Habilitationskolloquium

On Quality Issues in Complex Service-oriented Systems

Hong-Linh Truong
Distributed Systems Group, Institute of Information Systems
Vienna University of Technology

truong@dsg.tuwien.ac.at
http://dsg.tuwien.ac.at/Staff/truong

Outline

- Complex service-oriented systems - research challenges
- Selected works
  - Monitoring and analysis of performance metrics for workflows
  - Monitoring and analysis of data concerns and data contracts
  - Composable cost/performance evaluation for cloud applications
  - Programming human-based services in the cloud
- Conclusions and future work

Complex Service-oriented Systems

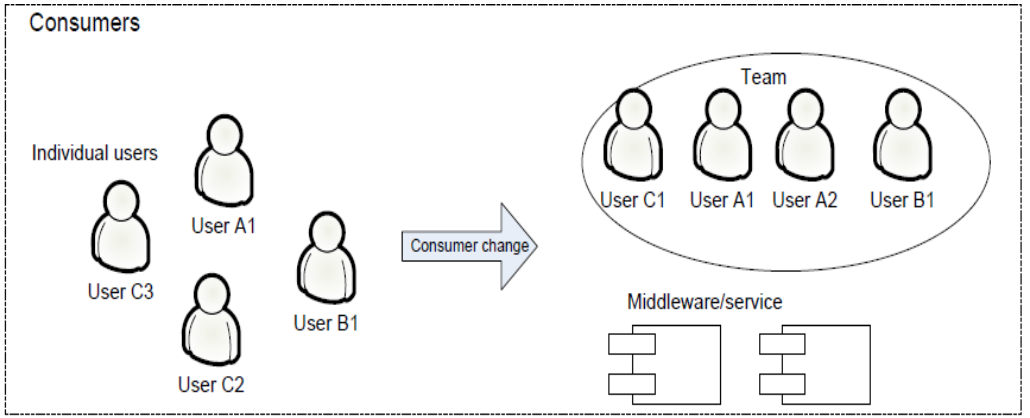

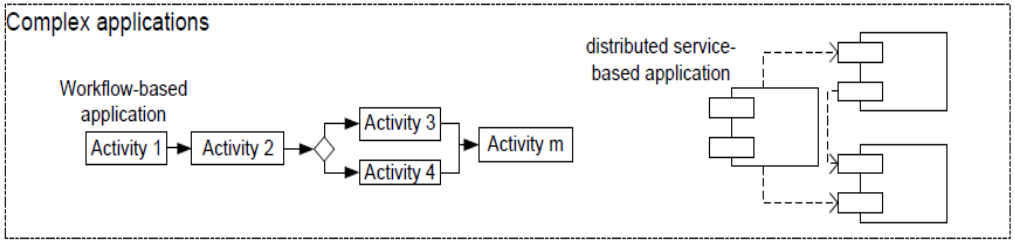

The complexity and diversity of computational, data and human resources as services, and rich service provisioning and utilization models lead to complex quality issues

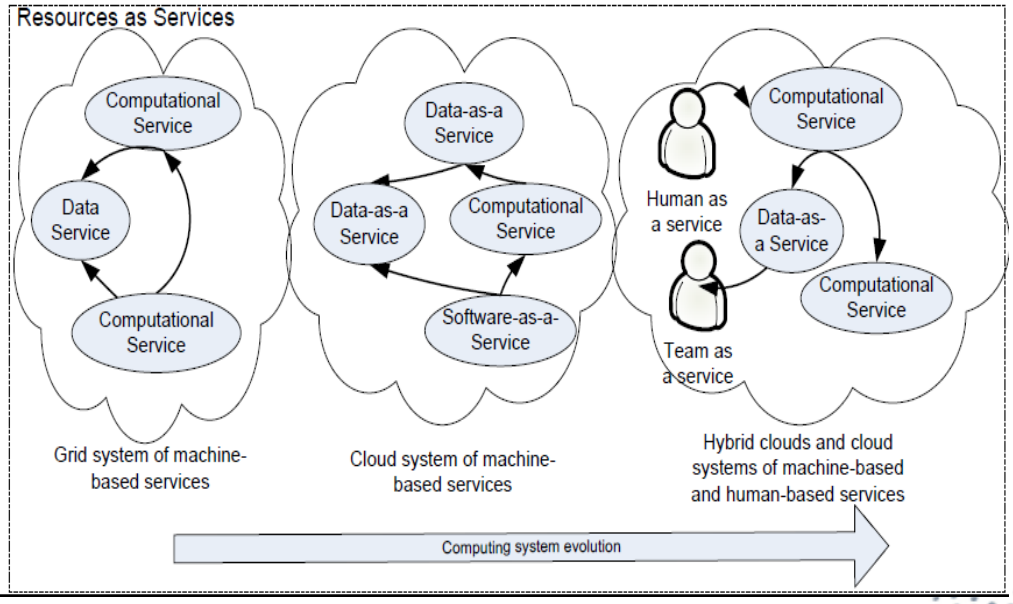

Example of complex service-oriented systems

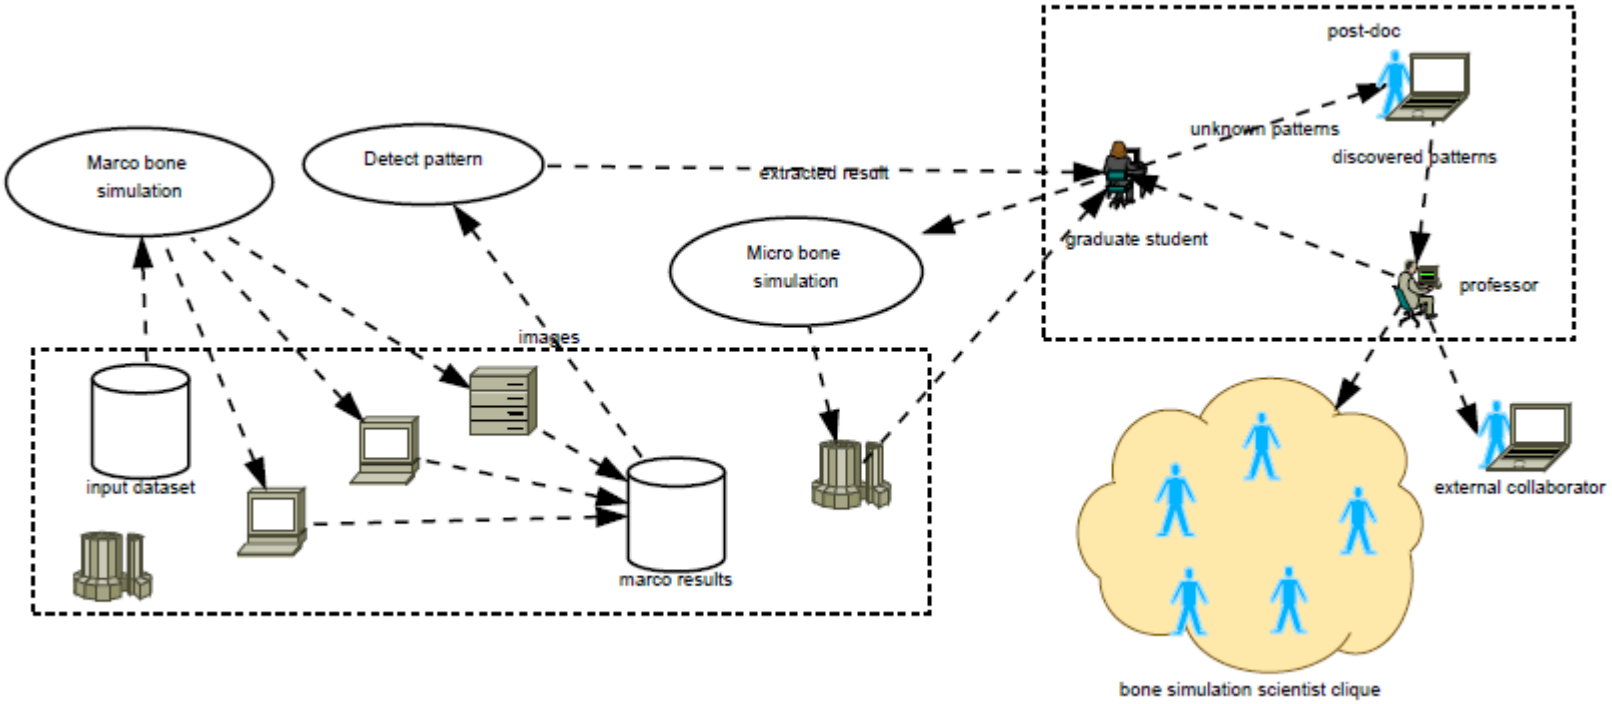

- Multiple types of resources
- Diverse types of interactions and roles
- Multiple perspectives in quality monitoring, analysis and utilization

Multiple types of resources

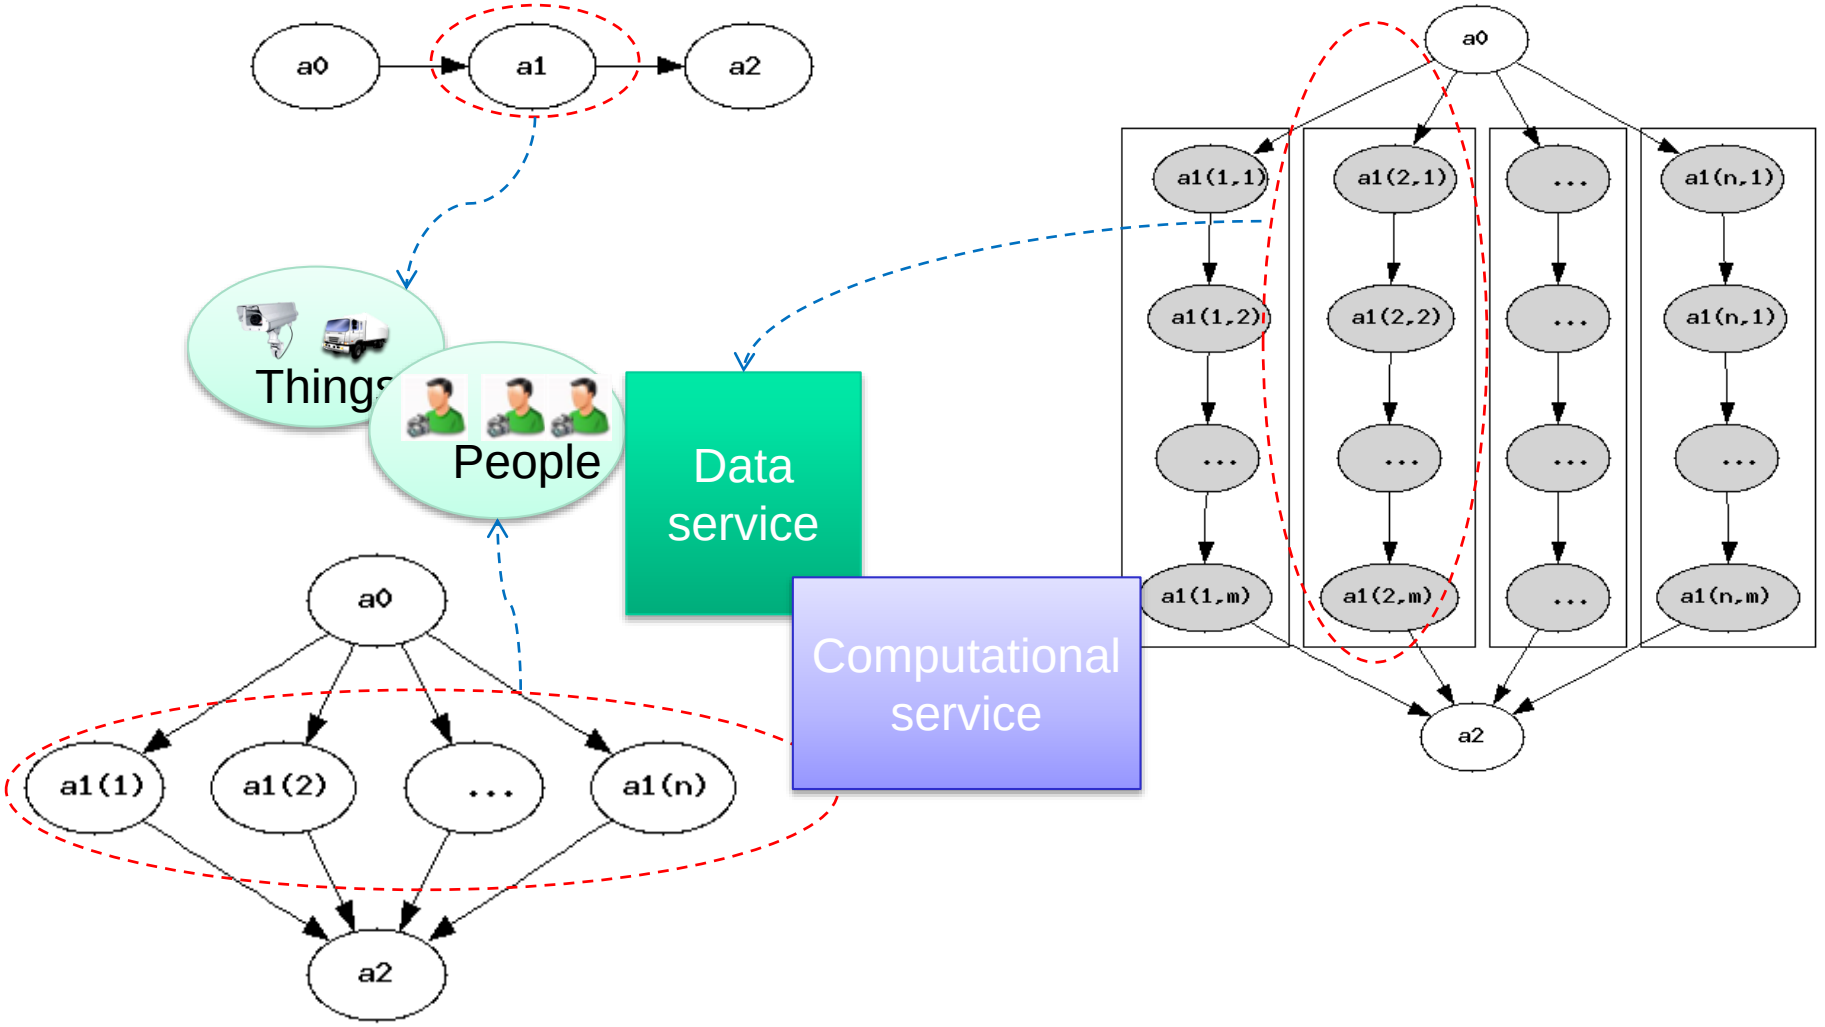

Multiple types of quality metrics

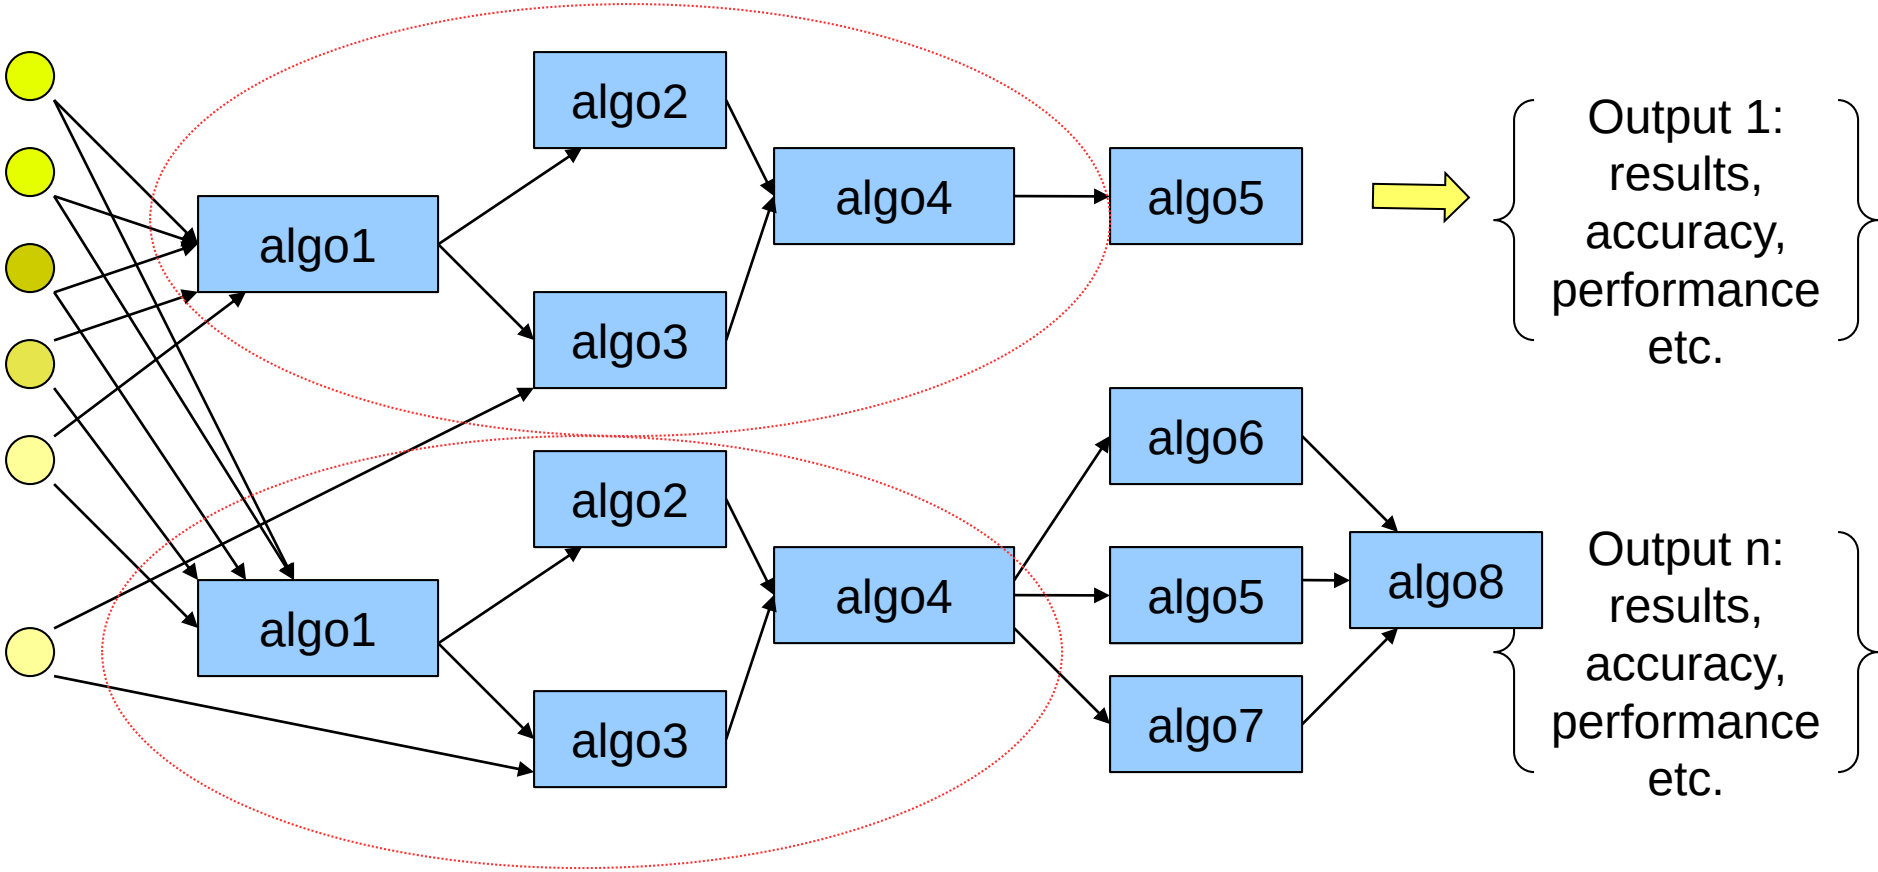

- N separated workflows for M consumers: common activities but different influence factors/expected quality

Multiple perspectives of elasticity in computing

1. Elastic demands from consumers
2. Output elasticity with different price and quality
3. Elastic data inputs, e.g., deal with opportunistic data
4. Elastic pricing and quality models associated resources

Multiple elasticity dimensions

Resource elasticity:
software/human-based computing elements, multiple clouds

Quality elasticity: performance, quality of data, service availability, human trust, ...

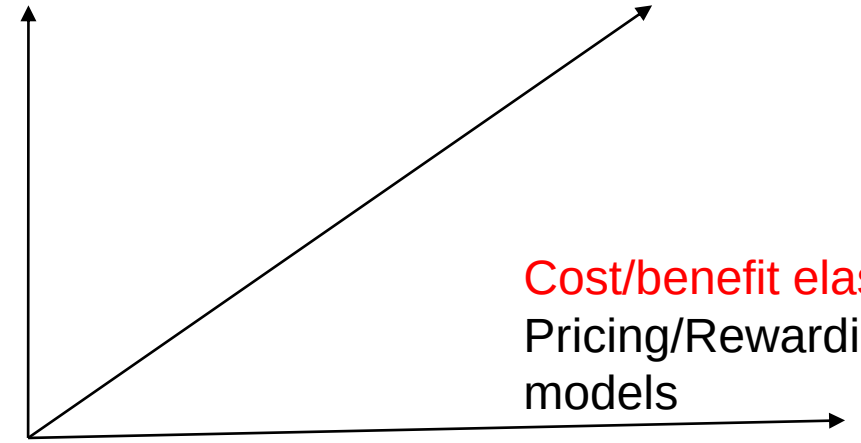

Cost/benefit elasticity:
Pricing/Rewarding/Incentive models

Elasticity

- Scaling software, services, and people in the same system
- But elasticity is not just “resource elasticity”: cost/benefit and quality are also important

Schahram Dustdar, Yike Guo, Benjamin Satzger, Hong Linh Truong: Principles of Elastic Processes. IEEE Internet Computing 15(5): 66-71 (2011

The Vienna Elastic Computing Model

- **Multi-dimensional elasticity**
- **Service computing models, and**
- **Cloud provisioning models**

Schahram Dustdar, Hong Linh Truong: Virtualizing Software and Humans for Elastic Processes in Multiple Clouds- a Service Management Perspective. IJNGC 3(2) (2012)

Focused research challenges

quality issues in elastic complex service-oriented systems

- How to characterize and evaluate
  - performance metrics, human interactions, quality of data *associated with software, humans and data* in complex systems
  - Quality of service, data concerns, and service contracts *associated with different types of services*
- What are the interdependencies among different types of quality metrics associated with different types of resources?
- How to utilize these metrics for service engineering and optimization to support elasticity?

Contribution area: monitoring and analysis of Grid workflow performance

Characterizing and evaluating quality for Grid workflows

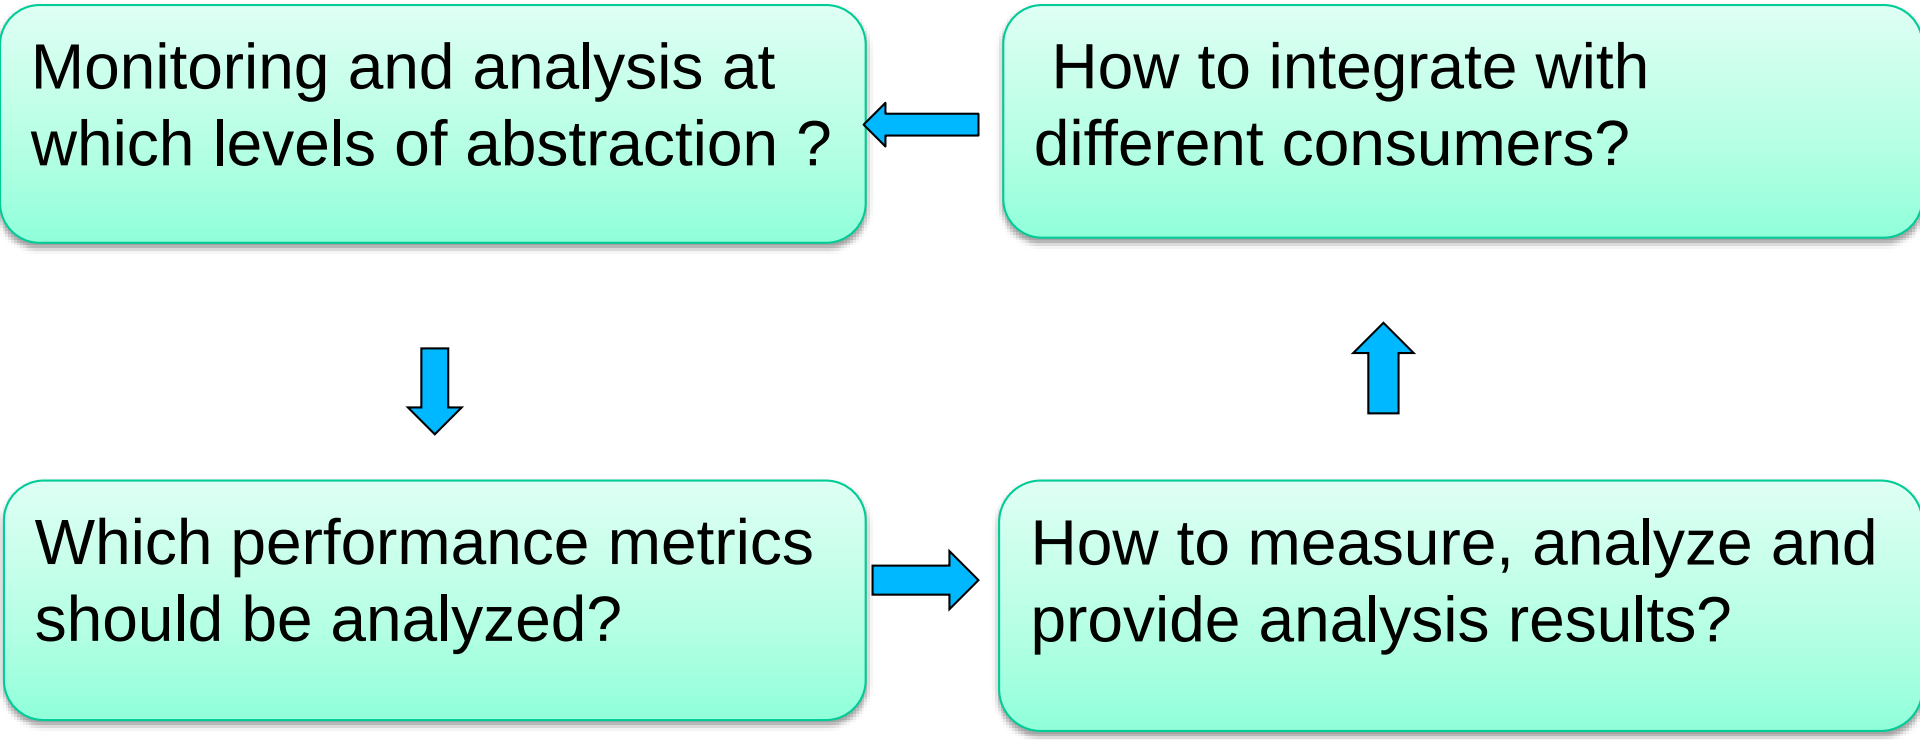

Contribution: performance metrics for Grid workflows

Abstracting workflows for quality monitoring, measurement and evaluation

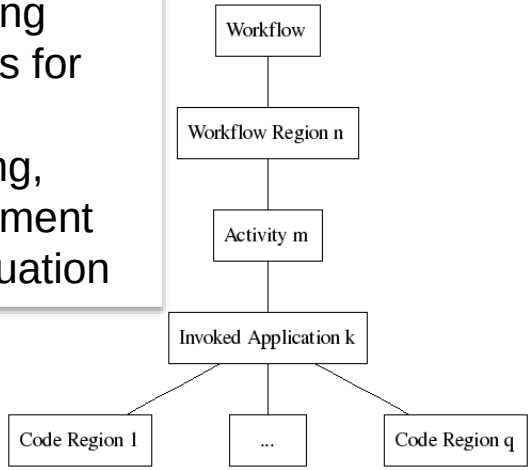

Multiple types of consumers

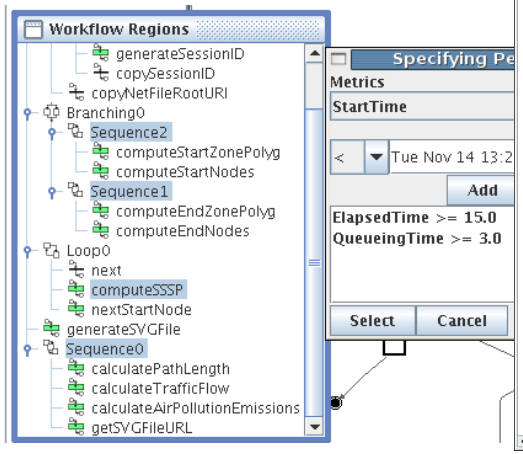

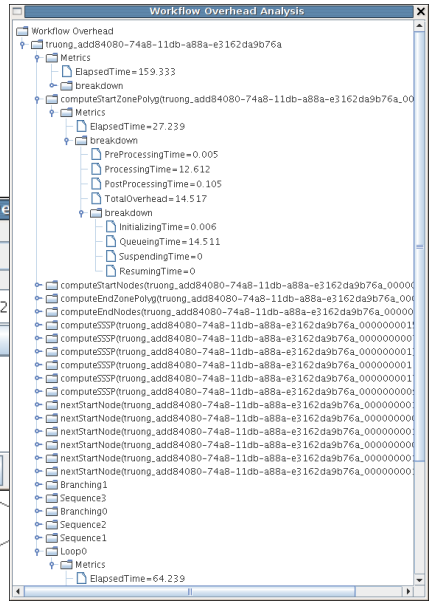

Characterizing performance metrics

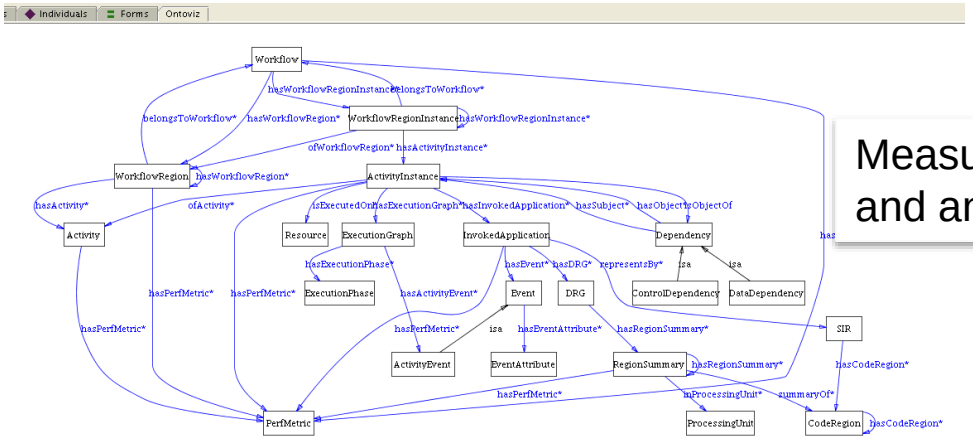

Measurement and analysis

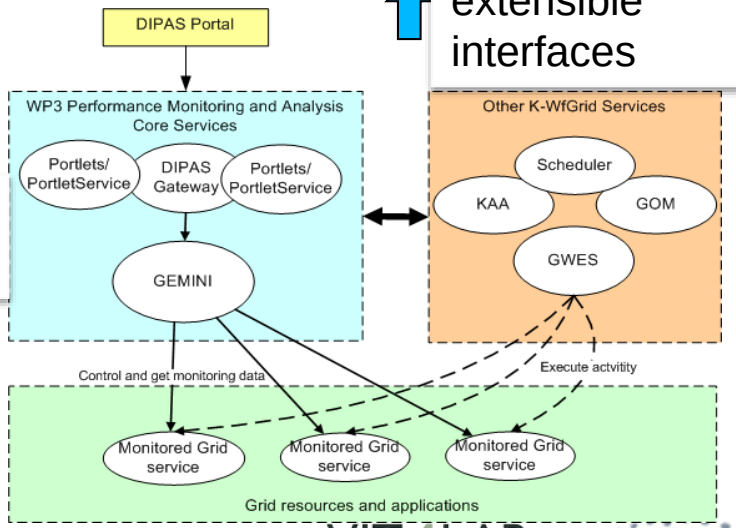

extensible interfaces

VITALAB Distributed Systems Group

Contribution: performance metrics for Grid workflows

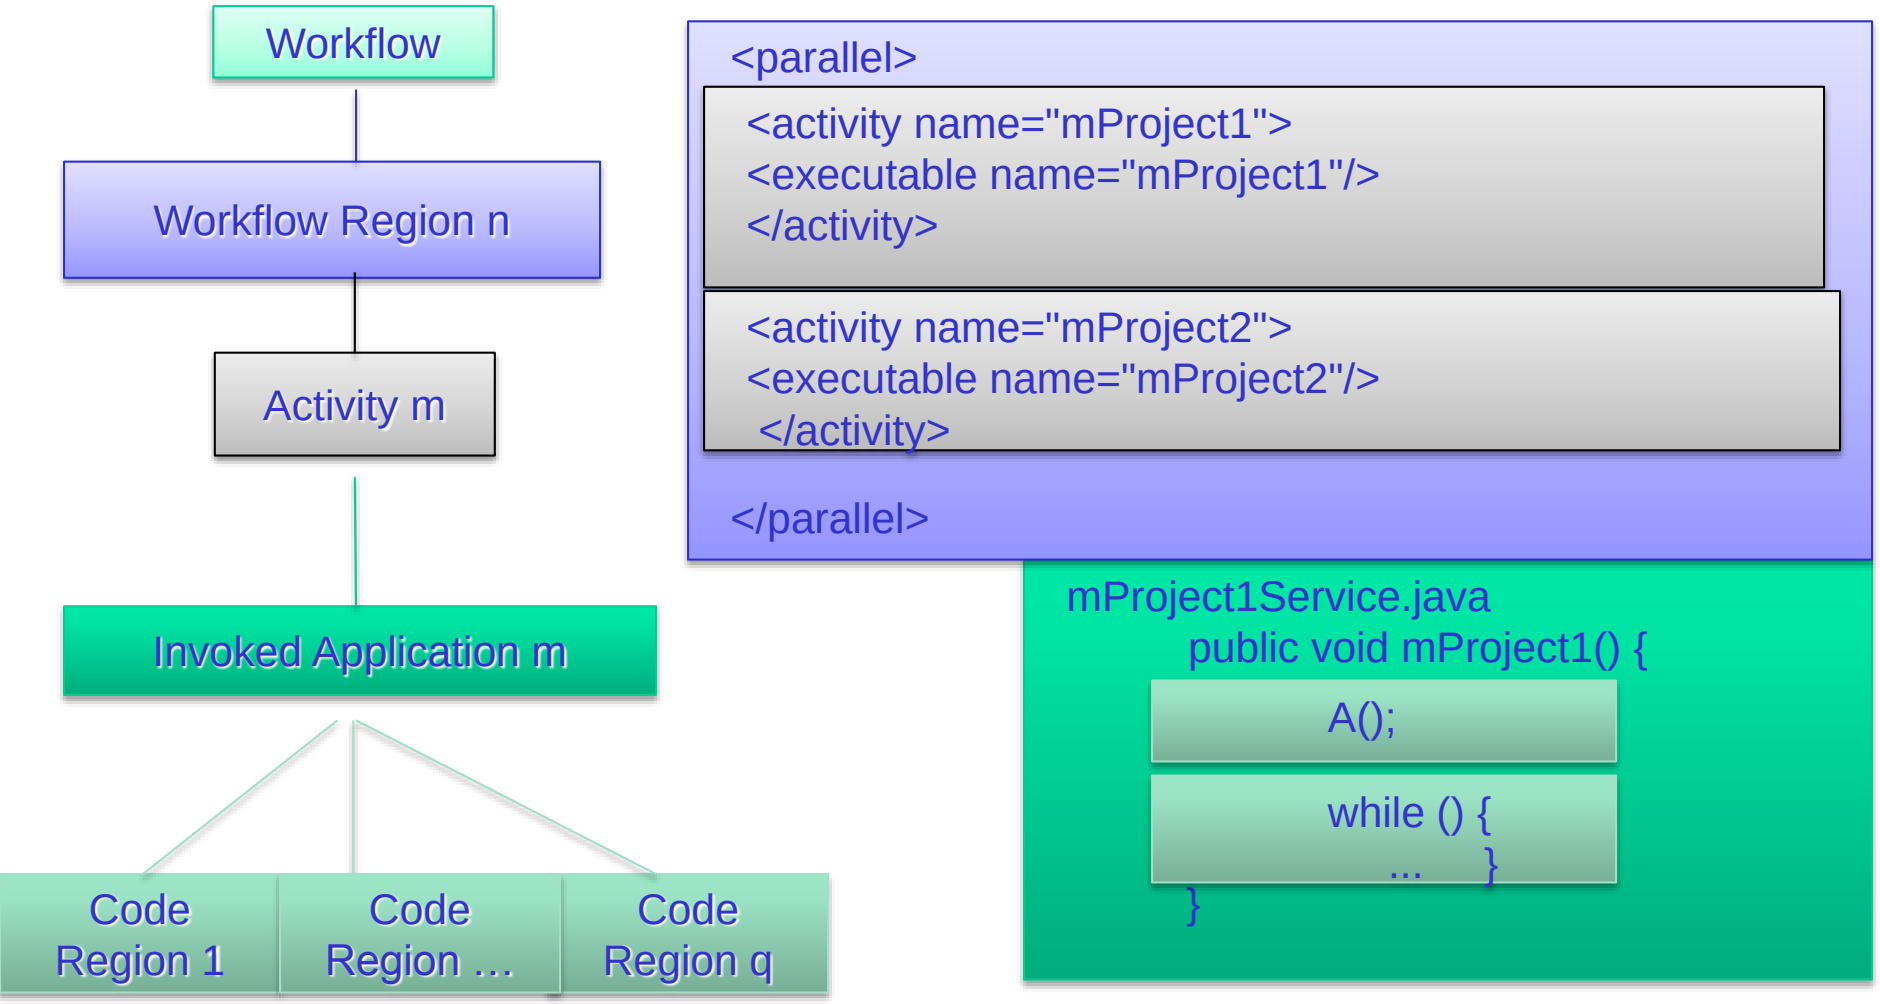

Hong Linh Truong, Schahram Dustdar, Thomas Fahringer: Performance metrics and ontologies for Grid workflows. Future Generation Comp. Syst. 23(6): 760-772 (2007)

Examples of metrics

Table 2
Performance metrics at code region level

Category	Metric name	Description

Table 4
Performance metrics at activity level

Category	Metric name	Description

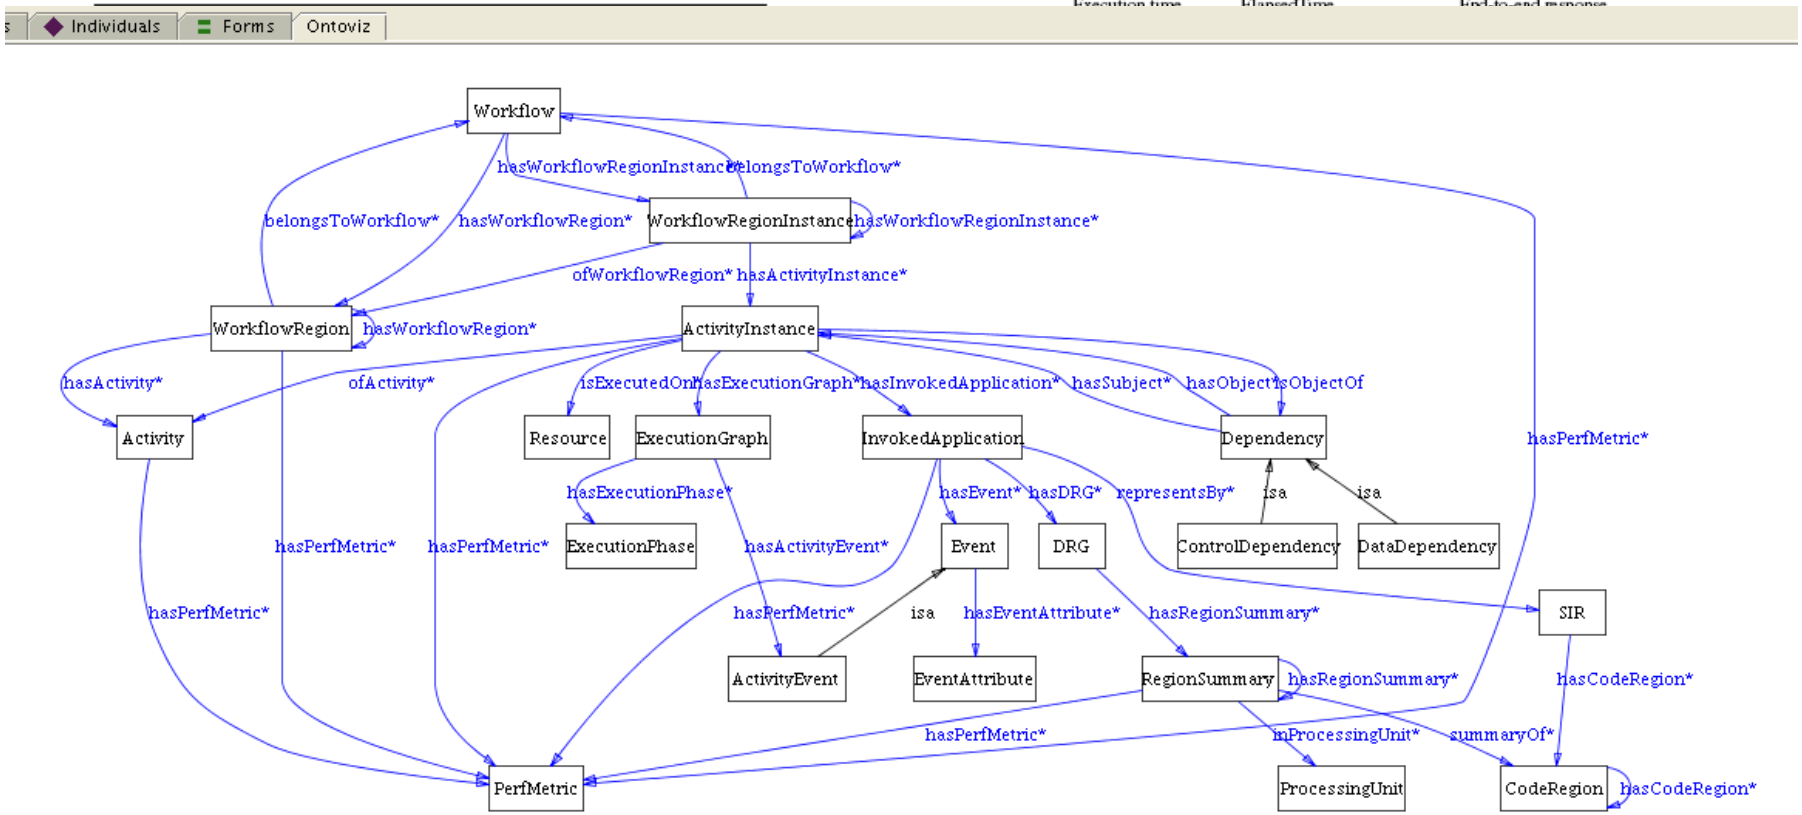

		hardware counters.
Temporal overhead	OCTRP, etc.	This type of metrics is defined only for parallel code regions.

		between a pair of activities.
Synchronization	SynDelay ExecDelay	Synchronization delay. Execution delay.

Contribution: distributed and online performance monitoring and analysis

Performance overhead analysis and search for performance problems

PDQS and WARL for performance service integration

PDQS, WIRL, and SIRWF for instrumentation and monitoring

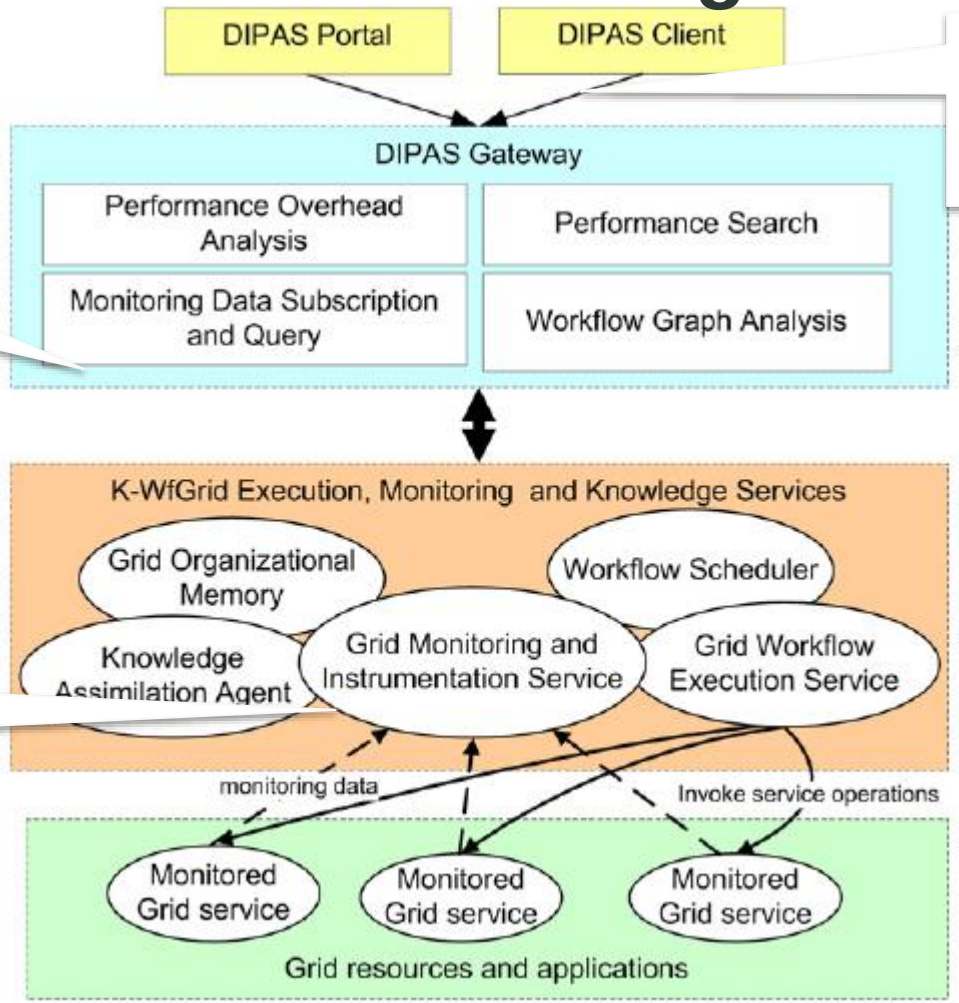

Hong Linh Truong, Peter Brunner, Vlad Nae, Thomas Fahringer: DIPAS: A distributed performance analysis service for grid service-based workflows. Future Generation Comp. Syst. 25(4): 385-398 (2009)

Online workflow tracing

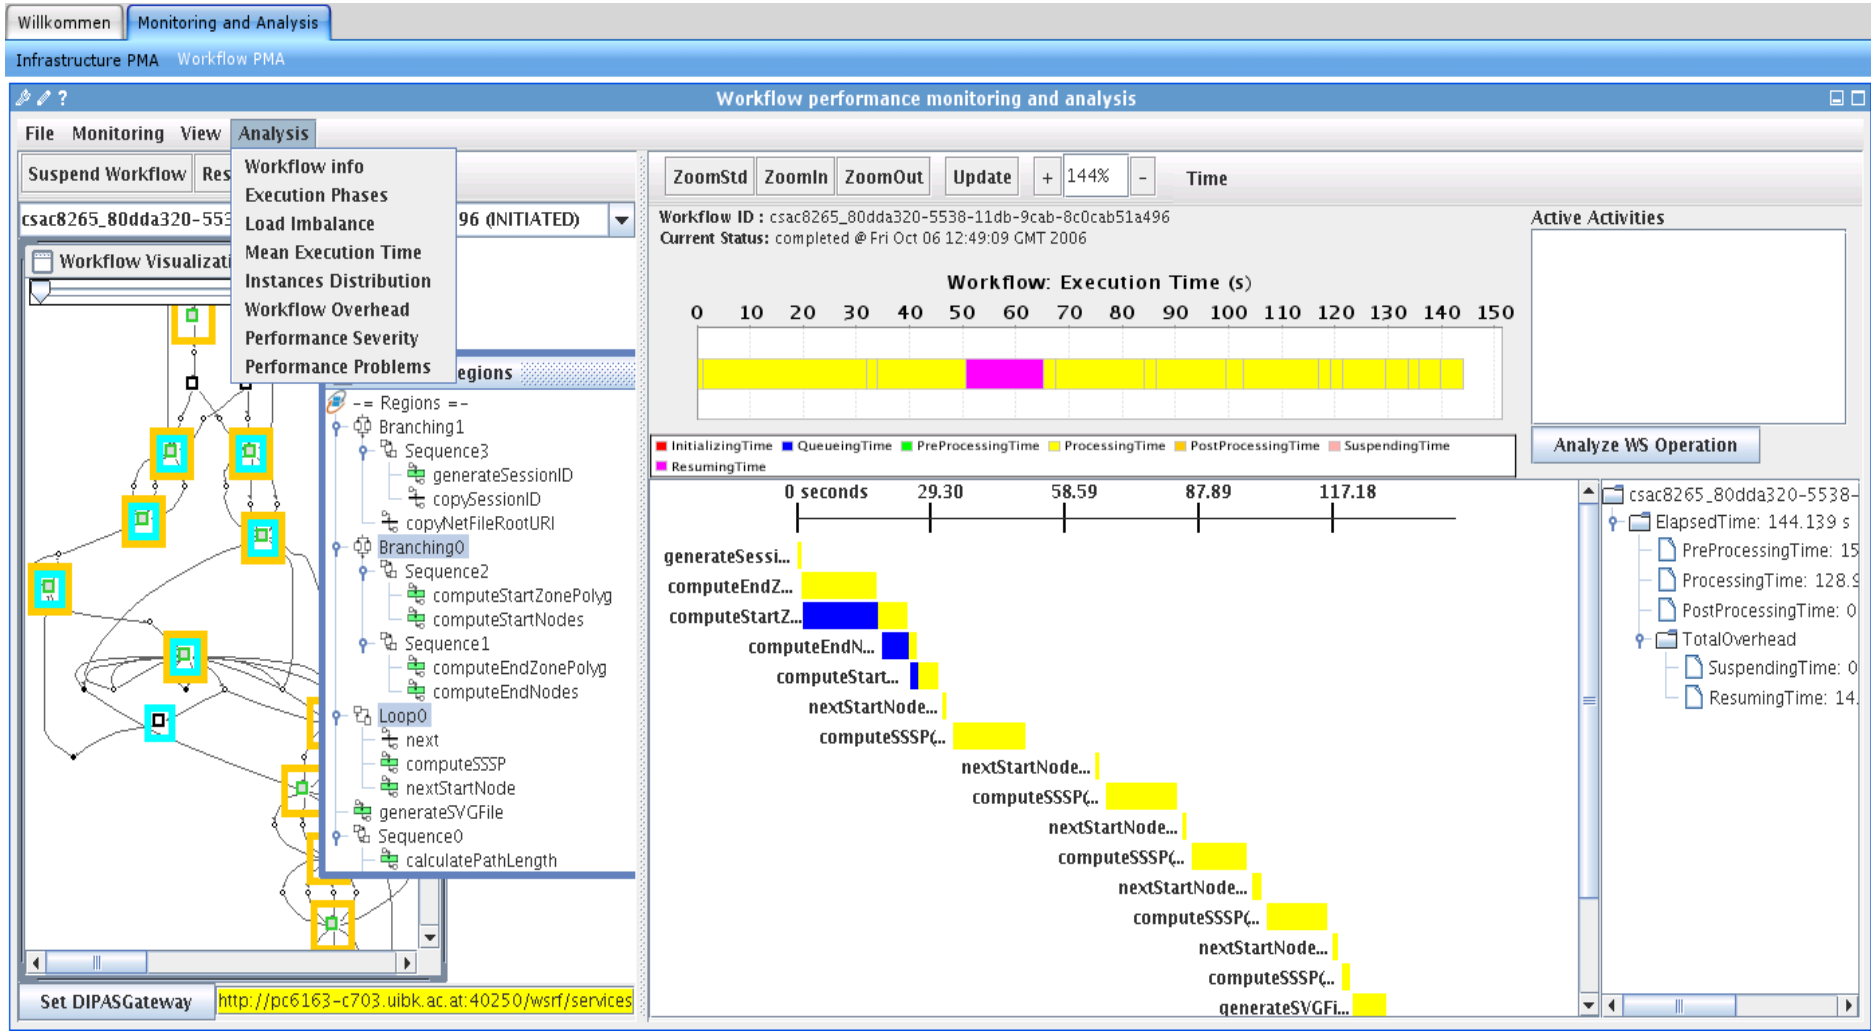

VITALAB Distributed Systems Group

Observing performance metrics and problems

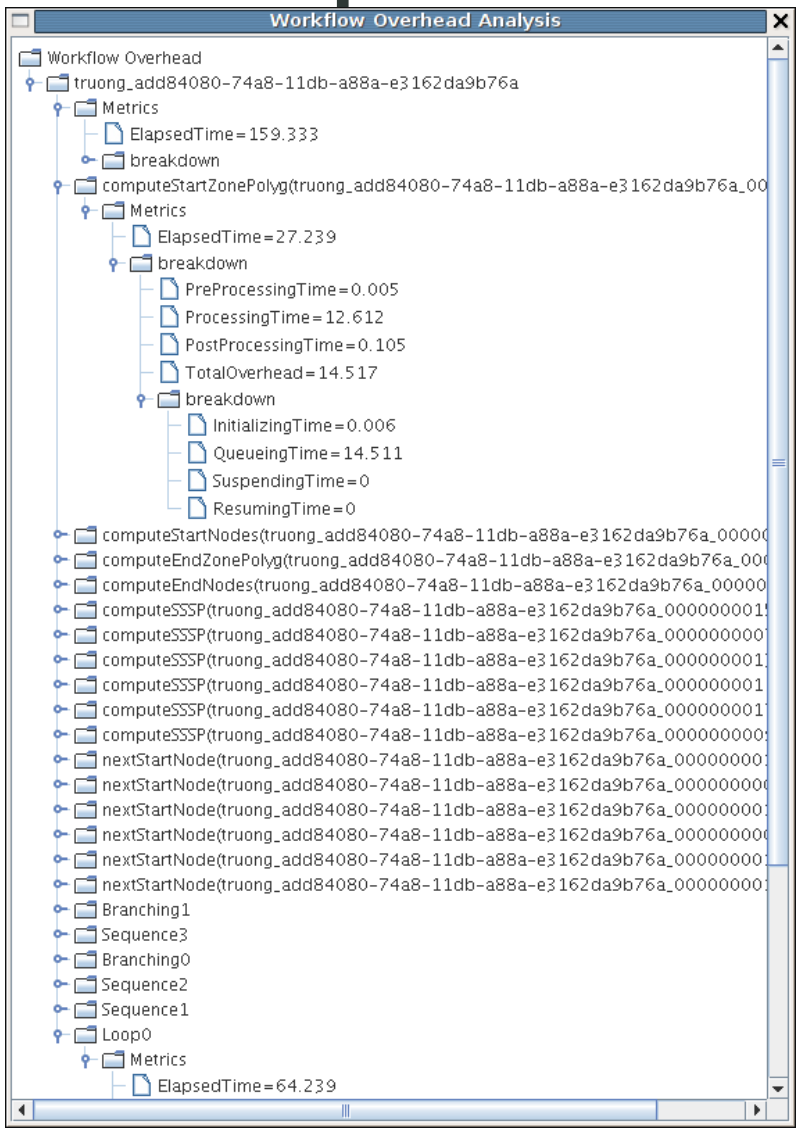

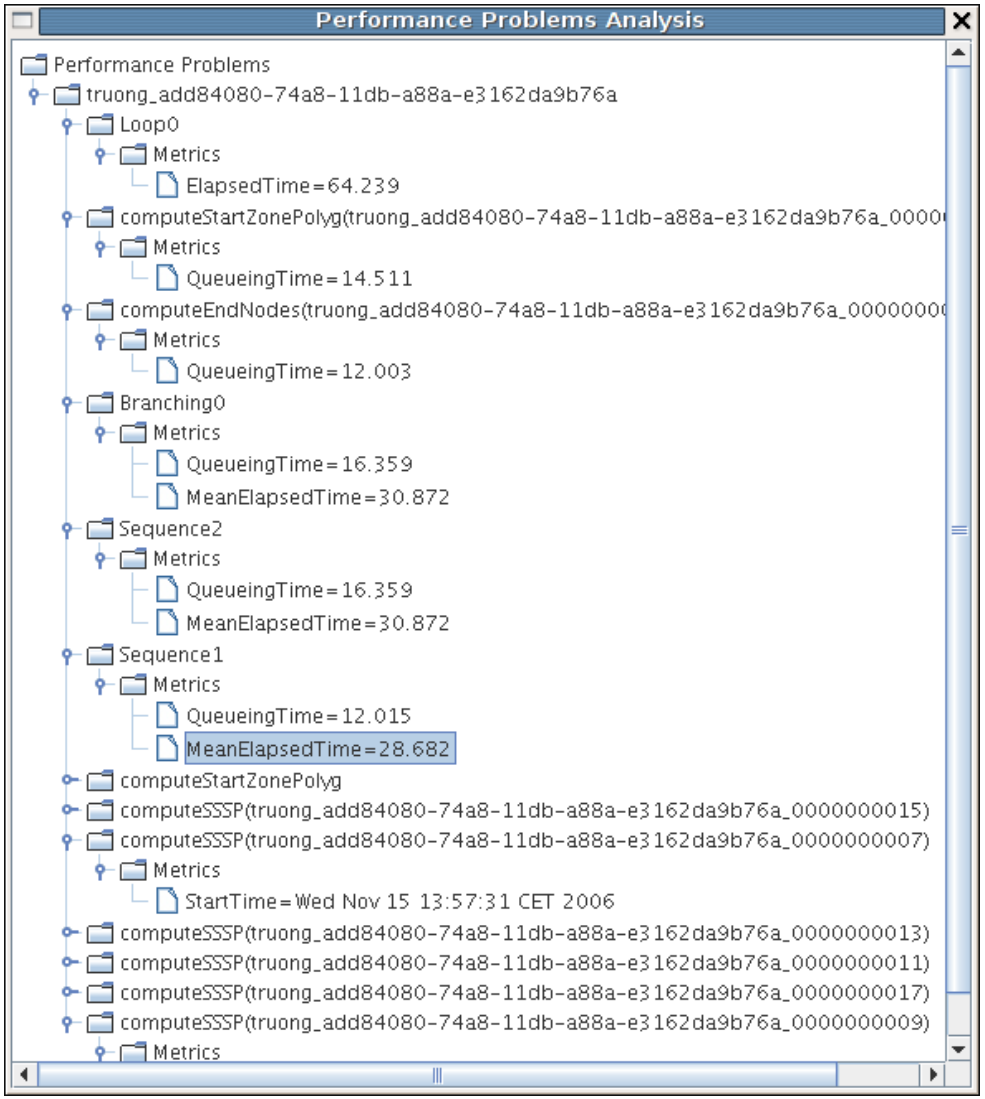

Summary of contributions

- Contributions
  - A comprehensive study of performance metrics associated with multiple levels of Grid workflow abstraction
  - Novel techniques and a framework for online and distributed monitoring and analysis of Grid workflows
  - Online demos/movies and software: http://www.dps.uibk.ac.at/projects/kwfgrid
- Related work
  - Most tools focus on activity levels and analysis at the workflow engine level, lack of comprehensive studies of performance metrics and target mainly to end users
- Selected publications
  - Hong-Linh Truong, Schahram Dustdar, Thomas Fahringer: Performance metrics and ontologies for Grid workflows. Future Generation Comp. Syst. 23(6): 760-772 (2007)
  - Hong-Linh Truong, Peter Brunner, Vlad Nae, Thomas Fahringer: DIPAS: A distributed performance analysis service for grid service-based workflows. Future Generation Comp. Syst. 25(4): 385-398 (2009)
  - Hong-Linh Truong, Schahram Dustdar: Online Interaction Analysis Framework for Ad-Hoc Collaborative Processes in SOA-Based Environments. T. Petri Nets and Other Models of Concurrency 2: 260-277 (2009)
  - Michael Reiter, Uwe Breitenbucher, Schahram Dustdar, Dimka Karastoyanova, Frank Leymann, Hong-Linh Truong: A Novel Framework for Monitoring and Analyzing Quality of Data in Simulation Workflows. eScience 2011: 105-112
  - Michael Reiter, Hong Linh Truong, Schahram Dustdar, Dimka Karastoyanova, Robert Krause, Frank Leymann, Dieter Pahr: On Analyzing Quality of Data Influences on Performance of Finite Elements Driven Computational Simulations. Euro-Par 2012: 793-804
  - Hong-Linh Truong, Robert Samborski, Thomas Fahringer: Towards a Framework for Monitoring and Analyzing QoS Metrics of Grid Services. e-Science 2006: 65

Contribution area: monitoring and analysis of data concerns and data/service contracts

Quality issues in data-as-a-service

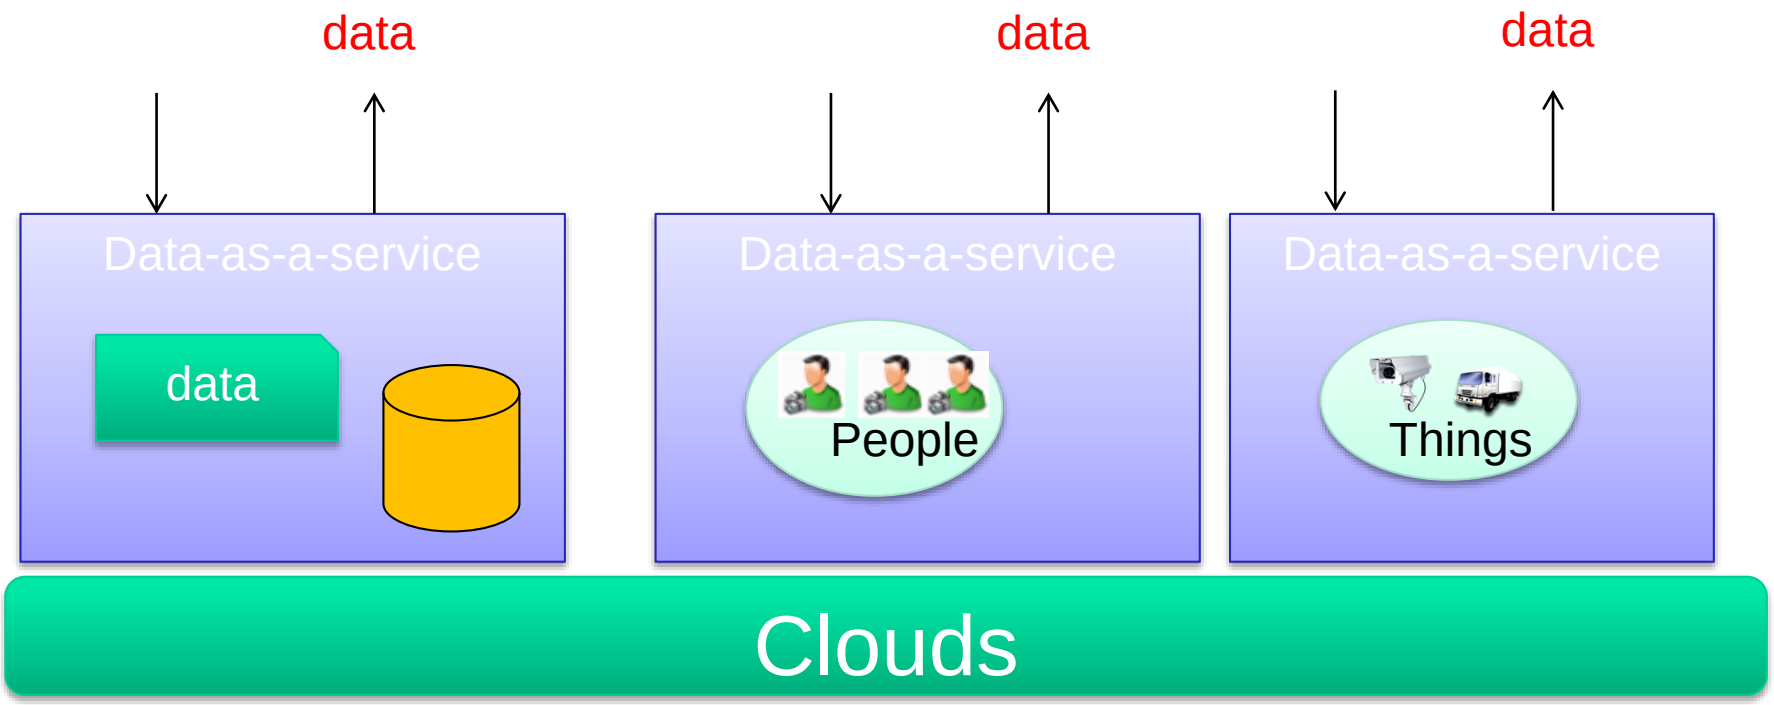

- *Multiple stakeholders, the rise of data provisioning, cloud models: → data-as-a-service*
- *Provide data capabilities* rather than provide computation or other software capabilities

DaaS design & implementation – not just „functional“ aspects

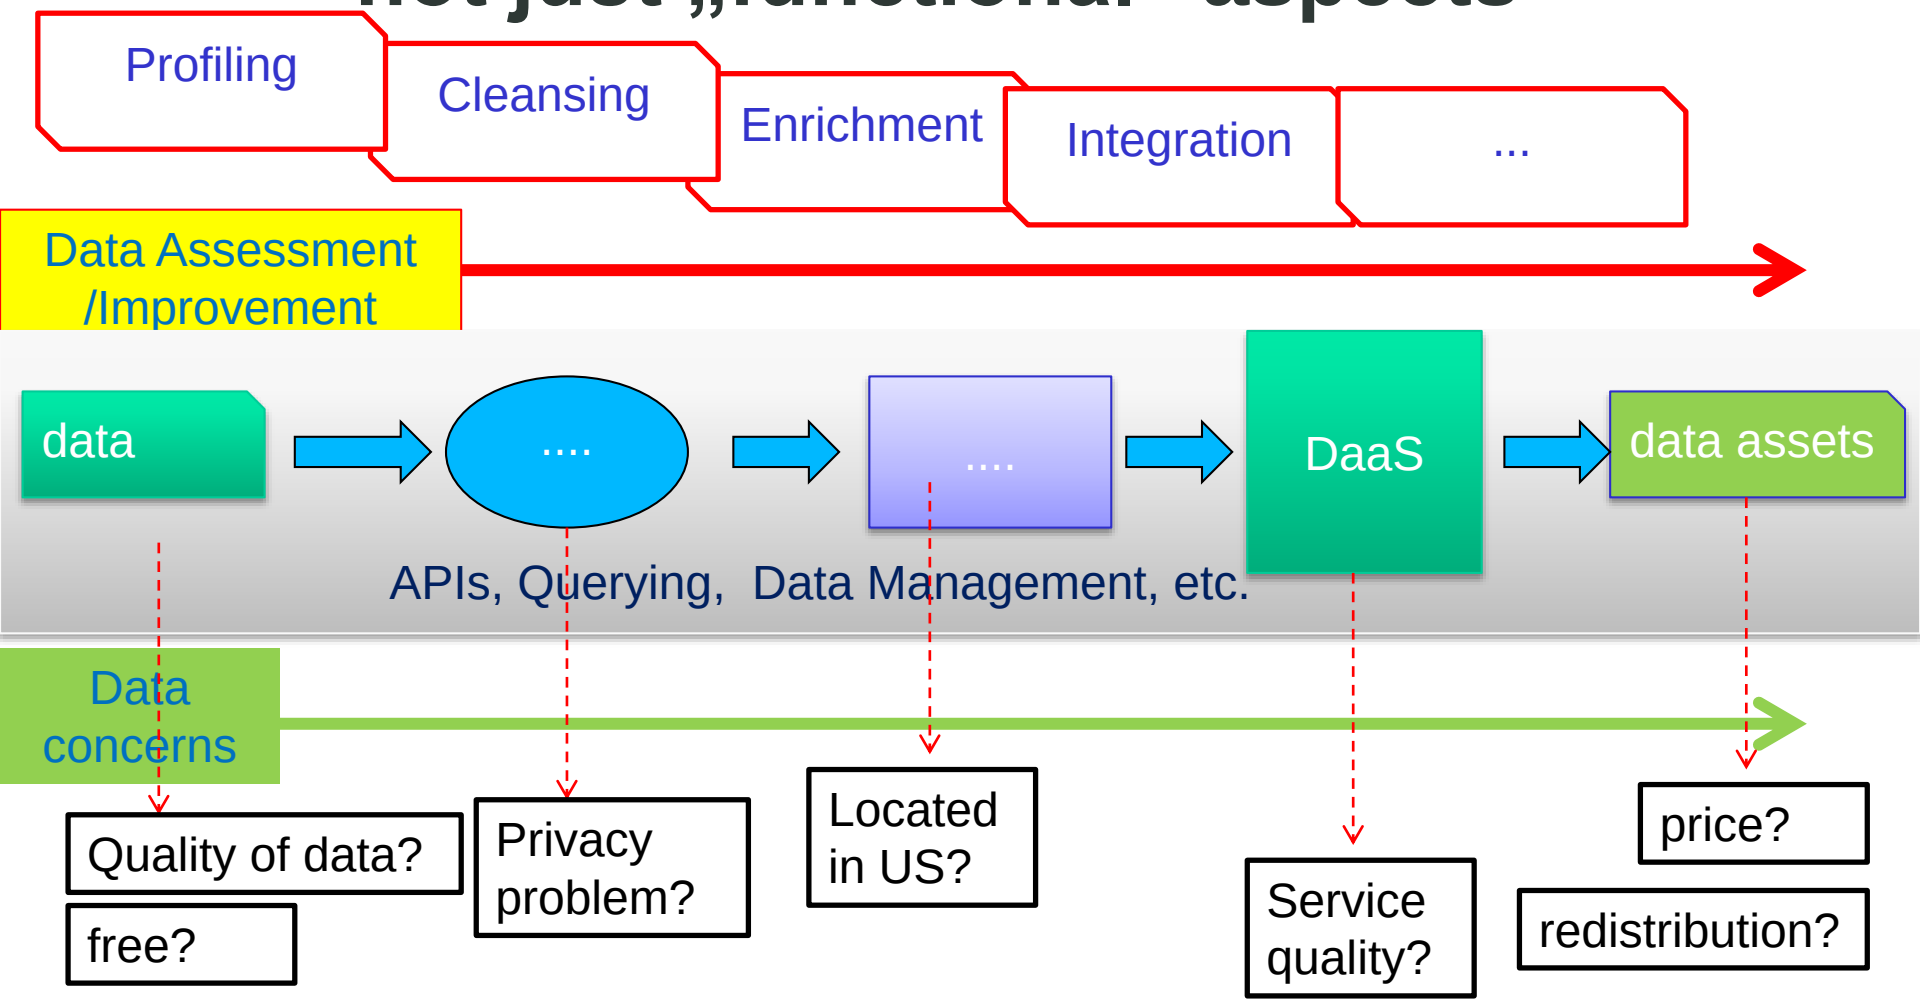

Schahram Dustdar, Reinhard Pichler, Vadim Savenkov, Hong Linh Truong: Quality-aware service-oriented data integration: requirements, state of the art and open challenges. SIGMOD Record 41(1): 11-19 (2012)

Contribution: analysis and conceptual model for DaaS concerns

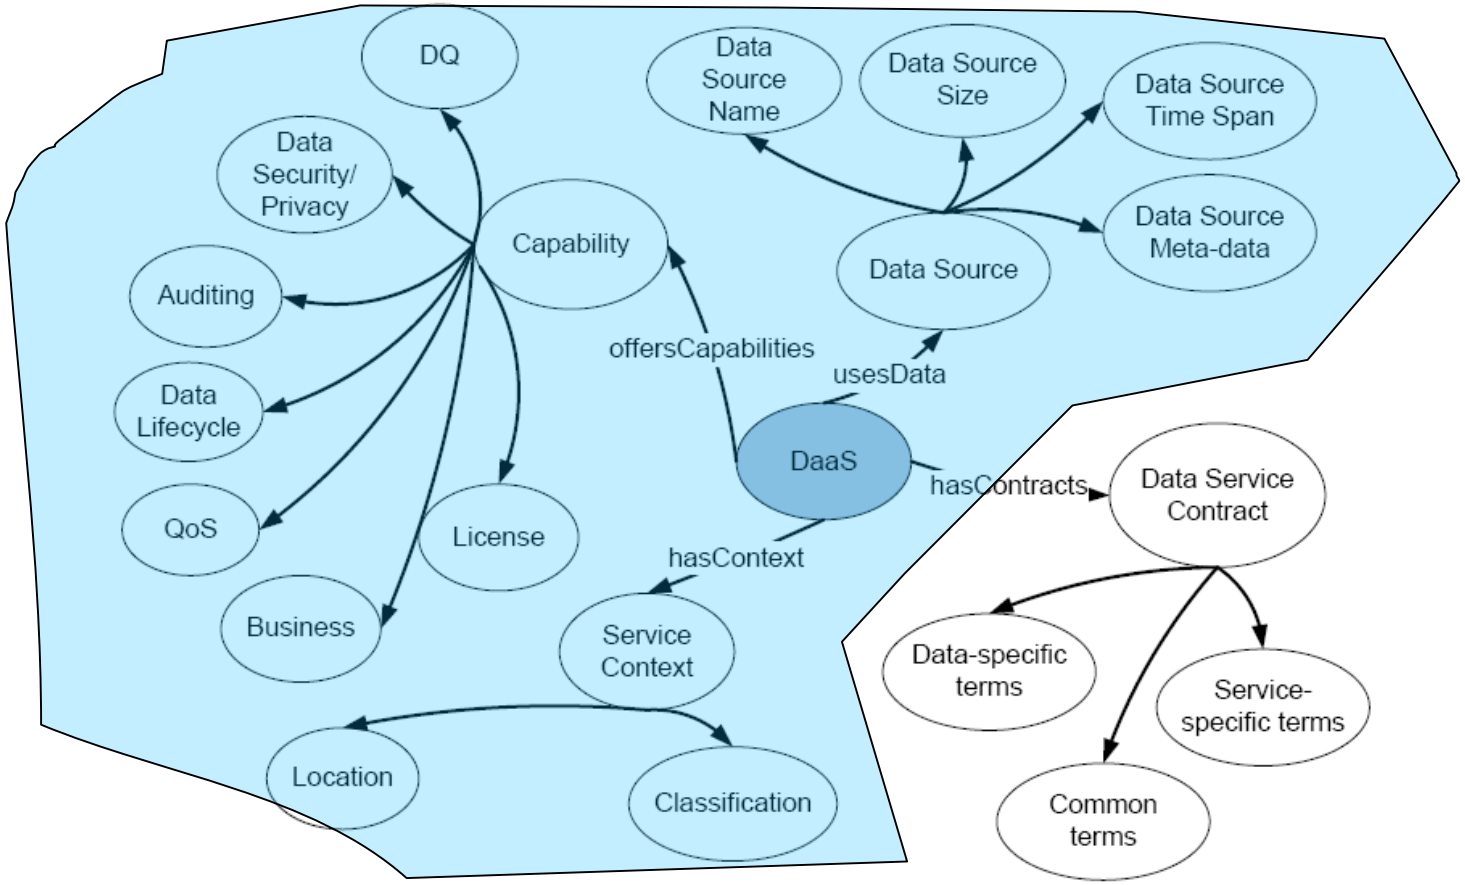

Hong Linh Truong, Schahram Dustdar On analyzing and specifying concerns for data as a service. APSCC 2009: 87-94

Contribution: techniques for evaluating data concerns in DaaS (2)

evaluation scope	• At which level the evaluation is performed?	*enable fine-grained evaluation with three scopes: data resources, DaaS operations, and DaaS*
evaluation modes	• When the evaluation is done?	*Offline and online are suitable for different types of data*
integration model	• How the evaluation tool is invoked?	*Push* and *pull* data concerns *Pass-by-value* versus *pass-by-reference data*

Hong Linh Truong, Schahram Dustdar: On Evaluating and Publishing Data Concerns for Data as a Service. APSCC 2010: 363-370

Contribution: modeling complex properties of data services utilized

A description model for DaaS

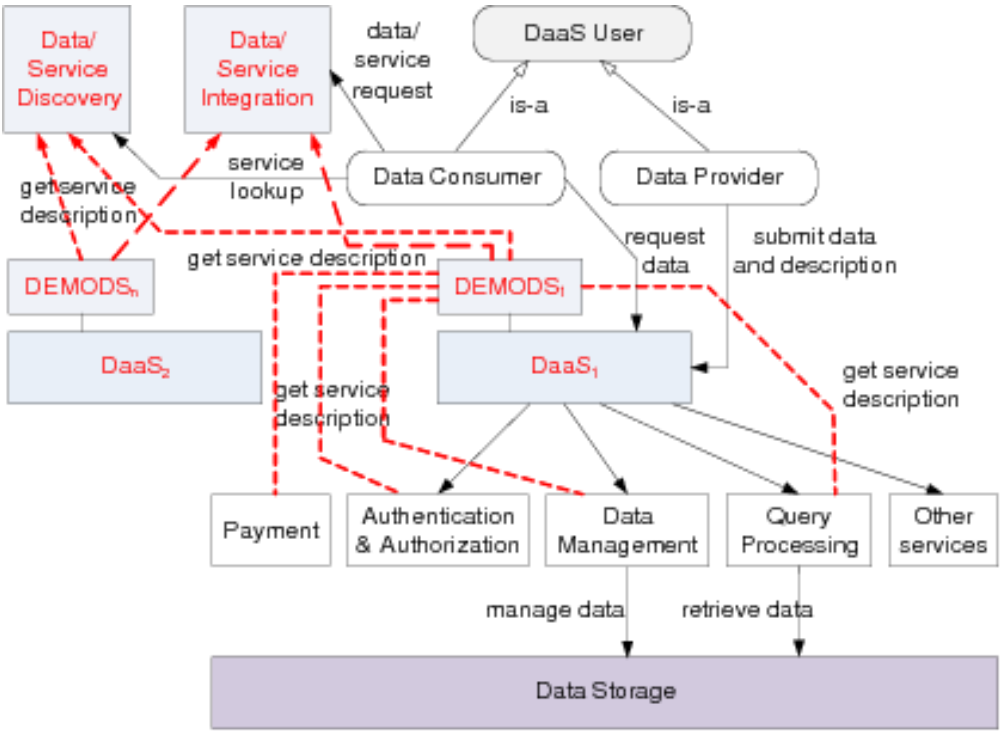

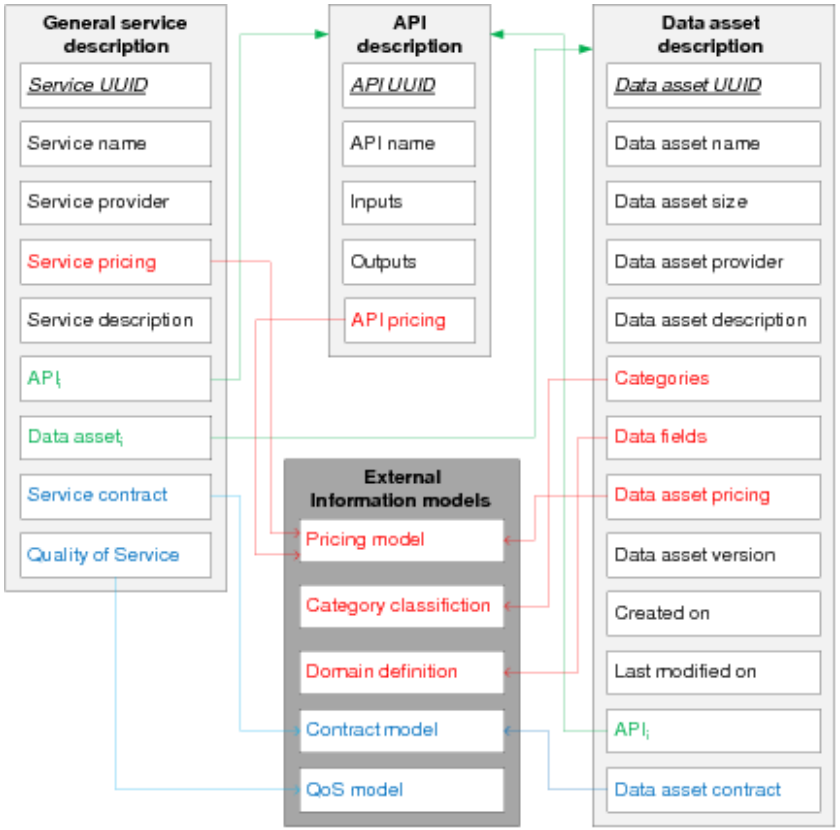

Quang Hieu Vu, Tran Vu Pham, Hong Linh Truong, Schahram Dustdar, Rasool Asal: DEMODS: A Description Model for Data-as-a-Service. AINA 2012: 605-612

Contribution: data contracts

Data contract development framework

Data contract specifications and evaluation algorithms

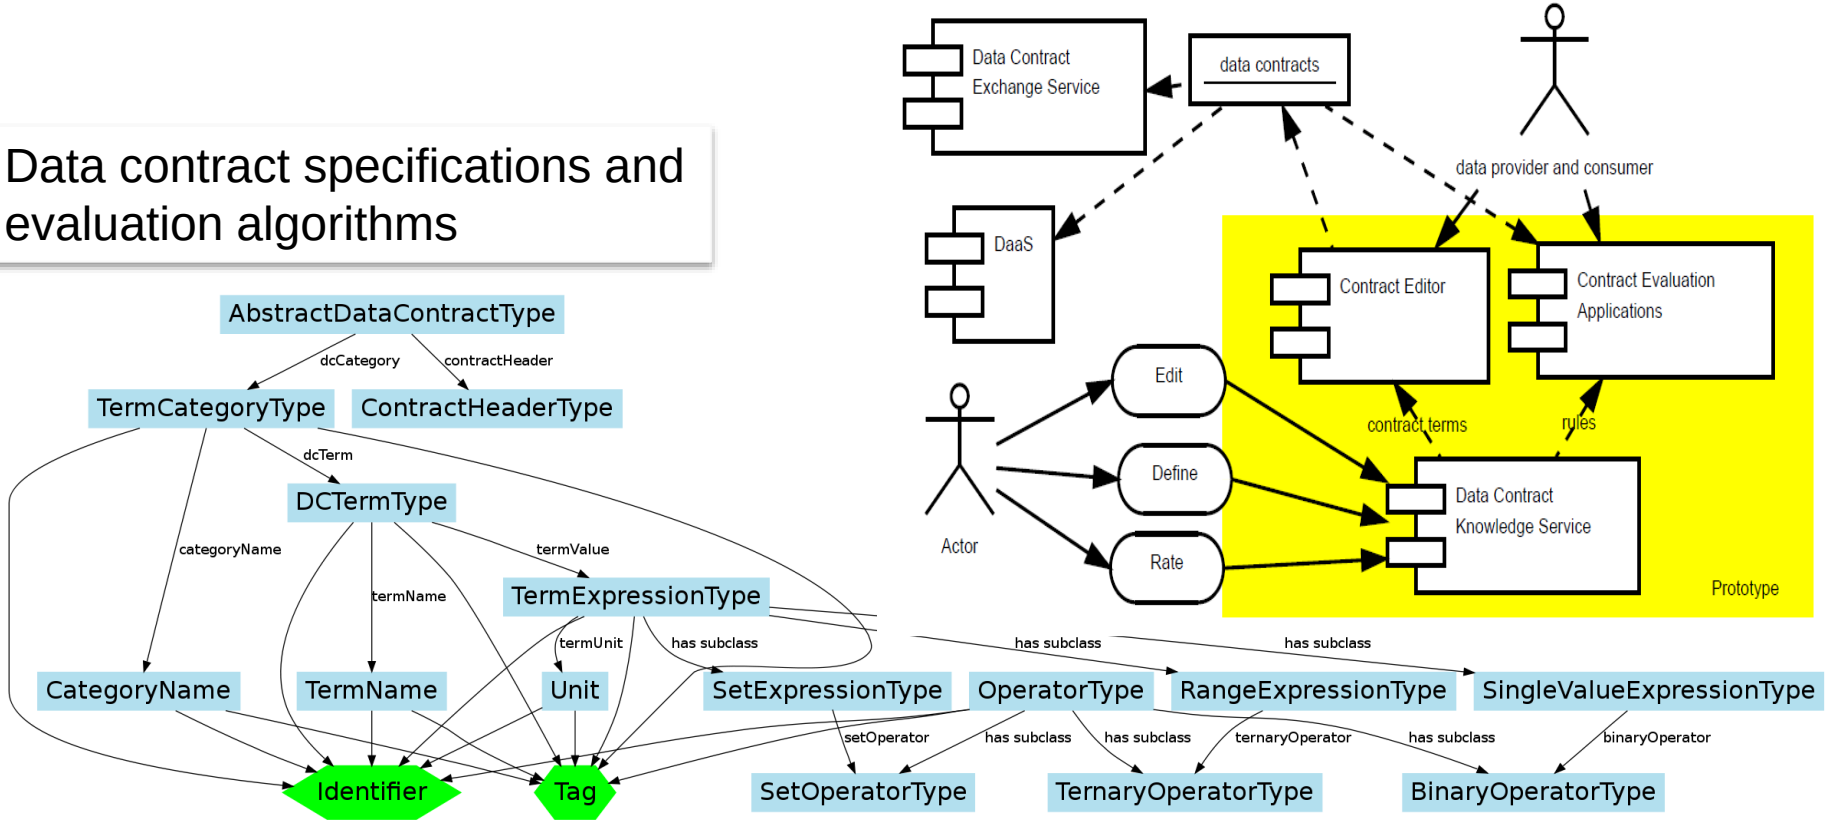

Hong Linh Truong, Marco Comerio, Flavio De Paoli, G. R. Gangadharan, Schahram Dustdar: Data contracts for cloud-based data marketplaces. IJCSE 7(4): 280-295 (2012

Contribution: reconcilation of service/data contracts

Composite service-based applications

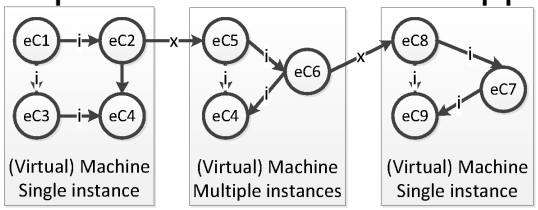

Multiple concerns at multiple levels

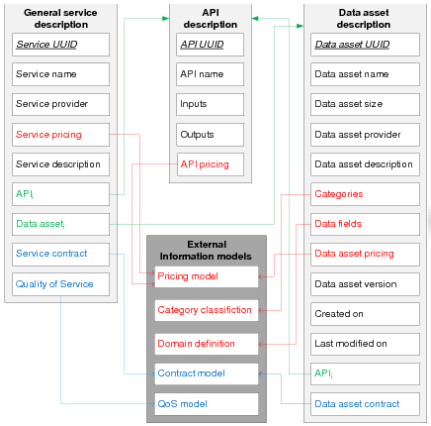

Service1: Derivation=„NotAllowed"
Service2: Collection="Allowed"

Compatibility?

Multiple language specifications

WSLA

WSOL

ODRL-S

....

Agorithms and framework for data/service contract compatibility evaluation

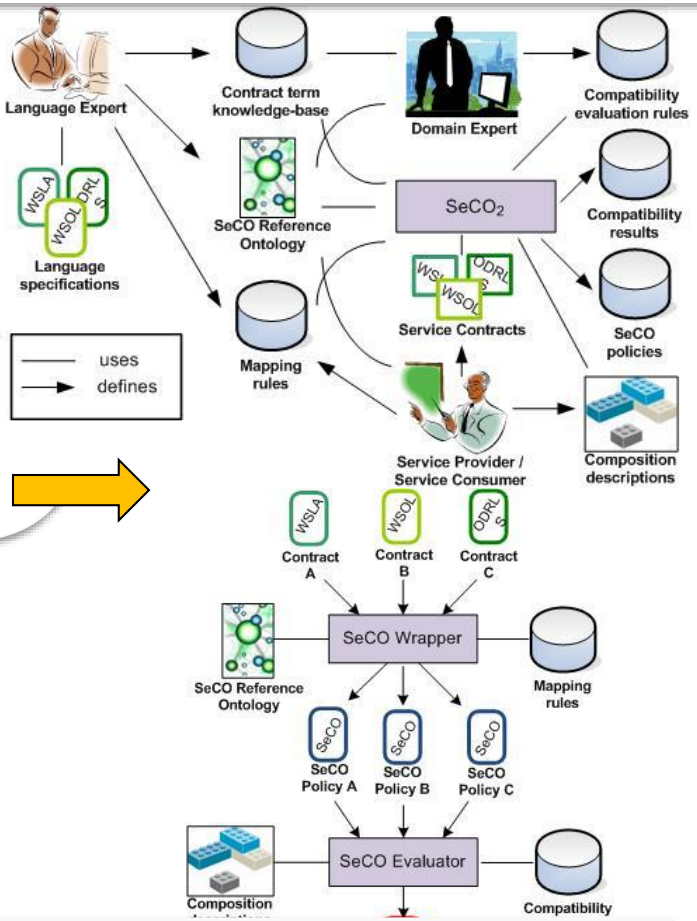

Marco Comerio, Hong Linh Truong, Flavio De Paoli, Schahram Dustdar: Evaluating Contract Compatibility for Service Composition in the SeCO2 Framework. ICSOC/ServiceWave 2009: 221-236

Hong-Linh Truong, G.R. Gangadharan, Marco Comerio, Vincenzo D' Andrea, Flavio De Paoli and Schahram Dustdar. "Reconciliation of Contractual Concerns of Web Services. In Handbook of Research on Service-Oriented Systems and Non-Functional Properties: Future Directions, 298-321 (2012)

Summary of contributions

- Contributions:
  - A data concern specification for DaaS, a framework for data concern evaluation and publishing, a data contract model, algorithms for data/service contract compatibilities, a description model for DaaS
  - Main contributions for the WWTF SODI project
  - Prototypes: http://www.infosys.tuwien.ac.at/prototyp/SOD1
- Related work
  - Service communities focus on QoS, database communites focus on the data level
- Selected publications:
  - Hong-Linh Truong, Schahram Dustdar: On analyzing and specifying concerns for data as a service. APSCC 2009:87-94
  - Hong-Linh Truong, Schahram Dustdar: On Evaluating and Publishing Data Concerns for Data as a Service. APSCC 2010:363-370
  - Hong-Linh Truong, Marco Comerio, Andrea Maurino, Schahram Dustdar, Flavio De Paoli, Luca Panziera: On Identifying and Reducing Irrelevant Information in Service Composition and Execution. WISE 2010: 52-66
  - Quang Hieu Vu, Tran Vu Pham, Hong-Linh Truong, Schahram Dustdar, Rasool Asal, DEMODS: A Description Model for Data-as-a-Service, (c)IEEE Computer Society, The 26th IEEE International Conference on Advanced Information Networking and Applications (AINA-2012), Fukuoka, Japan, March 26-29, 2012
  - Schahram Dustdar, Reinhard Pichler , Vadim Savenkov, Hong-Linh Truong, "Quality-aware Service-Oriented Data Integration: Requirements, State of the Art and Open Challenges", SIGMOD Record, Vol. 41, Number 1, March 2012
  - Hong-Linh Truong, G.R. Gangadharan, Marco Comerio, Vincenzo D' Andrea,Flavio De Paoli and Schahram Dustdar. "Reconciliation of Contractual Concerns of Web Services." In Handbook of Research on Service-Oriented Systems and Non-Functional Properties: Future Directions
  - Marco Comerio, Hong-Linh Truong, Flavio De Paoli, Schahram Dustdar: Evaluating Contract Compatibility for Service Composition in the SeCO2 Framework. ICSOC/ServiceWave 2009: 221-236
  - Hong-Linh Truong, Marco Comerio, Flavio De Paoli, G.R. Gangadharan, Schahram Dustdar, "Data Contracts for Cloud-based Data Marketplaces", International Journal of Computational Science and Engineering, Vol. 7, No. 4, 2012.

Contribution area: composable quality evaluation for cloud applications

Performance together other issues for complex service systems

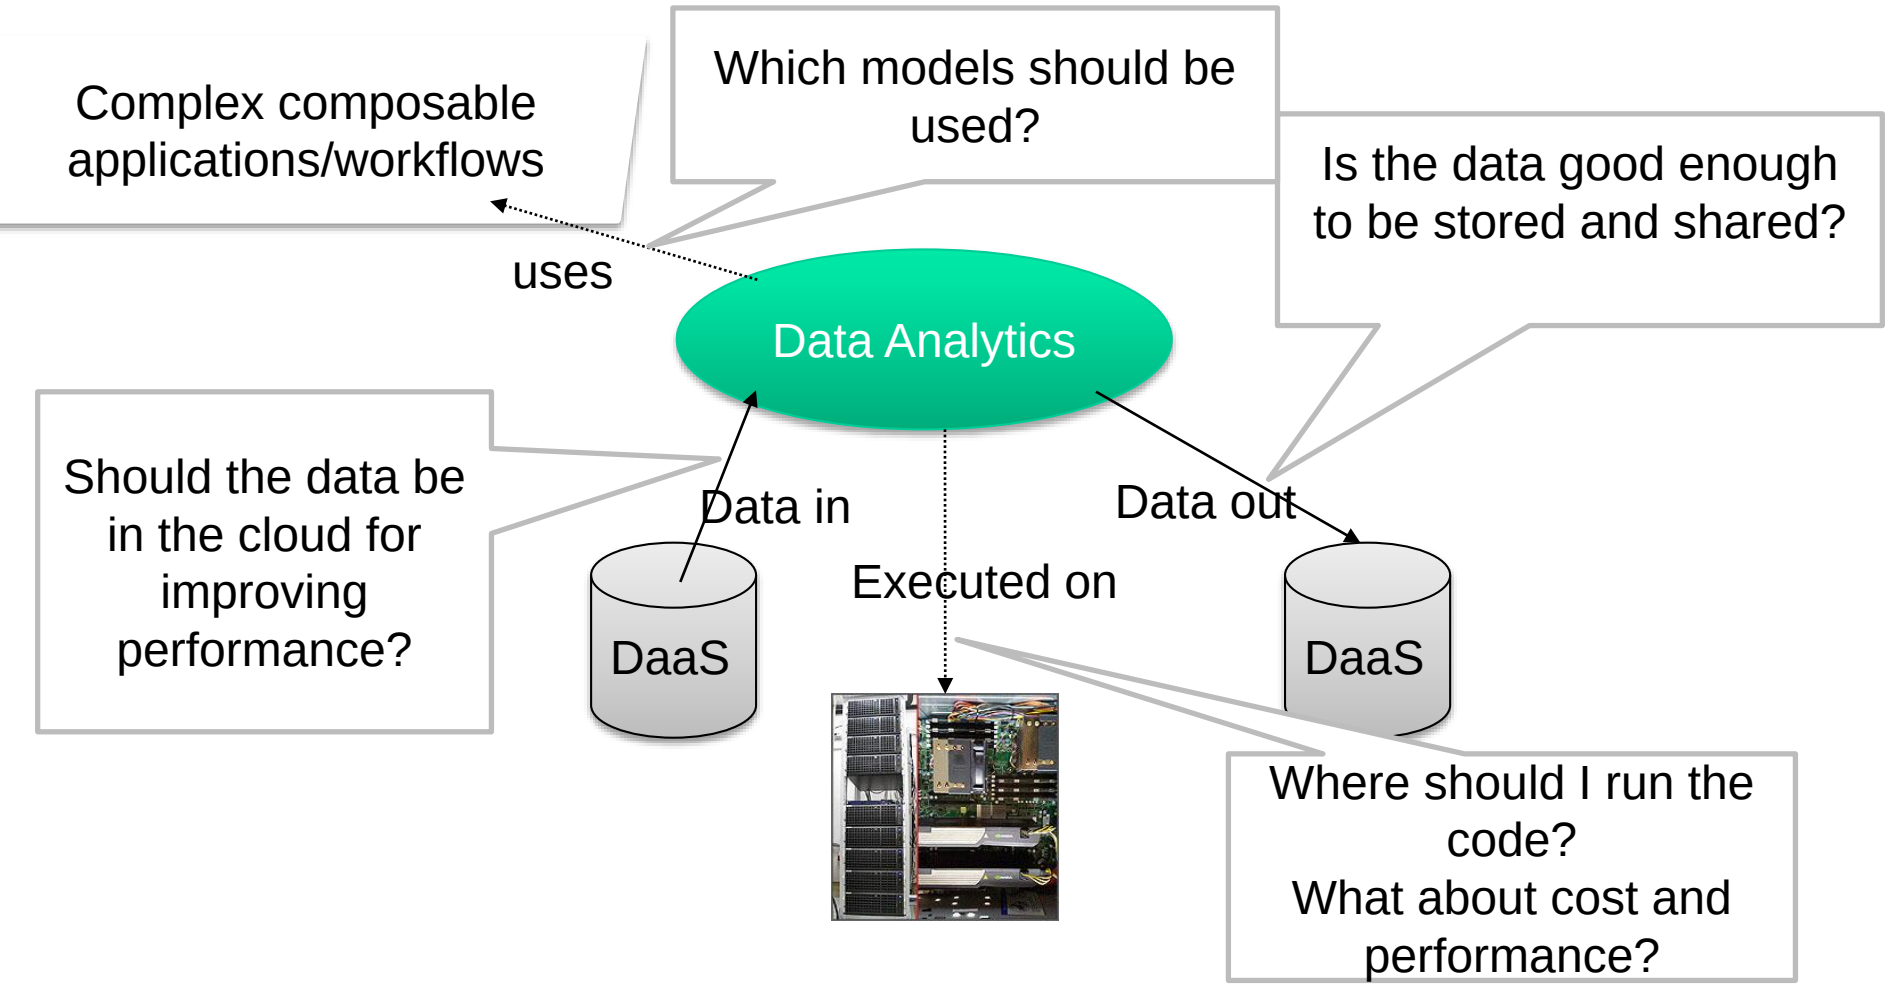

Dealing with performance and cost of complex applications in clouds

- Application complexity
  - Elastic high performance applications on multiple clouds: libraries, software services, virtual machines, etc.
  - Cost and performance are needed for determining which parts of the application should be excuted in the clouds and when
- Cost/performance model complexity
  - Coarse- and fine-grained cost models of clouds at different layers:
    - Too coarse-grained (networks, storages, machines) or too fine-grained (IO calls)
  - Software-, data-, human-specific cost/performance models
    - Cost models for individual parts (workflow, MPI, OpenMP, etc.)

Tran Vu Pham, Hong-Linh Truong, Schahram Dustdar "Elastic High Performance Applications - A Composition Framework", The 2011 Asia-Pacific Services Computing Conference (IEEE APSCC 2011), (c) IEEE Computer Society, December 12 - 15, 2011, Jeju, Korea

Hong Linh Truong, Schahram Dustdar: Cloud computing for small research groups in computational science and engineering: current status and outlook. Computing 91(1): 75-91 (2011)

Contribution: composable cost evaluation

Elastic high performance applications on multiple clouds: libraries, software services, virtual machines, etc.

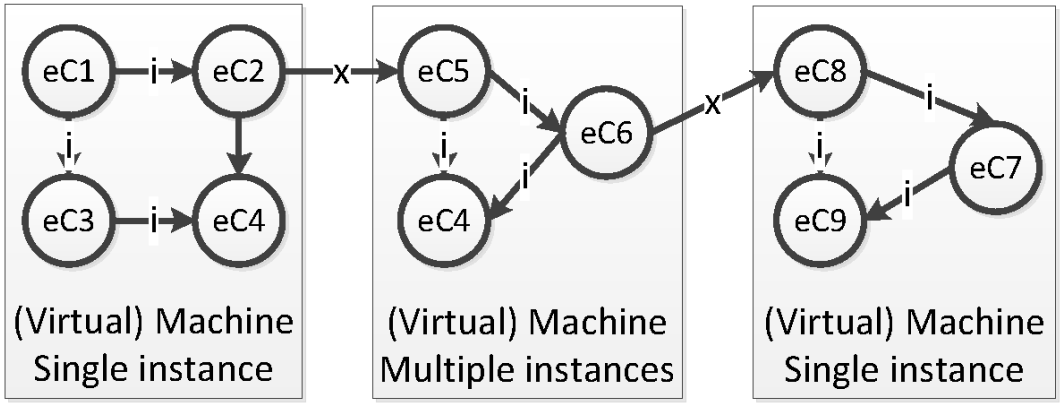

Utilize different performance and dependencies models for sequential, parallel, workflows, etc.

Runtime: Elastic processes

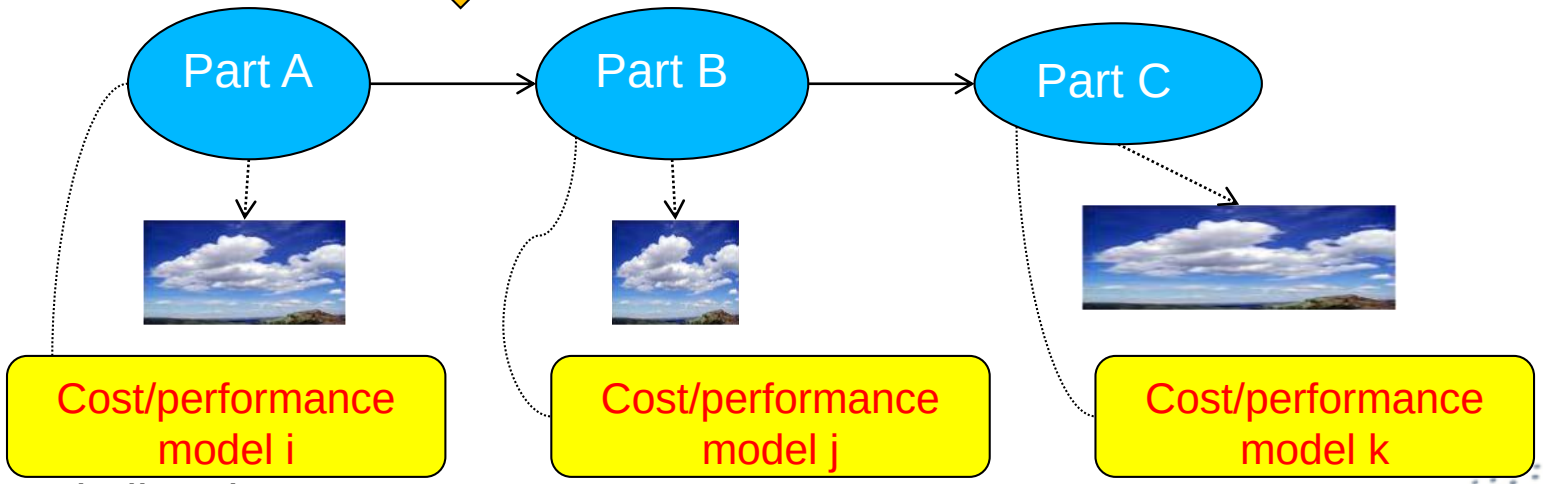

Examples of fine-grained composable cost models

Model	Activities	Cost
M_{ds}	Data storage	$size(total) \times t_{sub} \times cost(storage)$ where t_{sub} is the subscription time
M_{cm}	Computational machine	$cost(machine)$
M_{dfi}	Data transfer into the cloud	$cost(transfer_{in})$
M_{dfo}	Data transfer out to the cloud	$cost(transfer_{out})$
M_{sd}	Single data transfer without the cost for machines performing the transfer	$size(in) \times M_{dfi} + size(out) \times M_{dfo}$
M_{sm}	Sequential/multi-threaded program or single data transfer with the cost for machines performing the transfer (cost monitoring)	$t_e \times M_{cm} + size(out) \times M_{dfo} + size(in) \times M_{dfi}$
M_{se}	Sequential or multi-threaded program (cost estimation)	$f_{pi} \times M_{cm} + size(out) \times M_{dfo} + size(in) \times M_{dfi}$ where f_{pi} is an estimated performance improvement function when n expected threads to be used. f_{pi} can be provided by performance prediction tools or scientists. In our case, currently, we use an ideal parallel performance improvement $f_{pi} = \frac{p}{n} \times t_e(p)$ where p is the number of threads used to obtain $t_e(p)$. p and $t_e(p)$ are known knowledge.
M_{pm}	Parallel/MPI programs on multiple machines (cost monitoring)	$n \times M_{cm} \times t_e + size(out) \times M_{dfo} + size(in) \times M_{dfi}$
M_{pe}	Parallel/MPI programs on multiple machines (cost estimation)	$n \times M_{cm} \times f_{pi} + size(out) \times M_{dfo} + size(in) \times M_{dfi}$ where f_{pi} is an estimated performance improvement function when n processes are used.
M_{wm}	Workflows (cost monitoring)	$\sum_{i=1}^{k} (size(in_i) \times M_{dfi}) + \sum_{i=1}^{l} (size(out_i) \times M_{dfo}) + \sum_{i=1}^{n} (M_{cm} \times t_e(machine_i))$
M_{we}	Workflows (cost estimation)	$\sum_{i=1}^{nwr} cost(wr_i)$. For a workflow region wr_i, $cost(wr_i) = \sum_{j=1}^{q} (cost(activity_j))$ where $cost(activity_j)$ is determined based on M_{mp}, M_{sm}, and M_{sd}, when the activity $activity_j$ is a parallel activity, sequential activity, or a data transfer activity, respectively.

Online cost/performance evaluation

Application dependency model

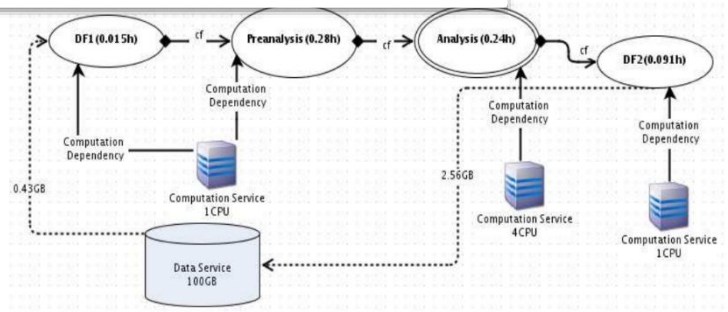

Application deployment

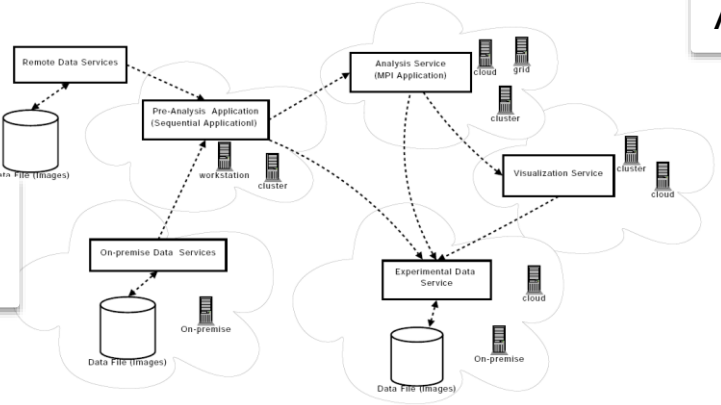

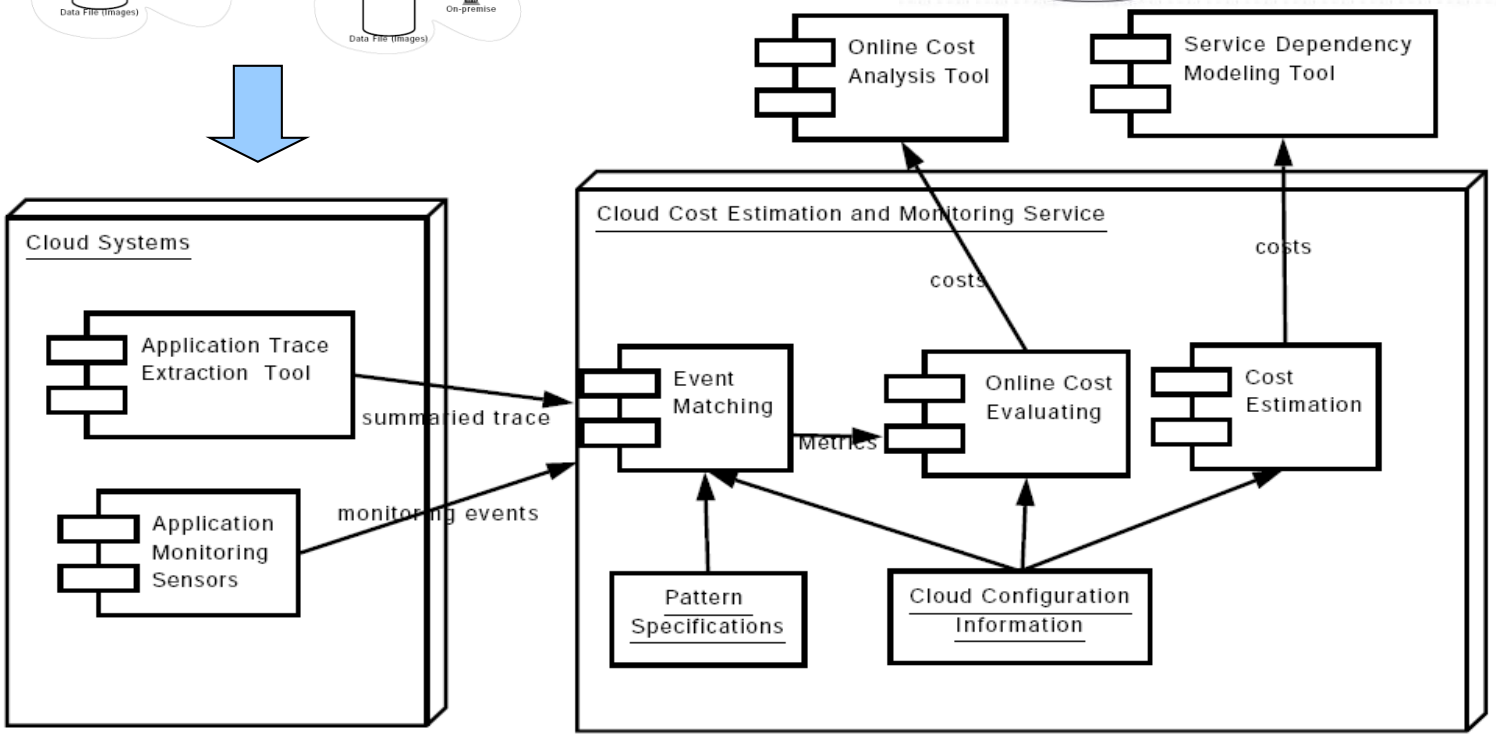

Hong Linh Truong, Schahram Dustdar: Composable cost estimation and monitoring for computational applications in cloud computing environments. Procedia CS 1(1): 2175-2184 (2010)

Examples of composable cost evaluation

Next generation sequencing analysis workflow (GSA)

Online analysis

Cloud provider bill

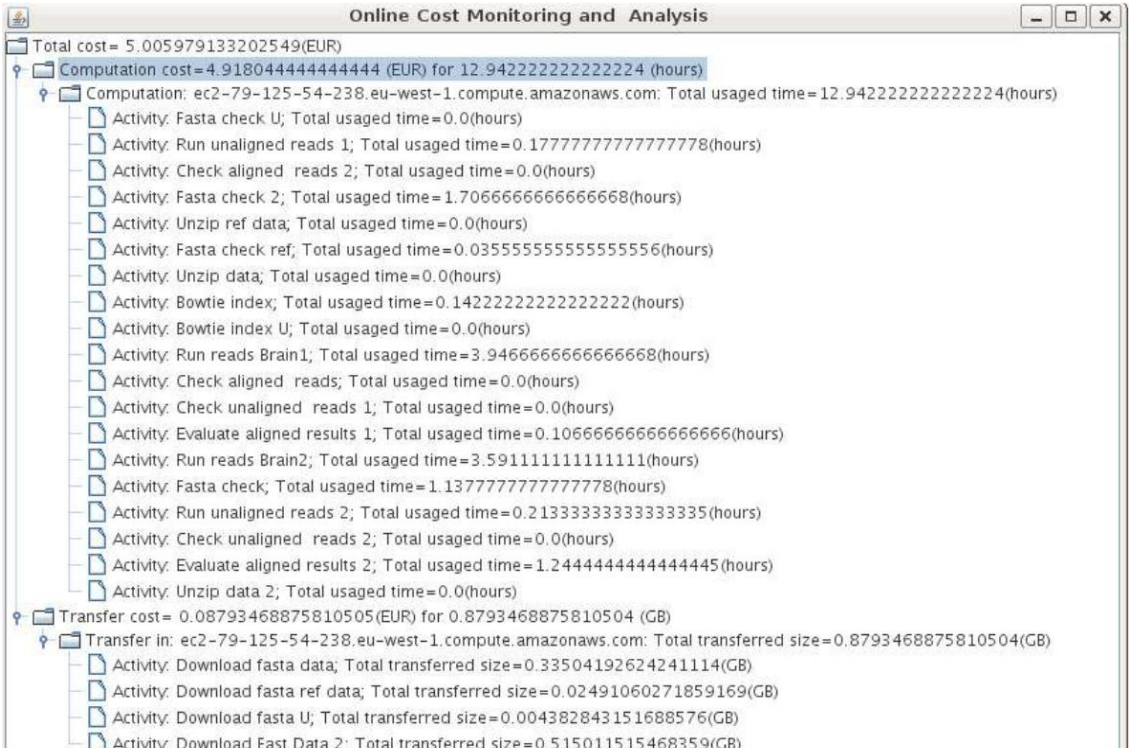

Operation Name	Usage Type	Usage Value	Cost (EUR)
RunInstances	EU-BoxUsage:m1.large	13 hours	4,94
RunInstances	EU-DataTransfer-In-Bytes	1,043 GB	0,104
RunInstances	EU-DataTransfer-Out-Bytes	4,48 GB	0,672
EBS:IO-Write	EU-EBS:VolumeIOUsage	2083828 I/O calls	0,229

Table 4: Operation name, usage types and usage values extracted from the Amazon billing report during the execution of the GSA workflow. We determined costs by using the following prices: 0.38 EUR per large instance (m1.large) instance-hour, 0.1 EUR per GB data transfer in, 0.150 EUR per GB data transfer out, 0.11 EUR per 1 million I/O requests.

Summary of contributions

- Contributions
  - A composable cost evaluation framework
  - Composation techniques for elastic high performance applications
  - Composable cost evaluation techniques play an important role for elasticity monitoring and controls in the ongoing FP7 CELAR (http://www.celarcloud.eu/)
- Related work
  - Simple cost model for particular types of clouds, focus on infrastructure not complex applications
- Scientific papers
  - Hong Linh Truong, Schahram Dustdar: Composable cost estimation and monitoring for computational applications in cloud computing environments. Procedia CS 1(1): 2175-2184 (2010)
  - Hong-Linh Truong, Schahram Dustdar: Cloud computing for small research groups in computational science and engineering: current status and outlook. Computing 91(1): 75-91 (2011)
  - Tran Vu Pham, Hong-Linh Truong, Schahram Dustdar: Elastic High Performance Applications - A Composition Framework. APSCC 2011: 416-423
  - Hong-Linh Truong, Tran-Vu Pham, Nam Thoai and Schahram Dustdar. "Cloud Computing for Education and Research in Developing Countries." In Cloud Computing for Teaching and Learning: Strategies for Design and Implementation, ed. Lee Chao, 64-80 (2012)

Contribution area: programming human-based services in the cloud

Contributions: incorporate humans into a programming paradigm

Programming languages

- Abstracting human compute units as program elements
- Extending programming languages to support human compute units
- Data/control flows via extensible APIs

Multiple programming models

- Shared memory (e.g., human –software – human), message passing (human-to-human), artifact-centric, etc., via APIs working atop the compute unit abstraction layer

Execution environment

- Computing capability /profile management: human computing power, reputation and incentive models
- Monitoring and enforcing incentives/rewards, quality of results, availability
- Communication between human-middleware, among Individual Compute Units (ICU)/Social Compute Units (SCU) for exchanging artifacts and comprehensing tasks

Our approach -- incorporate humans into a programming paradigm

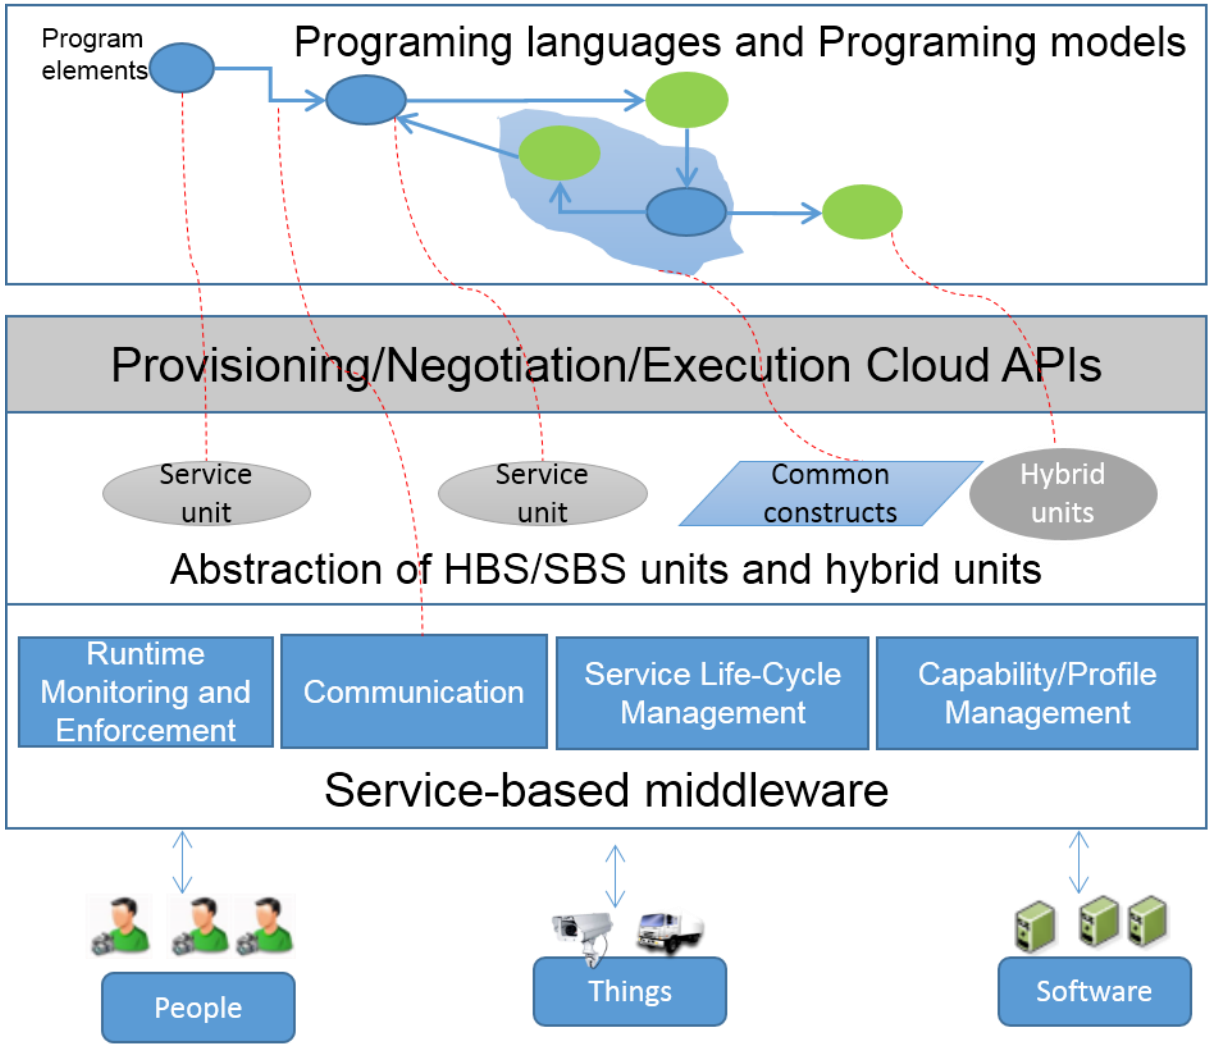

Programming constructs/elements

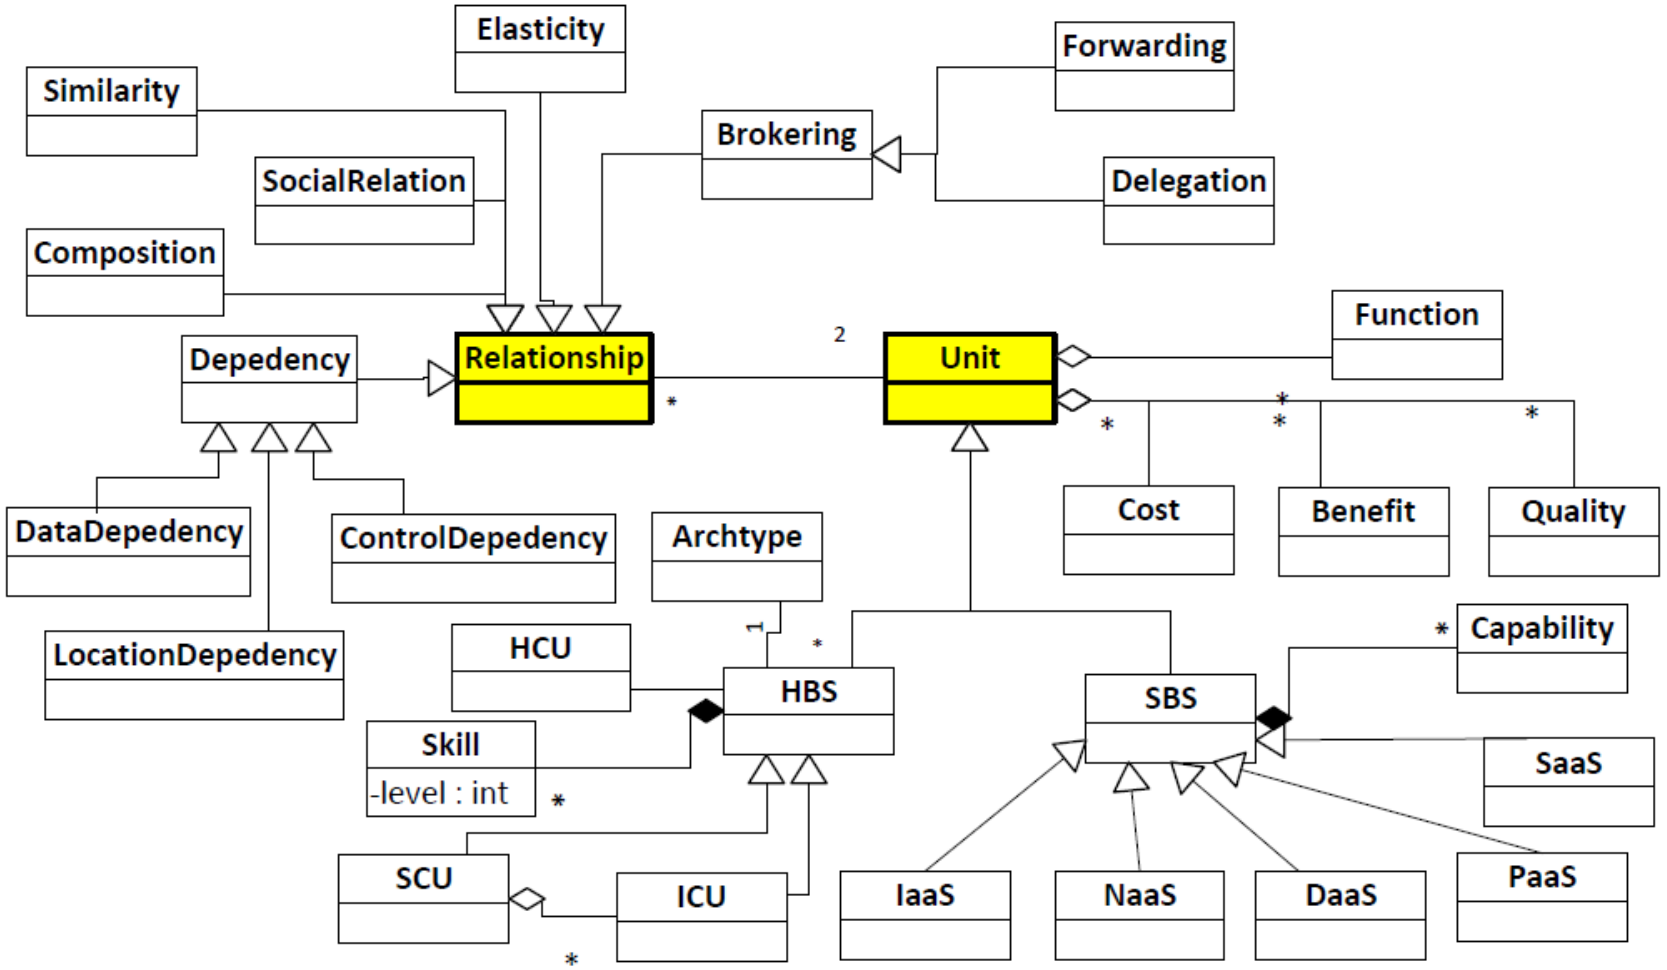

Contributions: programming hybrid services in the cloud

A cloud of hybrid services includes software-based services and human-based services that can be provisioned, deployed and utilized on-demand based on different pricing and incentive models

HPU Definition: *HPU is a value describing the computing power of an HBS measured in an abstract unit. A cloud of HBS has a pre-defined basic power unit, hpu_θ, corresponding to the baseline skill bs_θ of the cloud."*

HPU can be defined for different „archtypes"

An „archtype" characterizes the problem domain (e.g., bones simulation) that the ICU/SCU can solve (the type of tasks)

Hong Linh Truong, Schahram Dustdar, Kamal Bhattacharya: Programming Hybrid Services in the Cloud. ICSOC 2012: 96-110. Best paper award.

Programming a combination of HBS and SBS

e.g., preparing/managing inputs/outputs for HBS using SBS

```
//using JClouds APIs to store log file of web application server
BlobStoreContext context =
    new BlobStoreContextFactory().createContext("aws-s3","REMOVED
        ","REMOVED");
BlobStore blobStore = context.getBlobStore();
//.... and add file into Amazon S3
Blob blob = blobStore.blobBuilder("hbstest").build();
blob.setPayload(new File("was.log"));
blobStore.putBlob("hbstest", blob);
String uri = blob.getMetadata().getPublicUri().toString();
VieCOMHBS vieCOMHBS = new VieCOMHBSImpl();
//assume that WM6 is the HBS that can analyze the Web Middleware
    problem
vieCOMHBS.startHBS("WM6");
HBSRequest request = new HBSRequest();
request.setDescription("Find possible problems from " + uri);
vieCOMHBS.runRequestOnHBS("WM6", request);
```

Summary of contributions

- Contributions:
  - a novel model for clouds of HBS and hybrid services provisioning
  - a framework for solving complex problems using clouds of hybrid services
  - programming primitives for hybrid services
- Related work
  - Crowdsourcing techniques focus on simple models of human capabilities utilization, existing cloud techniques do not consider human-based services
- Selected publications
  - Schahram Dustdar, Hong Linh Truong: Virtualizing Software and Humans for Elastic Processes in Multiple Clouds- a Service Management Perspective. IJNGC 3(2) (2012)
  - Hong Linh Truong, Schahram Dustdar, Kamal Bhattacharya: Programming Hybrid Services in the Cloud. ICSOC 2012: 96-110. Best paper award
  - Schahram Dustdar, Yike Guo, Benjamin Satzger, Hong Linh Truong: Principles of Elastic Processes. IEEE Internet Computing 15(5): 66-71 (2011)
  - Lam-Son Lê, Hong Linh Truong, Aditya Ghose, Schahram Dustdar: On Elasticity and Constrainedness of Business Services Provisioning. IEEE SCC 2012: 384-391

Conclusions (1)

- The evolution of underlying systems and the utilization of different types of resources under different models for elasticity lead to complex quality issues
- We address several quality issues for software, data and people in an integrated manner for different perspectives
- Our contributions:
  - Online performance monitoring and analysis for workflows
  - Data concerns analysis and evaluation, data contracts, data and service contract compatibility for DaaS
  - Composable cost evaluation for cloud applications
  - Programming hybrid services in the cloud
- These contributions support the developent of multi-dimensional elasticity principles in complex service oriented systems

Conclusions (2)

- Research methodology:
  - Real-world systems with measurable metrics
  - Characterizing systems through the definition, measurement, and analysis of metrics using their dependencies and system structure dependencies
  - Composable evaluation techniques are a key to metrics evaluation for understanding complex systems
  - Validation in real projects with software prototypes
- Research collaboration
  - Complex service-oriented systems require intensive collaborations with other scientists in other disciplines

Future work

- Hybrid compute units
  - Modeling and provisioning things, software, and people under the same service models
- Programming hybrid compute units for elastic processes
- Service engineering analytics of elastic systems
  - Elasticity specifications and reasoning techniques
  - Elasticity spaces analytics
- Application domains
  - „Social computer“ and smart cities
  - Computational science and engineering

Thank you!

Acknowledgements

Kamal Bhattacharya, Peter Brunner, Marco Comerio, Vincenzo D'Andrea, Hoa Dam, Schahram Dustdar, Thomas Fahringer, G.R.Gangadharan, Yike Guo, Lam-Son Le, Frank Leymann, Andrea Maurino, Michael Mrissa, Vlad Nae, Flavio De Paoli, Tran-Vu Pham, Reinhard Pichler, Michael Reiter, Robert Samboski, Benjamin Satzger, Vadim Savenkov, Martin Treiber, and Quang-Hieu Vu