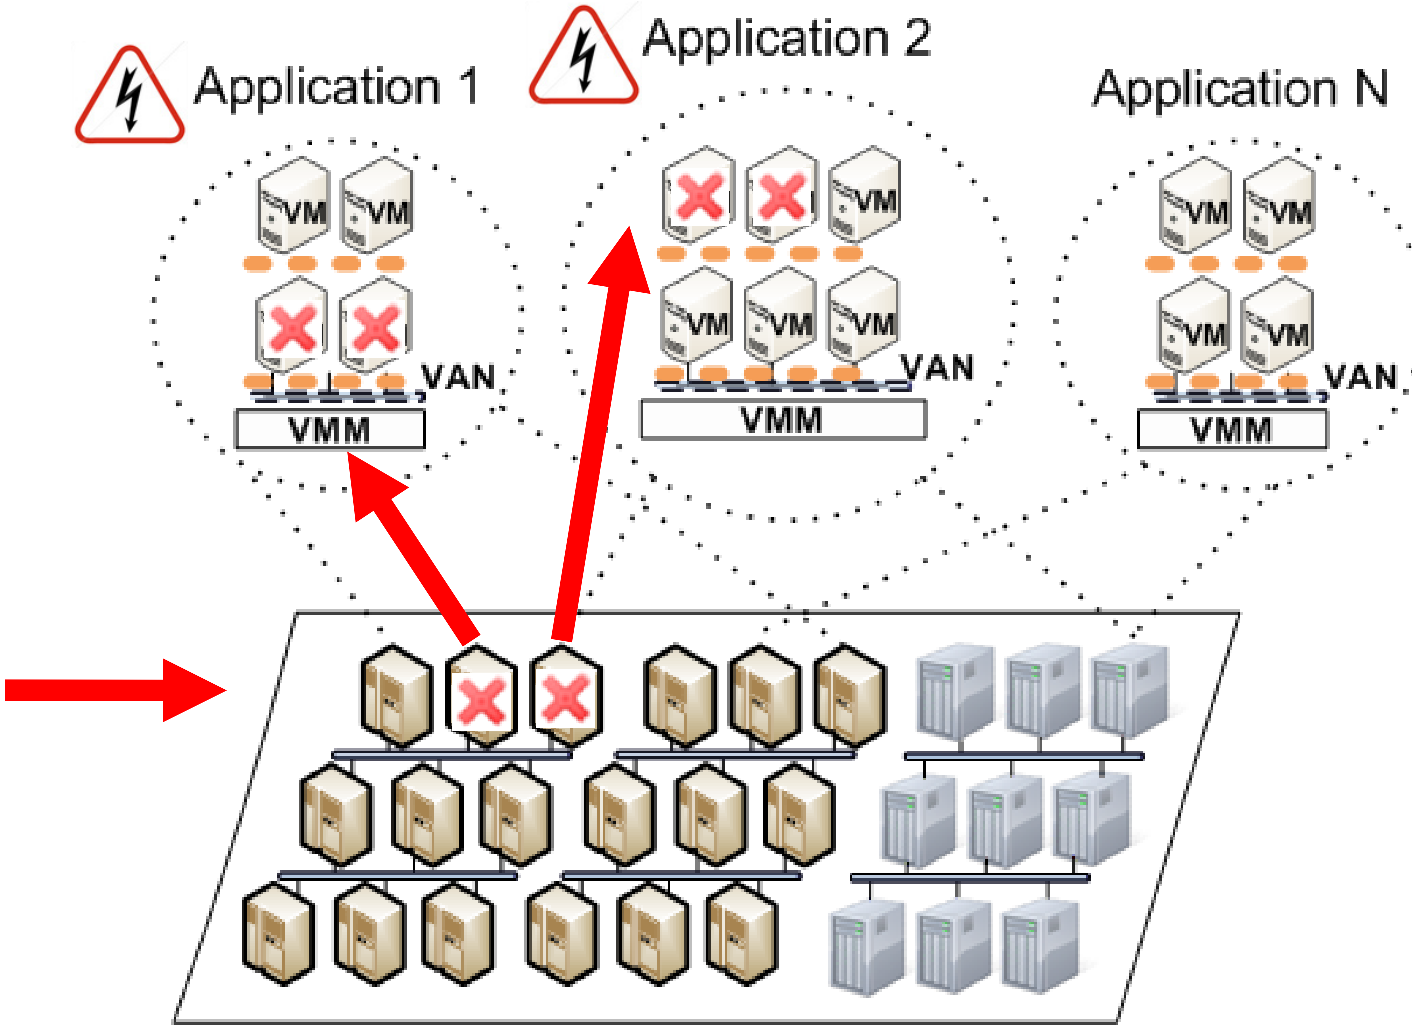

Hardware Failure

Challenges (4)

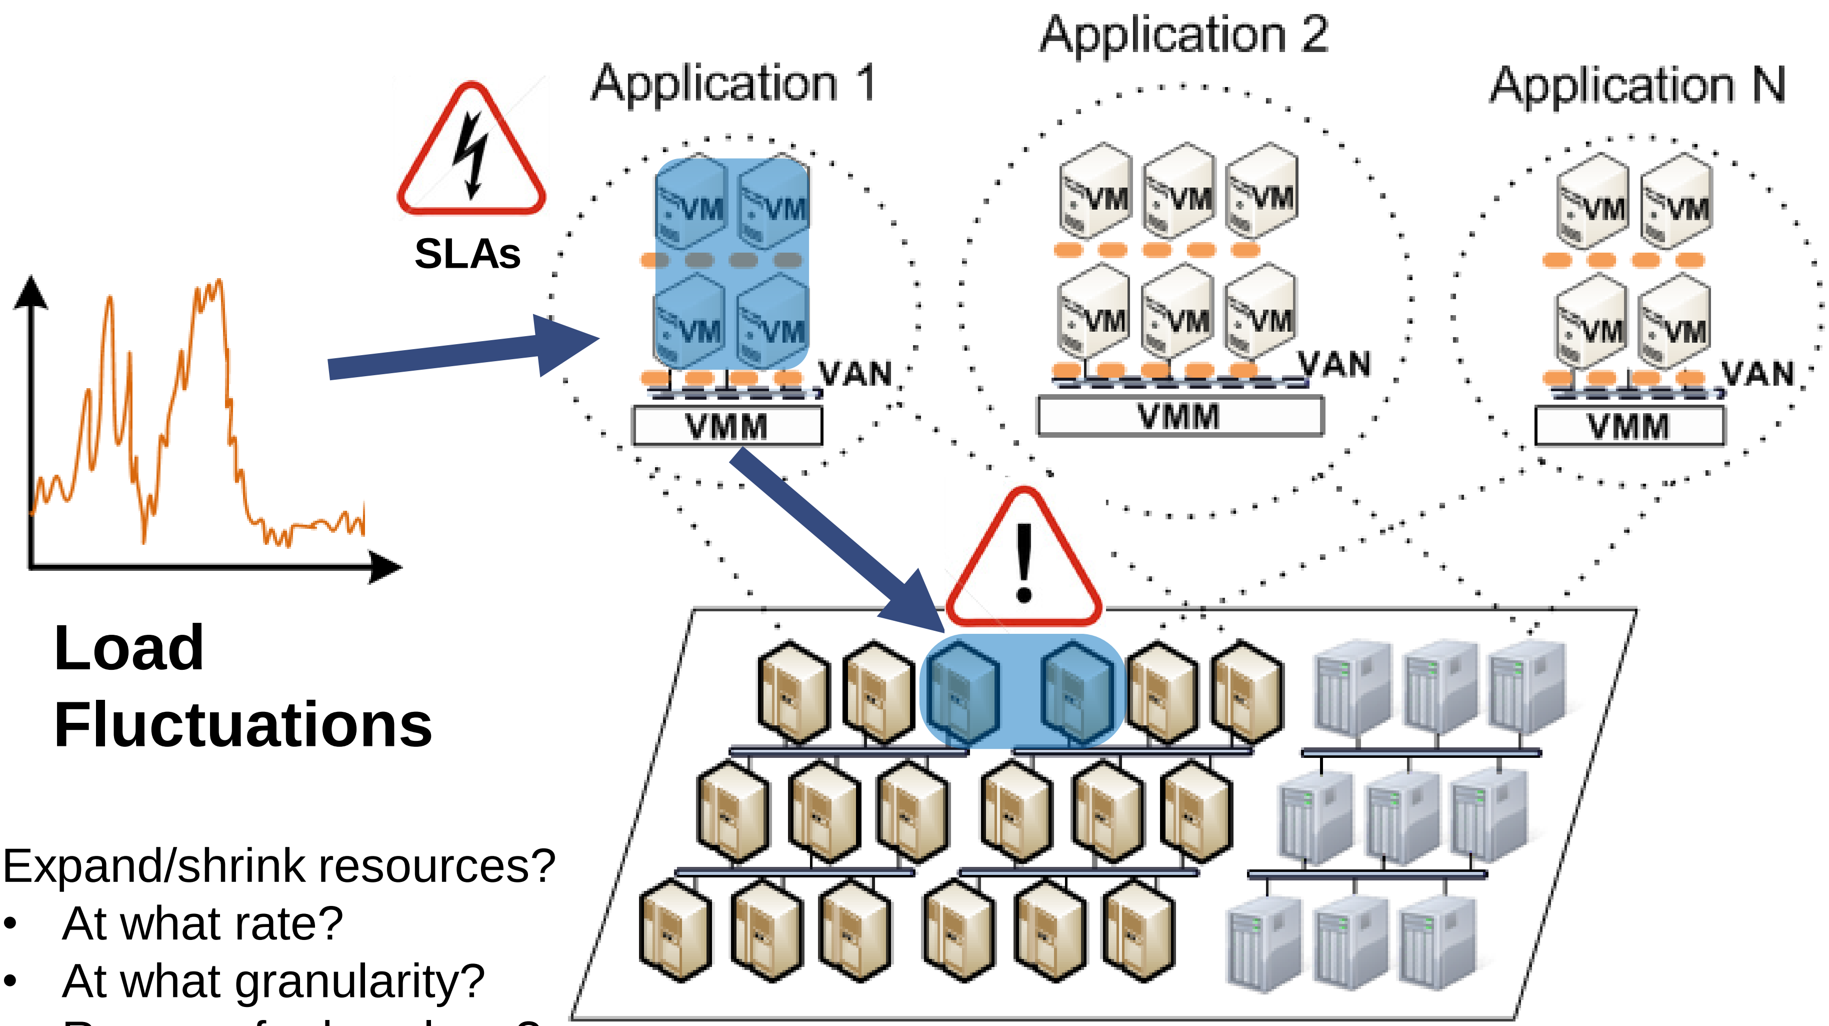

Load Fluctuations

Expand/shrink resources?
- At what rate?
- At what granularity?
- Reserve for how long?

Challenges

- Increased system complexity and dynamics
- Lack of direct control over underlying hardware
- New threats and vulnerabilities due to resource sharing
- Separation of service providers and infrastructure providers

↓

- Inability to provide dependability and QoS guarantees
- Lack of trust

High-Level Research Questions

- How to ***automatically predict*** vulnerabilities arising from varying workloads, network attacks or system failures?
- How to ***proactively adapt*** the system to avoid SLA violations or inefficient resource usage?
- How to **provide dependability / QoS guarantees** while ensuring high resource utilization and energy efficiency?

How to engineer **trustworthy** and **efficient** systems?

Further details in:

- S. Kounev, P. Reinecke, K. Joshi, J. Bradley, F. Brosig, V. Babka, S. Gilmore, and A. Stefanek. *Resilience Assessment and Evaluation of Computing Systems*, Chapter **Providing Dependability and Resilience in the Cloud: Challenges and Opportunities**. Dagstuhl Seminar 10292, Springer Verlag, 2012. *To appear.*

Agenda

- Motivation
- **Approach & Methodology**
- Exemplary Results
- Vision
- Conclusion

Proactive Self-Adaptive Systems Management

Online QoS prediction for problem anticipation

Online QoS prediction for reconfiguration impact analysis

Autonomic system adaptation

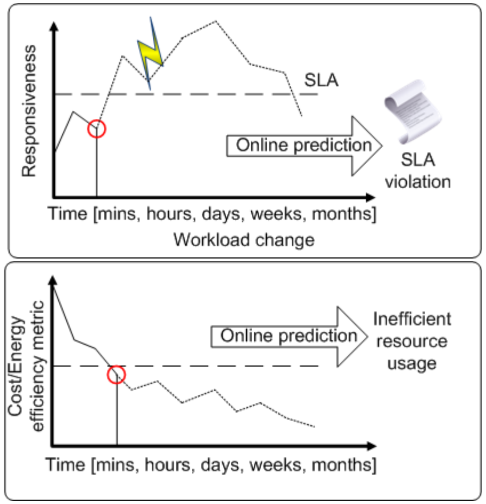

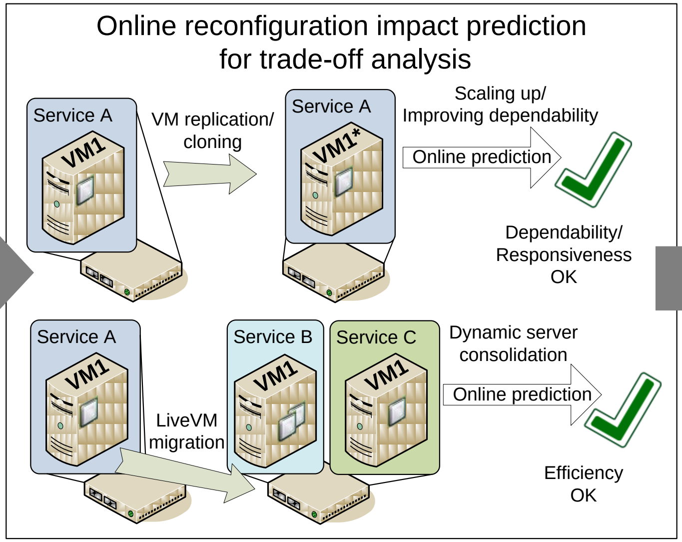

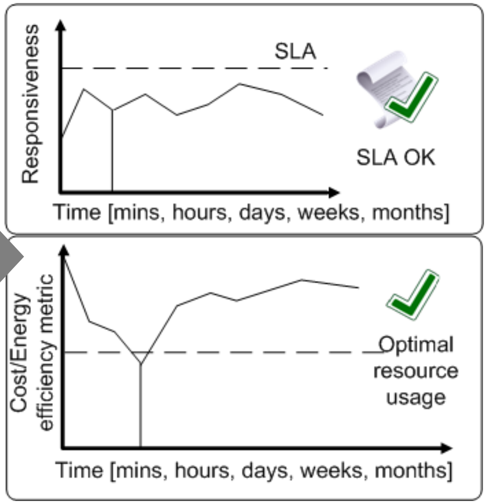

PART 1

PART 2

PART 3

Proactive Self-Adaptive Systems Management

System needs to be **explicitly** aware of its

Online QoS prediction for problem anticipation	Online QoS prediction for reconfiguration impact analysis	Autonomic system adaptation
Goals & objectives Dynamic changes Effect of changes	System architecture Resource landscape Possible adaptation actions Impact of adaptation (beforehand)	Dynamic adaptation strategies Impact of adaptation (afterwards)
PART 1	**PART 2**	**PART 3**

Self-Aware Software Systems

Self-Reflective

Aware of their software architecture, execution environment and hardware infrastructure, as well as of their operational goals

Self-Predictive

Able to anticipate and predict the effect of dynamic changes in the environment, as well as the effect of possible adaptation actions

Self-Adaptive

Proactively adapting as the environment evolves to ensure that their operational goals are continuously met

Details in:

- S. Kounev. **Self-Aware Software and Systems Engineering: A Vision and Research Roadmap**. In *Softwaretechnik-Trends 31(4), November 2011, ISSN 0720-8928.*

Examples of Performance-Influencing Factors

System workload and usage profile
- Number and type of clients
- Input parameters and input data
- Data formats used
- Service workflow

Software architecture
- Connections between components
- Flow of control and data
- Component resource demands
- Component usage profiles

Execution environment
- Number of component instances
- Server execution threads
- Amount of Java heap memory
- Size of database connection pools

Virtualization layer
- Physical resources allocated to VMs
  – number of physical CPUs
  – amount of physical memory
  – secondary storage devices

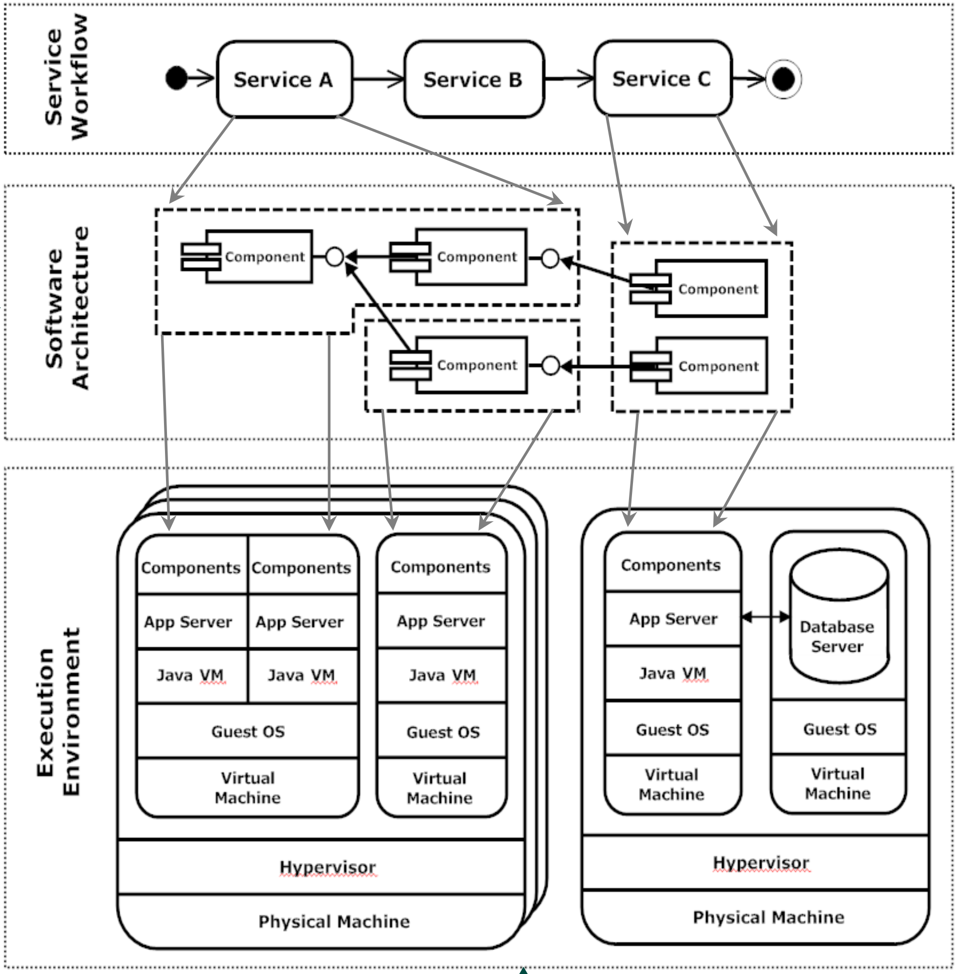

Network bandwidth between system nodes

State-of-the-Art: Model-based Approaches

Karlsruhe Institute of Technology

1. Models for QoS prediction at run-time

- Simple models used that abstract the system at very high level
- Many restrictive assumptions imposed
- Most of the mentioned aspects are not modeled explicitly

[G. Pacifici et al], [A. D'Ambrogio et al], [G. Tesauro et al], [D. Menasce et al], [C. Adam et al], [Rashid A. Ali et al], [I. Foster er al], [S. Bleul et al], [A. Othman et al], [P. Shivam et al], …

2. Models for QoS prediction at design & deployment time

- Overhead in building and analyzing models
- Models assume static system architecture
- Maintaining models at run-time prohibitively expensive

[M. Woodside et al], [D. Petriu et al], [R. Reussner et al], [C. Smith et al], [R. Mirandola et al], [K. Trivedi et al], [V. Cortellessa et al], [I. Gorton et al], [D. Menasce et al], [E. Eskenazi et al], …

Descartes Meta-Model (DMM)

- Architecture-level modeling language for **self-aware** run-time systems management of modern IT systems, infrastructures and services
- Main Goal: Provide Quality-of-Service (QoS) guarantees
  - Performance (current focus)
    - Response time, throughput, scalability and efficiency
  - Or more generally, dependability
    - Including also availability, reliability and security aspects

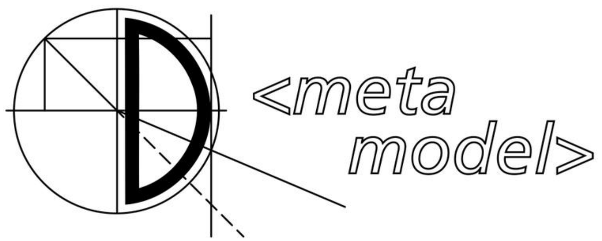

Descartes Meta-Model (DMM)

Collection of several meta-models each focusing on different system aspects

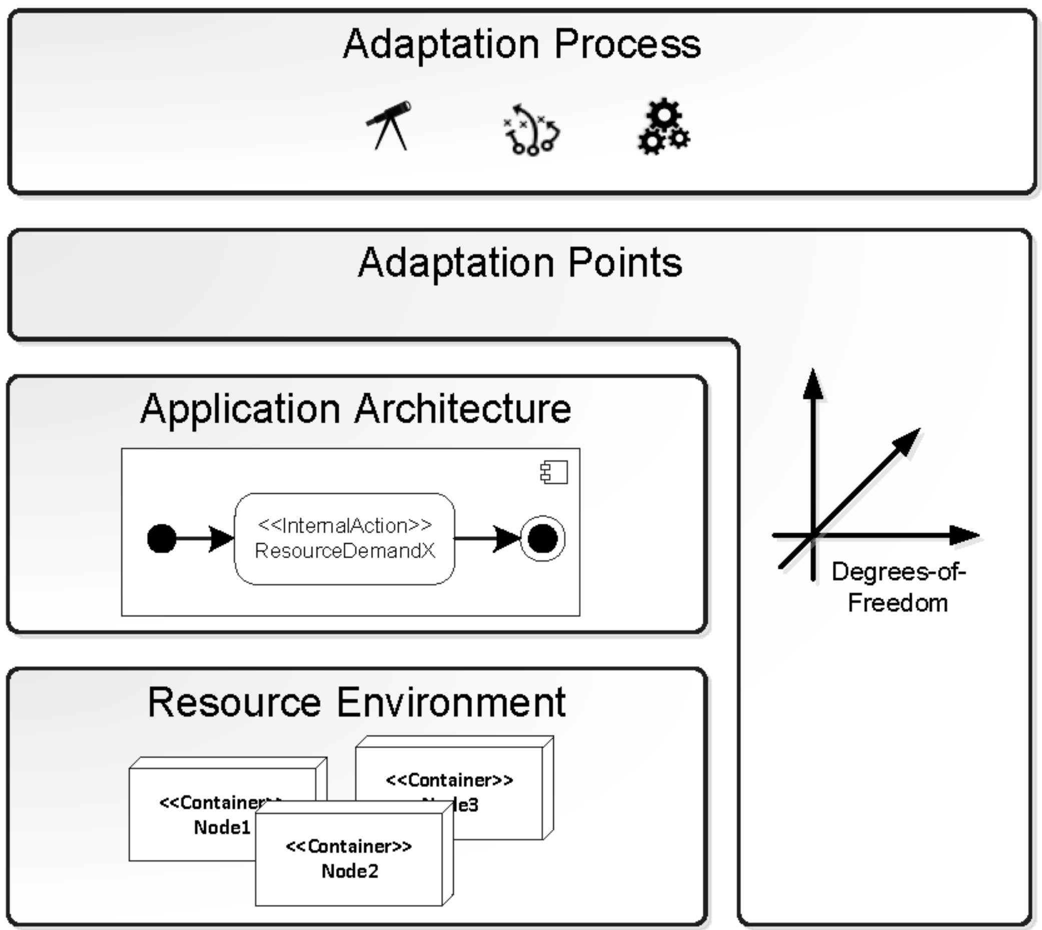

Dynamic Model Instance Composition

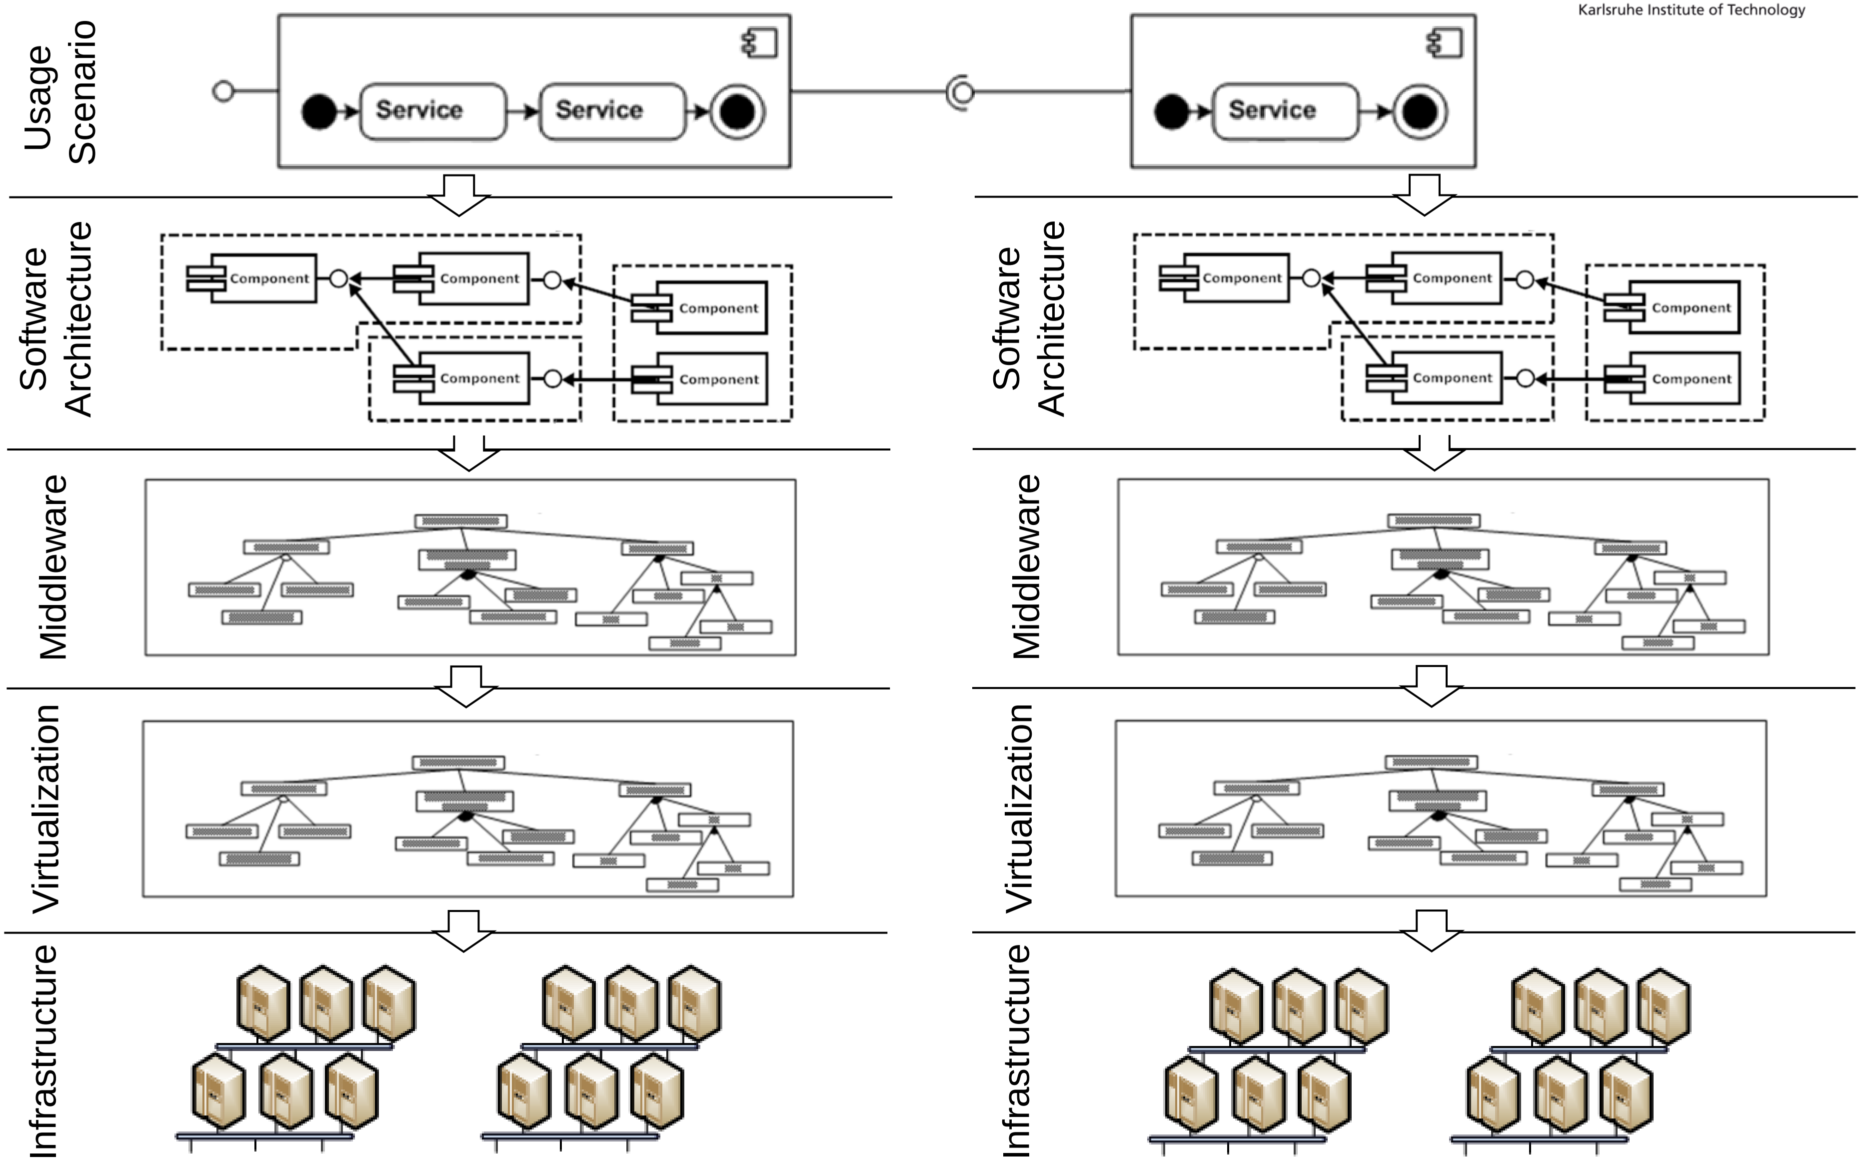

Tailored Model-to-Model Transformations

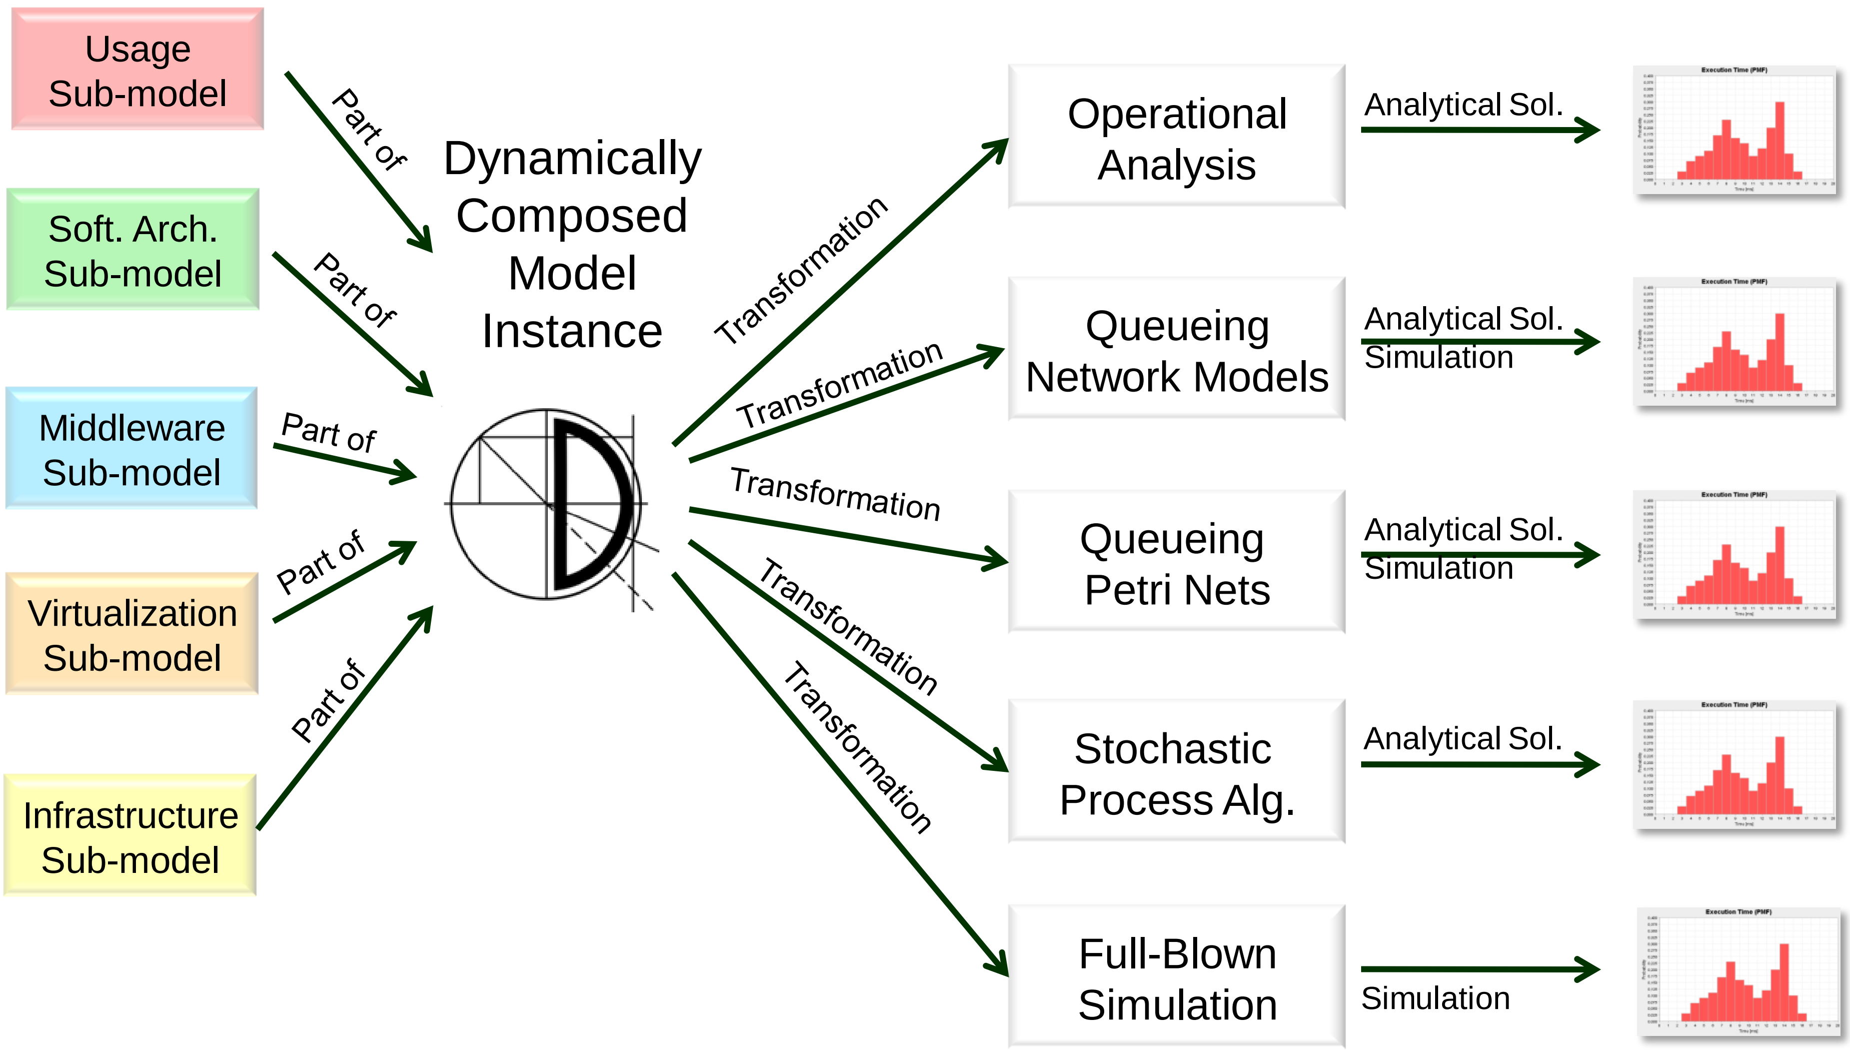

Example 1: Simple Bounds Analysis

$$R \geq \max\left[N \times \max\{D_i\}, \sum_{i=1}^{K} D_i \right]$$

$$X_0 \leq \min\left[\frac{1}{\max\{D_i\}}, \frac{N}{\sum_{i=1}^{K} D_i} \right]$$

$$\frac{N}{\max\{D_i\}[K+N-1]} \leq X_0 \leq \frac{N}{avg\{D_i\}[K+N-1]}$$

Example 2: Product-Form Queueing Network

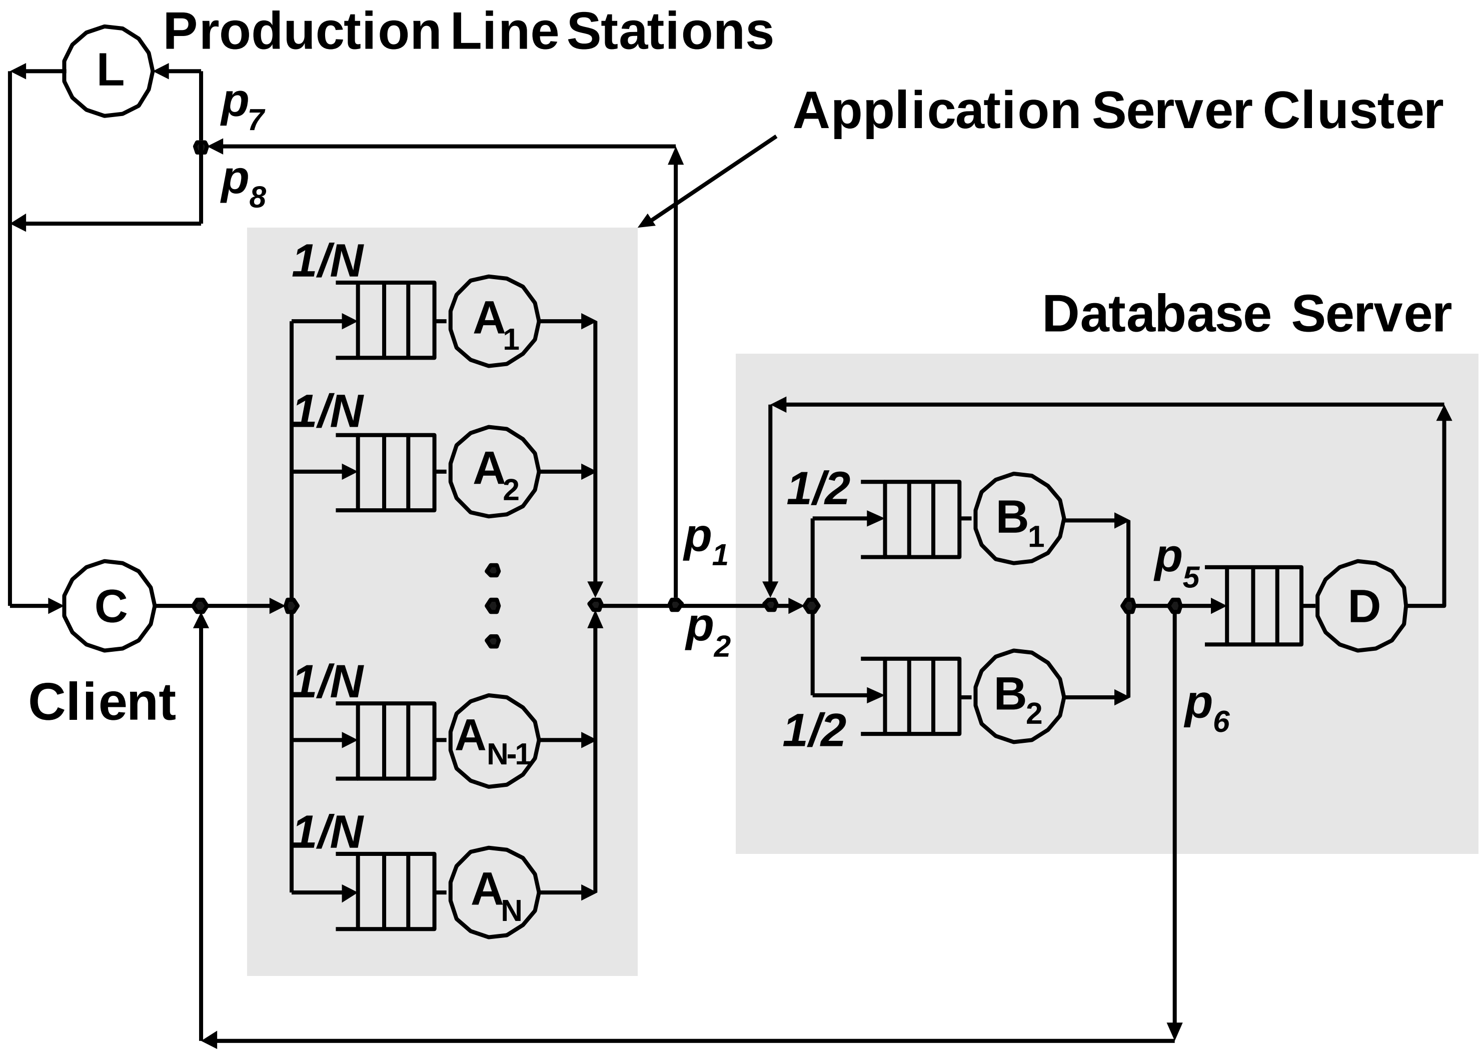

Example 3: Layered Queueing Network

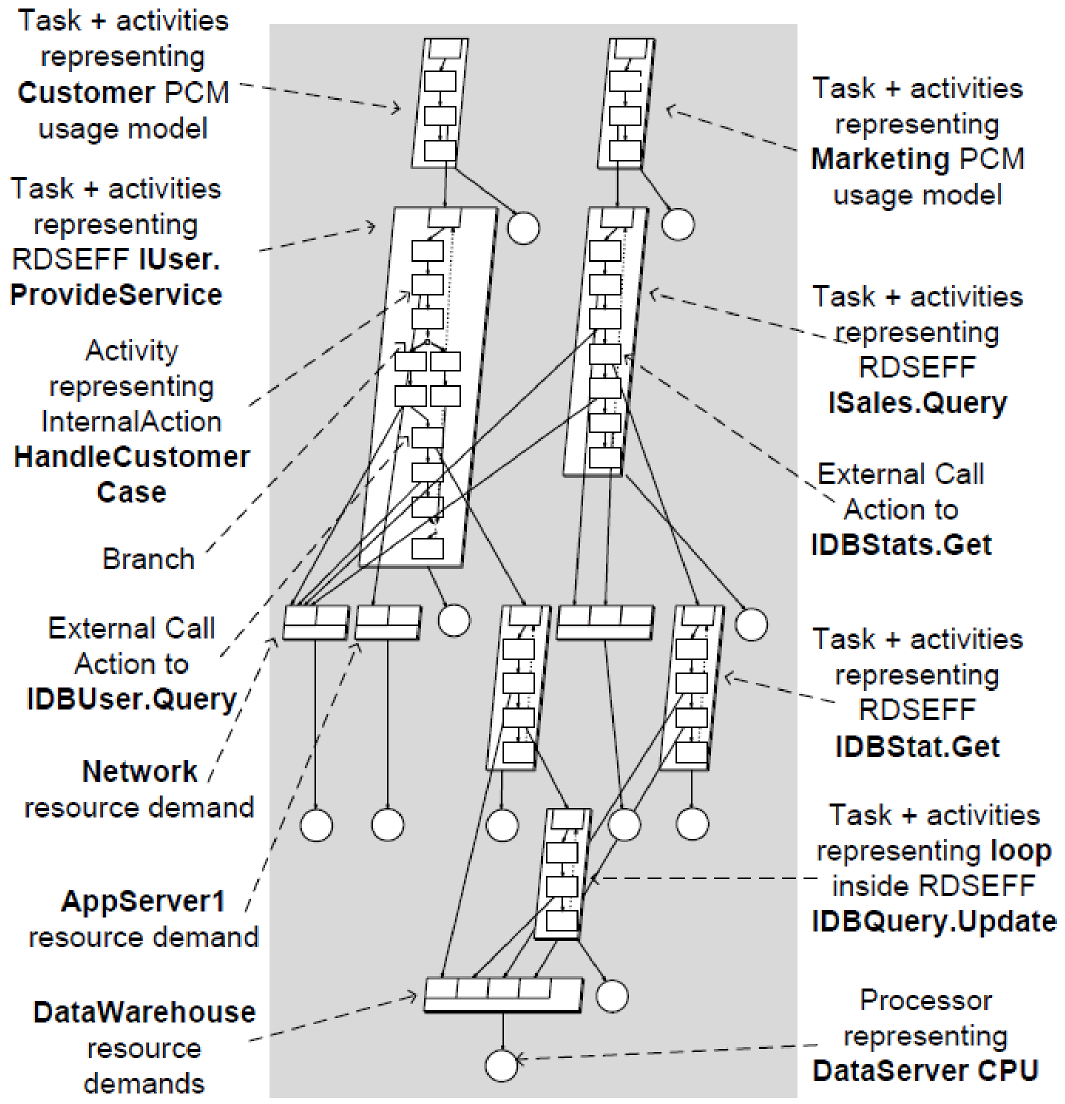

Example 3a: Layered Queueing Network

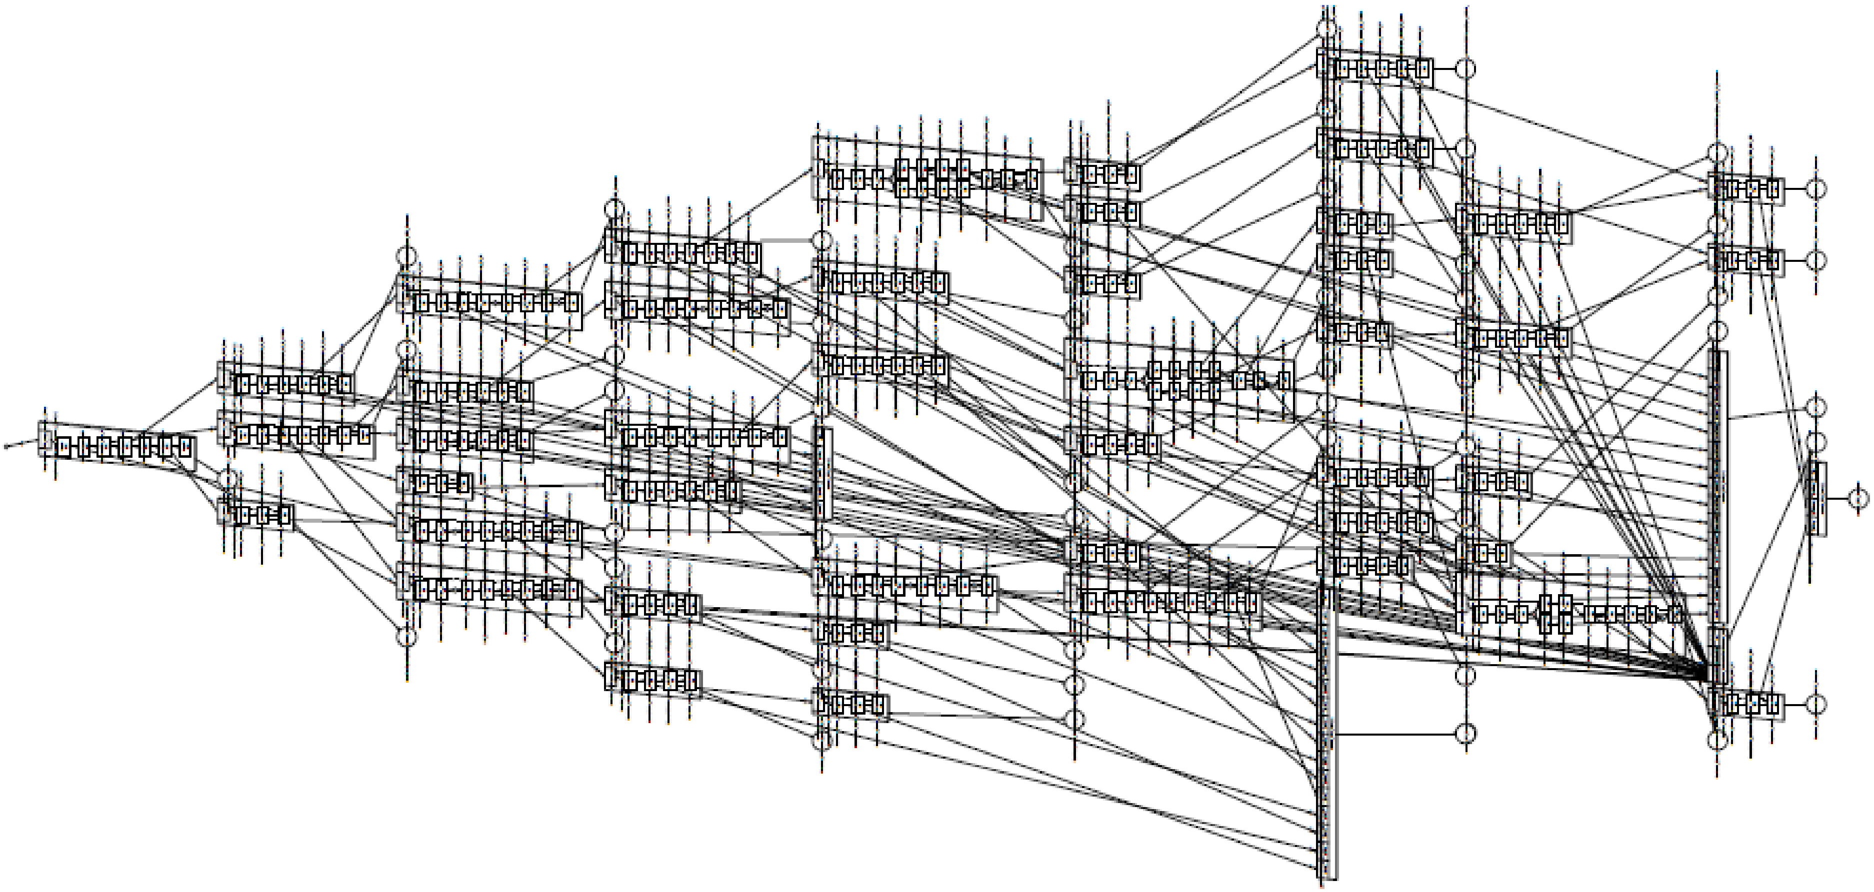

Example 4: Queueing Petri Net

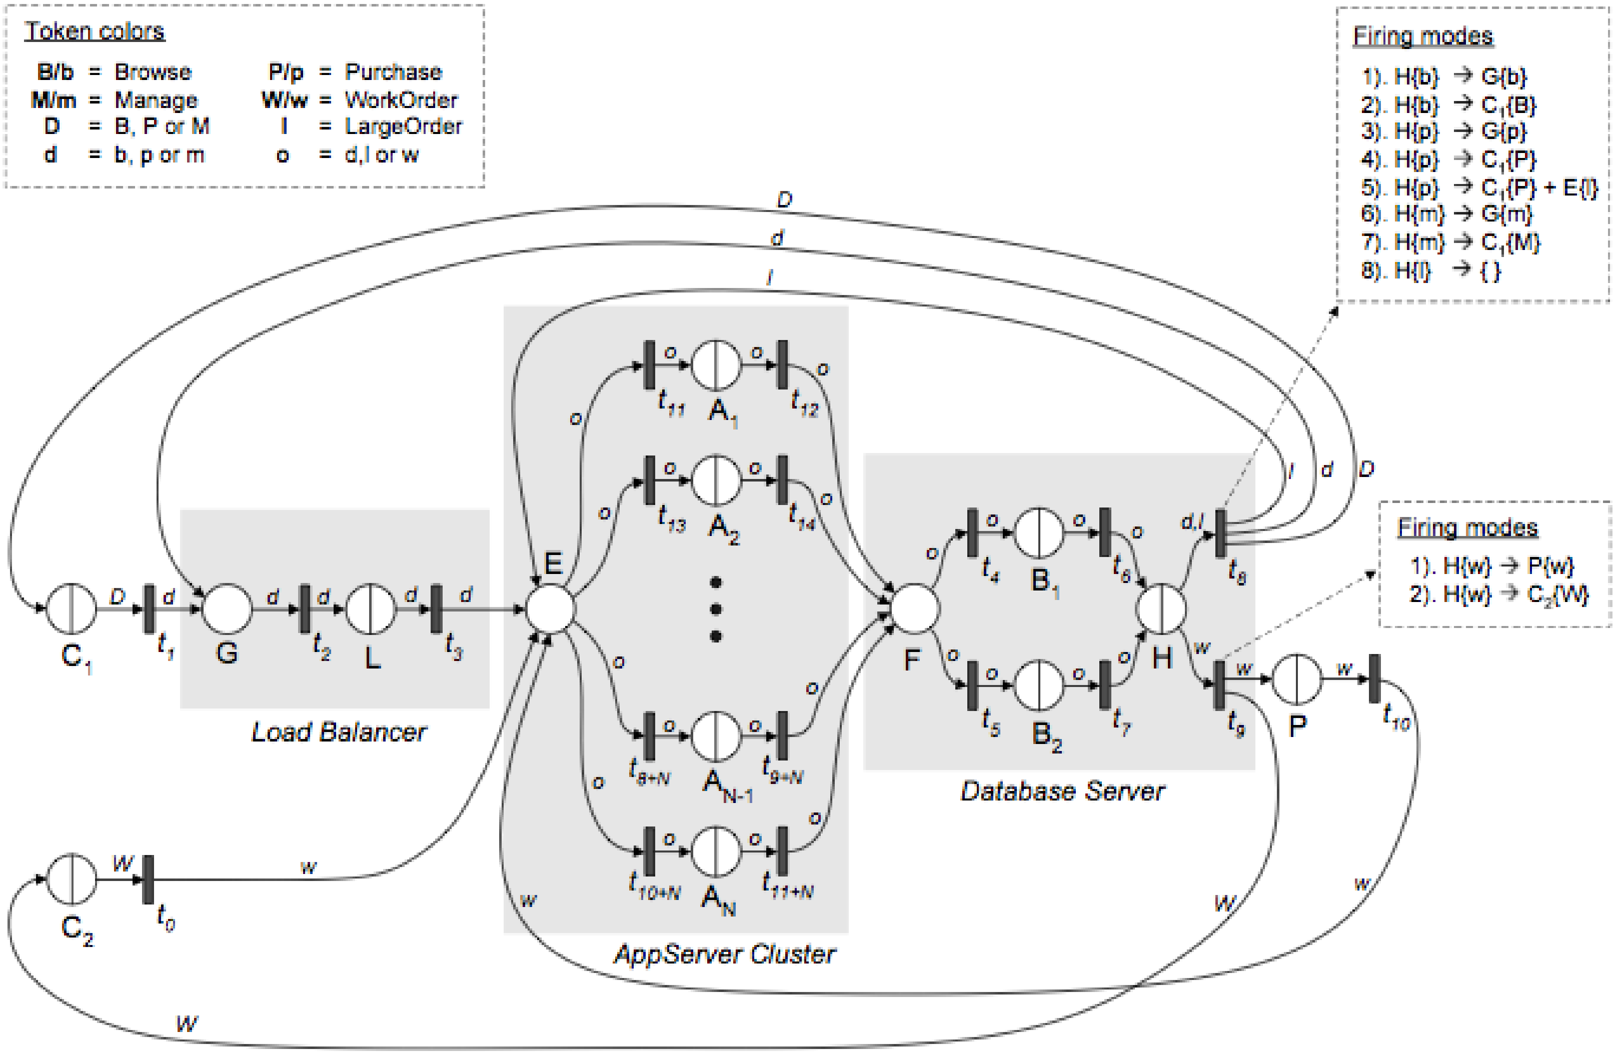

Example 4a: Queueing Petri Net

Further details in:

- P. Meier, S. Kounev and H. Koziolek. **Automated Transformation of Palladio Component Models to Queueing Petri Nets**. In *In 19th IEEE/ACM International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2011)*, Singapore, July 25-27 2011.

System Control Loop

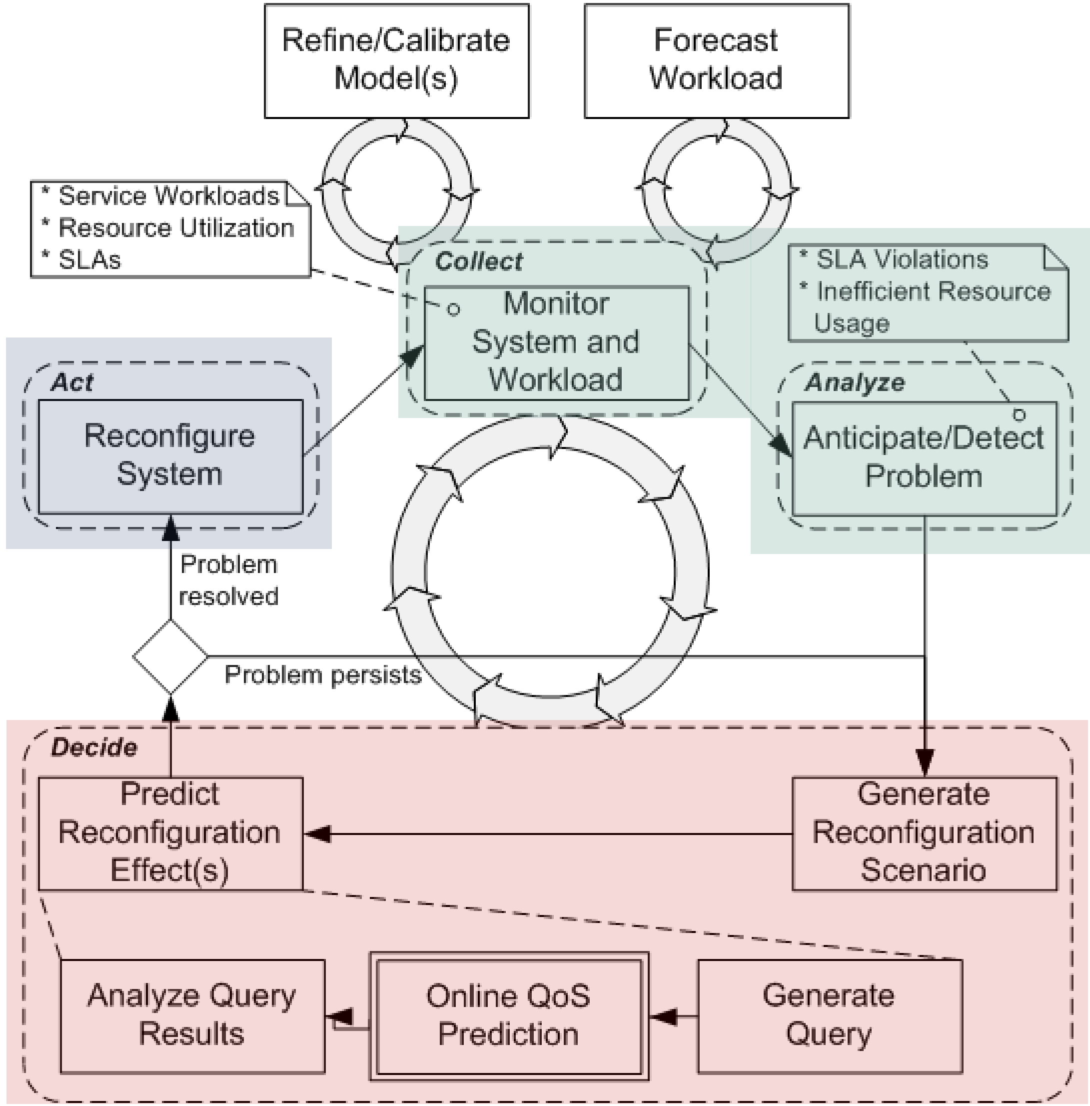

Part 1

Part 2

Part 3

DMM Technical Report

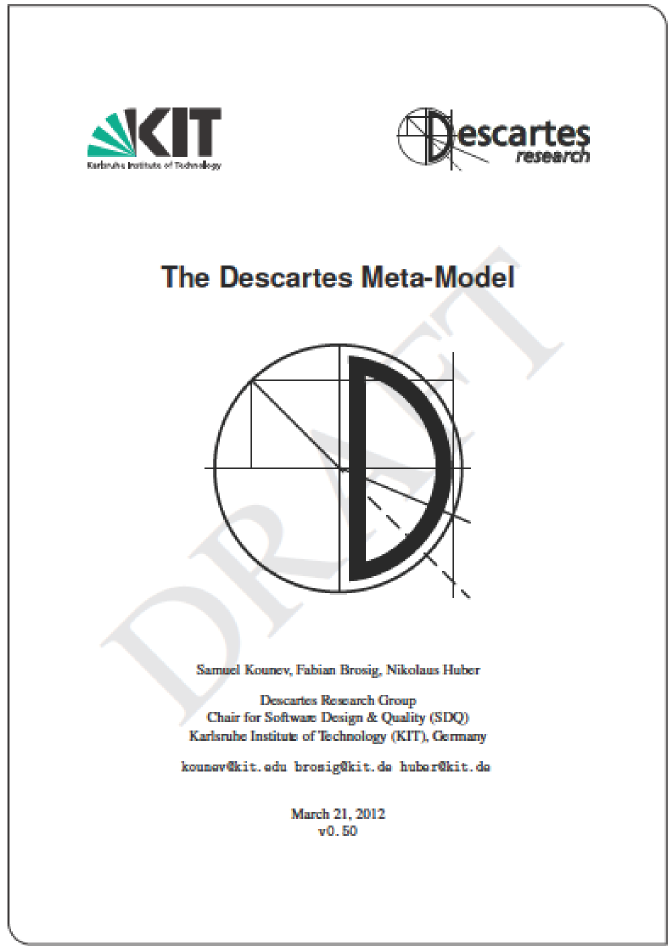
The Descartes Meta-Model

DRAFT

Samuel Kounev, Fabian Brosig, Nikolaus Huber

Descartes Research Group
Chair for Software Design & Quality (SDQ)
Karlsruhe Institute of Technology (KIT), Germany

kounev@kit.edu brosig@kit.de huber@kit.de

March 21, 2012
v0.50

www.descartes-research.net

Agenda

- Motivation
- Approach & Methodology
- **Exemplary Results**
- Vision
- Conclusion

DMM: Resource Environment

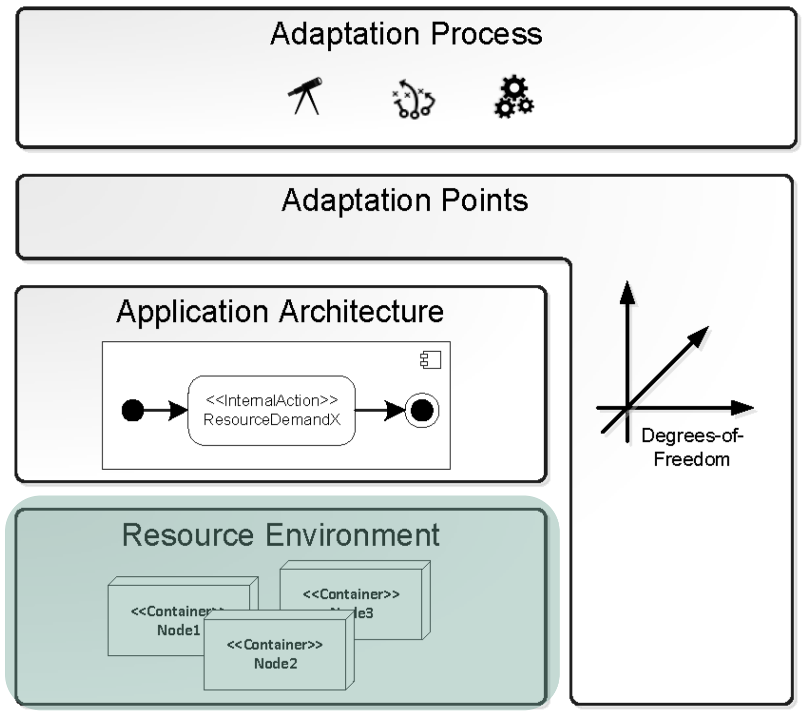

Further details in:

- N. Huber, F. Brosig and S. Kounev. **Modeling Dynamic Virtualized Resource Landscapes**. In Proceedings of the 8th ACM SIGSOFT International Conference on the Quality of Software Architectures (QoSA 2012), Bertinoro, Italy, June 25-28, 2012.

Resource Landscape: Example

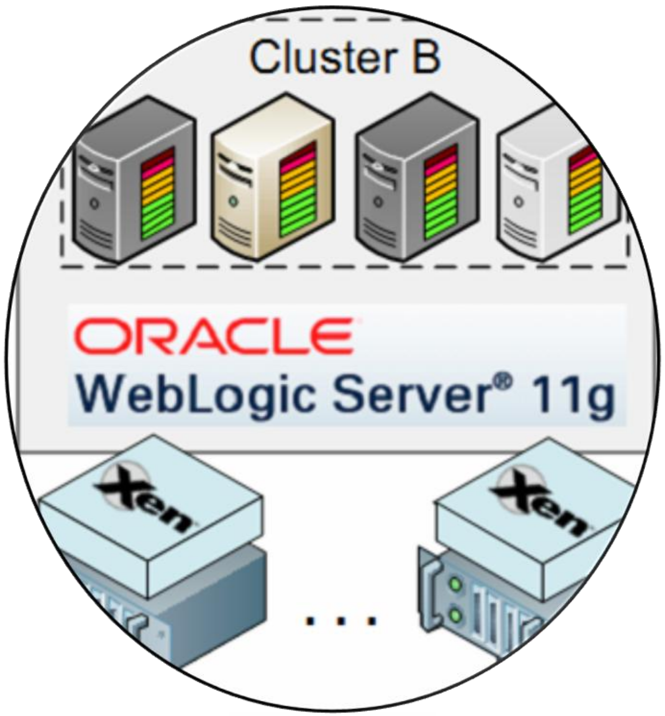

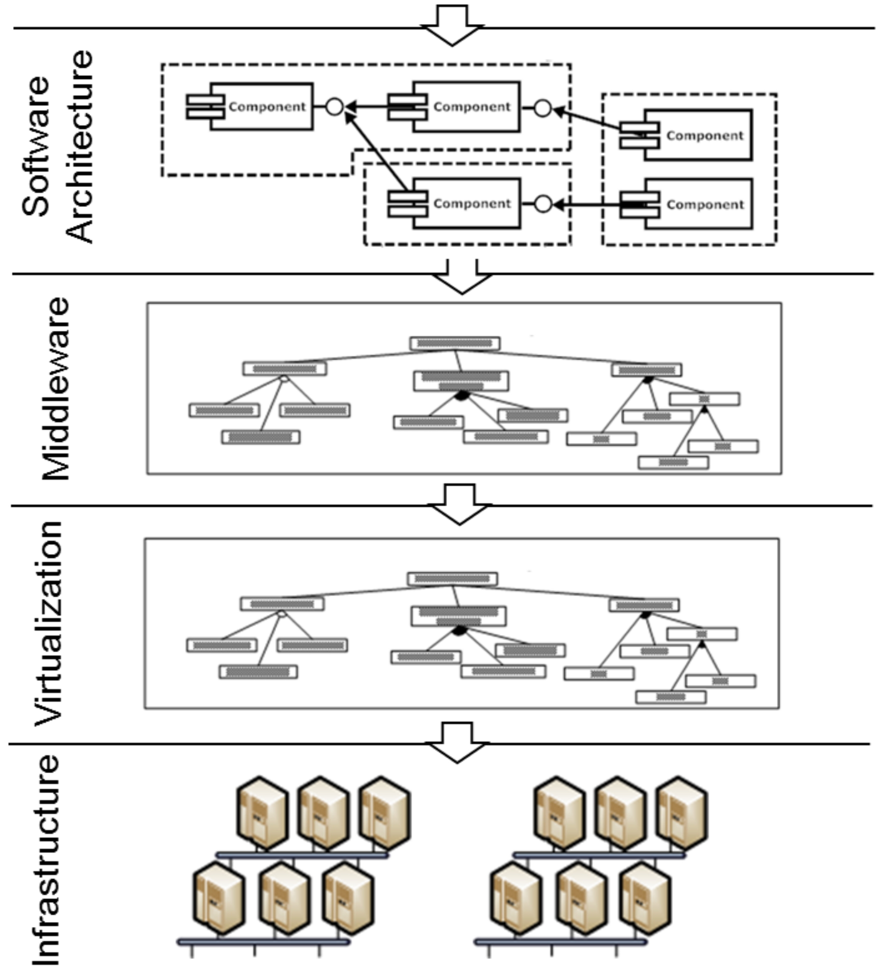

Resource Landscape: Layers

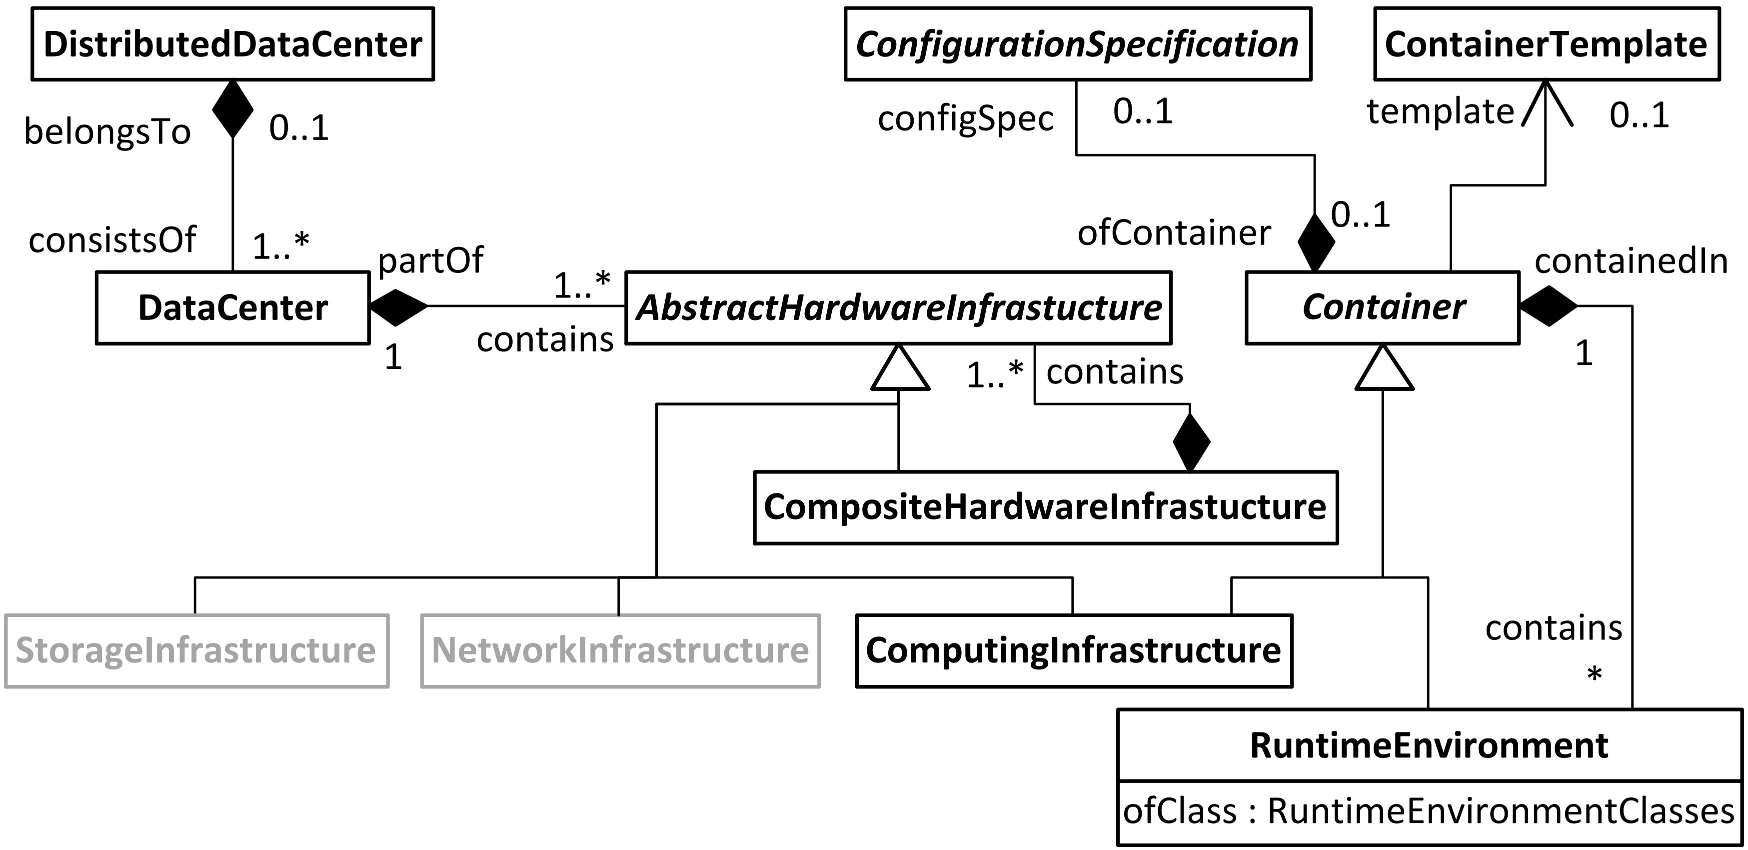

Resource Landscape: Container Templates and Runtime Environment Classes

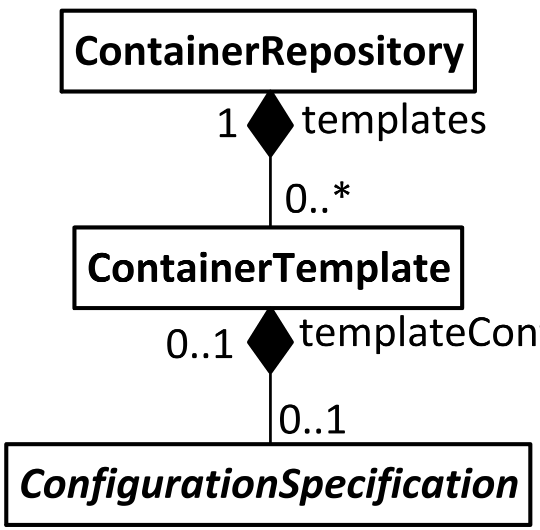

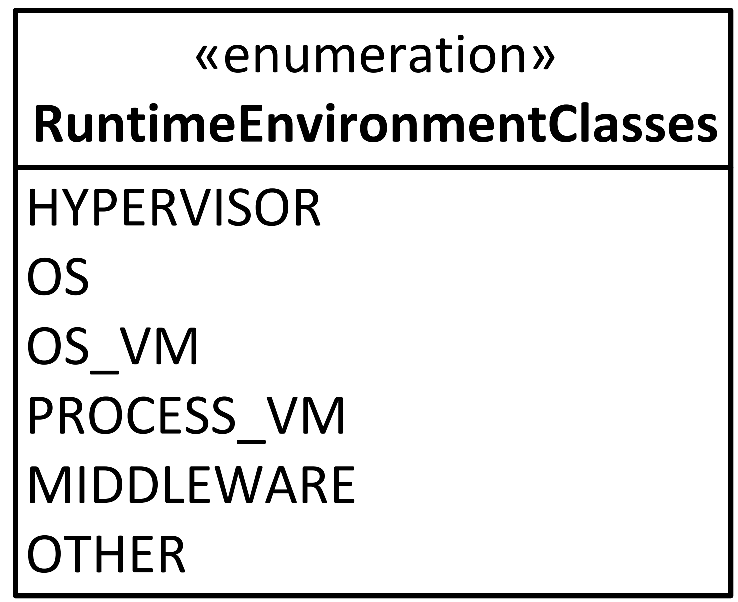

Resource Landscape: Configuration Specification

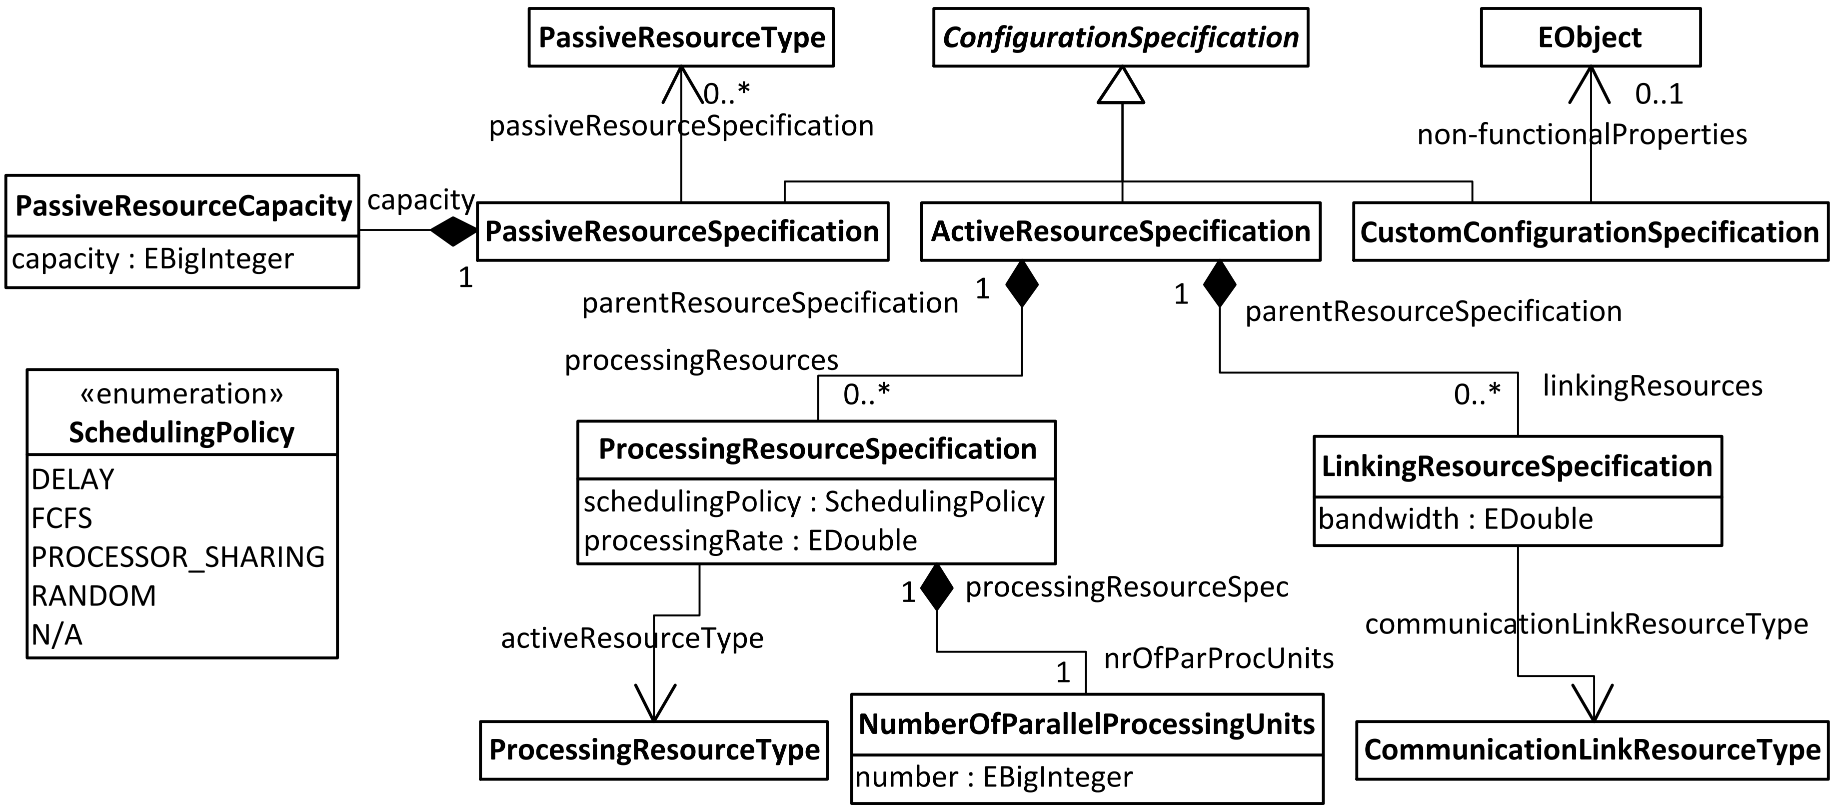

Example: Virtualization Layer

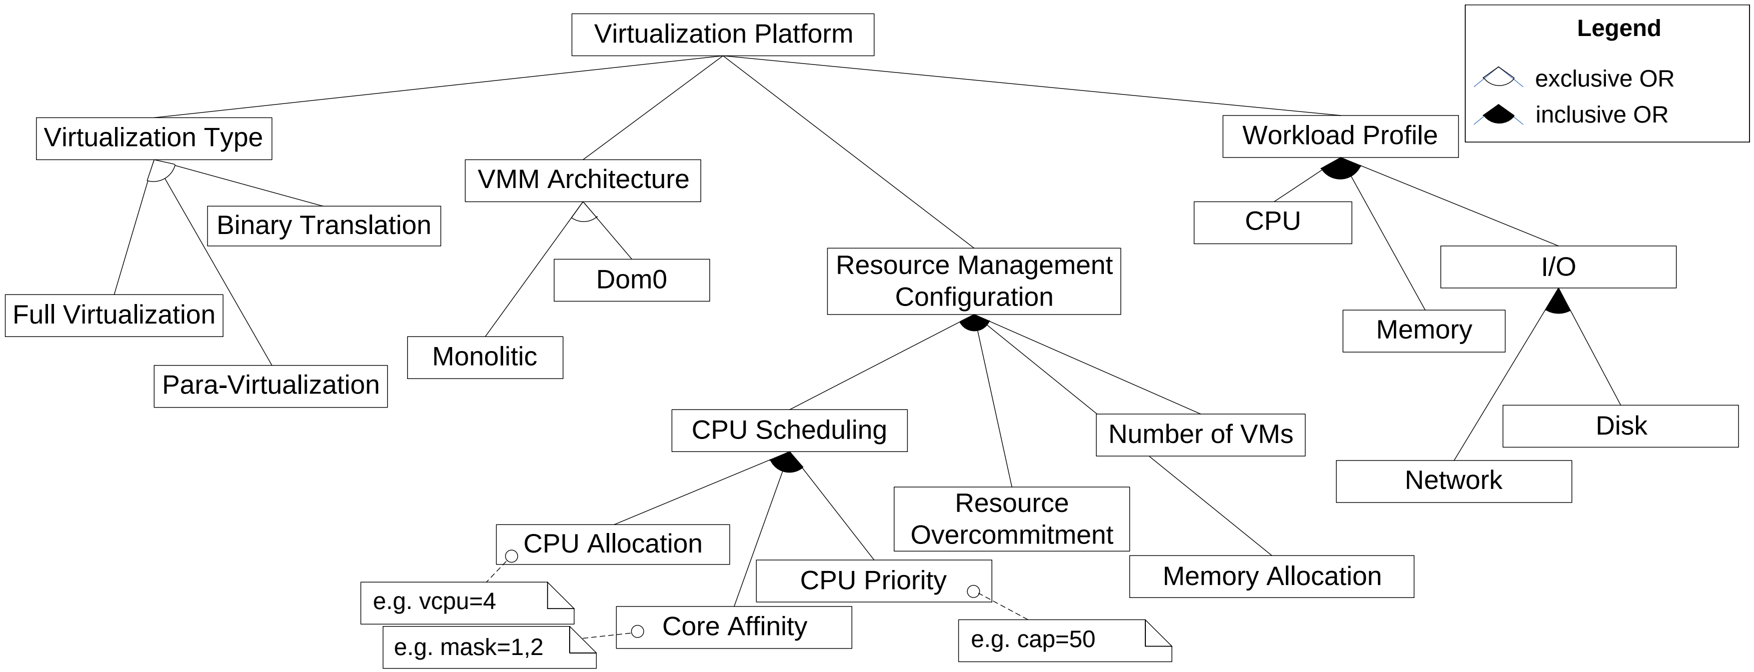

DMM: Adaptation Points

Karlsruhe Institute of Technology

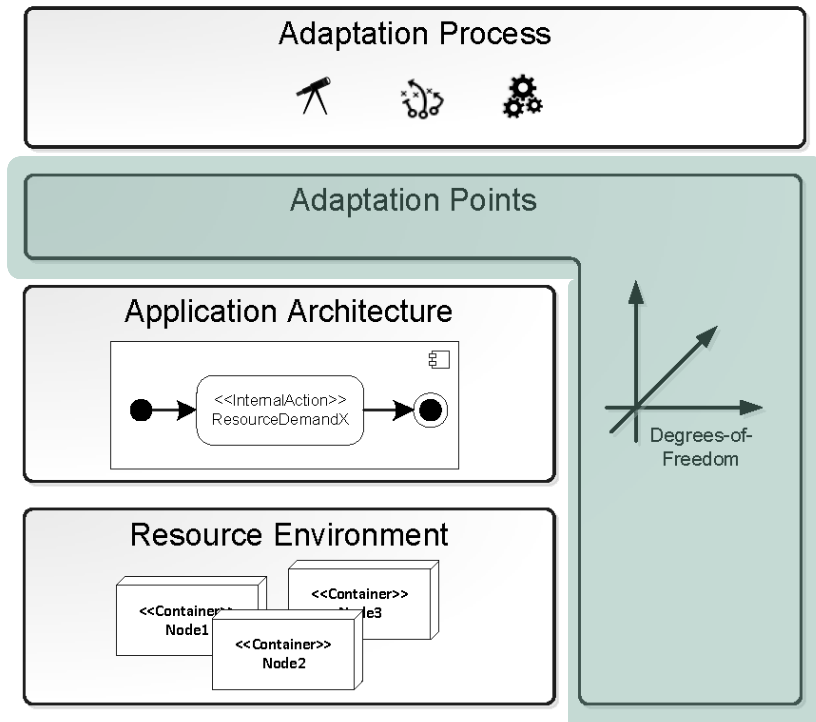

Further details in:

- N. Huber, F. Brosig and S. Kounev. **Modeling Dynamic Virtualized Resource Landscapes**. In Proceedings of the 8th ACM SIGSOFT International Conference on the Quality of Software Architectures (QoSA 2012), Bertinoro, Italy, June 25-28, 2012.

Adaptation Points: Examples

- Scaling Resources

- Replicating VMs, Migrating VMs

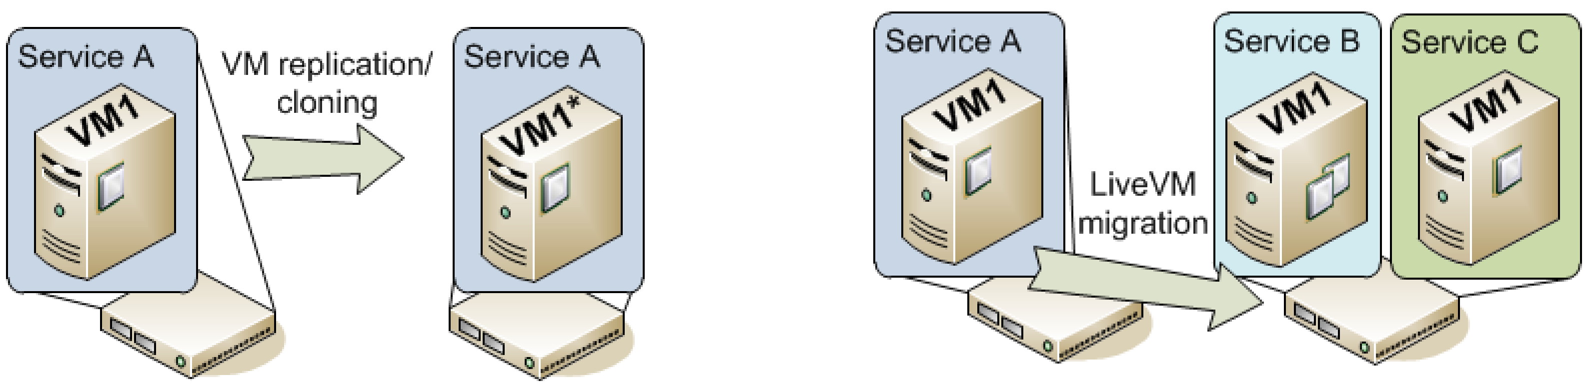

Adaptation Points

- Specification of valid system configurations
- “Decorator” model of static view

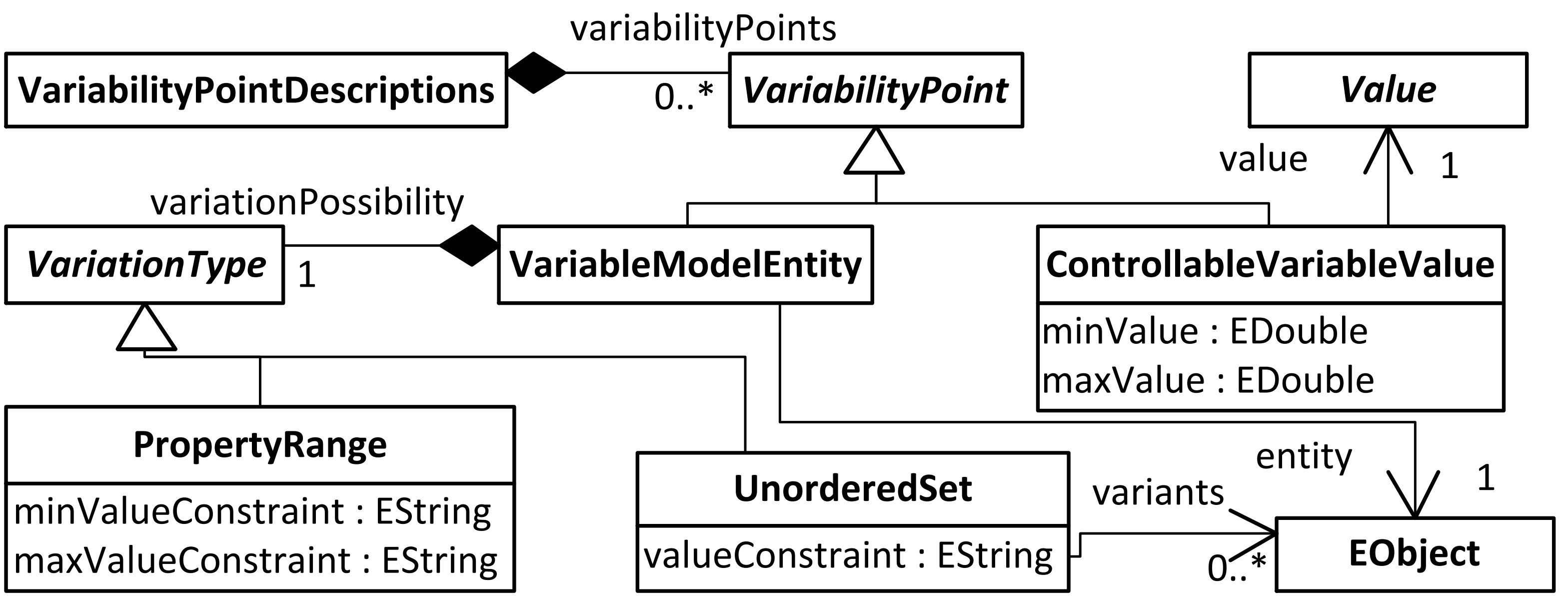

[HBK12]

DMM: Application Architecture

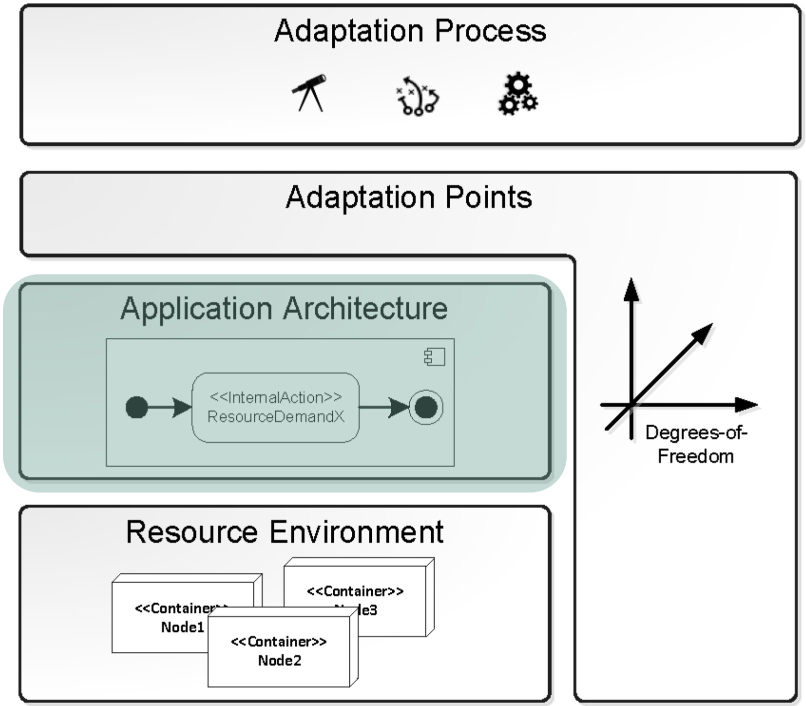

Further details in:

- F. Brosig, N. Huber, and S. Kounev. **Modeling Parameter and Context Dependencies in Online Architecture-Level Performance Models**. In *Proceedings of the 15th ACM SIGSOFT International Symposium on Component Based Software Engineering (CBSE 2012), Bertinoro, Italy*, June 26-28, 2012.

Modeling the Application Level

- Service Behavior Abstractions
  for Different Levels of Granularity

- Probabilistic Parameter Dependencies

- Deployment-Specific Resource Demands /
  Response Times

Online Performance Prediction Scenario

Karlsruhe Institute of Technology

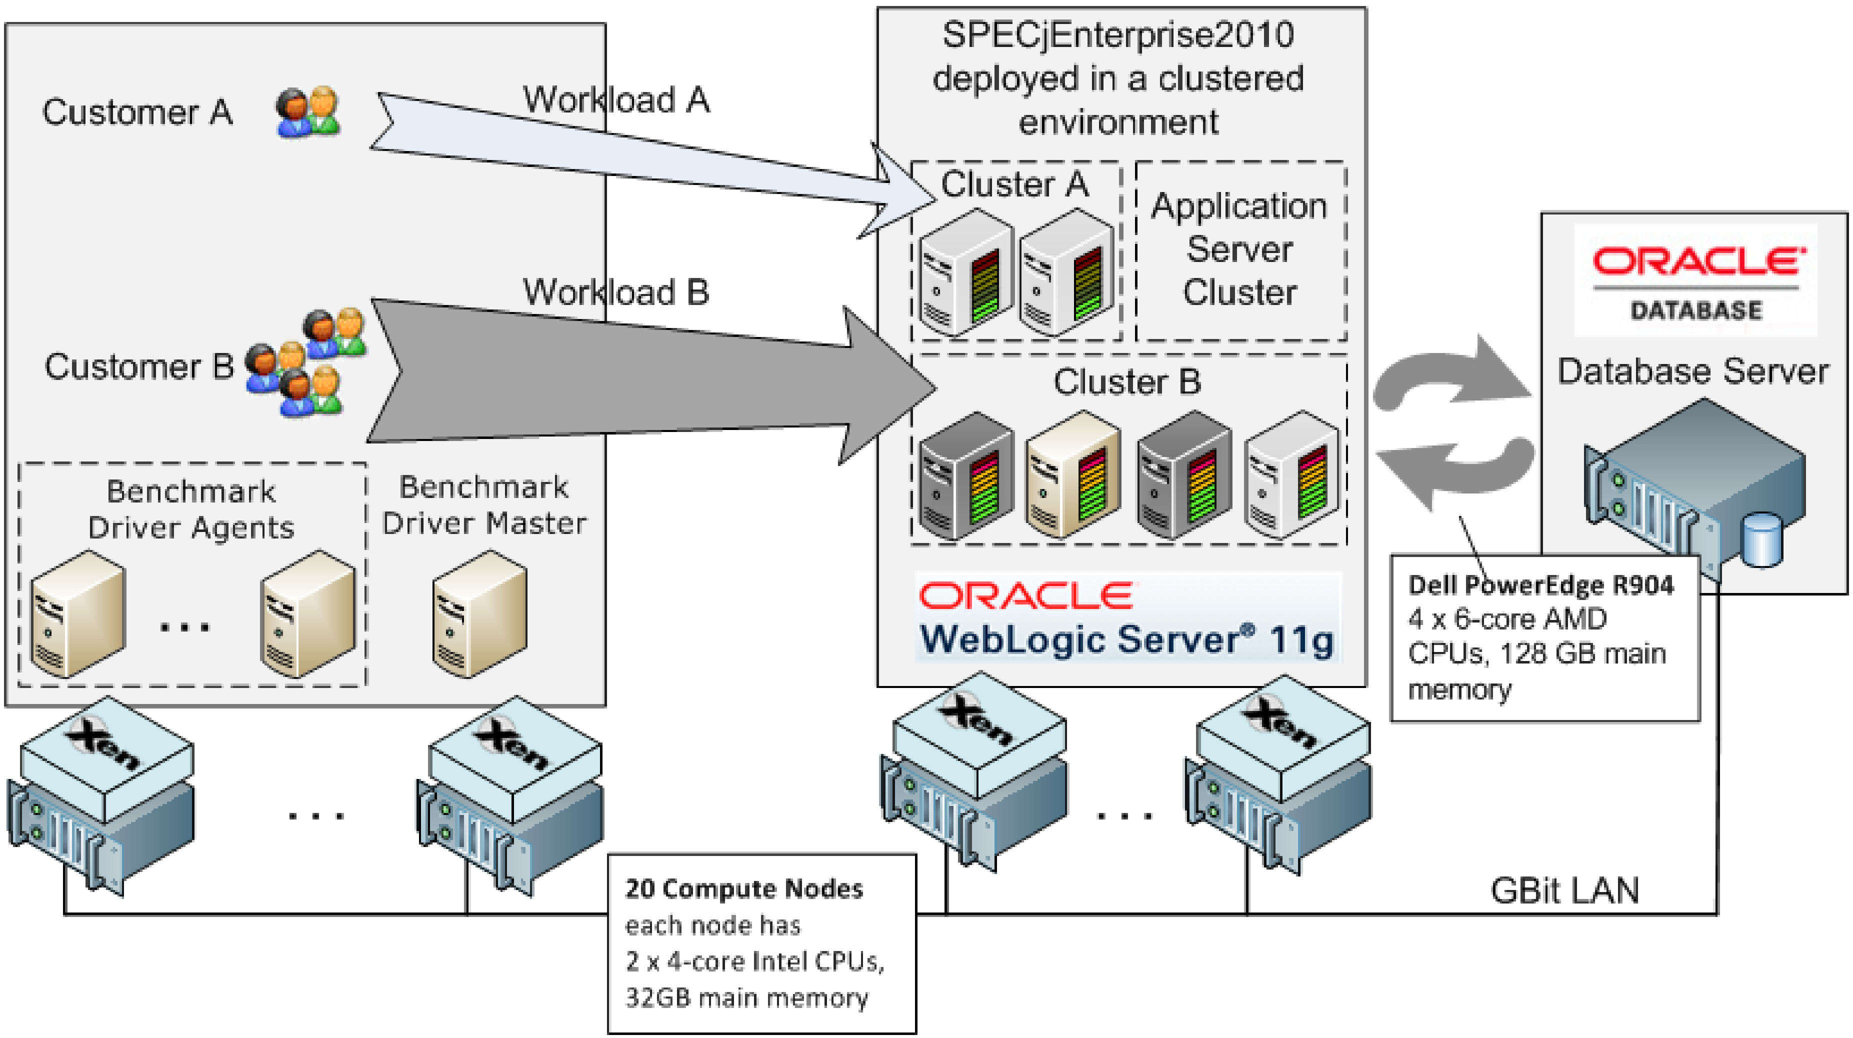

Service Behavior Abstractions

- BlackBoxBehavior
  - No information about resources, resource demands, control flow, call frequencies,…
- CoarseGrainedBehavior
  - Information at component boundary level (external services, resource consumption,...)
- FineGrainedBehavior
  - Information about component-internals (control flow, resource demands, parametric dependencies,...)

Service Behavior Abstractions

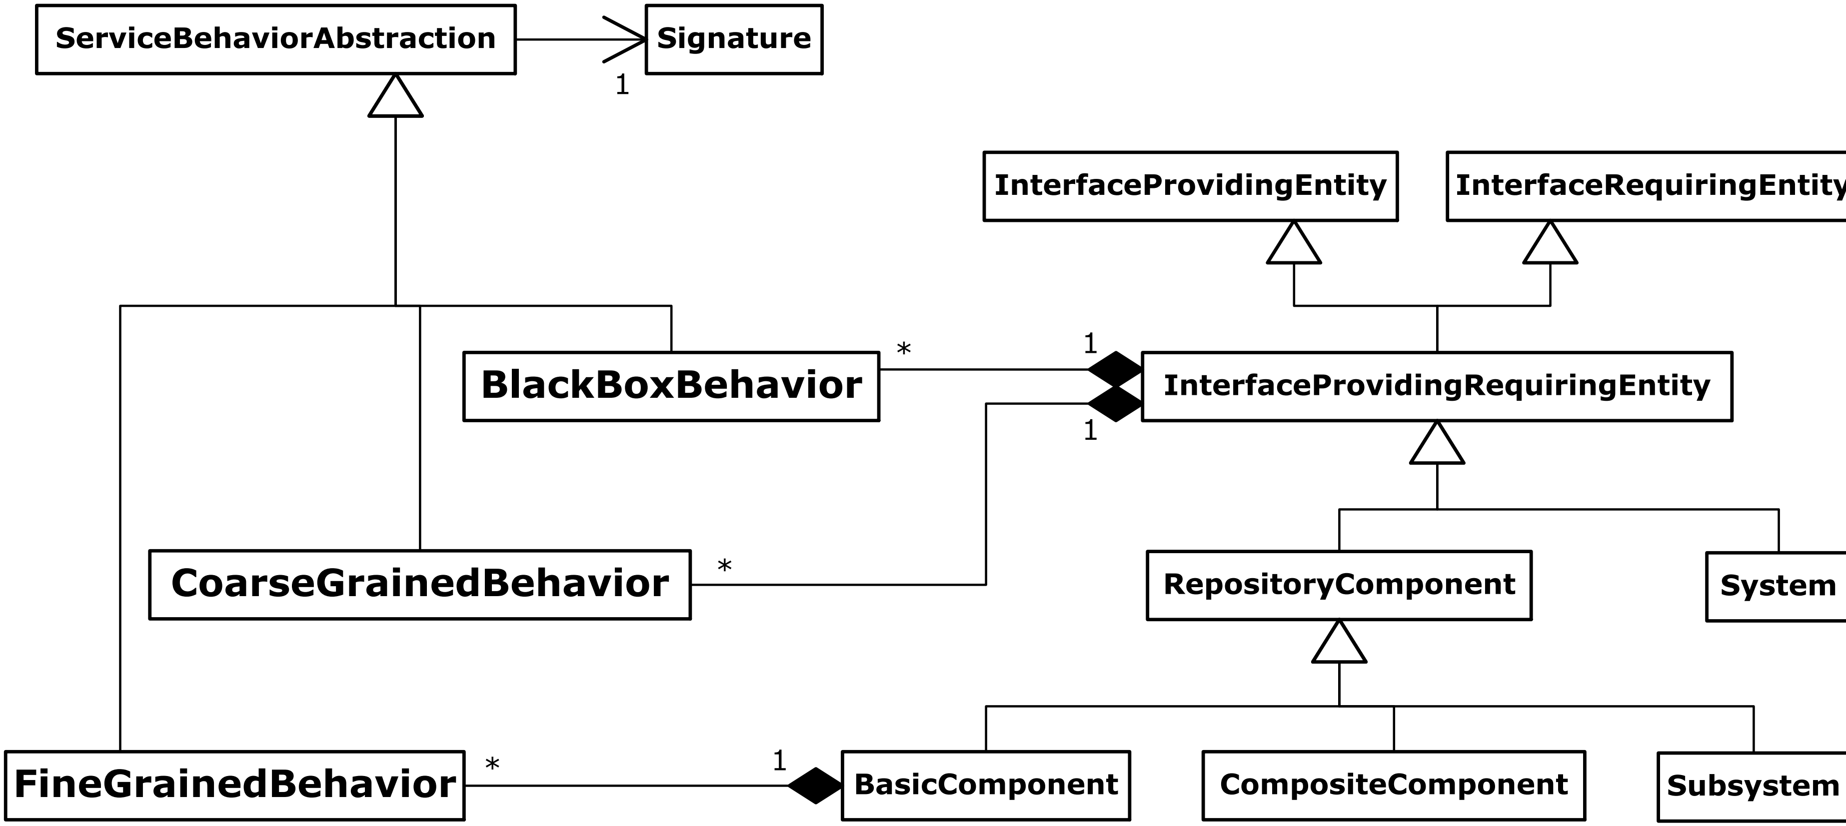

Service Behavior Abstractions

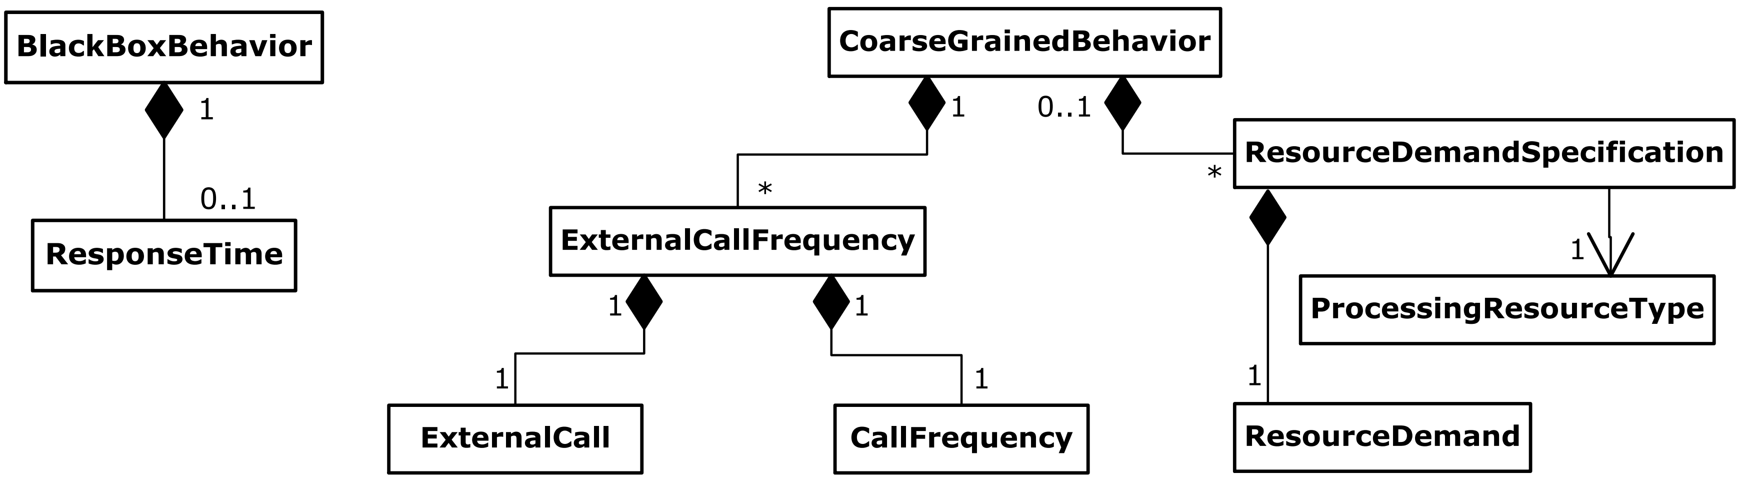

Service Behavior Abstractions

- Control flow abstraction

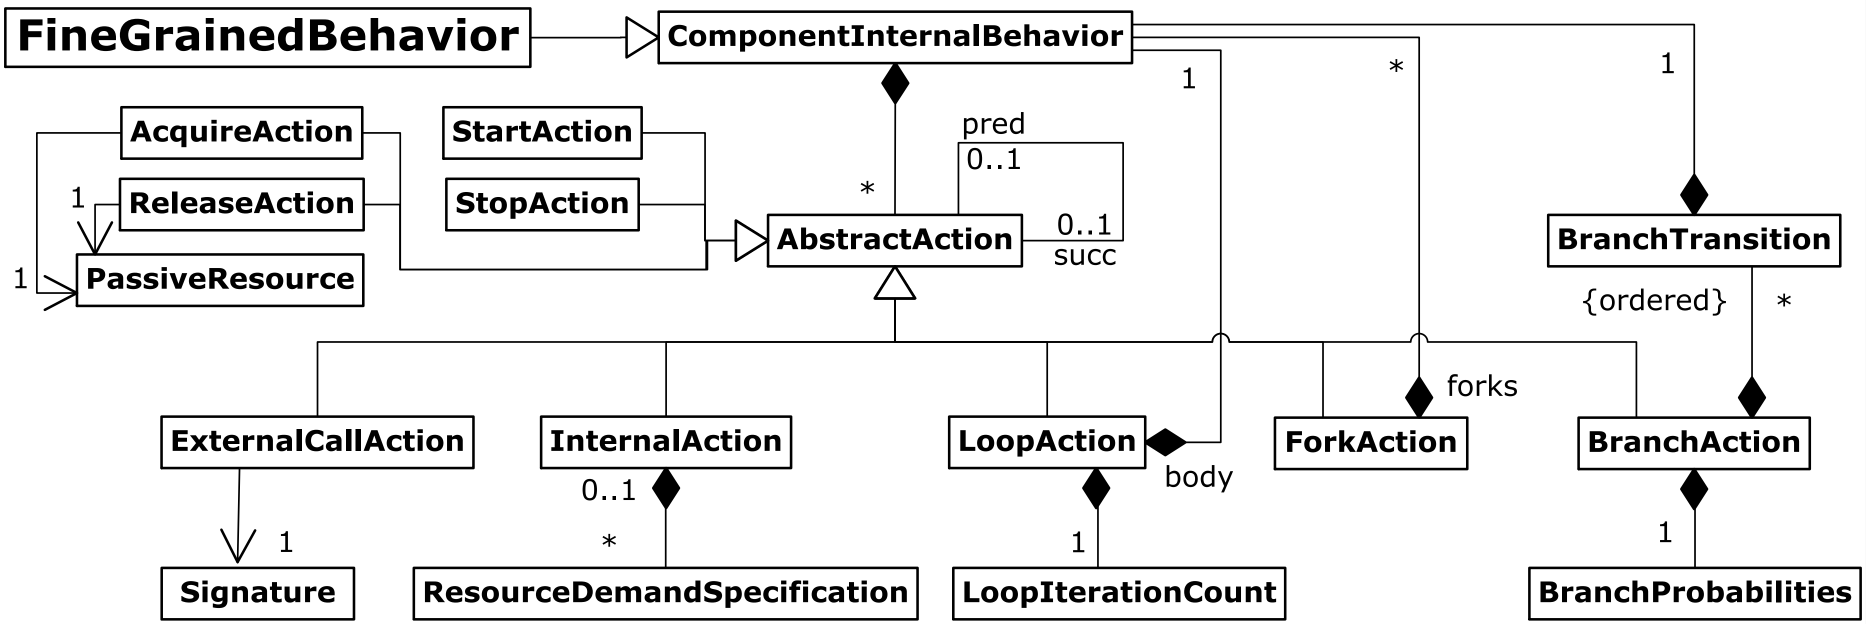

Probabilistic Parameter Dependencies

- Characterize dependencies statistically

- Influencing parameters
  - Service input parameters
  - Return parameters of external service calls

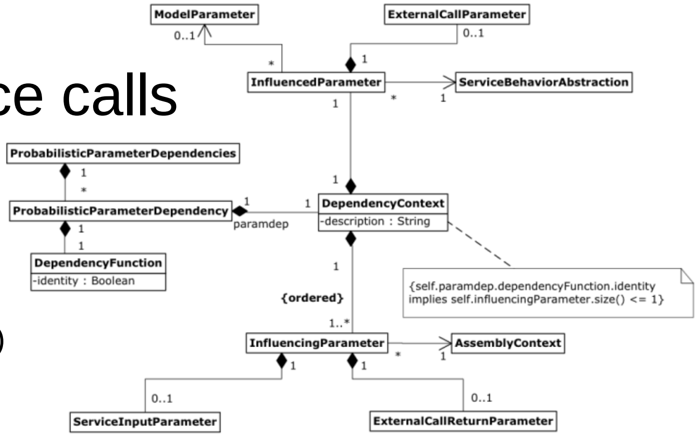

- Influenced quantities
  - Loop iteration numbers (FineGrainedBehavior)
  - Branch probabilities (FineGrainedBehavior)
  - Call frequencies (CoarseGrainedBehavior)
  - Resource demands (FineGrainedBehavior, CoarseGrainedBehavior)
  - Response times (BlackBoxBehavior)
  - Input parameter of ext. service call (FineGrainedBehavior, CoarseGrainedBehavior)

Model Variables - Metamodel

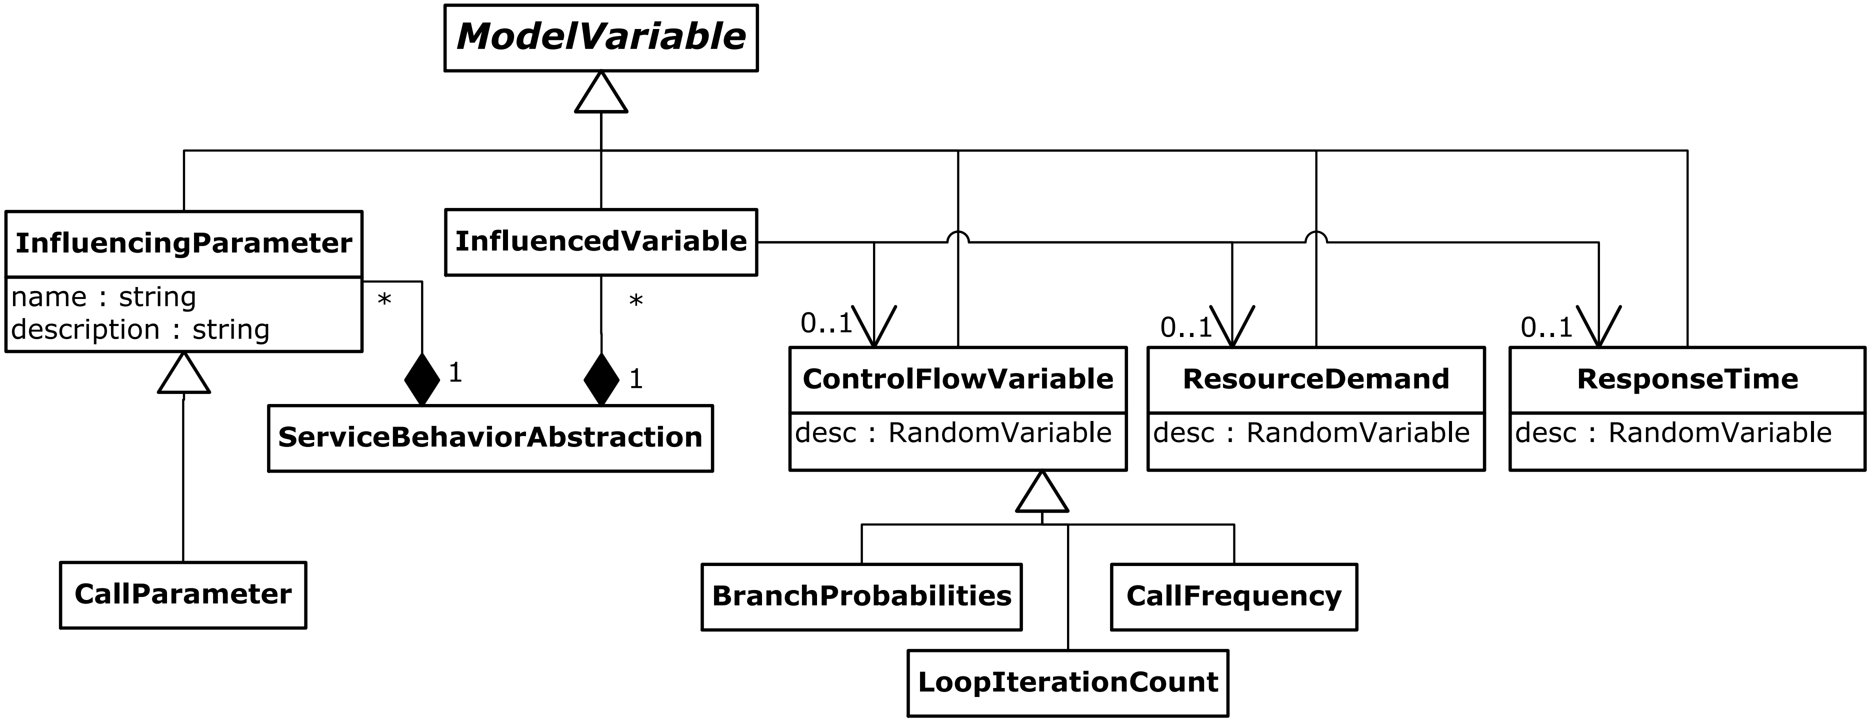

Probabilistic Parameter Dependencies

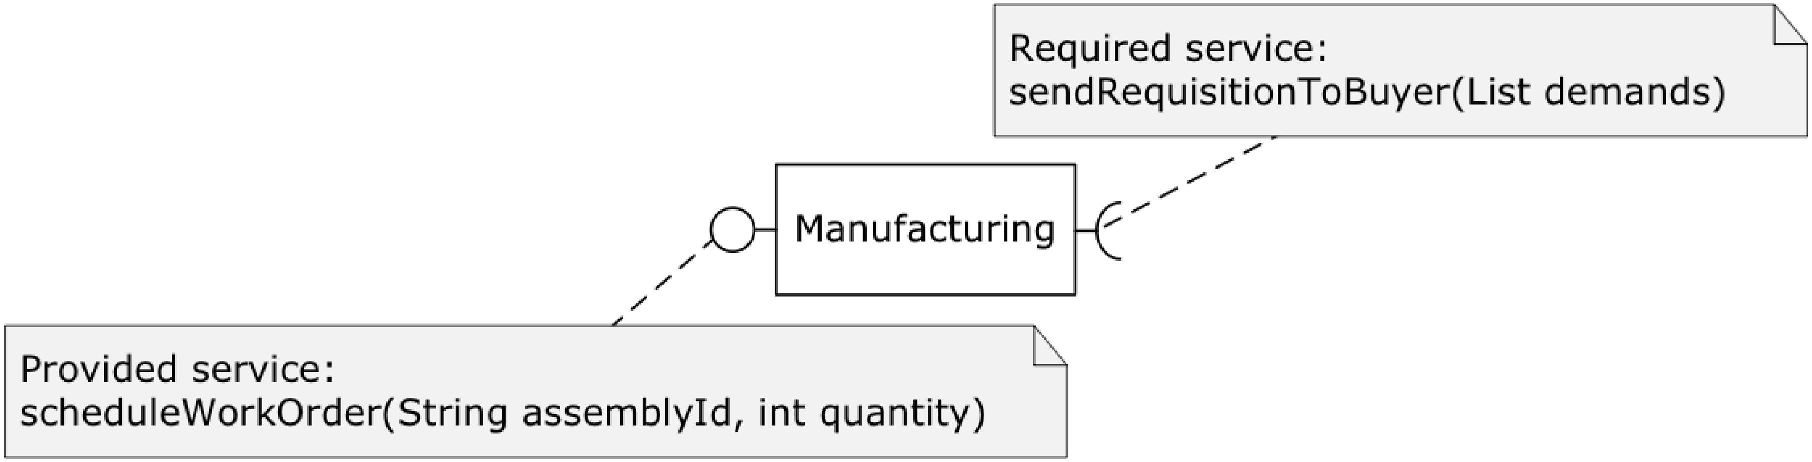

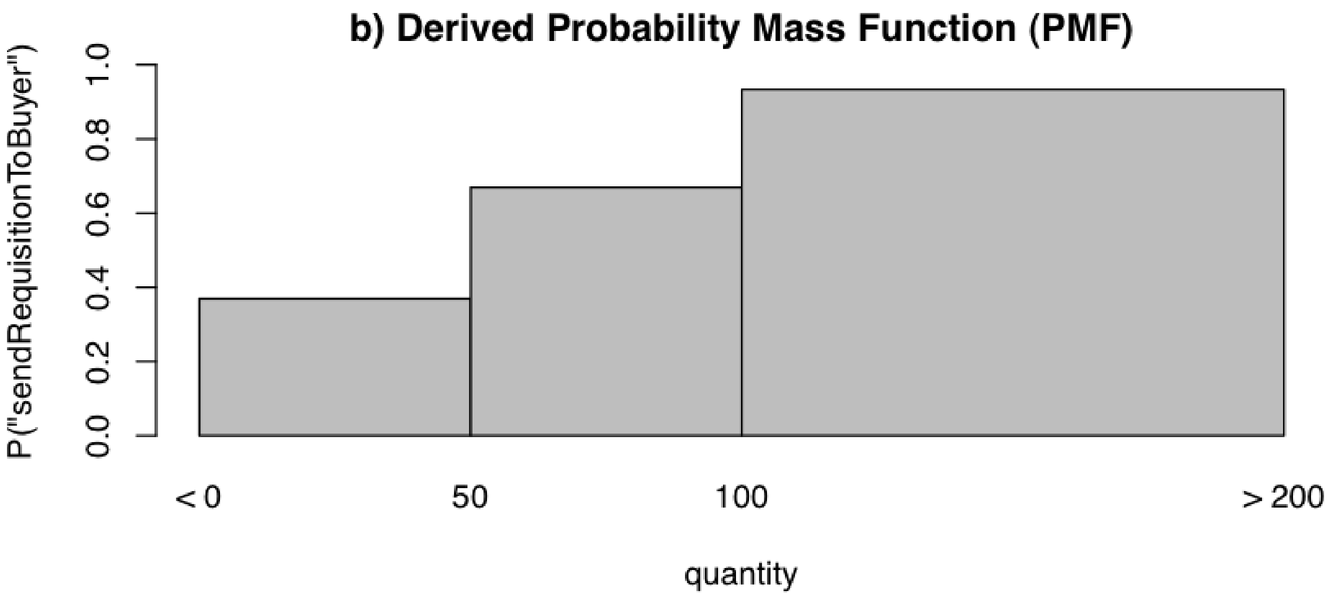

Scope - Motivation

- Customer-specific application server cluster
  - Scenario: Replicate Server Instance

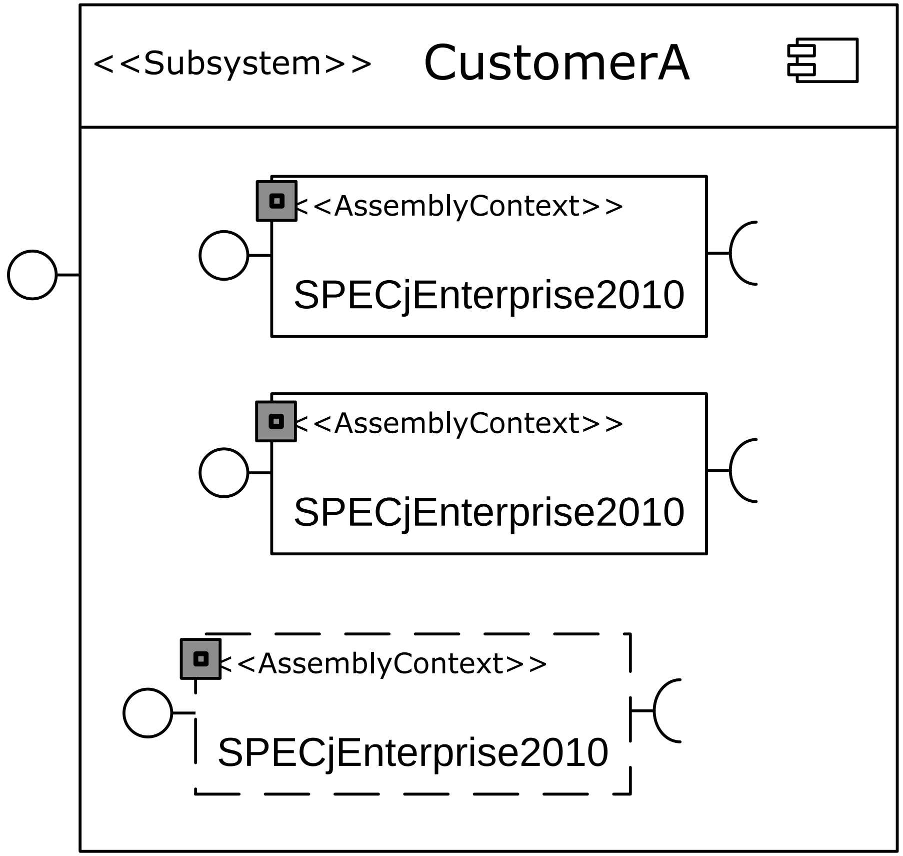

Scope - Motivation

- Customer-specific application server cluster

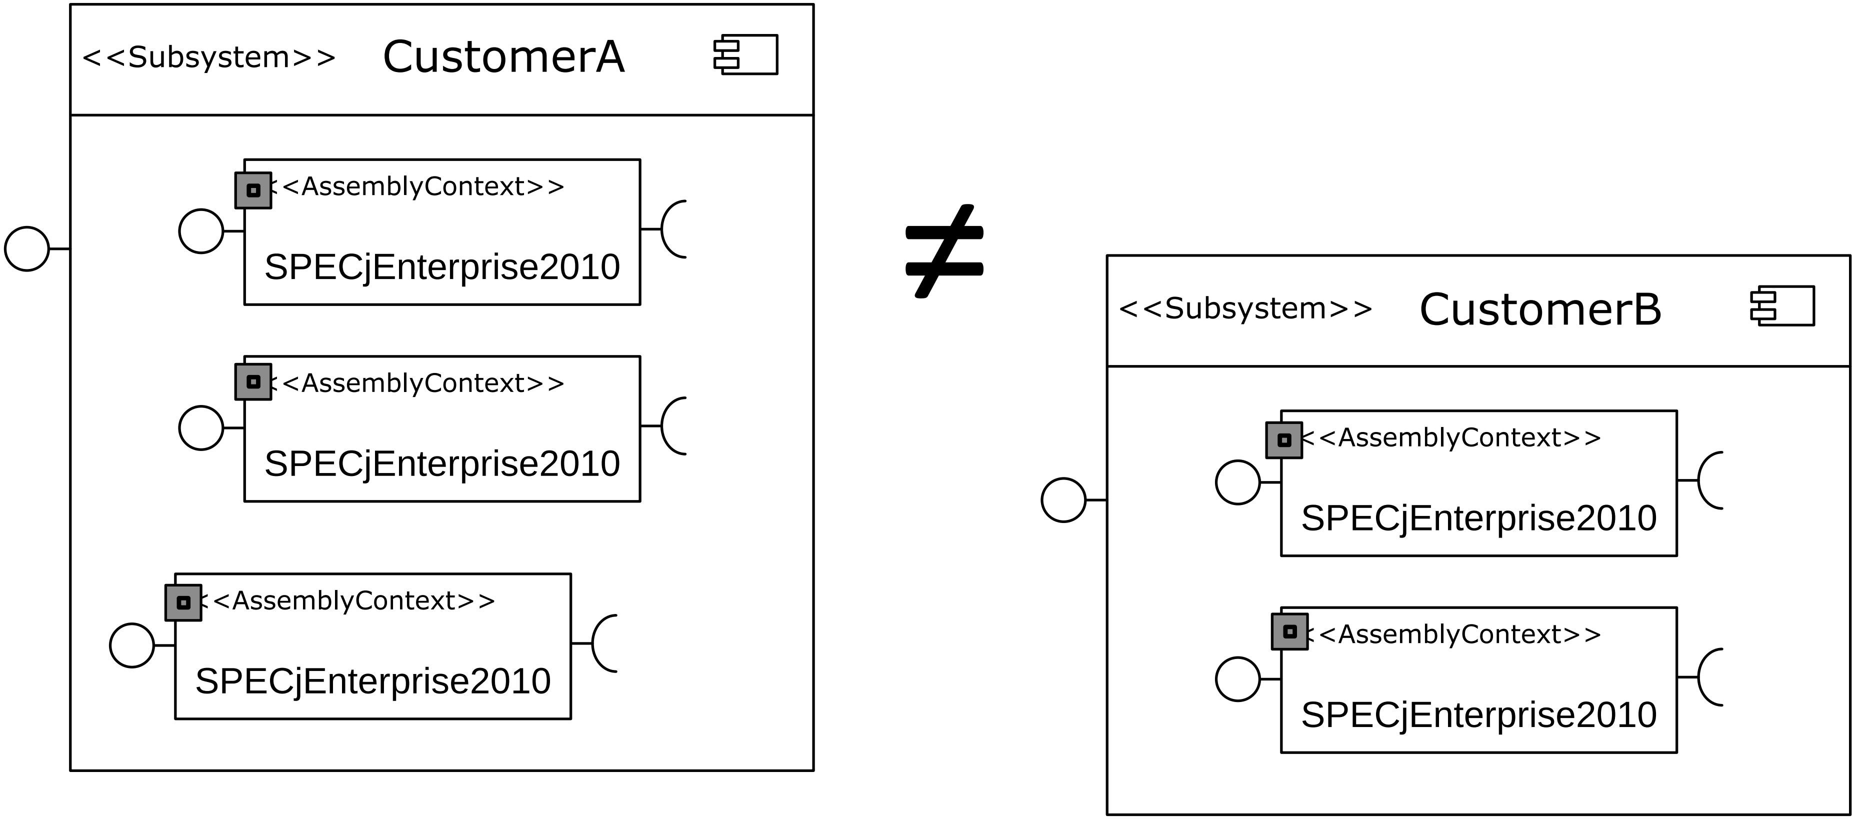

Relationship/Scope - Metamodel

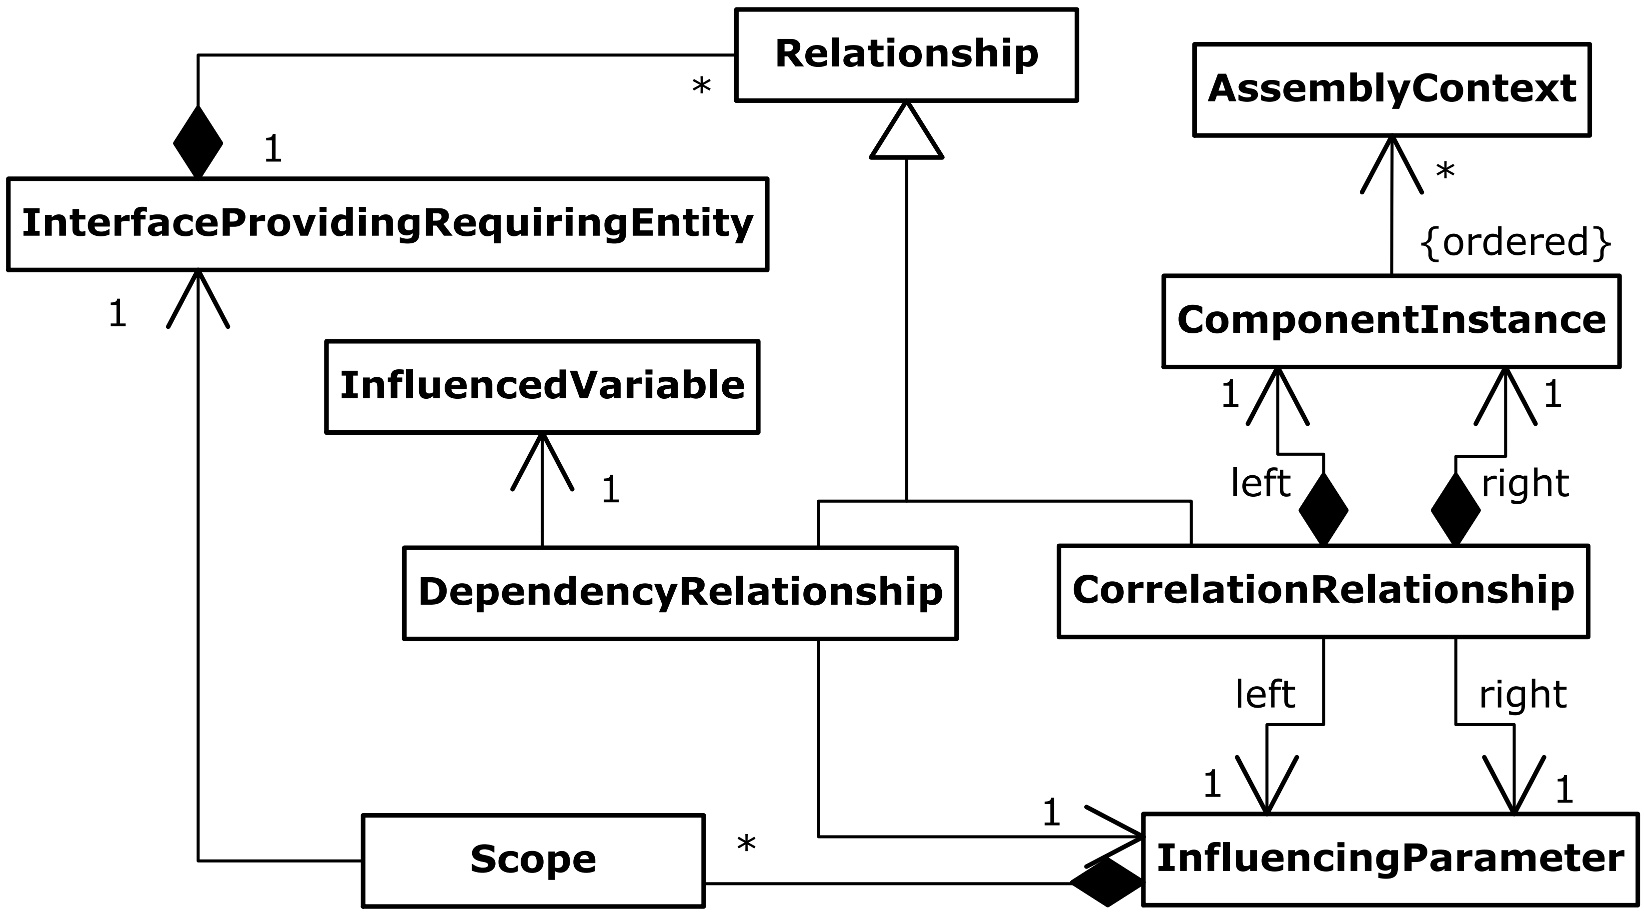

Relationship Characterization - Metamodel

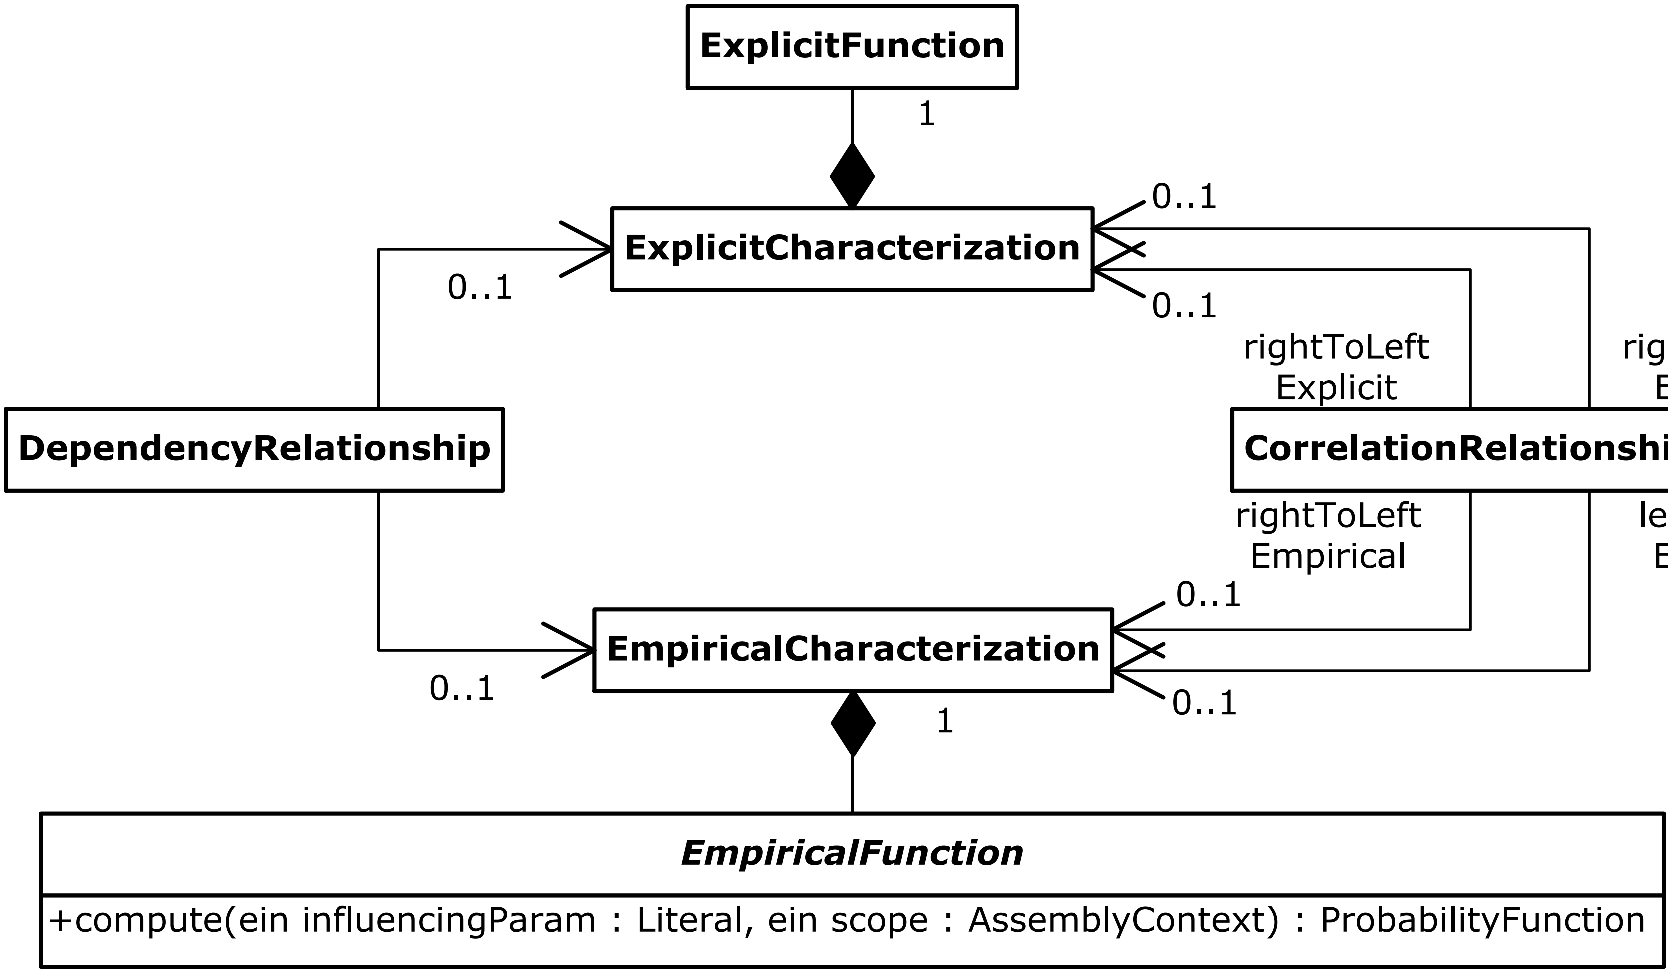

Deployment-Specific Resource Demands

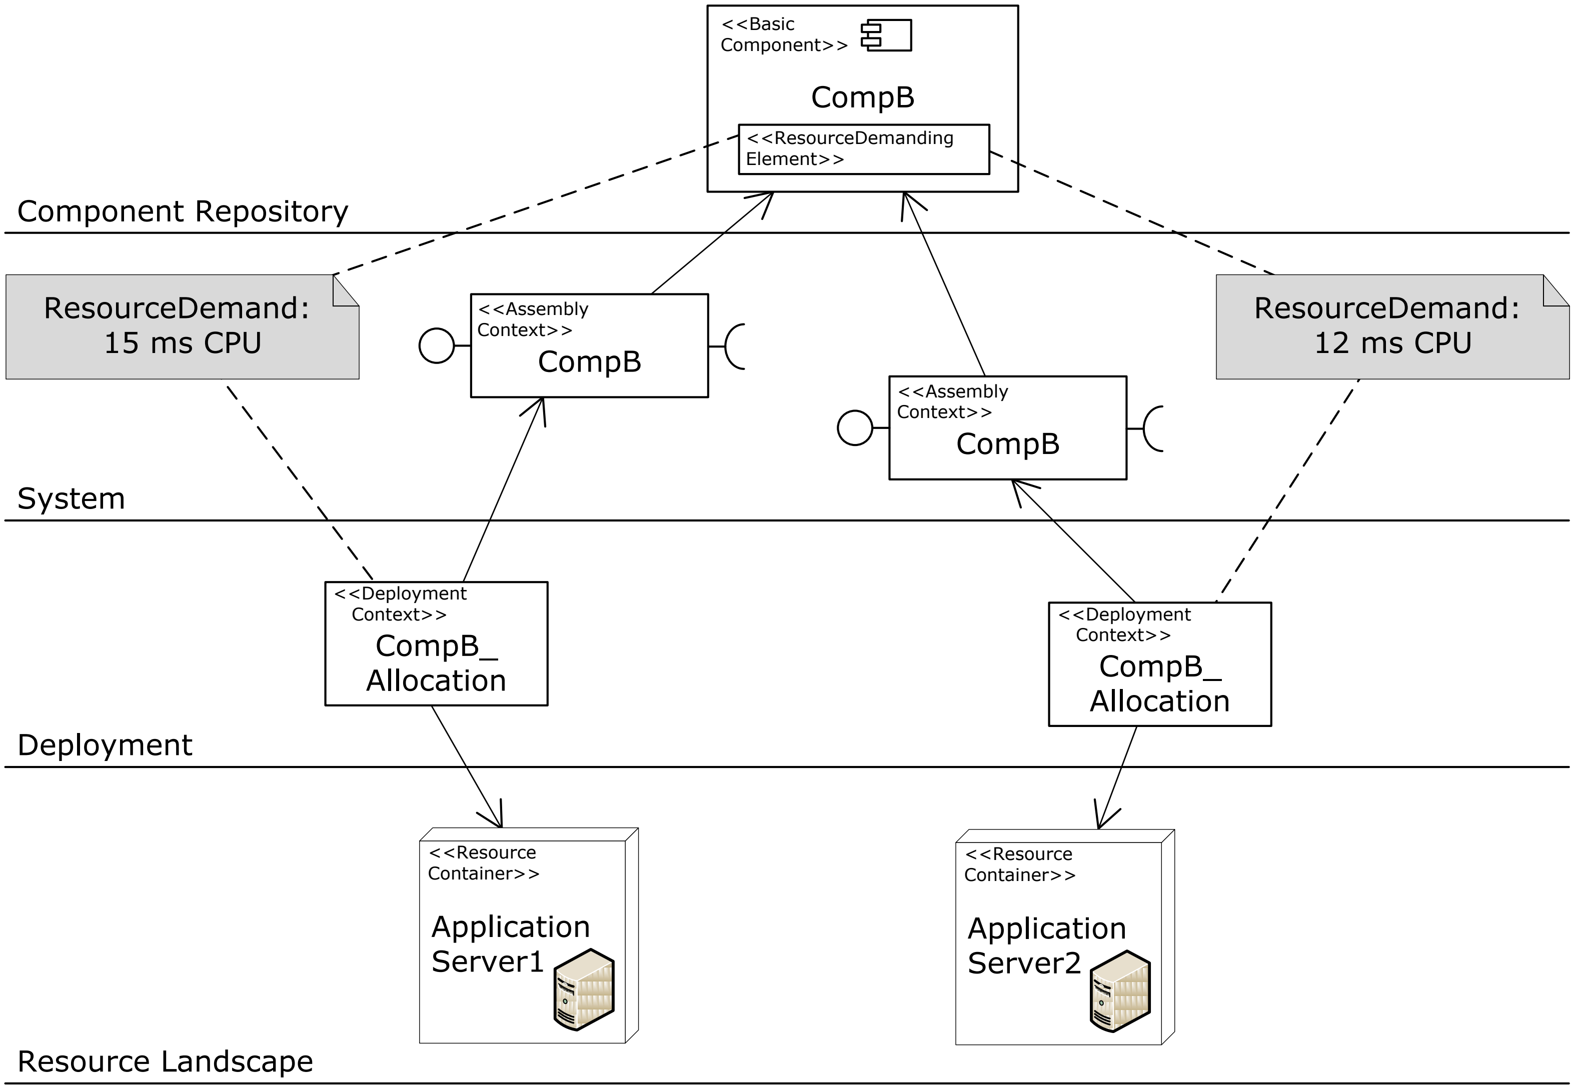

DMM: Modeling Adaptation Processes

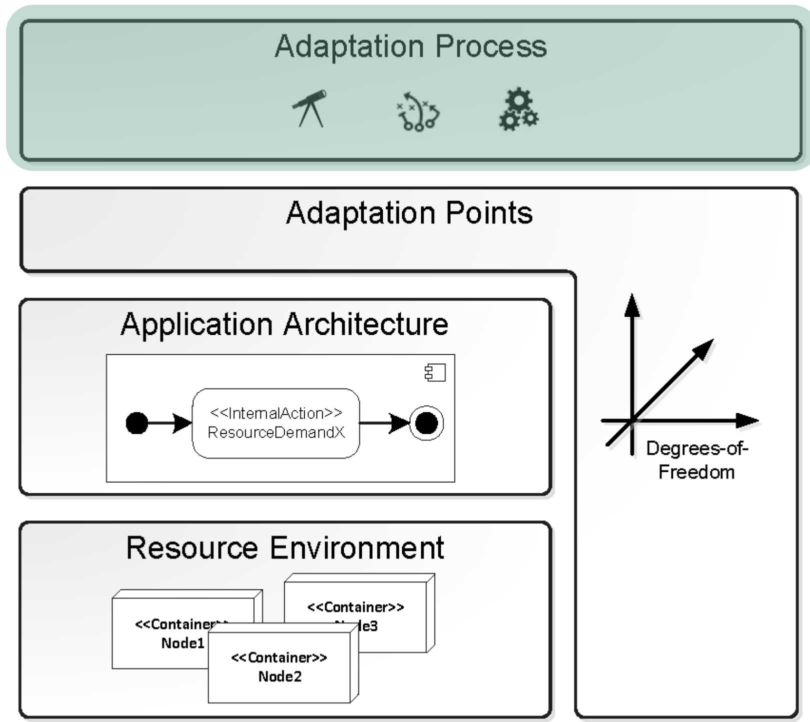

Details in:

- N. Huber, A. van Hoorn, A. Koziolek, F. Brosig, and S. Kounev. **S/T/A: Meta-modeling Run-time Adaptation in Component-based System Architectures**. Under review, 2012.

Motivation

- Rapid growth of autonomic computing and self-adaptive systems engineering
- Open challenges
  - System-specific reconfiguration techniques typically hard-coded in the system's implementation
  - How to separate software design and implementation from system reconfiguration logic?
- Main issue:
  - How to abstract from system-specific details?
  - How to enable the reuse of adaptation strategies?

Holistic Model-based Approach

- Describe system adaptation processes at the system architecture level
  - Distinguish high-level reconfiguration objectives from low-level implementation details
  - Explicitly separate technical from logical aspects
  - Capture reconfiguration logic in a generic, human-understandable, machine-processable and reusable way
  - Provide intuitive modeling concepts that can be employed by system architects and software developers
  - Facilitate maintenance and reuse

Autonomic Resource Management

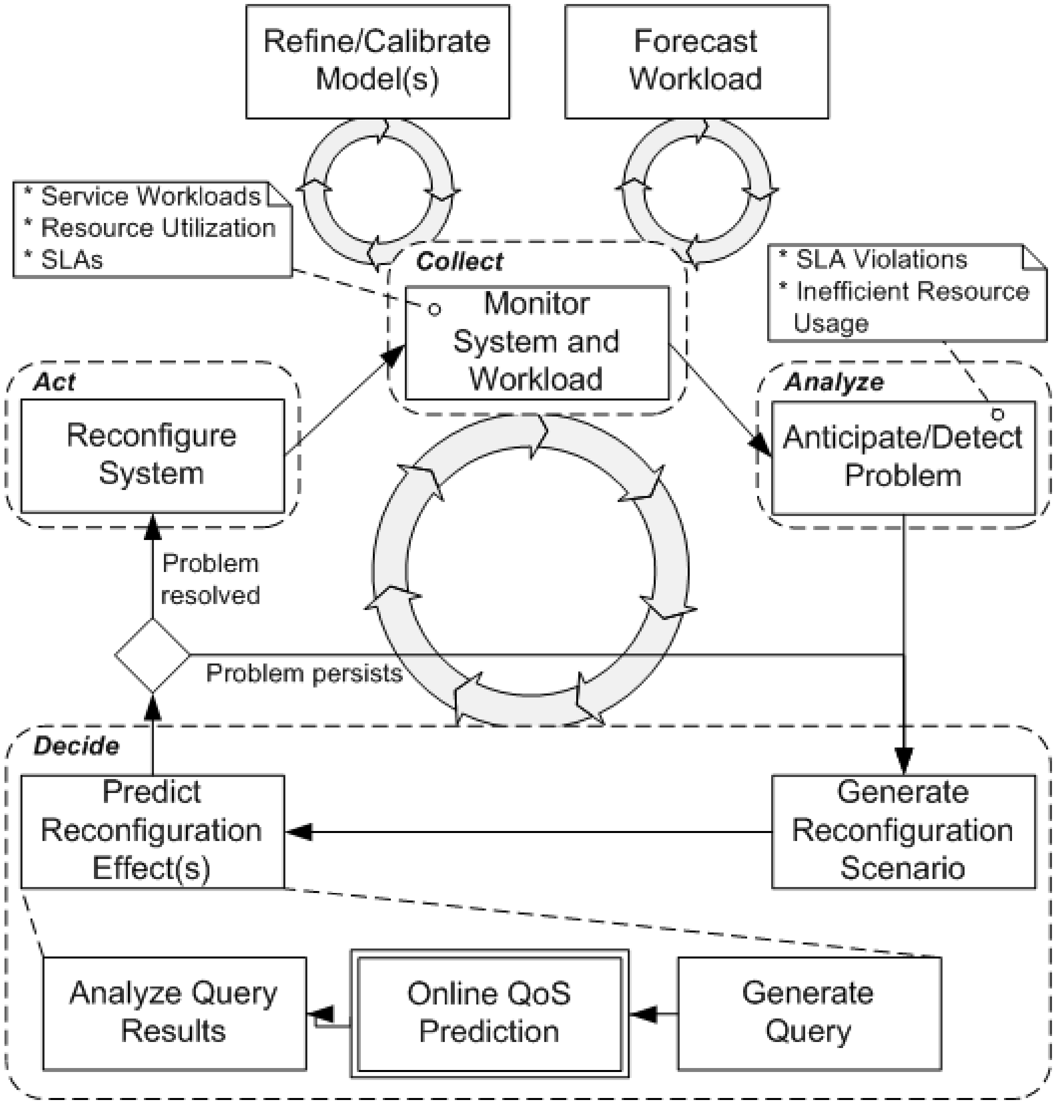

Model-based System Reconfiguration

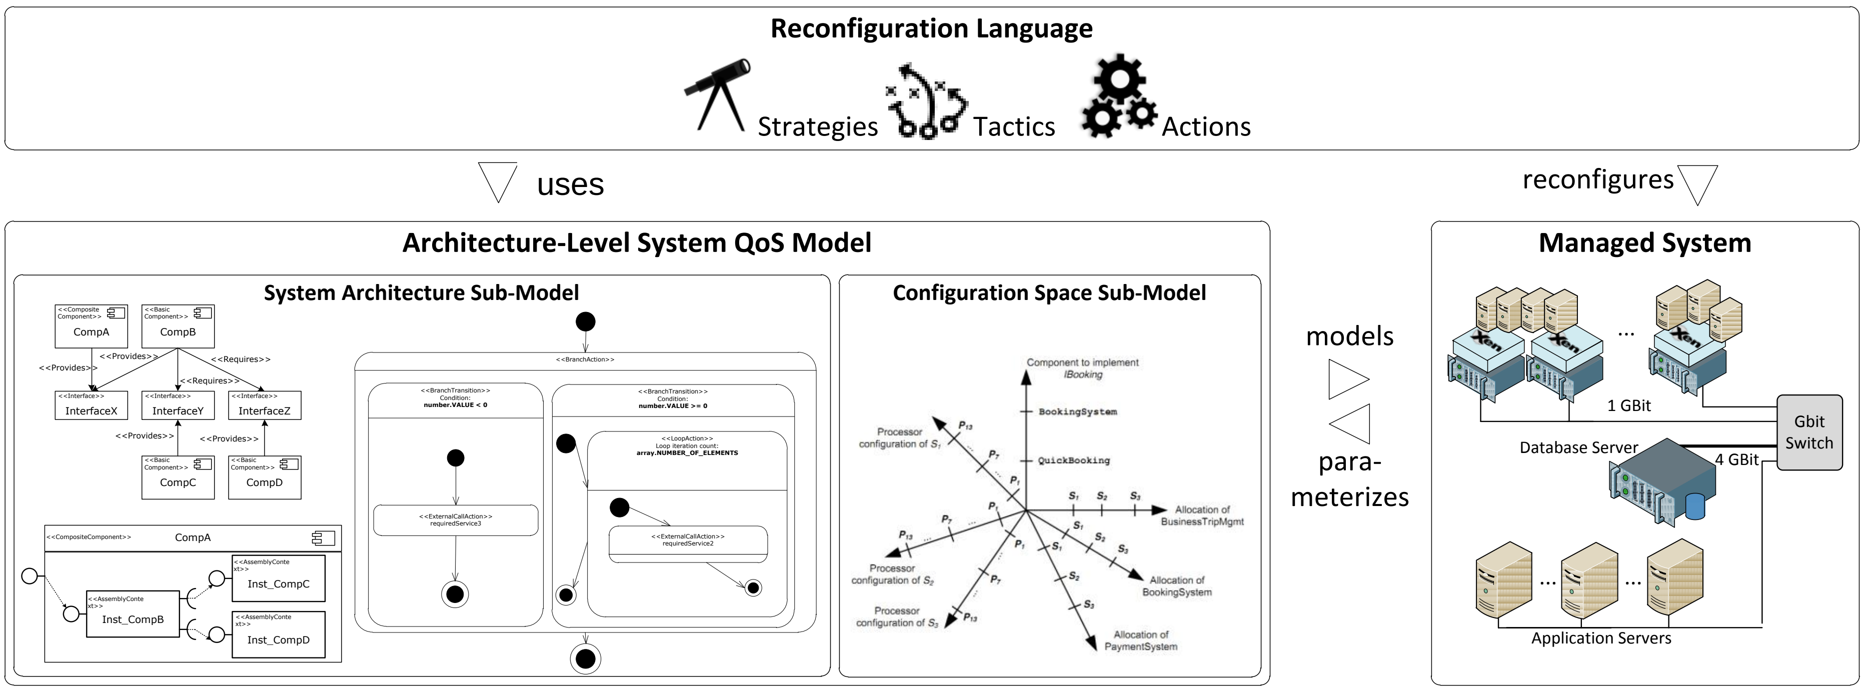

Motivation

- Detailed models of the environment
  - Technical: Structuring, Migration
  - QoS property influences
  - Resource layers
  - → Improved decision making
- Dynamic Reconfiguration
  - On the model level (end-to-end)
  - Decouple from technical, system-specific details

S/T/A Reconfiguration Language

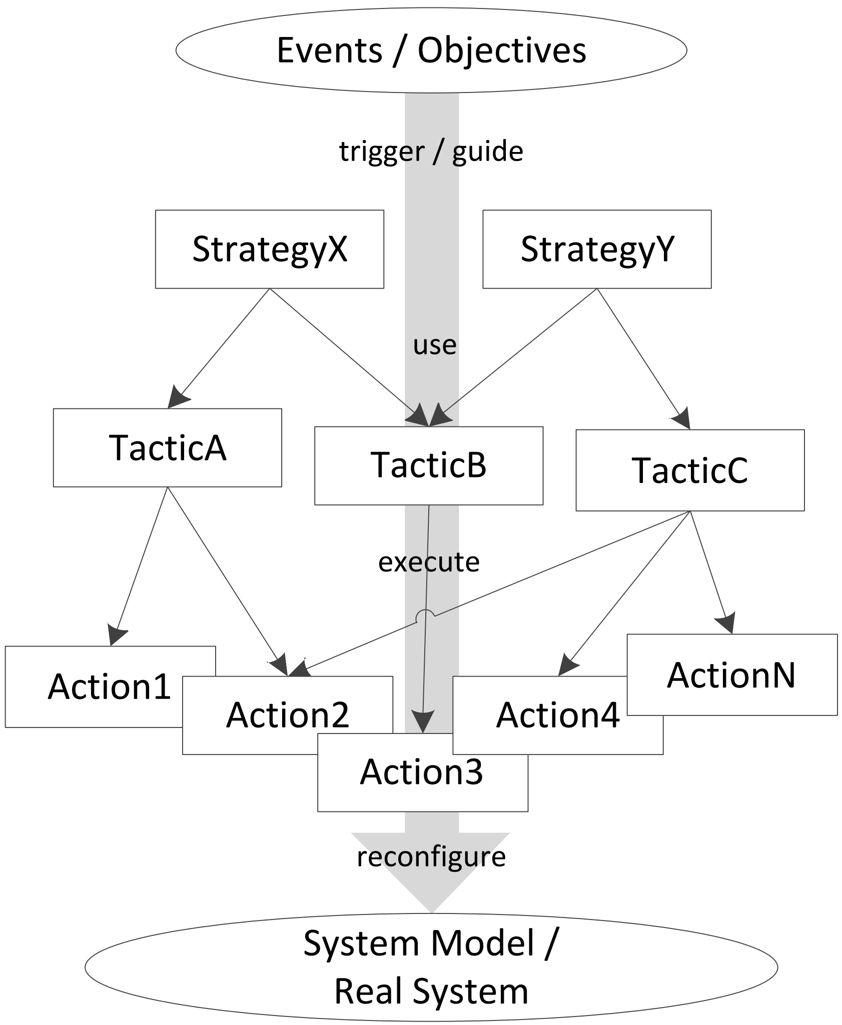

Separate

- Logical view, high-level process
- Technical view, low-level operations

[HHK+12]

Separation of Concerns

Karlsruhe Institute of Technology

Strategies

- High-level
- Independent of system specific details
- Describe process view
- Indeterminism

Tactics & Actions

- Low-level
- System specific
- Reconfiguration operations
- Deterministic

S/T/A Meta-Model

Karlsruhe Institute of Technology

- Actions refer to adaptation points / DoF Model
- Tactics execute Actions in Reconfiguration Plans
- Strategies use weighted Tactics

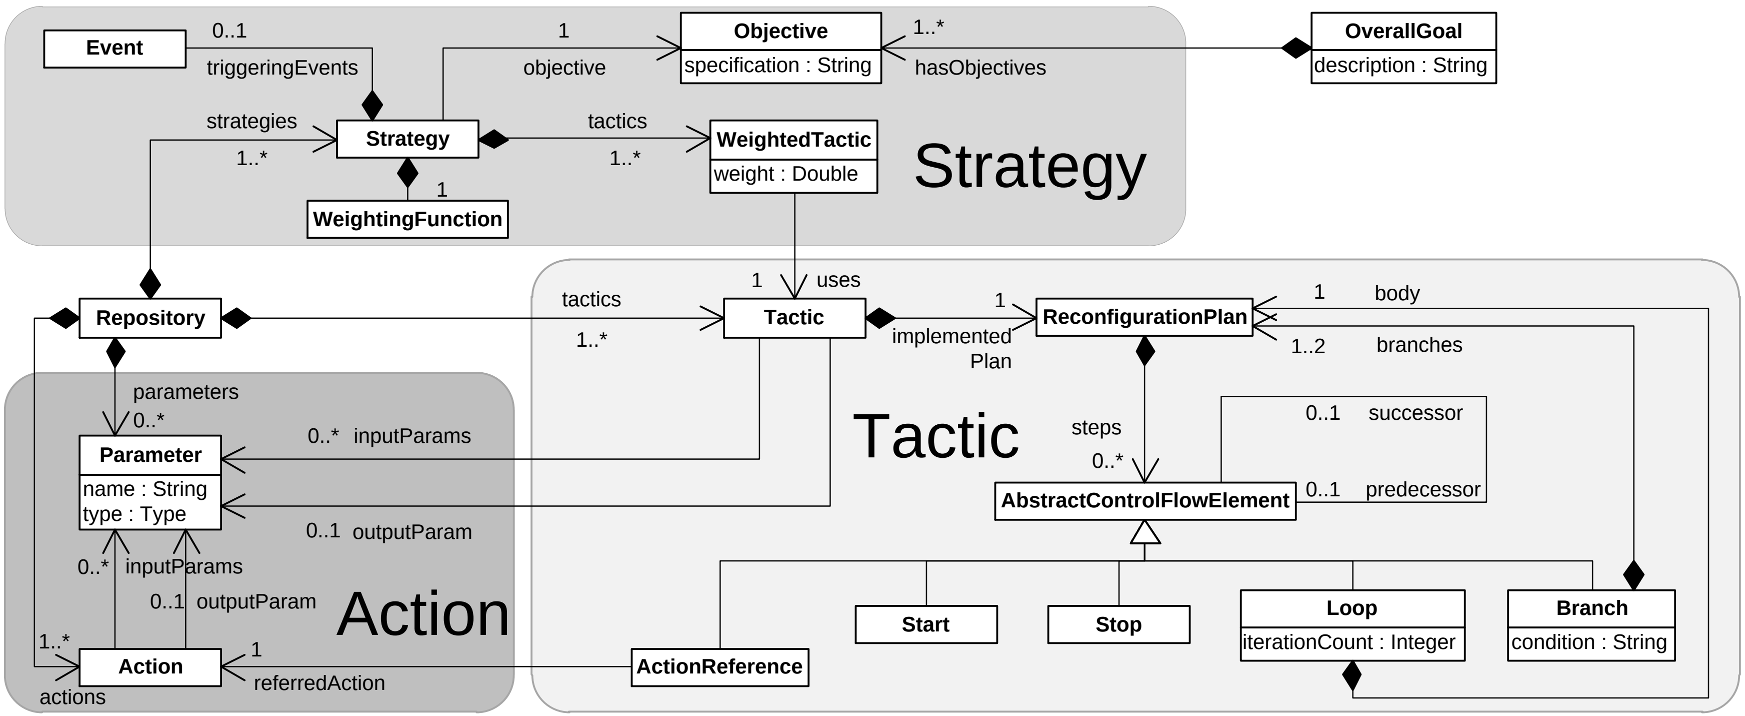

Example Strategies

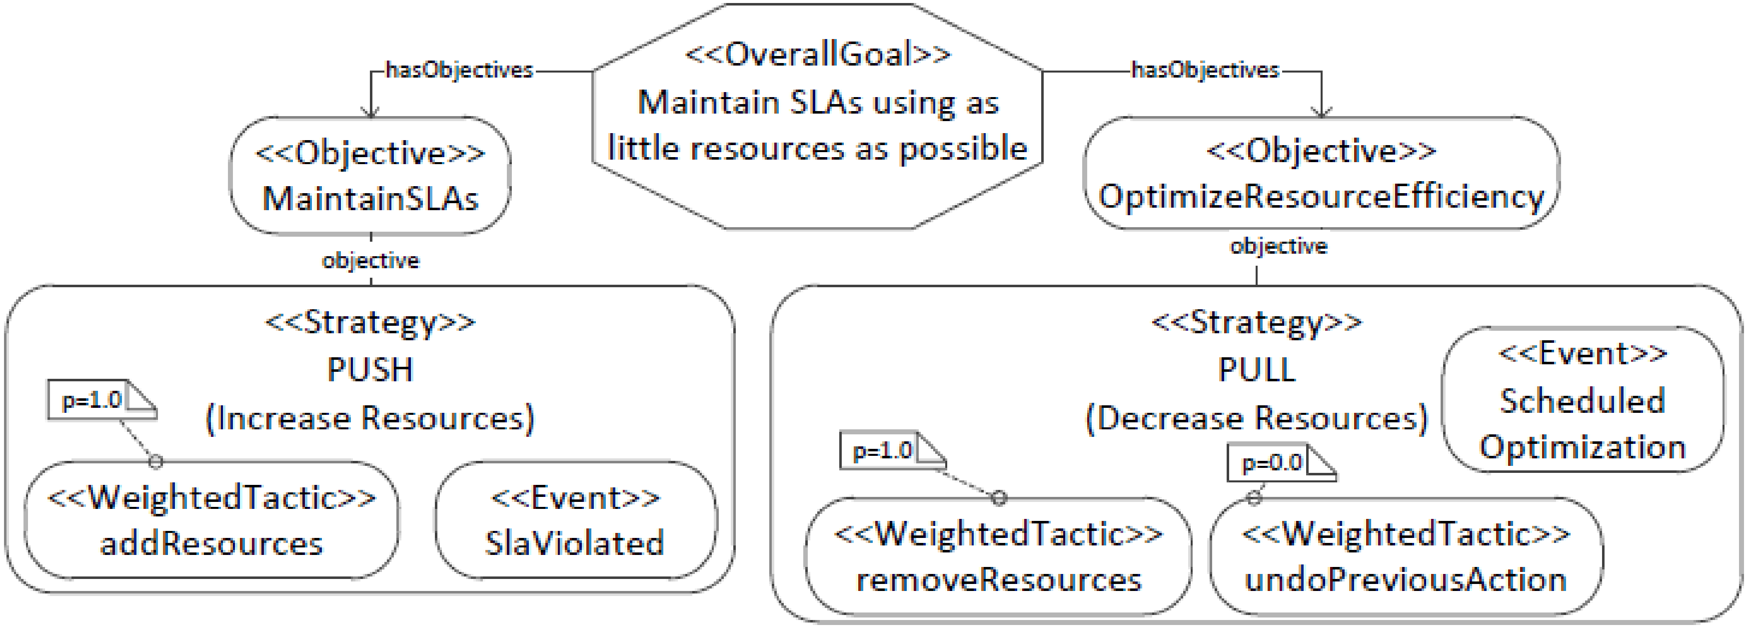

Example Tactics & Actions

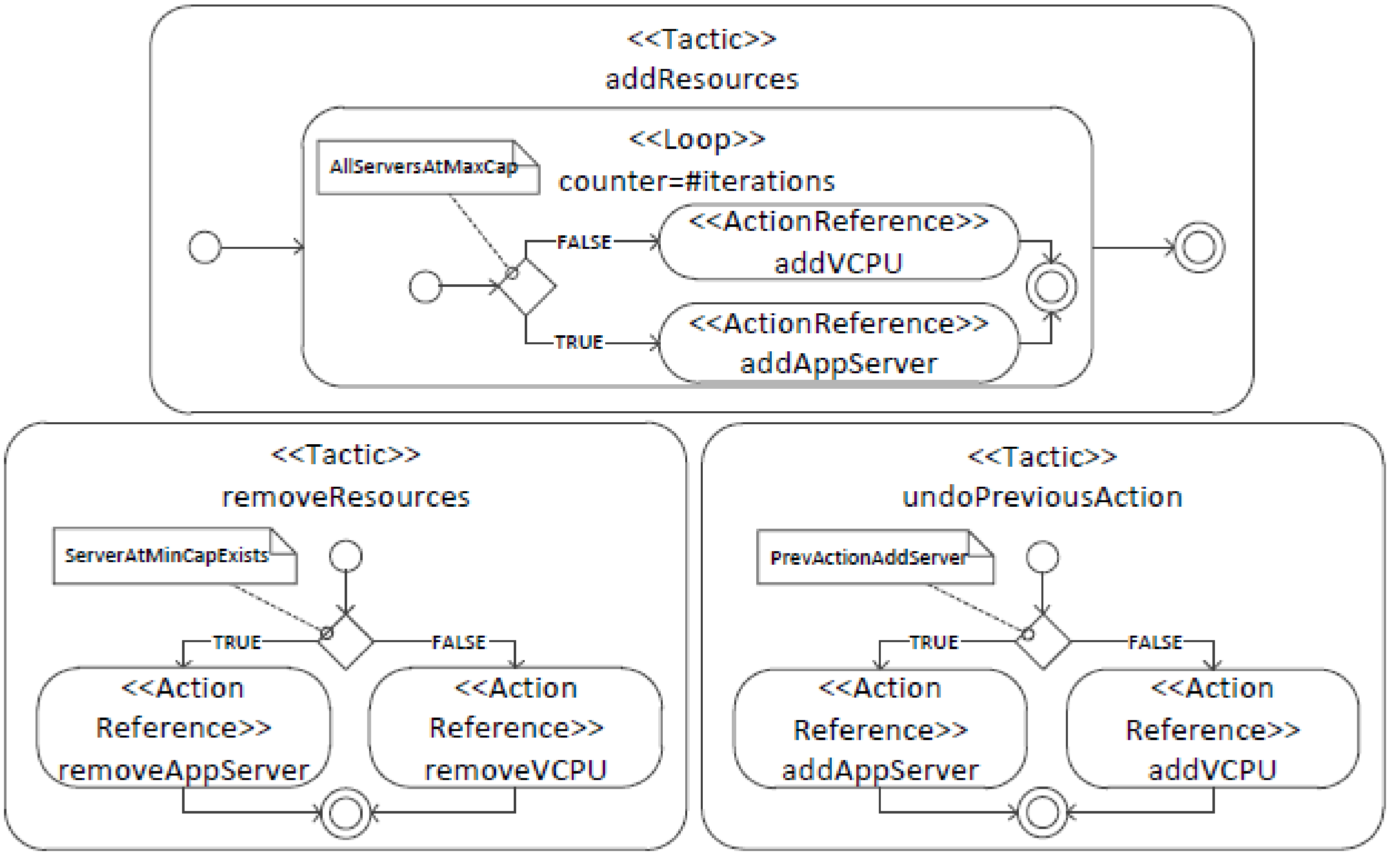

Evaluation

- S/T/A implemented in PerOpteryx

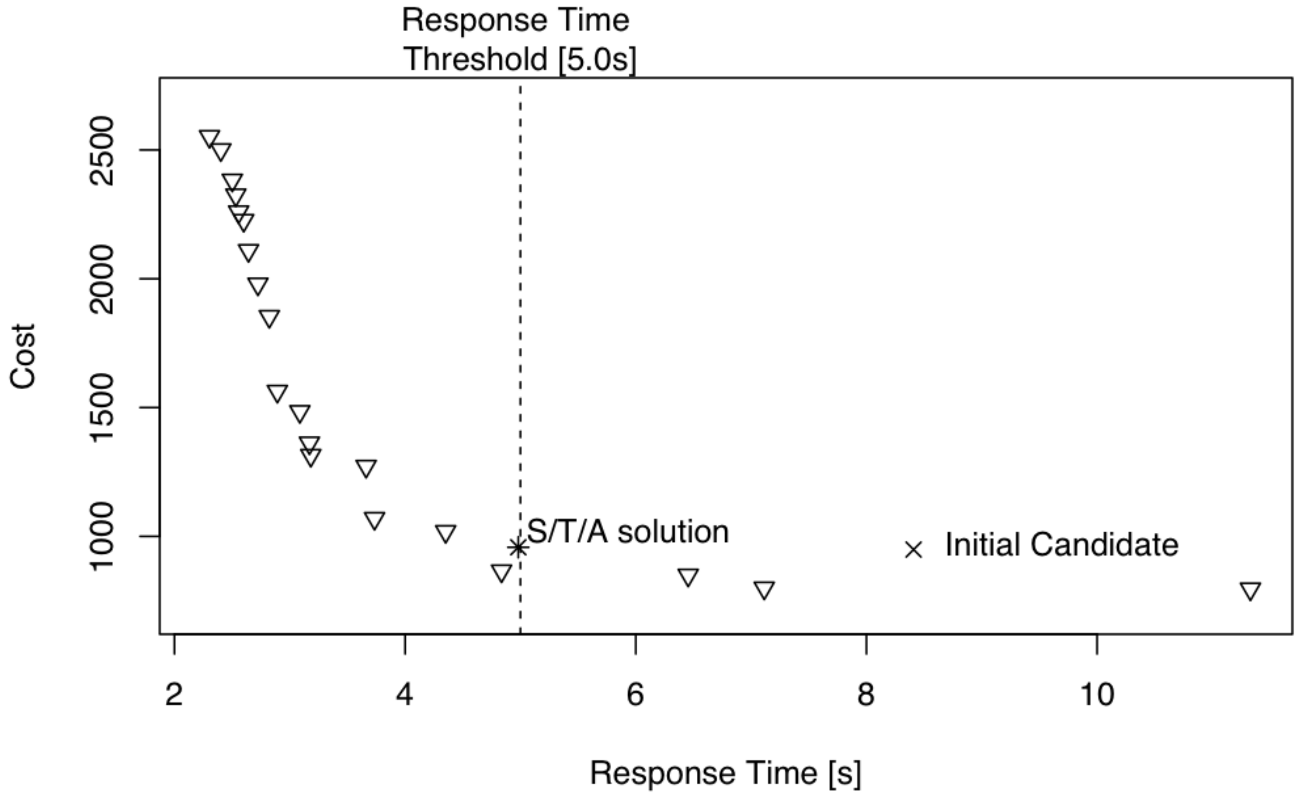

Ongoing and Future Work

- Efficient heuristics and optimization algorithms specifically tailored for use in S/T/A
- Graphical tools for modeling reconfiguration processes
- Model-based evaluation of reconfiguration processes/algorithms at system design time
  - Efficiency, Self-stabilization, Scalability, Elasticity, …
- Standard metrics and benchmarks

Self-Adaptive Resource Management

Karlsruhe Institute of Technology

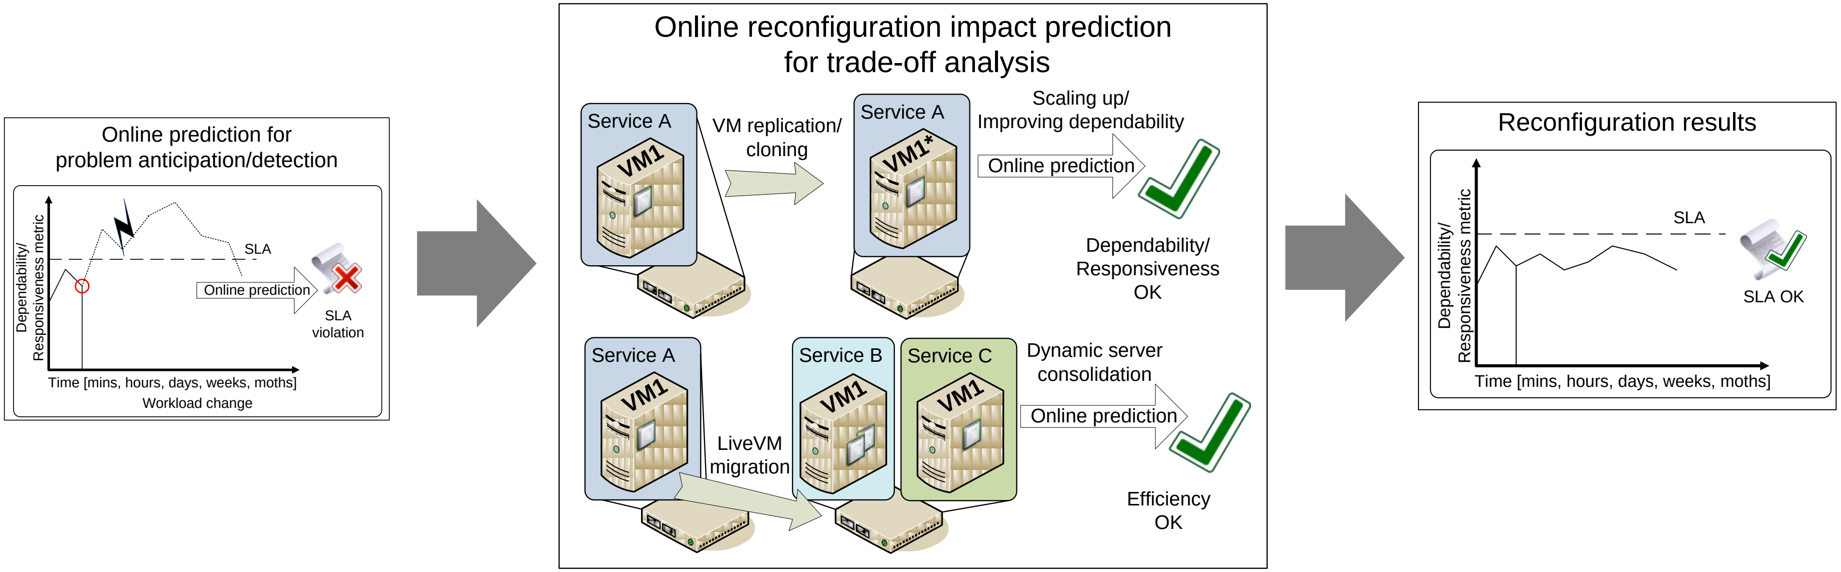

Further details in:

- N. Huber, F. Brosig, and S. Kounev. **Model-based Self-Adaptive Resource Allocation in Virtualized Environments**. In 6th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS 2011), Honolulu, HI, USA, May 23-24, 2011.

Self-Adaptive Resource Allocation

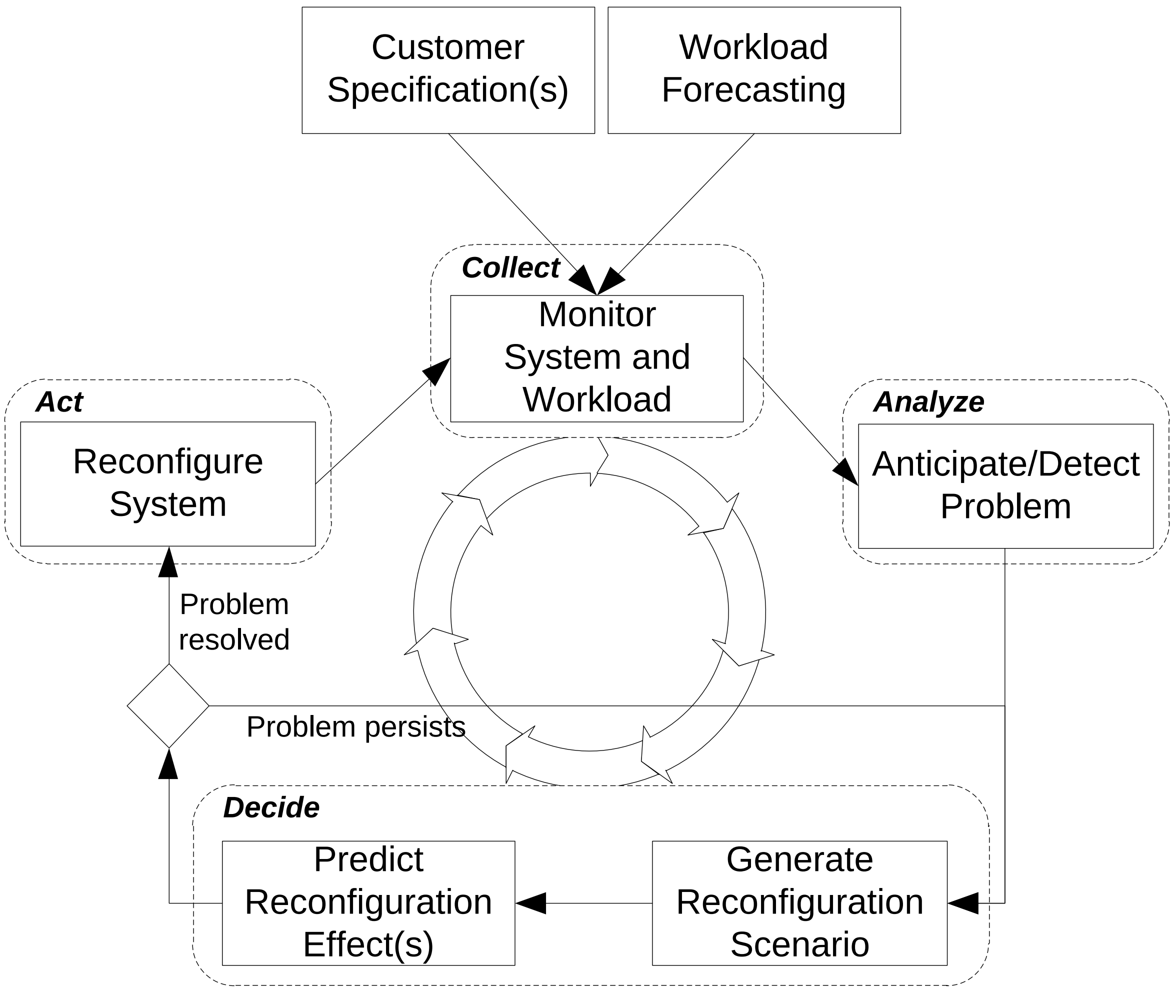

Reconfiguration Algorithm

Decide

PUSH Phase

- Add resources
  - vCPUs (if available)
  - Application server nodes
  - until

$$\overline{cap}(c,t) = \left\lceil \frac{\sum_{c\in\widetilde{C}} c[\lambda]\cdot D(c[s])}{\sum_{c\in C} c[\lambda]\cdot D(c[s])} \right\rceil \cdot cap(c,t)$$

PULL Phase

- Remove underutilized resources as long as no SLAs are violated

Reconfiguration Algorithm

PUSH

---

while $\exists c \in \widetilde{C} : \neg P_R(c)$ **do**
    for all $t \in V(c[s]) : \neg P_U(t)$ **do**
        while $cap(c,t) \leq \overline{cap}(c,t)$ **do**
            if $\exists i \in F(c[s],t) : i[\kappa] < i[\overline{\kappa}]$ **then**
                $i[\kappa] \leftarrow i[\kappa] + 1$
            else
                $F(c[s],t) \leftarrow F(c[s],t) \cup \{\widehat{i}\}$
            end if
        end while
    end for
end while

PULL

---

for all $c \in C$ **do**
    while $\exists t \in V(c[s]) : \overline{U}(t) - U(t) \geq \epsilon$ **do**
        if $\exists i \in F(c[s],t) : i[\kappa] > 0$ **then**
            $i[\kappa] \leftarrow i[\kappa] - 1$
            if $\neg P_R(c)$ **then**
                $i[\kappa] \leftarrow i[\kappa] + 1$
            end if
            if $i[\kappa] = 0$ **then**
                $F(c[s],t) \leftarrow F(c[s],t) \setminus \{i\}$
            end if
        end if
    end while
end for

Case Study: SPECjEnterprise2010

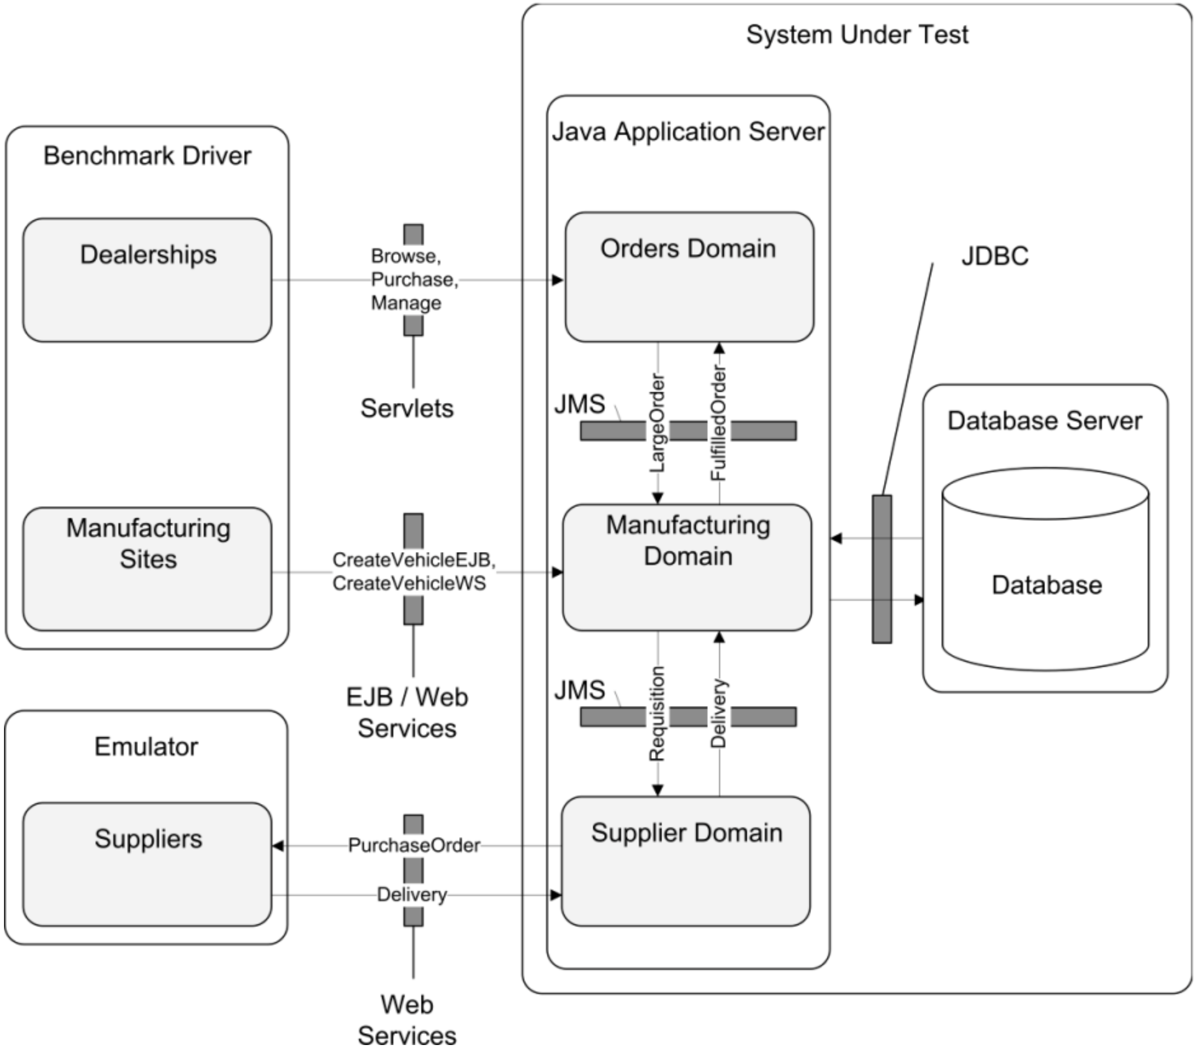

Architecture-Level Performance Model

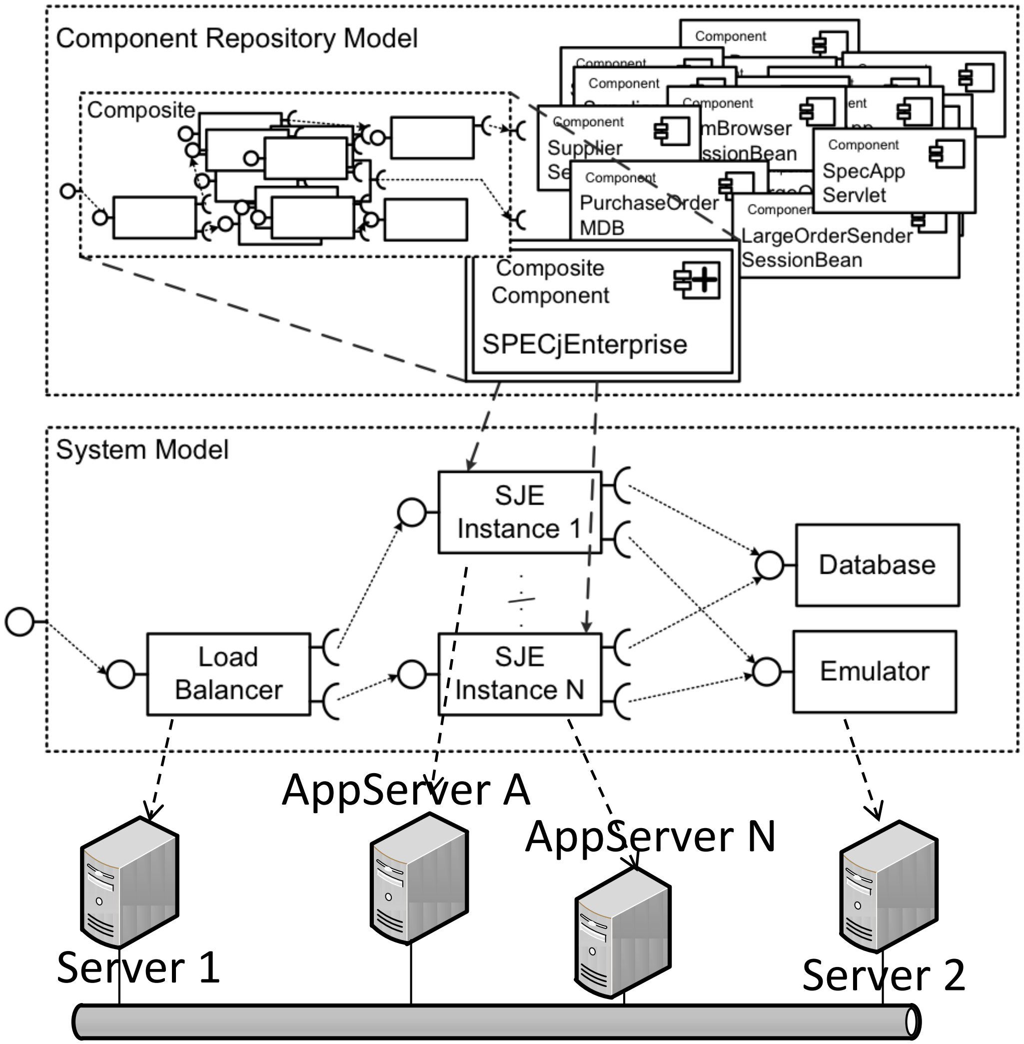

Architecture-Level Performance Model

- Semi-automatic extraction [Brosig11]
- 28 components, 63 behavior specs
- Example control flow specification

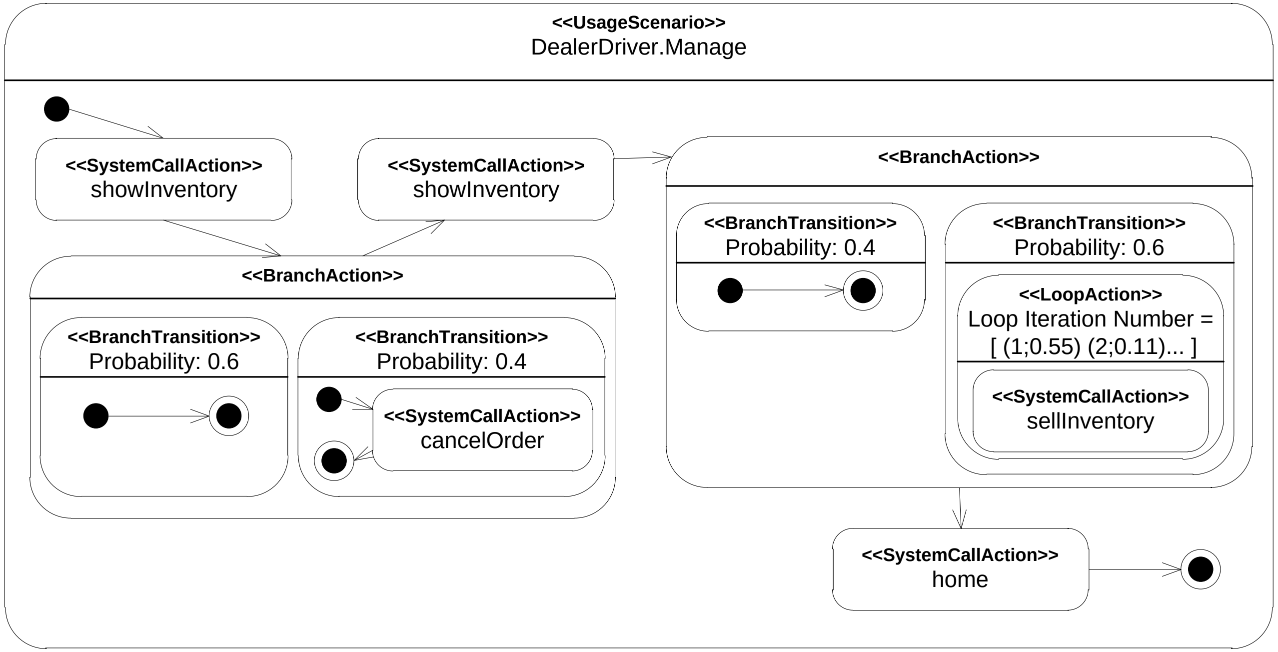

Experimental Setup

- Six blade servers
  - 2 Xeon E5430 4-core CPUs
  - 32 GB of main memory
  - Citrix XenServer 5.5
  - Oracle WebLogic 10.3.3
  - Oracle Database 11g

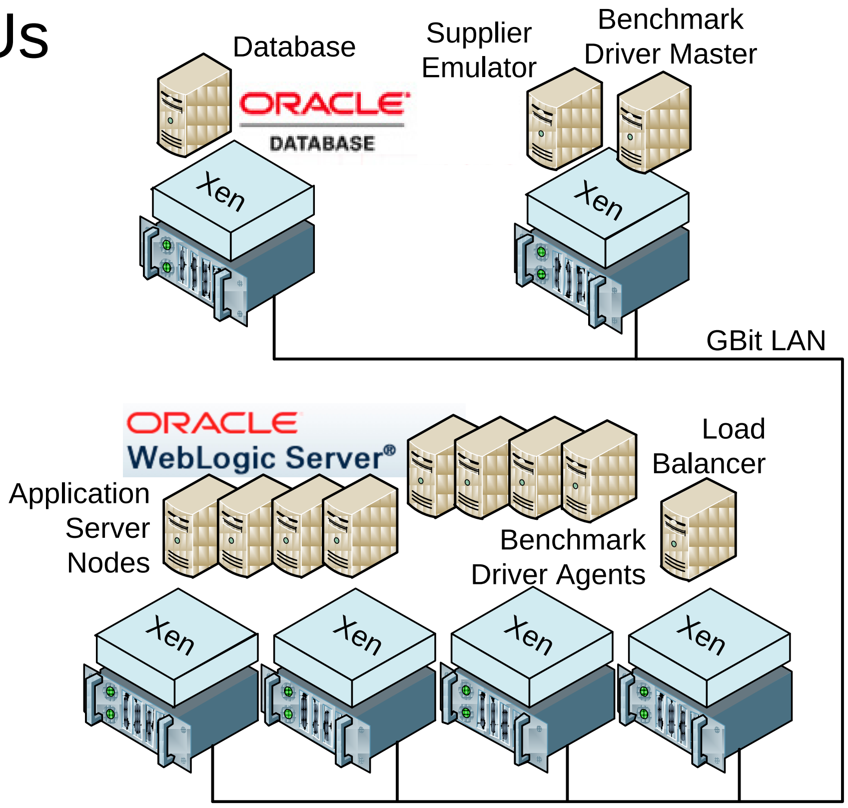

What If: New Service Added?

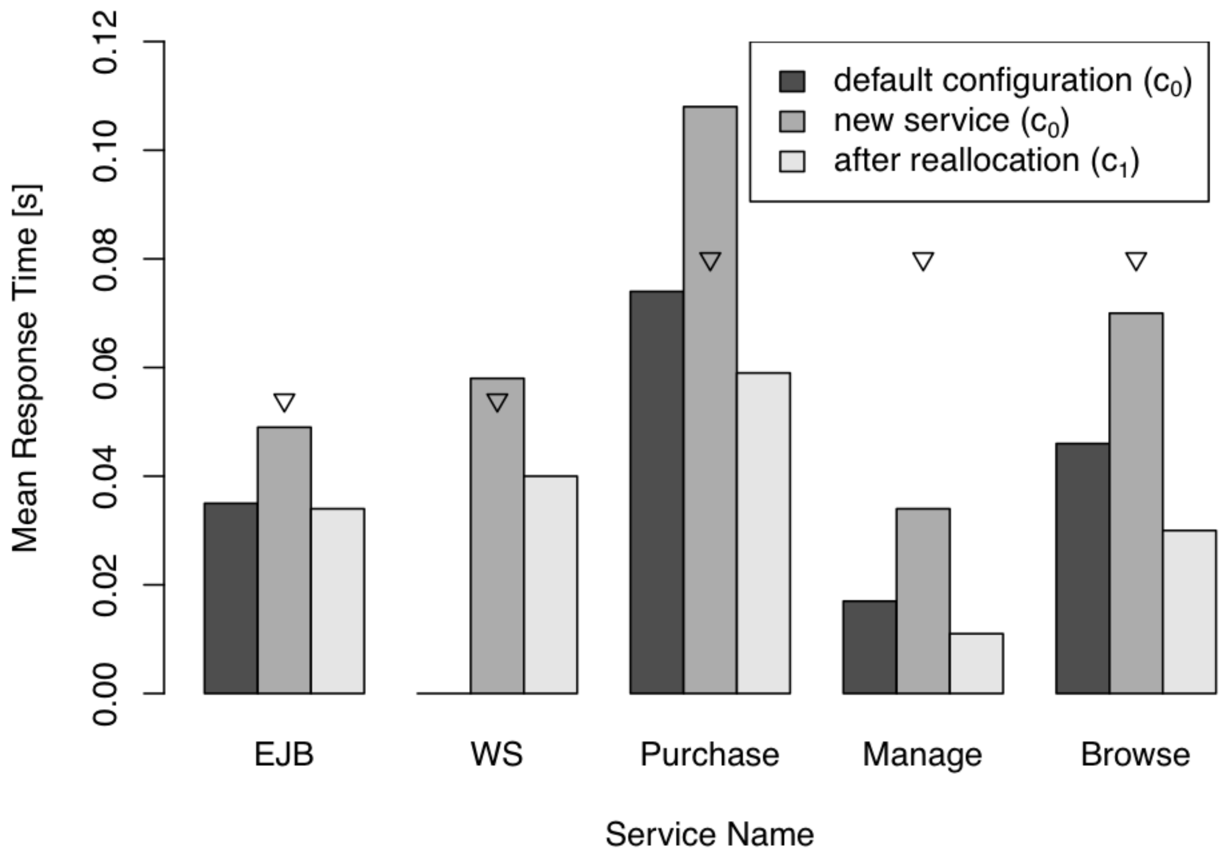

What If: Workload Changes?

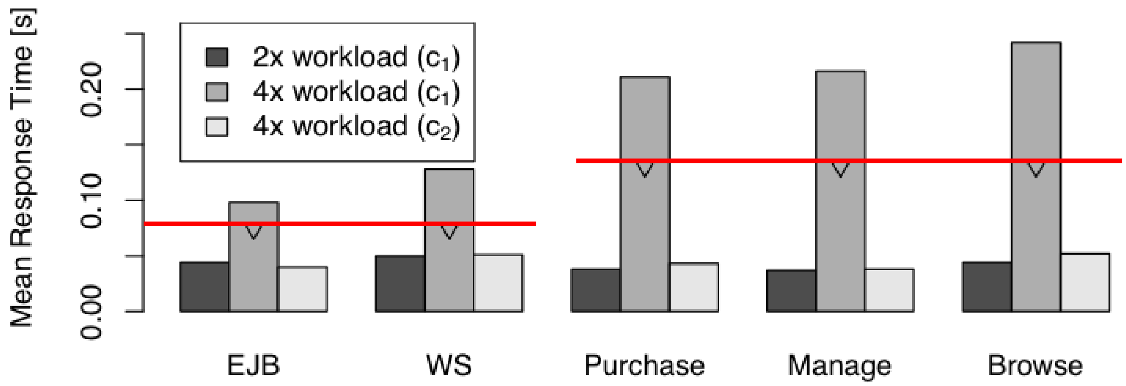

a) increasing workload from 2x to 4x

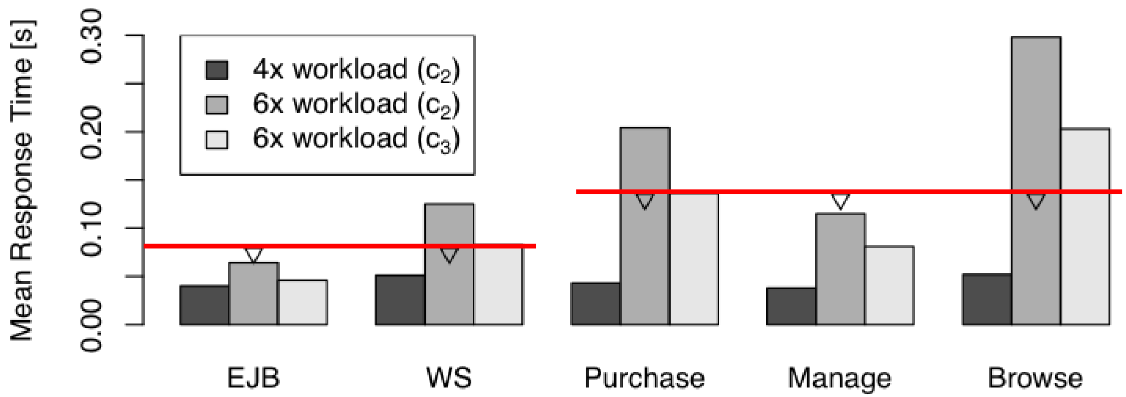

b) increasing workload from 4x to 6x

Benefits in Cost Savings

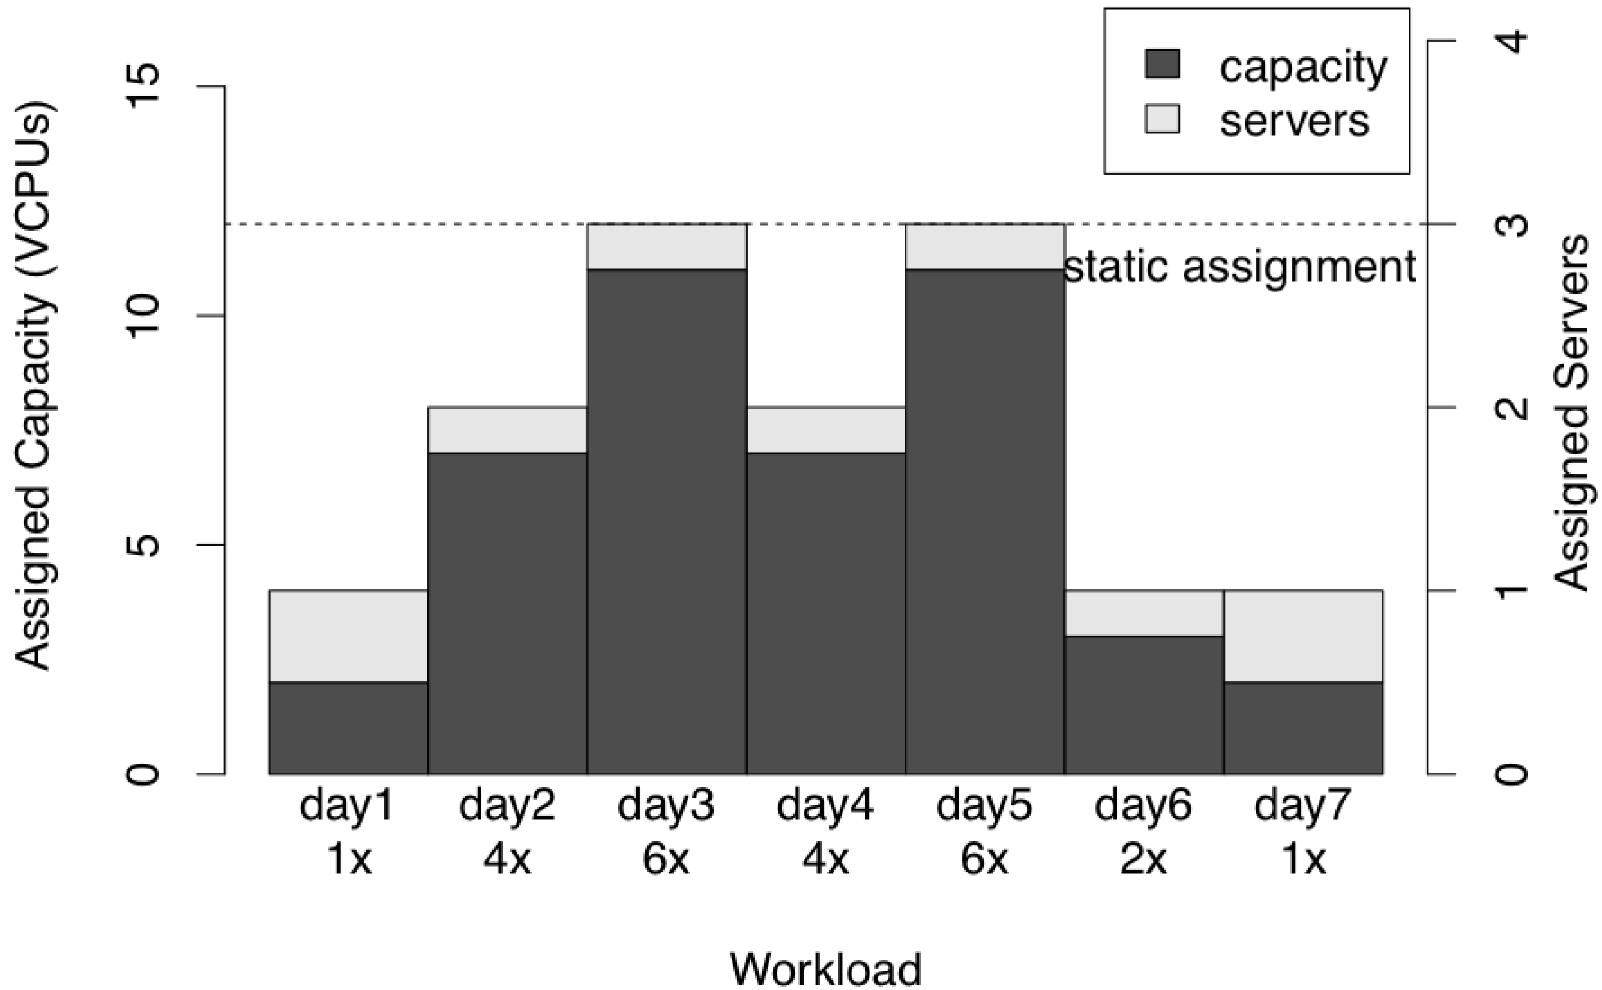

Modeling Virtualization Platforms

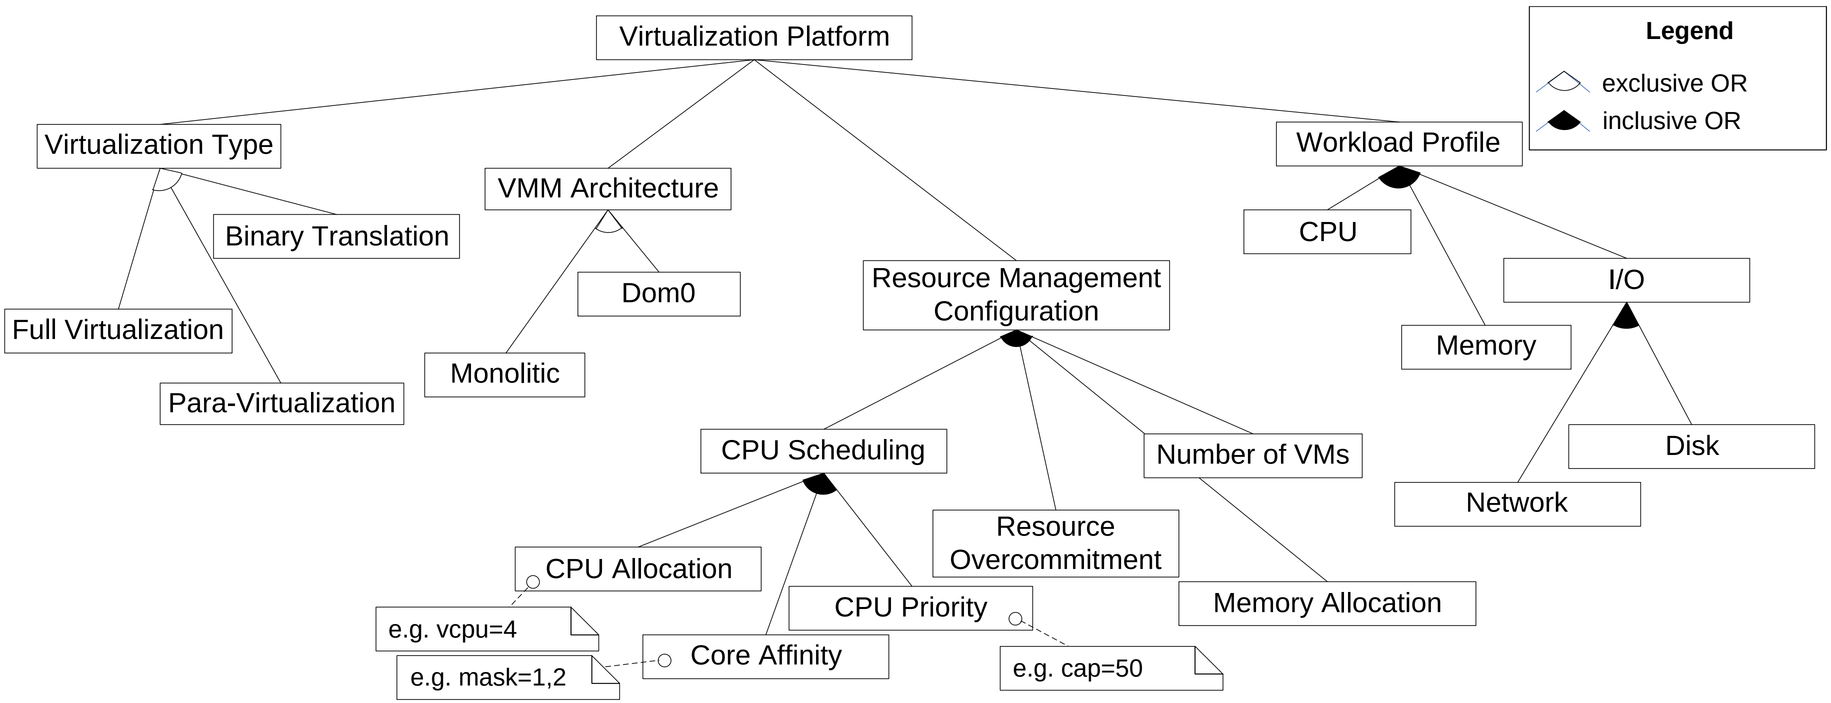

Further details in:

- N. Huber, M. Quast, M. Hauck, and S. Kounev. **Evaluating and Modeling Virtualization Performance Overhead for Cloud Environments**. In *Proceedings of the 1st International Conference on Cloud Computing and Services Science (CLOSER 2011), Noordwijkerhout, The Netherlands*, May 7-9 2011. Best Paper Award.

Automated Experimental Analysis

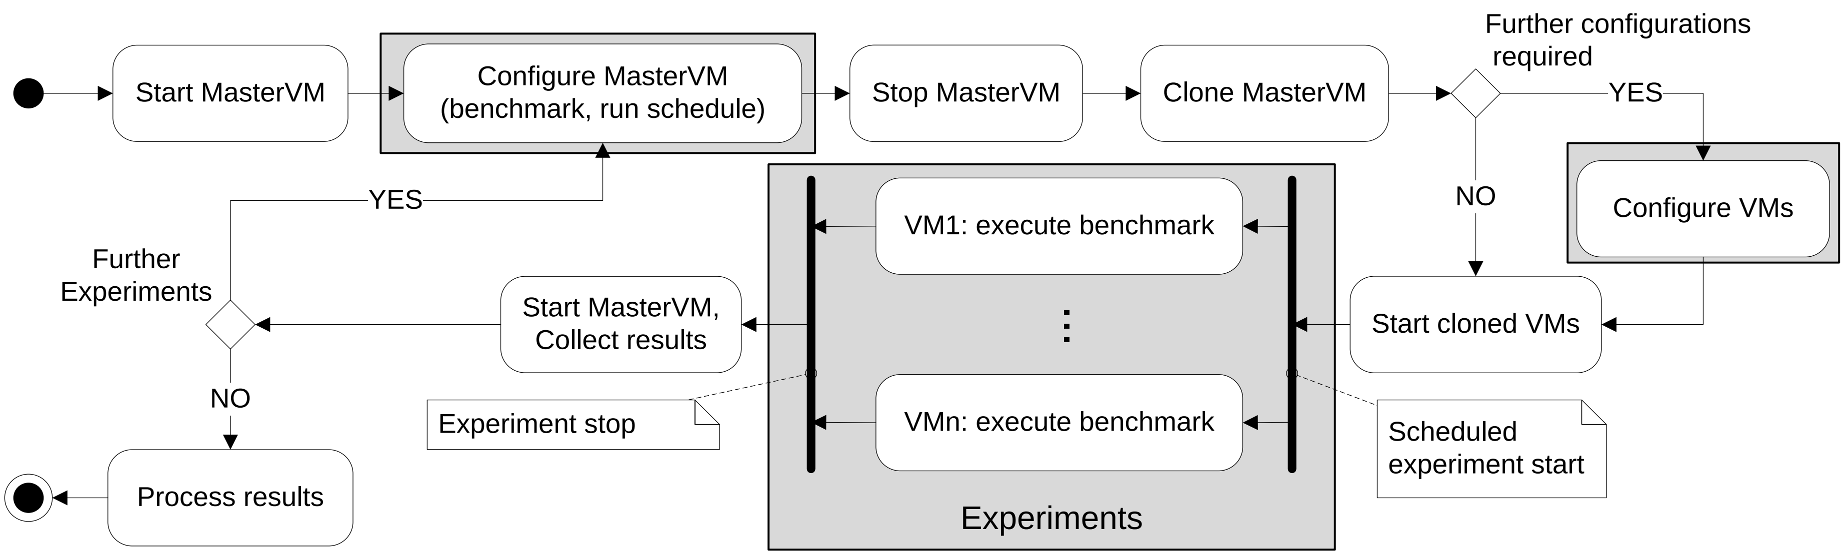

Further details in:

- N. Huber, M. von Quast, F. Brosig and S. Kounev. **Analysis of the Performance-Influencing Factors of Virtualization Platforms**. In 12th International Symposium on Distributed Objects, Middleware, and Applications (DOA 2010), Crete, Greece, October 2010. Springer Verlag.

Experiment Setup

- Virtualization Platforms
  - Citrix XenServer 5.5
  - VMware ESX 4.0
- Experimental environment
  - SunFire X4440 Server, AMD Opteron 24*2.4 GHz, 128 GB DDR2 RAM

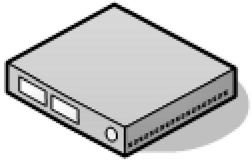

- Different benchmark types
  - Passmark PerformanceTest v7 (CPU, Memory, HDD)
  - SPECcpu (CPU + Memory)
  - Iperf (Network)

Virtualization Overhead

XenServer 5.5

Throughput metric: higher values are better

Benchmark	native	virtualized	Delta (abs.)	Delta (rel.)
Passmark CPU, 1 core	639.3	634.0	5.3	0,83%
Passmark CPU, 2 cores	1232.0	1223.0	9.0	0.97%
SPECint(R)_base2006				3.61%
SPECfb(R)_base2006				3.15%
Passmark Memory, 1 core	492.9	297.0	195.9	39.74%
Passmark Memory, 2 cores	501.7	317.5	184.2	36.72%
Iperf (send)	527.0	393.0	134.0	25,43%
Iperf (receive)	528.0	370.0	158.0	29,92%

Virtualization Overhead (2)

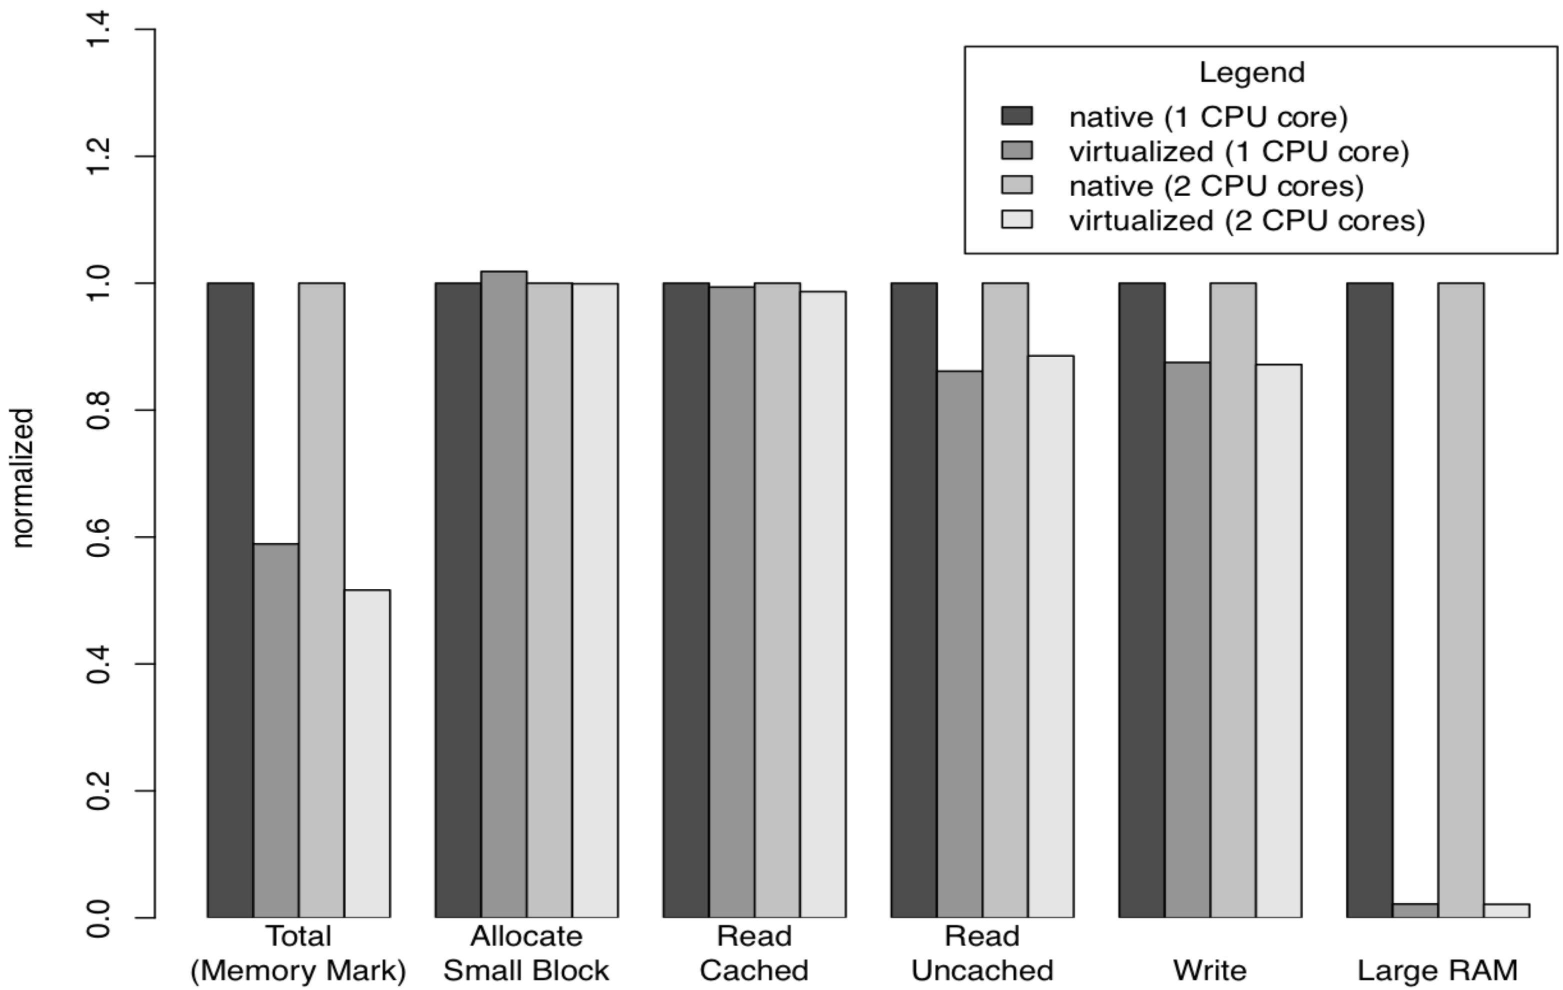

Scalability

- Scaling CPU resource
- Performance impact of affinity

Affinity OFF

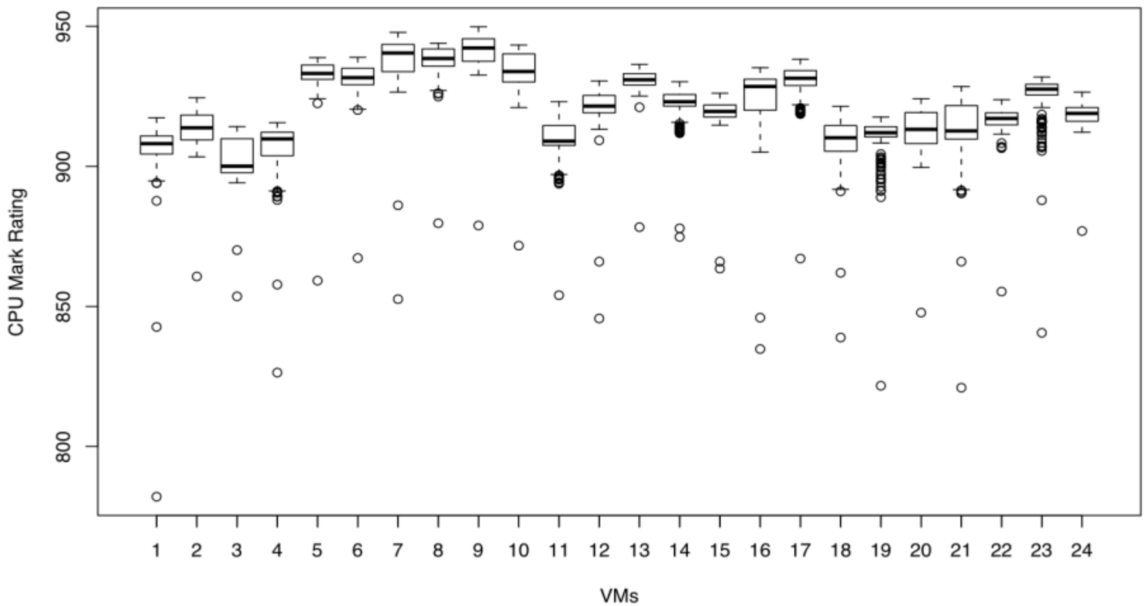

Mutual Performance Influences

- Citrix XenServer 5.5
- VM_A and VM_B pinned on the same core > $core_0$
- r_A, r_B performance drop compared to isolation

VM_A	CPU	CPU	Mem	CPU	Mem	Disk	CPU	Mem	Disk
VM_B	CPU	Mem	Mem	Disk	Disk	Disk	Net	Net	Net
r_A	46,71%	50.64%	50.33%	23.35%	24.82%	31.16%	52.88%	52.85%	3.07%
r_B	52.44%	45.93%	49.04%	1.49%	-0.99%	45.99%	40.46%	42.18%	33.31%

Automated Model Extraction @ Run-Time

KIT
Karlsruhe Institute of Technology

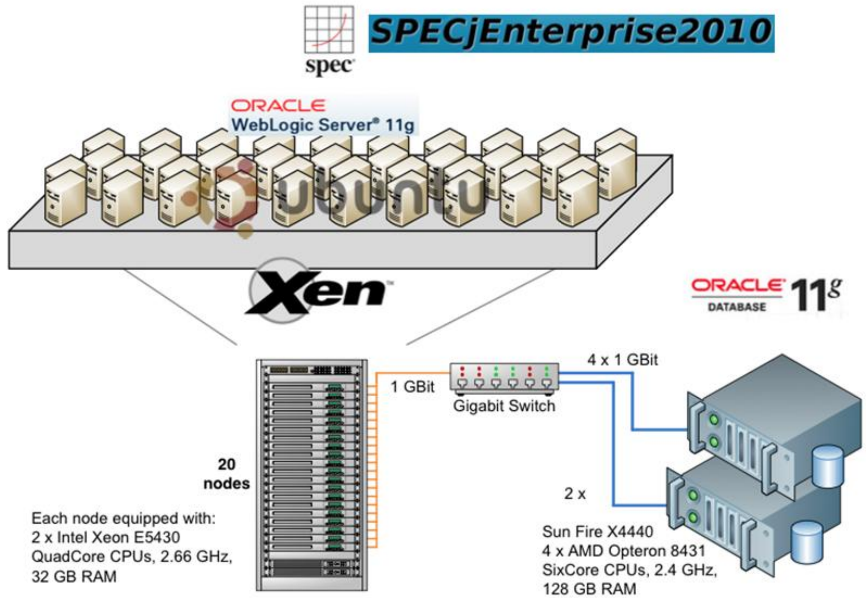

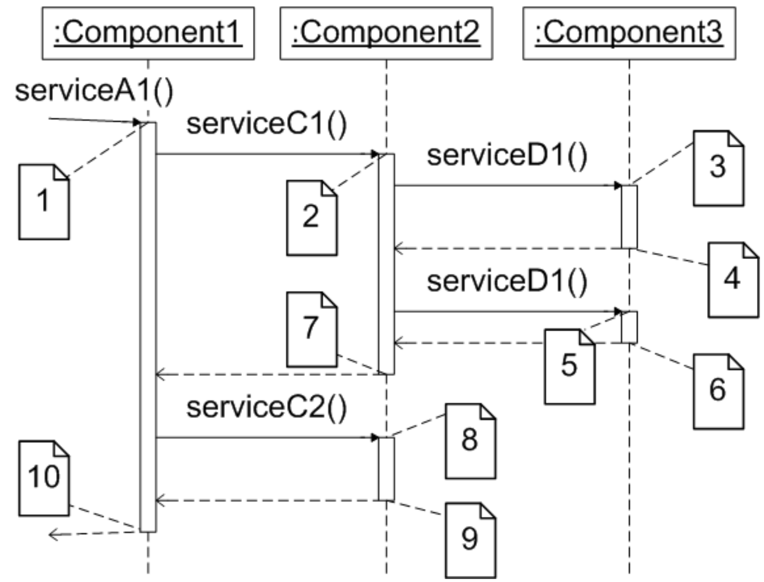

Details in:

- F. Brosig, N. Huber and S. Kounev. **Automated Extraction of Architecture-Level Performance Models of Distributed Component-Based Systems**. 26th IEEE/ACM International Conference On Automated Software Engineering (ASE 2011), Oread, Lawrence, Kansas, November 2011.
- S. Kounev, K. Bender, F. Brosig, N. Huber, and R. Okamoto. **Automated Simulation-Based Capacity Planning for Enterprise Data Fabrics**. In 4th International ICST Conference on Simulation Tools and Techniques (SIMUTools 2011), Barcelona, Spain, 2011. Best Paper Award.

Automated Model Extraction @ Run-Time

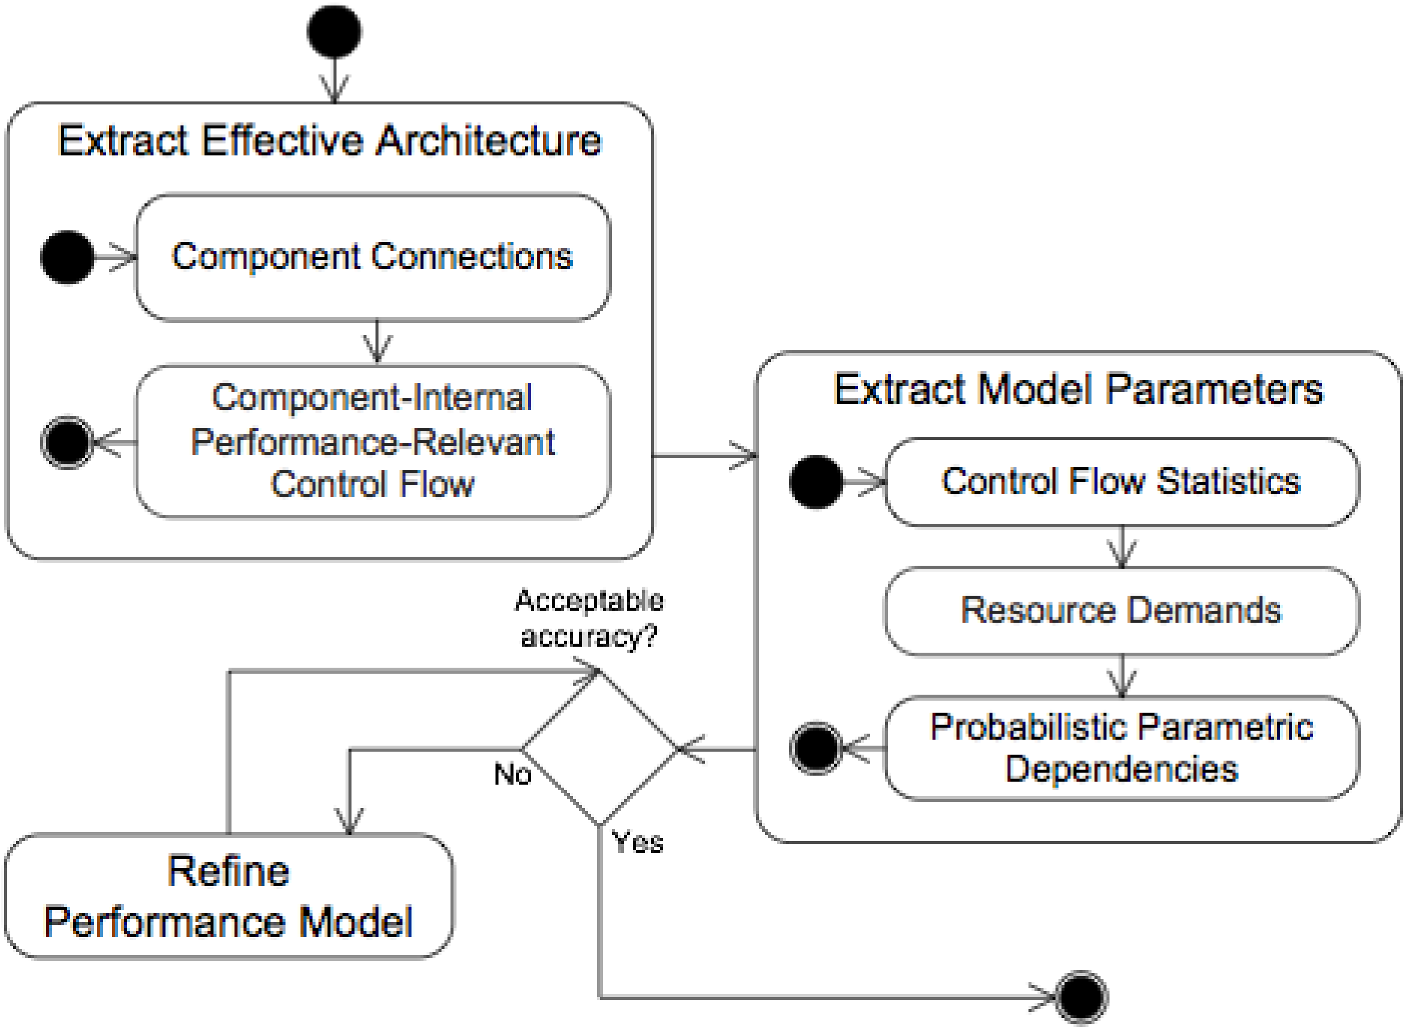

Automated Model Extraction @ Run-Time

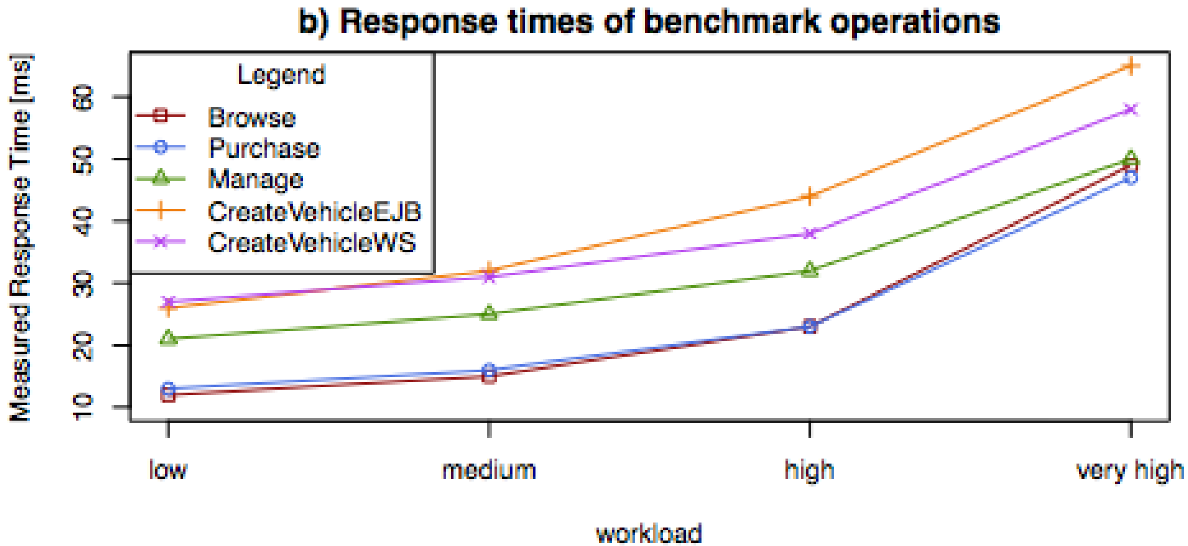

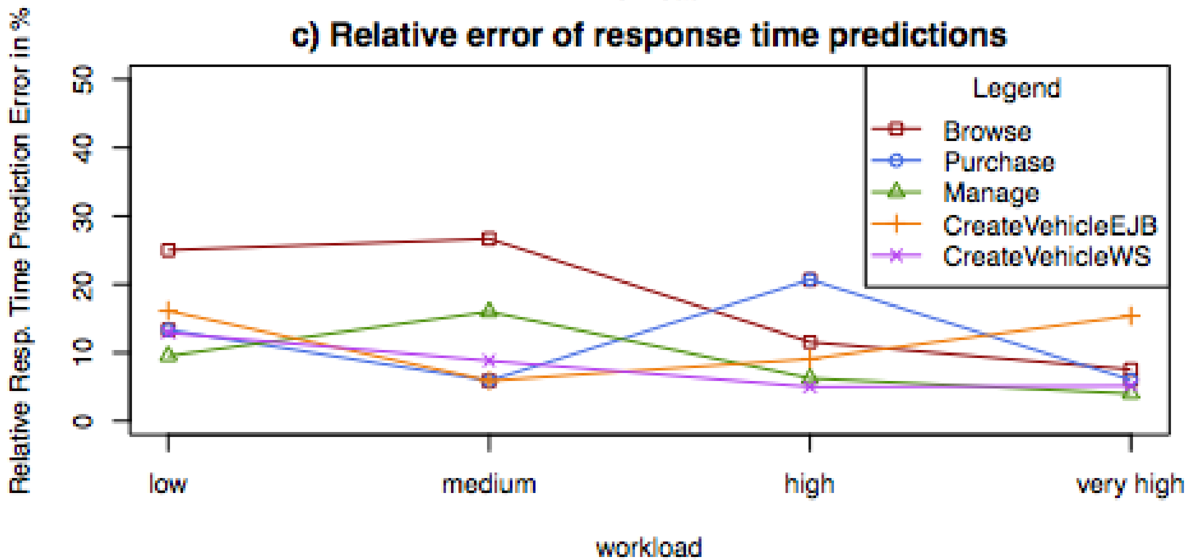

Automated Model Extraction @ Run-Time

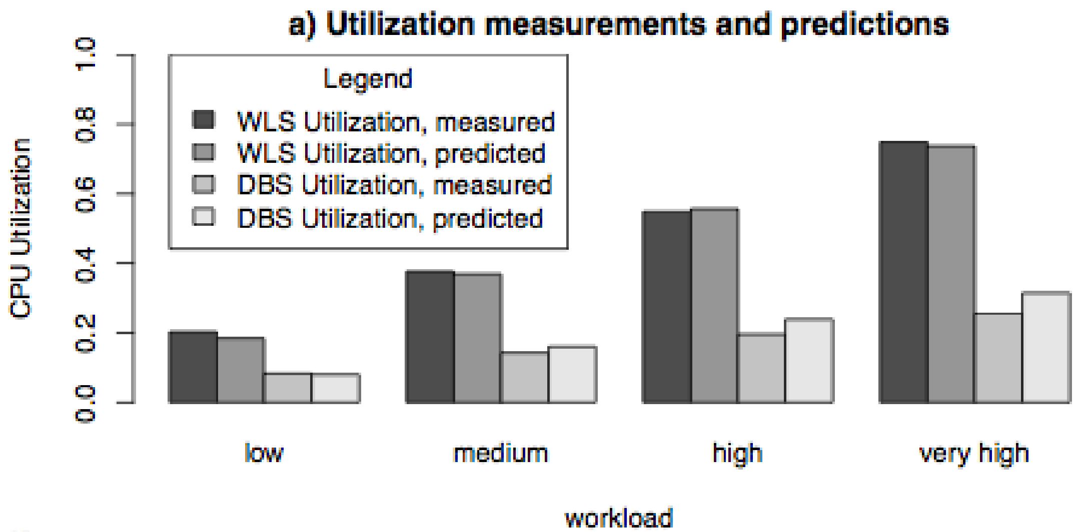

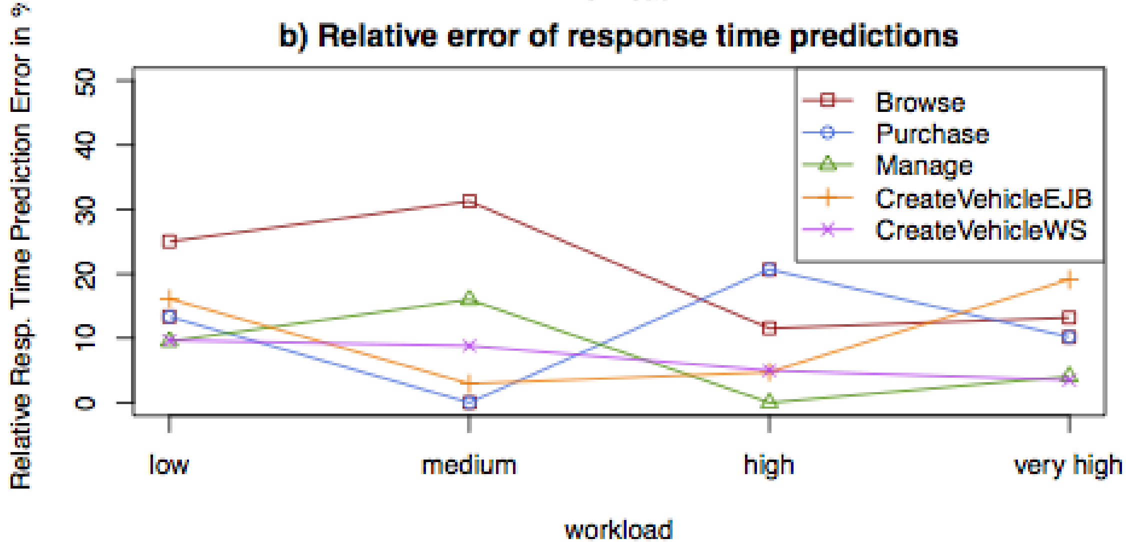

Agenda

- Motivation
- Approach & Methodology
- Exemplary Results
- **Vision**
- Conclusion

Today

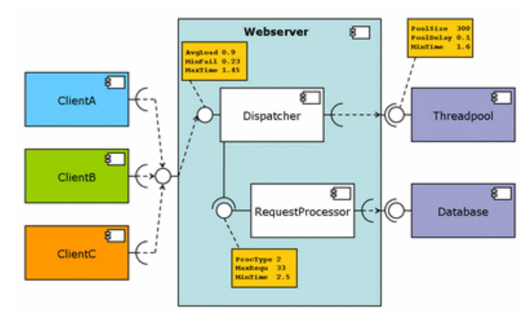

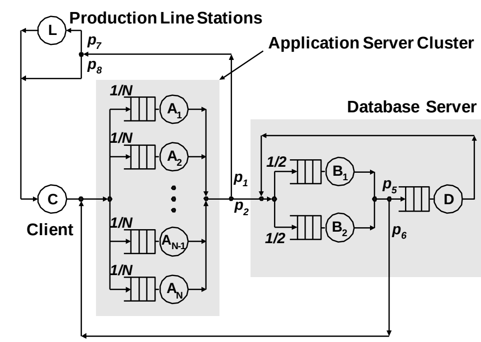

Systems

Models

The Future

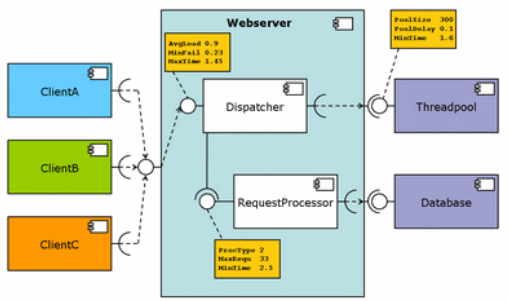

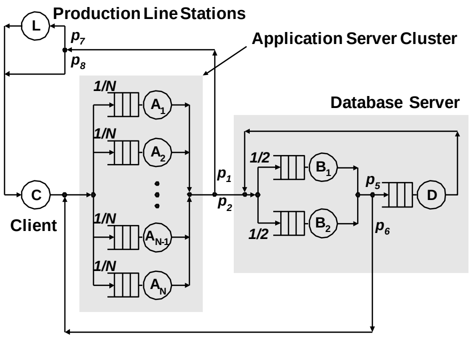

Systems **Models**

The Future

"I think, therefore I am …"
-- René Descartes

Self-Aware Software & Systems Engineering

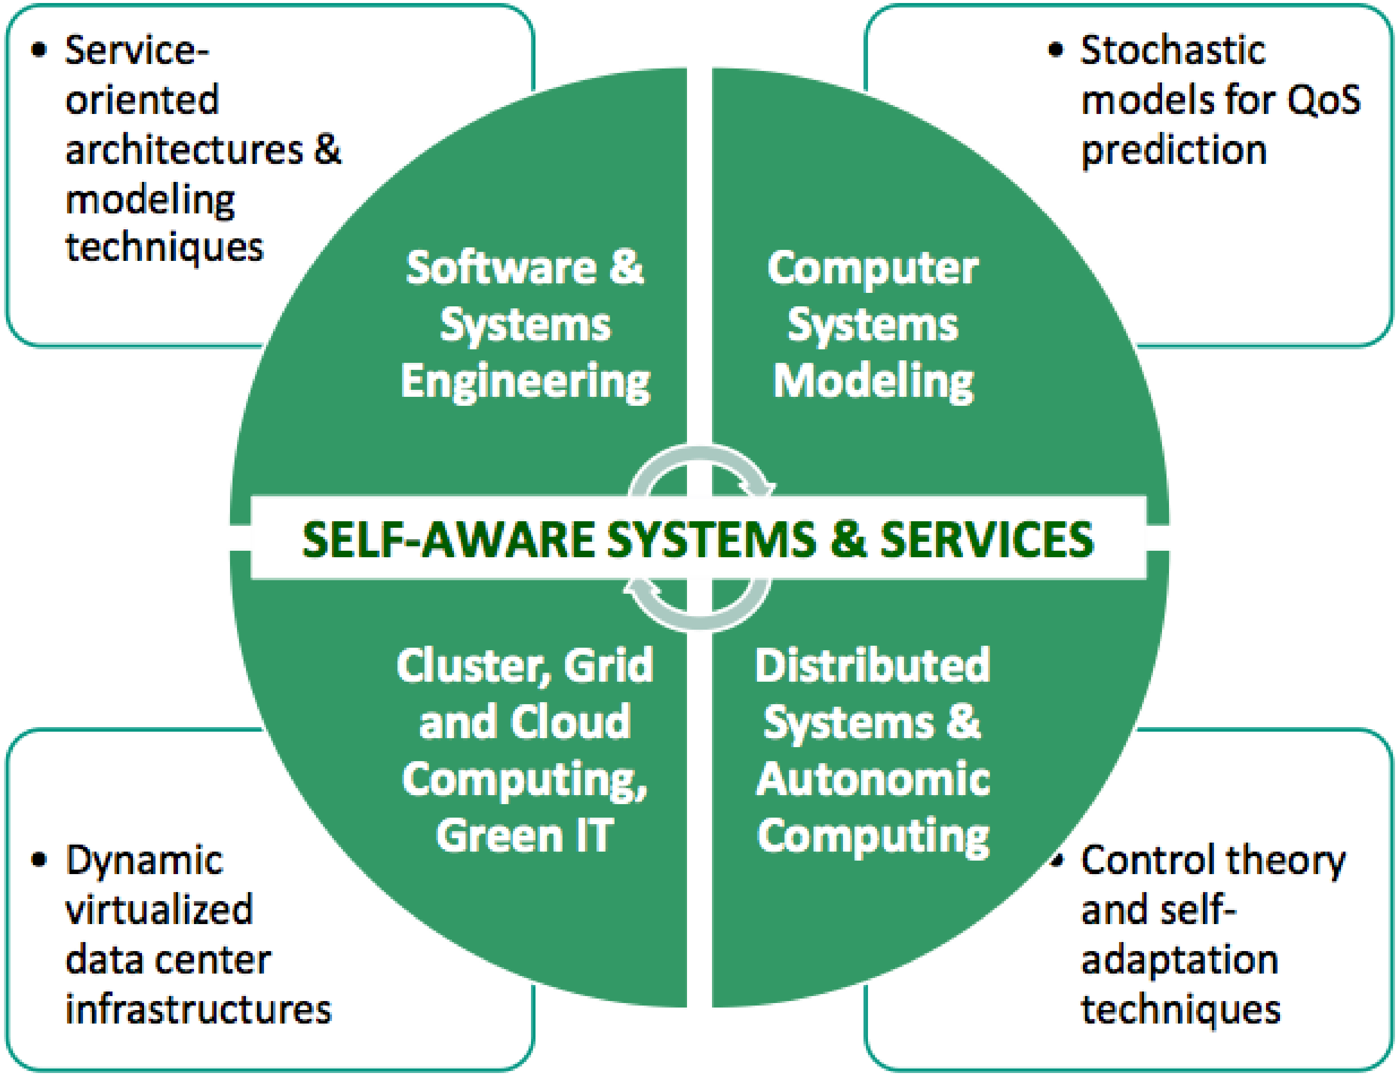

SPEC Research Group (SPEC RG) http://research.spec.org

- Founded in March 2011
- Platform for collaborative research in the area of quantitative system evaluation and analysis
- Foster interaction between industry and academia
- Wider scope
  - Metrics and benchmarking methodologies
  - Methods and tools for experimental system analysis
  - Covering all stages of the system lifecycle
  - Both existing and newly emerging technologies
  - Evaluation of early prototypes and research results
- Classical performance metrics: response time, throughput, scalability, resource/cost/energy efficiency, elasticity
- Plus dependability in general: Availability, reliability and security

Members (April 2011)

KIT Karlsruhe Institute of Technology

KIT Karlsruhe Institute of Technology

PURDUE UNIVERSITY

IBM

SAP

hp invent

Berkeley University of California

NORTHEASTERN UNIVERSITY BOSTON, MASSACHUSETTS

AMD

intel

CISCO SYSTEMS

TECHNISCHE UNIVERSITÄT DARMSTADT

vmware

ORACLE

RWTH AACHEN UNIVERSITY

THE UNIVERSITY OF TEXAS AT AUSTIN

UNIVERSITEIT GENT

SIEMENS

YAHOO!

UNIVERSITÄT MANNHEIM

UNIVERSITY OF ALBERTA

redhat

NICTA

IIT Bombay

SAN FRANCISCO STATE UNIVERSITY

浙江大學 ZHEJIANG UNIVERSITY

UNIVERSITY OF MINNESOTA Driven to Discover℠

· U C ·

Thank You!

http://www.descartes-research.net
http://www.relate-itn.eu
http://research.spec.org