

Chapter 20

Self-Awareness of Cloud Applications

Alex Iosup, Xiaoyun Zhu, Arif Merchant, Eva Kalyvianaki, Martina Maggio, Simon Spinner, Tarek Abdelzaher, Ole Mengshoel, Sara Bouchenak

Abstract Cloud applications today deliver an increasingly larger portion of the Information and Communication Technology (ICT) services. To address the scale, growth, and reliability of cloud applications, self-aware management and scheduling are becoming commonplace. How are they used in practice? In this chapter, we propose a conceptual framework for analyzing state-of-the-art self-awareness approaches used in the context of cloud applications. We map important applications corresponding to popular and emerging application domains to this conceptual framework, and compare the practical characteristics, benefits, and drawbacks of self-awareness approaches. Last, we propose a roadmap for addressing open challenges in self-aware cloud and datacenter applications.

Alexandru Iosup
Delft University of Technology, the Netherlands, e-mail: A.Iosup@tudelft.nl

Xiaoyun Zhu
Futurewei Technologies, CA, USA, e-mail: xiaoyzhu@yahoo.com

Arif Merchant
Google, Inc., CA, USA, e-mail: aamerchant@google.com

Eva Kalyvianaki
Imperial College of London, UK, e-mail: Evangelia.Kalyvianaki.1@city.ac.uk

Martina Maggio
Lund University, Sweden, e-mail: mmartimay@gmail.com

Simon Spinner
University of Wuertzburg, Germany, e-mail: simon.spinner@uni-wuerzburg.de

Tarek Abdelzaher
University of Illinois at Urbana Champaign, IL, USA, e-mail: zaher@illinois.edu

Ole Mengshoel
CMU Silicon Valley at the NASA Ames Research Center, PA, USA, e-mail: ole.mengshoel@sv.cmu.edu

Sara Bouchenak
INSA Lyon, France, e-mail: sara.bouchenak@insa-lyon.fr

20.1 Introduction

Cloud computing is the Information and Communication Technology (ICT) paradigm under which services are provisioned by their users on demand, with payment covering only what is actually used. Cloud users can today lease infrastructure, platform, software, and others “as a service”, from commercial clouds such as Amazon, Azure, Google, and SAP. Governments and entire industries are building large-scale datacenters that are and will increasingly host cloud computing applications. At the same time, data becomes an integral part of cloud computing: by 2017, over three-quarters of our personal and business data will reside in datacenters, according to a recent IDC report [32]. Cloud applications, often consumed by users as services, already represent over 10% of the entire ICT market in Europe [15]. Netflix, whose users consume a large fraction of the global Internet traffic, relies on ICT services from Amazon Web Services (AWS)¹. The market, growing in size, diversity of applications, and sophistication, already exceeds hundreds of millions of users and, as a consequence, \$100 billion world-wide [10]; the cloud market will likely contribute over 100 billion Euro to the European GDP in 2020 [15]. At this scale and with this importance, human management of datacenter resources is prohibitively expensive and, often, too error-prone. Thus, the use of self-awareness techniques to manage cloud applications is increasingly more present. In this chapter, we analyze the use of self-awareness in cloud computing and its applications.

Cloud applications raise a complex management challenge, derived from the goals of three main stakeholders: application users, application operators, and cloud operators. Each of these stakeholders has different requirements, which are often conflicting. For example, application users could demand that an interactive application is always responsive, even under bursty arrivals of user-issued requests. To meet this demand, application operators could require that enough capacity is always provided by cloud operators, yet only want to pay for what is actually consumed. Tension arises between performance and other requirements, including the cost of operations. As a consequence, the management challenge is to optimize non-trivial efficiency metrics and to meet complex service level agreements (SLAs), to an extent that already exceeds the capabilities of human management.

We investigate in this chapter the current state of self-awareness in cloud computing, and in particular datacenter-based cloud applications. Our goals are to introduce practical cases of self-awareness in such applications; to present a conceptual framework for analyzing state-of-the-art self-awareness approaches used in practice; to map already important and emerging application domains to the conceptual framework of self-awareness approaches used in practice, and analyze similarities and differences, benefits and costs of self-awareness approaches; and to identify and analyze open challenges in self-aware cloud applications, and propose a roadmap for advancing the state-of-the-art. The main contribution is structured as follows.

In Section 20.2, we introduce a framework for the analysis of self-awareness techniques used in cloud computing and its applications. Our framework consists of

¹ Details: <https://aws.amazon.com/solutions/case-studies/netflix/>.

a structured way to analyze the types of applications, of problems, and of approaches for which self-awareness is relevant in practice. The framework also structures the analysis of directions for future research. Although the framework is currently built for the analysis of self-awareness in cloud applications, and thus is adapted to the operational conditions in cloud computing (metrics, stakeholders, etc.), the framework could be extended to other domains. We show the usefulness of this framework by applying it in practice, with the results presented in the next sections.

In Section 20.3, we focus on eight popular or emerging application domains, with important commercial, scientific, governance, and other societal impact. Although any application domain is applicable, the market volume and the number of users, today or in the foreseeable future, are important criteria for selecting the application domains for this chapter. Among the selected applications are business applications, compute-intensive and data-intensive batch processing, data-stream processing, online gaming, partial processing, and cyber-physical applications. Some of these applications, such as online gaming, partial processing, and cyber-physical applications are emerging in terms of number of users and adoption of cloud technology. We also include in this section the workloads generated by the datacenters themselves, which can be seen as overhead, but are already consuming large amounts of resources and must meet complex, albeit internal, SLAs.

In Section 20.4, we identify ten types of problems that are already addressed by self-awareness techniques, including recovery planning, resource autoscaling, runtime architectural reconfiguration and load balancing, fault-tolerance in distributed systems, energy-proportionality, workload prediction, performance isolation, diagnosis and troubleshooting, discovery of application topology, and intrusion detection and prevention. Some of these problems, such as autoscaling, energy-proportionality, performance isolation, and intrusion detection and prevention, have developed a new form or even appeared specifically in the context of cloud computing.

In Section 20.5, we identify and analyze seven types of self-awareness approaches used in practice: feedback control based techniques, metric optimization with constraints, machine learning based techniques, portfolio scheduling, self-aware architecture reconfiguration, and stochastic performance models. Although none of these approaches are unique to cloud computing, their adaptation to cloud computing and its applications is non-trivial.

In Section 20.6, we identify and analyze four directions for future use of self-awareness approaches for cloud computing and its applications. We focus on directions that are not only needed for practical applications, but for which we can envision the next research steps and that the results of this research can be put in practice in the following three to five years.

Our survey of applications, problems, self-awareness approaches, and open challenges in self-awareness is by far not exhaustive. However, we study for each broad types with existing popularity and likely future impact. Moreover, we envision that the approach we take in this work will also be useful for studying other types.

This chapter is the result of original work by the authors, and in particular the survey started during the Dagstuhl Seminar 15041, “Model-driven Algorithms and Architectures for Self-Aware Computing Systems”.

20.2 Overview of the Framework

We propose a framework for understanding the practice of using self-awareness techniques in managing and scheduling of cloud applications. This framework follows the structure of a natural discussion about the field, with three main questions and a format for answering them that is conducive to surveying the field. The first question focuses on specific applications: *Which cloud applications raise the challenges that self-awareness techniques are particularly good in addressing?* Answering this question requires an understanding of the nature and characteristics of self-awareness challenges that affect such cloud applications. Thus, the second question is *Which are the important self-awareness challenges for cloud applications?* As a third question, *Which are the self-awareness approaches that address the self-awareness challenges in this context?* Last, a fourth question focuses on the future: *Assuming a research horizon of 3-5 years, what are the most promising directions for future research in enabling self-aware cloud applications?*

Answering the first three questions is sufficient to yield a survey of cloud applications whose self-awareness challenges are addressed or resolved in practice by self-awareness approaches (techniques, methods, best-practices, or even entire methodologies). For example, mapping all the different application-challenge-approach paths_i can create a survey of the entire space. Then, it is easy to group together applications raising the same challenge; similarly, self-awareness approaches that address the same challenge can be grouped together. The survey is practical, in that a complete “application-challenge-approach” path can be directly considered by the practitioner.

The first three questions in our framework are also necessary. There are hundreds of application *types* commonly used in software engineering practice, as indicated for example by the extensive taxonomy of Forward and Lethbridge [21]. Thus, surveying without the guidance of specific applications provided by the first question could lead to a variety of self-awareness challenges and approaches, all with the merit of being applicable, but without much proof of use in cloud context. Without the specific problems provided by the second question, the self-awareness techniques could be used in a variety of cases, limited only by the creativity of the designer and by the difficulty of proving their benefit for practical use. Thus, limiting the survey to the set of challenges that are currently addressed in the context of cloud applications is necessary; we address this through the combined expertise of the authors regarding the field. Last, although many self-awareness techniques already exist, not all have yet been applied to cloud settings. Thus, surveying could go well beyond the scope of the third question, and generic techniques that may not work well (enough) in practice will also be surveyed.

The framework also structures the analysis of directions for future research. We leverage here the collective expertise of the authors, in which important questions are proposed by individuals and discussed by the community. By iterating this process, we believe the community can propose and refine its own most important goals. The results of the first iteration are shared with everyone interested to help the community make progress, through the text in Section 20.6.

20.3 Types of Applications

In this section, we present the following eight types of applications that already benefit from the use of self-awareness techniques.

1. Enterprise applications
2. Computing-intensive batch processing
3. Data-intensive batch processing
4. Data-stream processing
5. Workloads generated by datacenter operations
6. Online gaming
7. Partial and delayed processing
8. Cyber-physical applications

Each of these selected applications is already popular, or generates a significant amount of revenue, or is critical to the operation of many businesses, or uses a significant amount of resources, or is promising to emerge as such; often, the applications we select have a combination of these characteristics.

20.3.1 *Enterprise Applications*

Description: We reuse the definition of enterprise applications proposed by Shen et al.: “the user-facing and backend services, generally supporting business decisions and operations and commonly contracted under strict SLA requirements, whose downtime or even just low performance will lead to reduced productivity, loss of revenue, customer departure, or even legal actions. These workloads include enterprise multi-tier applications, and business-critical workloads that often include applications in the solvency domain or other decision-making tools. Other applications that characterize business-critical workloads are email, collaboration, database, ERP, CRM, and management services, when used in conjunction with other workloads.” [71]

20.3.1.1 Multi-Tier Enterprise Applications

Application components: Multi-tier enterprise applications refer to those web-based business applications that each comprises a collection of cooperating components, organized as multiple logical tiers. The most common three-tier architecture consists of a presentation tier, an application tier, and a database tier. The presentation tier receives requests from the clients, the application tier handles the business logic, and in turn interacts with the database tier to obtain and store persistent data. Examples of such applications including Enterprise Resource Planning (ERP) and Customer Relationship Management (CRM) are often subject to an interactive workload. Different types of requests often incur different loads on the system (e.g., read- vs. write-intensive transactions, compute- vs. data-intensive workloads). The multi-tier architecture makes it challenging to implement self-awareness schemes for such applications, as there may be complex control flows between the different tiers, and each tier may have different resource requirements and performance bottlenecks. With the trend towards service-oriented architectures, the different tiers are often split into different services, making the control flow even more complex.

Metrics of interest: *Availability* - the percentage of time the application service remains up and running; *reliability* - the ratio of successful requests to the total number of requests; *performance* - characterized by such metrics as throughput (requests/sec) and request response times; *resource* (CPU, memory, I/O) *utilization* of the underlying system; and *cost* - metrics related to financial or energy cost.

Typical problems: Performance isolation and service differentiation, tradeoff between multiple metrics of interest, diagnosis and troubleshooting, dynamic load balancing, end-to-end service level assurance, and autoscaling of resources.

Typical self-aware elements: Metric optimization [50], machine learning [57, 77], feedback control [36, 50, 67], stochastic performance models [5, 73], and statistical estimation [47] are the main approaches researchers have applied to help create self-aware multi-tier applications.

20.3.1.2 Business-Critical Applications

Application components: Business-critical workloads often include applications that provide support for decision-making, such as Monte Carlo simulations, financial and other types of modeling applications programmed as tightly coupled parallel jobs of relatively small size, but also the regular management services described in Section 20.3.1.1. It is typical for the system user to not specify the applications, due to privacy and business secrecy. Instead, users request service expressed only in SLA terms, e.g., number and size of virtual machines, generally provisioned for long periods of time and operated by the users IT team. The current practice in the datacenter is to require engineering confirmation for the most important provisioning and allocation decisions, especially at the initial installation of the long-running virtual machines. Self-aware resource management and scheduling tools [75] provide advice that engineers may take into account.

Metrics of interest: Various traditional metrics, including latency and throughput, and reliability and system load. Risk-related metrics, such as the risk score [75], which expresses the risk of significant under-performance and thus penalties paid by the service operator to the service user, and loss of trust.

Typical problems: Reduce the risk of performance degradation; use resources efficiently; avoid system overload and unavailability.

Typical self-aware elements: Portfolio scheduling [75], topology-aware resource management [75]², prediction of runtimes and resource occupancy, bin-packing-based optimization [27].

20.3.2 Compute-Intensive Batch Processing

Description: Compute-intensive batch processing includes workloads where computation, rather than data I/O and movement, consumes the largest portion of the runtime and of the consumed resources, and is thus the main focus of resource management and scheduling. This type of applications has evolved much over the past few decades, from select few users running large parallel jobs on supercomputers (late 1980s – early 1990s), to practically every research and engineering lab running in multi-cluster grids many small, independent tasks [33], integrated through scripts into mostly compute-intensive jobs (mid-1990s to today). These bags of tasks (BoTs), which are effectively conveniently parallel implementations of scientific and engineering workloads (e.g., simulations), have emerged as a response to the shift from expensive supercomputers that offered high performance and availability, to commodity hardware that crashes often. Since the early 2000s, workflows of inter-dependent tasks, where dependencies are expressed programmatically and inter-task data transfers occur through batch transfers of (typically POSIX) files, have also become increasingly more common in practice [33, 35].

20.3.2.1 Compute-Intensive Batch Processing In Clusters

Application components: Workloads include bags of predominantly sequential tasks and small-scale parallel jobs (in engineering and research labs).

Metrics of interest: Throughput, response time/bounded slowdown, makespan for BoTs and normalized schedule length for workflows.

Typical problems: Increase throughput, reduce response time, and in particular the (bounded) slowdown/makespan for BoTs and the (normalized) schedule length for workflows. Balance performance and cost.

Typical self-aware elements: Traditional techniques for dynamic and adaptive scheduling and resource management. Flagship projects include Condor, Globus, and, more recently, Mesos.

² Commercial products in this domain are scarce. Notable products include VMware's open-source Project Serengeti <http://www.vmware.com/hadoop/serengeti>

20.3.2.2 Compute-Intensive Batch Processing In and Across Datacenters

Application components: At the scale of entire datacenters and in multi-datacenter environments, that is, on the order of 10,000 to over 100,000 machines, load is submitted by thousands of users. Workloads come from scientific computing, financial, engineering, and other domains, and are dominated by bags of tasks of highly diverse sizes and resource demands [33].

Metrics of interest: User metrics are similar to those in the cluster context, but also include aggregate measures of the fraction of deadlines and throughput goals satisfied under extreme conditions such as large-scale failures and flash crowds. Energy costs. Metrics interesting to datacenter administrators include scalability, availability, load balance, and achievable utilization.

Typical problems: Load balancing, particularly across data centers. Handling dynamic load variations, likely due to normal bursty behavior, time-of-day effects, and failures. Enabling high resource utilization without adverse impact on performance and with performance isolation across the workloads of many diverse users.

Typical self-aware elements: Resource management that is aware of dynamic loads, service-level objectives, and failure/maintenance issues. Examples include automatic job queue reconfiguration [17], self-aware job managers [33], and large-scale datacenter management systems [66].

20.3.3 Data-Intensive Batch Processing

Description: There is a plethora of application domains including commercial applications, retail and science domains, that generate big data (large volume, high variety, low veracity, etc.) Data-intensive batch processing involves systems to process sets of big data without interactive data-processing sessions. Such processing is performed by a cluster of compute nodes of data that are typically stored on distributed storage, with intermediary results stored in-memory or on disks local to each node. The nodes collectively execute software that coordinates the distribution and computation of the data sets across the cluster according to the processing semantics. We classify these systems into the following two categories.

20.3.3.1 MapReduce-Based Data-Intensive Batch Processing

Application components: MapReduce is a popular programming model for developing and executing distributed data-intensive and compute-intensive applications on clusters of commodity computers. A MapReduce job is an instance of a running MapReduce program and is comprised of Map and Reduce tasks. Tasks are executed according to the programming model, but embed functions (code) provided by the user. High performance and fault-tolerance are two key features of typical MapRe-

duce runtime environments. They are achieved by automatic task scheduling; data placement, partitioning and replication; and failure detection and task re-execution.

The MapReduce model has proven to be versatile in industry, where it is used for many Big Data tasks including log processing, image processing, and machine learning. For example, MapReduce has been used to learn conditional probability tables of Bayesian Networks (BNs). Both traditional parameter learning (complete data) and the classical Expectation Maximization algorithm (incomplete data) can be implemented within the MapReduce model [6].

Metrics of interest: Performance metrics such as job response time and throughput (jobs/minute), and data input and output (IOPS); reliability - measured as the ratio of successful MapReduce job requests to the total number of requests; various cost metrics; other low-level MapReduce metrics related to the number, length, and status (i.e., success or failure) of MapReduce jobs and tasks.

Typical problems: Performance and dependability guarantees, trading-off between multiple metrics of interest. Chains and workflows of MapReduce jobs are useful, but could be difficult to manage and troubleshoot. Vicissitude-workflows of MapReduce jobs lead to diverse challenges, by stressing different system resources at different or even the same time. Workloads can be dominated by a few MapReduce jobs, used periodically or in bursts.

Typical self-aware elements: Performance models and management [68], self-aware architecture reconfiguration [25].

20.3.3.2 Other Data-Intensive Batch Processing

Application components: Many programming models for data-intensive batch processing, and significantly different from MapReduce, exist today: Spark, Naiad, Ciel, etc. Such systems rely on functional, imperative, or dataflow models to express computations, and may be general or highly specialized.

Metrics of interest: Response time, recovery time, and cost. Some these systems, such as Naiad, are also designed to support low-latency results, to provide support to process data and deliver results in near-real-time.

Typical problems: Handling workload variations and data recovery after machine failures. Synchronization of data state when processing spans many nodes. For graph processing, the algorithm but also its input data set affect performance significantly, but predicting how is challenging. Auto-scaling is very challenging, due to the possible need to transfer large state.

Typical self-aware elements: Recovering from failures. Scaling to additional nodes to handle workload variations and arbitrary computation.

20.3.4 Data-Stream Processing

Description: We observe an avalanche of data continuously generated from various sources such as sensor networks, business operations, web applications, and social networks. There is a pressing need to process such data in real-time. For example, several companies like Facebook and LinkedIn used to analyze their daily web logs to better support their operations [64], but are shifting to real-time analysis.

Application Components: Data stream processing (DSP) involves the real-time processing of data that are continuously generated from several distributed sources at time-varying rates. Data analysis is represented via user queries that describe the type of processing users wish to operate over source data. DSP queries are typically represented by directed dataflow graphs where vertices correspond to operators and directed edges indicate the flow of data among operators. Each operator typically corresponds to certain parts of the query processing—often associated with well defined semantics such as joins, aggregates, filters, etc. Finally, queries are deployed on a cluster of nodes, referred to as data stream processing systems (DSPSs).

Metrics of interest: Primarily, query performance on a DSP is measured via the delivery of low-latency and high-throughput results regardless of the workload demands and their time-based variations.

Typical problems: Dynamic workload and operator diversity (e.g., different semantics), load balancing in datacenters (e.g., System S, SPQR [38], Soda), architectural reconfiguration and performance isolation in cloud environments [7, 28, 60].

Typical self-aware elements: Optimization models for resource allocation and placement [38] and CPU-based heuristics to trigger scale-out operations [7].

20.3.5 Workloads Generated by Datacenter Operations

Description: Unlike the other applications described in this section, this workload is created by the system itself, as response to real-input workloads, in particular to give probabilistic operational guarantees. Typical workloads here are the product of backup, logging, checkpointing, and recovery systems. Although needed to meet declarative specifications of the availability, durability, and recovery time requirements, these workloads cause significant reliability-related overheads that need minimization. For example, on the order of 20% of the resources are currently wasted on failures and spent for recovery in large-scale infrastructure [14].

20.3.5.1 Addressing Failures in the Datacenter

Application components: Implementing data redundancy for availability and disaster tolerance results in some of the largest workloads in datacenters, especially regarding data transfer and storage. Resources consumed may include disk and tape

capacity and bandwidth; CPU and memory for redundancy operations (such as encoding); and network bandwidth within and across datacenters.

Metrics of interest: Availability (expected fraction of time that a desired data object will not be accessible); durability (annual data loss rates, MTDL); amount of data lost during failures (e.g., how far a checkpointed system will need to rewind); recovery time (how long it will take to recover back to a normal operating state).

Typical problems: Recovery planning by selecting the combination of redundancy techniques to meet reliability requirements; designing a schedule of backup operations to fit resource availability and limit interference with other workloads; designing a schedule of recovery operations after a large-scale failure while enabling diagnosis (see Section 20.4.1).

Typical self-aware elements: Automated designers combine multiple resilience techniques to meet reliability and recovery requirements within resource limitations. Both deterministic and stochastic models [42] are used to represent the monitored workload levels and changes, component failure rates, and the reliability provided by the resilient system. Design methods include machine-learning-based techniques, mathematical optimization and meta-heuristic techniques [24, 40].

20.3.5.2 Addressing Failures at Exascale

Application components: The future exascale machines, which will exceed 1 exaflop sustained performance (so, 2-3 orders of magnitude larger than today's Top500 machines) exacerbate the problems and pose important scale challenges to the aspects observed in Section 20.3.5.1. Checkpointing and other mechanisms designed for this scale operate with waves or hierarchies of periodic or triggered operations, either partial or for the entire system, that affect CPU, memory, network, and storage resources.

Metrics of interest: same as for Section 20.3.5.1; also energy and human resource costs (e.g., does recovery require human resources?).

Typical problems: Recovery planning and general automatic recovery approaches are the key challenge in the field, possibly aided by advanced workload prediction, with current approaches leading to poor energy proportionality (high energetic cost). Containment, including performance isolation is important, because correlated (e.g., cascading) failures can cause significant problems to other components and applications than affected by the original failure. Diagnosis and troubleshooting pose important challenges, because at this scale applications can have over a million concurrent threads of execution and are very difficult to debug; even error identification and reporting are important challenges in exascale systems.

Typical self-aware elements: Various methods, surveyed in a recent overview of the field [72], among which: stochastic performance models, to trade-off re-computation of results for stored backups and checkpoints; and, from another research community, redundancy of execution [2], of all or of critical tasks, and on all or a selection of more reliable resources, to reduce the effects of failures efficiently.

20.3.6 Online Gaming

Description: Hundreds of online games (OGs) entertain over 250,000,000 online players, in a global market that generates over 30 billion Euros yearly. Massivizing, which means to scale efficiently while meeting strict SLAs, is the biggest challenge of massively multiplayer online games (MMOGs). We consider here only the resource management for the in-game virtual world, excluding external processes such as gaming analytics (similar to Sections 20.3.3 and 20.3.4) and game-content generation (not standardized). There are many types of online games, among which the most popular are online social games (OSG); First-Person Shooters (FPS) and Real-Time Strategy games (RTS); and Massively Multiplayer Online Role-Playing Games (MMORPG).

Metrics of interest: In general, response time, cost of operation, performance variability impact as aggregate performance penalty (own metric), time- and space-varying reliability or availability.

Typical problems: Reduce cost. Balance cost-performance. Impact of performance variability. Unavailability at critical time or for the critical component.

20.3.6.1 Datacenter-based Approaches

Application components: The most popular OSGs (e.g., the FarmVille series and Clash of Clans) have over 100 million daily active users, and hundreds of OSGs attract over 1 million daily active users. OSGs use multi-tier web applications (described earlier) with hundreds of thousands to millions of concurrent yet short-lived user sessions. Their populations can fluctuate significantly over time, especially during initial deployment and after their peak popularity is gone [55]. MMORPGs (e.g., World of Warcraft, Destiny) commonly use geographically distributed clusters of servers to support (multi-)hour game sessions. FPS games, e.g., the Call of Duty series, and RTS games, e.g., the StarCraft and DotA series, typically use servers to run independent game instance that run for a few tens of minutes; often, these servers are hosted by gaming-friendly datacenters.

Typical self-aware elements: Self-aware provisioning of resources from datacenters, especially in hybrid clouds, using workload prediction and modeling. Cost-aware operation. Portfolio scheduling [69]. Availability on demand [70].

20.3.6.2 Offloading of Mobile Interactive Applications

Application components: An emerging application in this space is that of mobile games that use clouds as offloading target. The structure of such an application is typically a workflow or dataflow, where tasks have inter-dependencies and typically execute iteratively (the input-update-synchronize cycle common in game design). Some or all tasks can be offloaded to the cloud. For example, cloud gaming appli-

cations could offload all update tasks to the cloud and stream back to the mobile device a video rendering of the current game status.

Metrics of interest: The general metrics, plus energy costs, and metrics relating to more complex costs of operation (e.g., roaming and other special rates).

Typical problems: the general online gaming problems, plus focus on energy.

Typical self-aware elements: Applications may offload only (a part of) computation, data acquisition, or another resource-consuming part. Several, but not many, feedback and reconfiguration techniques, and stochastic performance models to address which part to offload, where, and how already exist [56, 76].

20.3.7 *Partial and Delayed Processing*

Description: Cloud applications have functional requirements on the computations they should perform and non-functional requirements on the additional properties they should have. Some of these requirements can be expressed as incremental requirements on application behavior. For example, the accuracy of the answer and the reliability of the operation could be bounded in stages. In partial processing, applications can gradually downgrade user experience to avoid system saturation.

Partial and *delayed* processing addresses such requirements. For example, recommendation engines used by online shops to offer end users suggestions of similar products are often highly demanding on computing resources [46]. Using a partial input instead of the entire database of choices can lead to acceptable recommendations for both users and system. With delayed processing, applications can extend their response time to cope with overload conditions [37] or dynamically vary resource allocations to batch processing jobs to increase overall utilization. [51].

Metrics of interest: Typically, (bounded) response time: the latency experienced by the users or the latency to deliver results in real-time data stream processing. Also, reliability, and cost to produce the results.

Typical problems: In addition to traditional requirements, dynamic loads and variable number of users, making dynamic resource provisioning necessary [62]. Unexpected events and potential failures, such as large load spikes, software and hardware failures, and lack of performance isolation during workload consolidation.

Typical self-aware elements: For partial processing, enabling and disabling optional components on the fly, which leads to bounded response time [13, 44] and ability to address multiple failures [45]. Stochastic performance models can help in analyzing the tradeoffs in use of content vs. capacity requirements, in use of content vs. response time, and in (iterative) response accuracy vs. response time.

In some cases, extending the response time slightly (e.g., a few seconds every hour) is typically unnoticeable for users, but solving this optimization problem with flexible constraints can significantly reduce the impact of overloads [37, 51].

20.3.8 Cyber-Physical Applications

Description: In cyber-physical system (CPS) applications [61], a computing system interacts with the physical world in some non-trivial manner. While originally exemplified by closed embedded systems, such as industrial automation, the scope of CPS applications has grown over the years to include larger open systems such as disaster response, medical applications, energy management, and vehicular control.

Metrics of interest: There are two often conflicting evaluation metrics for CPS applications: safety and performance. Attaining higher performance often requires greater coupling between components, but such coupling introduces complexity and pathways for failure propagation which may compromise safety.

Typical problems: A self-aware architecture should offer ways to meet both safety and performance requirements despite the conflict between them. Adaptation is needed to attain good tradeoffs, as discussed in the examples below.

20.3.8.1 Medical Applications

Application components: Consider a future implanted smart insulin pump for a diabetic patient. Multiple control mechanisms may be present that take measurements of activity levels from accelerometers and measurements of sugar levels in the bloodstream to modulate insulin delivery. While inner control loops will run locally, there is an opportunity to perform some predictive optimization in the cloud based on context derived from user location, synched calendar, and other factors. Decisions on where to execute which functionality at what time may be revisited dynamically to adapt to different network and patient conditions, as well the current control objectives (see Section 20.5.1 for details).

Metrics of interest: The safety requirement may specify that the pump shall never overdose the patient (which could be fatal). A performance requirement may specify that the pump should predictively adapt its output depending on the person's activity and food intake. The predictive aspect is key, because the human body has too large of a response time for purely reactive (feedback) control schemes to offer tight sugar regulation.

Typical Problems: Attaining better predictions requires exploitation of more information. Acquisition of this additional information comes at the cost of having to connect to other less reliable subsystems (e.g., the cloud), creating dependencies that may act as pathways for failure propagation and hence safety violations. Hence, a conflict manifests between performance and safety.

Typical self-aware elements: The need to reconcile safety and performance gives rise to a new type of adaptation, where the system toggles between meeting performance objectives and meeting safety objectives, depending on current state. In the nominal (normal operation) state, the system optimizes performance. However, when system approaches boundaries of safety, the objective switches to enforcement of safety, even if performance is affected. This approach is commonly referred to as the Simplex architectural pattern. For example, when the insulin pump

observes large deviations in patient's blood sugar levels, it may switch to a simple PID control mode based only on trusted local sensors and disconnect itself from the less reliable cloud inputs that might be offering bad predictions.

20.3.8.2 Vehicular Applications

Application components: Vehicles are an interesting and emerging case study for CPSs. Up until now, development has focused on making the individual more autonomous. While this trend continues, we believe there will be an increasing emphasis on (i) communication between vehicles of different degrees of autonomy; (ii) careful sharing of CPS information resources; and (iii) system health management [8].

Metrics of interest: Safety requirements might be that a vehicle shall avoid hitting other vehicles or pedestrians; traffic signs shall be obeyed. Performance requirements could be to minimize fuel consumption and travel time.

Typical problems: Better performance may be achieved by exploiting global information, typically from a cloud service. The service might include a database of all traffic signs, estimates of current traffic conditions, and a schedule of traffic signals. This information can be used to plan itineraries and driving speeds such that fuel, trip time, and other passenger preferences are optimized. To ensure safety, however, only reliable sensors that are local to the car should be used in making decisions. However, these sensors have only a local view and may miss various performance optimization opportunities.

Typical self-aware elements: To attain a good tradeoff between safety and performance, a local override mechanism is needed to take control when a safety requirement is about to be violated. For example, if the perceived state at an approaching intersection is different from that reported by the cloud (e.g., light is red, not green), the local override should take over and manage the car based on local sensors only. The idea is to use the subset of most reliable information only, when the system state is close to a safety violation boundary, while exploiting additional less reliable sources for optimization in other states.

20.4 Types of Problems

In this section, we identify the following ten types of problems that affect cloud applications, and that can benefit from the use of self-awareness techniques.

1. Recovery planning
2. Autoscaling of resources
3. Runtime architectural reconfiguration and load balancing
4. Fault-tolerance in distributed systems
5. Energy-proportionality and energy-efficient operation
6. Workload prediction

7. Performance isolation
8. Diagnosis and troubleshooting
9. Discovery of application topology
10. Intrusion detection and prevention

20.4.1 Recovery Planning

Context: Enterprise storage systems are designed to be resilient to failures, but when a large failure occurs—for example, a datacenter level failure—recovery is a complex process, and frequently involves some application downtime. It is important to recover the most important applications quickly; but *What is the sequence of recovery operations that will minimize the damage?* [40] and *What can prevent the failure from reactivating?* [72]

Problem: When an enterprise storage system experiences a large failure, critical applications must be recovered quickly, even at the expense of additional downtime for less important applications. Administrators are under stress, and have little time to design the best sequence of recovery operations, and the default methods may be far from optimal, since each failure can be different. Preventing the re-occurrence of the failure, for example by isolating, reconfiguring, or micro-rebooting the failed component, also requires careful planning. A self-aware recovery system that understands the failures, the criticality of the applications to be restored, and the possible recovery options, can propose or automatically execute a customized recovery plan.

Expected Advancement: By codifying the recovery operations, the cost of downtime for applications, and what is needed to bring each application back up so that future failures are also avoided, it is possible to model many possible recovery plans. An automated system can select an optimal schedule that balances recovery targets with cost and resource waste.

Expected Impact on Application Types: A recovery planning optimizer should be integrated with the overall mechanisms for managing backups, failures, and related workloads in datacenters (see Section 20.3.5.1). Clearly, the availability of resources (e.g., inter-datacenter bandwidth) impacts the recovery process, and should be planned accordingly. In the other direction, the design of the backup and restore mechanisms should be visible to the recovery planner in a way that enables introspection, modeling, and updates when the mechanisms or the applications change.

20.4.2 Autoscaling of Resources

Context: Many applications are subject to time-varying workloads. For instance, workloads of internet and enterprise applications typically contain time patterns (e.g., day vs. night, seasonal effects), long-term trends (e.g., increasing customer base), and bursts (e.g., flashcrowds of interest for content).

Problem: As a consequence of time-varying workloads, sizing a system for the expected peak workload is very inefficient (by orders-of-magnitude!) and may be infeasible if the workloads are hardly predictable. Therefore, a system should be able to dynamically acquire and release resources (e.g., number of replicated VMs) as required for serving the current workload with a certain level of performance. Target levels for the application performance may be specified in Service-level Agreements (SLAs). Cloud computing provides the required flexibility to dynamically change the amount of resources allocated to applications. However, scaling controllers in state-of-the-art cloud platforms are following simple trigger-based approaches (e.g., if “utilization is above a given threshold, add one VM”), lacking knowledge about the structure and behavior of the application. Moreover, few are able to respond to rapid, bursty load transitions [23]. Scaling controllers when data processes require non-trivial management, for example for big data processing (see Section 20.3.3) are even more difficult to design [25].

Expected Advancement: A certain self-awareness of the system is required to take acceptable or even optimal decisions about when to scale an application vertically or horizontally and by what amount of resources; about which part of an application to scale if the application is comprised of multiple tiers, components or tasks, or concurrent threads of execution; about which (part of the) application to ensure against the risk of under-performance or failure; etc. Ideally, the self-aware system will be able to enforce high-level objectives as specified in SLAs (e.g., end-to-end response times and maximum risk of SLA breaches).

Expected Impact on Application Types: Autoscaling affects resource provisioning in both multi-tier enterprise applications (Section 20.3.1), and component-based applications such as data-intensive batch processing (Section 20.3.3), data stream processing (Section 20.3.4), and datacenter-based online gaming (Section 20.3.6.1). This allows the auto-scaled applications to meet their quality of service goals in spite of increased workload demands.

20.4.3 Runtime Architectural Reconfiguration and Load Balancing

Context: For many of the applications we describe in Section 20.3, the conditions they are operating under can change at runtime. These include workloads from users or other systems interacting with the application, resource usage patterns of the application itself, component failures from the hosting platform, as well as competing demands from other applications sharing the same infrastructure.

Problem: In contrast to autoscaling, which keeps the same architecture while changing the scale of the system, architectural reconfiguration is a family of techniques that reconfigure some architectural components of the overall system. When workload changes occur, the current architecture of the system may become obsolete. For example, even if individual virtual clusters can auto-scale, the overall architecture of how the virtual clusters are deployed on the physical infrastructure also

need self-aware capabilities. A self-aware architecture should reconfigure and manage its components, for example by resizing its queues and changing their scheduling policies, or by changing the paths for sharing loads between different queues.

Another common problem in managing a cluster of resources (e.g., servers or storage devices) is how to automatically balance the load across the cluster. This involves migrating workloads from one server or device to another quickly and without any downtime to the applications. Load blancing can also happen among multiple clusters running similar or different workloads.

A further complication can arise in data centers with thousands of hosts and services where some services have affinity or anti-affinity constraints. For example, a service and its backup service should not reside on the same host (*anti-affinity*), while it is preferred to have a user interface service and its corresponding DB service on the same host (*affinity*). It is a challenge to maximize the utilization of the servers while providing a high-quality user experience and not violating these constraints.

Expected Advancement: Runtime reconfiguration and load balancing typically require solving an online optimization problem, whose objective involves specific performance or availability metrics. First, for each type of applications, such metrics need to be identified, collected, and calculated in real time. Second, we need to develop techniques for quantifying the cost of each reconfiguration step (e.g., an architectural change or a VM migration) and for weighing the cost against the benefit of the reconfiguration. Thrid, we need frameworks for dealing with the fundamental tradeoff between faster response and stability. Finally, advances on load rebalancing optimization under affinity or anti-affinity constraints are expected, especially with the ever increasing scale and complexity of such problems.

Expected Impact on Application Types: Runtime reconfiguration benefits especially applications who are not negatively impacted by the duration and other costs of reconfiguration. Among such applications are batch processing applications, either compute-intensive (Section 20.3.2) or data-intensive (Section 20.3.3), and some data-stream processing applications (Section 20.3.4).

20.4.4 *Fault-Tolerance in Distributed Systems*

Context: Due to ever-increasing scale and complexity, hardware and software faults (which lead to errors, which *may* lead to a failure) in cloud computing infrastructures are the norm rather than the exception [4]. This is why many from the application classes introduced in Section 20.3 include fault-tolerance techniques, such as replication, early in their design.

Problem: Failures in cloud infrastructures are often correlated in time and space [22, 79], which means they a failure may affect tens to hundreds of nodes, or even entire datacenters. Therefore, it may be economically inefficient for the service provider to provision enough spare capacity for dealing with all failures in a satisfactory manner. When correlated failures occur, the service may *saturate*, i.e., it can no longer serve users in a timely manner. This in turn leads to dissatisfied users,

that may abandon the service, thus incurring long-term revenue loss to the service provider. Note that the saturated service causes infrastructure overload, which by itself may trigger additional failures, thus aggravating the initial situation. Hence, a mechanism is required to deal with rare, cascading failures, that feature temporary capacity shortage. The main problem is to maintain bounded response times in the presence of failure, while wasting an acceptable amount of resources (today, about 20% of the entire capacity, but the goal for exascale systems is to waste under 2% [72]). The problem of fault handling also includes a component about preventing fault re-occurrence, which includes elements of diagnosis, troubleshooting, isolation, and (micro-)rebooting. [72]

Expected Advancement: Advances are expected in the use of control theory with the brownout approach, in smart load-balancing using knowledge gained with control, in self-checking and self-diagnosing, and in self-reconfiguration and in smart decisions about micro-reboots. Using these techniques, the applications perform better at hiding faults from the user, as measured in the number of timeouts a user would observe.

Expected Impact on Application Types: This problem affects request-response applications (Section 20.3.7).

20.4.5 Energy-Proportionality and Energy-Efficient Operation

Context: A problem tightly related to autoscaling is one of energy proportionality. Workloads in many applications are becoming more data-centric. In other words, the data volume, and not the algorithmic complexity, is becoming the primary contributor to resource consumption. Single pass algorithms are used on most streaming data, and their complexity is largely linear in the data size. Moving data across machines is therefore very expensive, compared to the cost of data processing.

Problem: It is hard to design systems where resources operate at capacity all the time. Necessarily, some resources will not be fully utilized. Solution such as autoscaling could become prohibitively expensive if they involve frequent data movement. On the other hand, in the absence of autoscaling, some machines will not be fully utilized. This operating mode exposes a problem with most current data center hardware; namely, energy proportionality (or, rather, lack thereof). A server that is only 30% utilized may be using 80% of the energy needed at full load. One needs to design solutions where energy consumption shrinks proportionally to load.

Expected Advancement: Attaining energy proportionality in data-centric applications is challenging. It requires algorithms that minimize unnecessary data movement, while performing autoscaling. These algorithms must amortize cost of data movement over time [49]. The latter may require a prediction of future data access patterns [39].

Expected Impact on Application Types: Energy proportionality will benefit both data-centric multi-tier enterprise applications and stream processing applica-

tions by allowing them to operate in a more energy efficient manner while minimizing the need for data movement.

20.4.6 Workload Prediction

Context: The increased volume of data involved in modern cloud applications suggests that initial data placement will play a big role in application performance and energy consumption. Improper placement will create future load imbalance (e.g., if many popular items are collocated) or needlessly increase energy consumption (e.g., if infrequently accessed items are placed together with some frequently accessed ones thus preventing machines from being turned off). Chapter 18 discusses these issues in further detail, from both a time-series and a machine learning perspective.

Problem: In (nearly) stateless services, such as those serving static Web pages, a load balancer can rectify imbalance simply by distributing future requests in a more equitable fashion. In applications where moving data is costly, it is harder to predict computing load because such load has a substantial data affinity. Hence, data placement dictates where computation runs. Getting the placement right in the first place becomes important. This motivates techniques for predicting future access patterns to data items at the time these items first enter the system and are stored [39].

Expected Advancement: Self-awareness techniques that can understand and represent efficiently the state of the system. Collecting and summarizing monitoring data at the scales expected for cloud computing infrastructures and workloads is challenging, yet needed. Proper data access pattern prediction techniques will minimize the need for moving data unnecessarily, and hence improve both performance and energy consumption of data centers. For example, data predicted to be of no further interest could be moved proactively to servers that operate in more aggressive energy saving modes, or are in places that are harder to cool, hence saving energy. Similarly, data predicted to be popular could be partitioned among a sufficient number of servers, reducing the chances of developing hotspots and needing to relocate some of the data to other machines. To conclude, increasing awareness for the lowest possible cost is an important trade-off that remains largely unexplored.

Expected Impact on Application Types: Most cloud computing applications will benefit from some form of workload prediction.

20.4.7 Performance Isolation

Context: Support for multi-tenancy is an important feature of clouds. For example, SaaS offerings are typically implemented by multi-tenant application architectures. Multi-tenancy means that different tenants from separate organizations are sharing the same application instance and see their own tenant-specific view of the data

and functionality. Thus, the operator of a SaaS provider can increase the efficiency compared to running separate application instances.

Problem: The tenants of a cloud service may, unwillingly or even willingly, affect the operation of the system and thus each other. If one tenant exceeds its shared portion, or if the services are oversubscribed and the rightful tenants access the service simultaneously, the performance as experienced by the other tenants can fluctuate or even depreciate significantly. Likely because of (lack of) performance isolation, SaaS, but also PaaS and IaaS clouds, can experience high variability in the performance of their service [34].

Expected Advancement: To ensure performance isolation in such disruptive scenarios, per-request admission control is required, that automatically throttles users exceeding their quota to avoid breaking the SLAs of other tenants.

Expected Impact on Application Types: Multi-tier enterprise applications in a cloud environment need both performance isolation among different application instances (e.g., on IaaS) [50,58] or different tenants of the same application instance (e.g., on SaaS) [47].

20.4.8 Diagnosis and Troubleshooting

Context: When an application or service goes down or fails to reach the service level objective (SLO) regarding its end-to-end performance, one needs to engage in the process of diagnosis and troubleshooting.

Problem: The key problem is to identify the faulty component that has caused the failure, or the associated performance bottleneck that has led to the service degradation. This can be challenging due to the increasingly more complex and distributed nature of modern applications, their growing space of configurations, the typically time-varying workload demands, and the applications' dependency on a variety of hardware resources, such as processors, memory, storage and network I/O devices, as well as software resources, such as locks, threads, and connection pools.

Expected Advancement: Traditionally, maintenance personnel, system administrators or datacenter operators perform diagnosis and troubleshooting manually, using a combination of logs, performance charts, best practices menus, and their domain knowledge, which is time consuming and error prone. With the utilization of self-awareness techniques, we should build management services that can automatically determine the likely causes of failures or performance problems [43, 77, 78].

Expected Impact on Application Types: Applications from enterprise multi-tier systems to networked cyber-physical front-ends can benefit from automatic diagnosis and troubleshooting, resulting in shorter durations of failures or service-level violations, and reduced cost in management and operations. Prior work addresses system health management, including diagnostics and prognostics capabilities [52,63,65], application troubleshooting [43], and troubleshooting uncoordinated self-aware managers [30].

20.4.9 Discovery of Application Topology

Context: Automatic discovery of application topology or runtime architecture is a required feature for any mature application performance monitoring or management solutions, according to the Gartner APM Conceptual Framework.

Problem: The problem here is to automate the process of identifying the relationship and dependency among individual application components at runtime, as well as how they are mapped to the physical infrastructure (e.g., servers), with no or only minimum input from human operators.

Expected Advancement: To solve the above problem, one needs to implement real-time, fine-grained tracing of individual transactions as they traverse through the execution paths of the application. Such monitoring solutions can be passive or require instrumentation at the kernel, middleware, or application level. Statistical techniques for inferring correlations or discovering dependencies are often needed.

Expected Impact on Application Types: Having access to an accurate application topology can help diagnose or debug performance degradations and discover hidden performance bottlenecks in multi-tier enterprise applications (Section 20.3.1.1) during their operation or help identify potential root causes of observed failures through event correlation.

20.4.10 Intrusion Detection and Prevention

Context: In cloud environments, different applications coming from diverse organizations may share the same physical resources. Depending on the cloud service model (IaaS, PaaS, or SaaS), the data-center owner has different levels of control over the executed application and their configuration. Vulnerabilities in the infrastructure software (e.g., hypervisors) or in shared services (e.g., storage services) can be exploited to widen an attack from any application to any other virtual machine in the same infrastructure. For instance, attackers may rent virtual machines (in case of public clouds) or exploit vulnerabilities of applications (e.g., private website) to get access to sensitive data (e.g., e-commerce system with credit card data) in other virtual machines in the same cloud environment.

Problem: The detection and prevention of attacks in a cloud environment requires a classification of cloud workloads (either using application or network probes) into benign and malicious ones. False positives can result in unnecessary actions countering attacks, which may have negative effects on the applications.

Expected Advancement: A self-aware system automatically learns to distinguish between benign and malicious workloads and can filter out false positives. Furthermore, the system is able to react to attacks and adapt itself to ensure its self-protection capabilities. For more details on quantifying the self-protection capabilities of self-aware systems we refer the reader to Chapter 14.

Expected Impact on Application Types: All application types running in a cloud are potential targets for attacks. Prior work in this area aimed at automatic workload anomaly detection and at enabling self-healing capabilities of systems.

20.5 Types of Approaches

In this section, we identify and analyze the following seven types of self-awareness approaches used in practice to address the problems identified in the previous section.

1. Feedback control-based techniques
2. Metric optimization with constraints
3. Machine learning-based techniques
4. Portfolio scheduling
5. Self-aware architecture reconfiguration
6. Stochastic performance models
7. Other approaches

20.5.1 Feedback Control-Based Techniques

Description: Control theory is a branch of mathematics that studies how to influence the behavior of dynamical systems [48]. Based on a formal model of the target system in the form of equations (for time-based control) or automaton (for event-based control), control theory provides principles for how to synthesize a controller that would regulate the behavior of the system and obtain prescribed properties.

Expected Impact: Although control theory was invented to deal with physical systems, the same principles have been successfully applied to many different application domains in computing systems [29], including resource allocation [36,50], application performance via bounded response times [37,44], reliability [19], and fault tolerance [7,45], stream processing [38], and big data [18].

Details: A controller should provide the following properties [12,20]:

- *Setpoint Tracking.* A setpoint refers to the goals to be achieved. For example, a system is considered responsive when its user-perceived latency is sub-second.
- *Transient behavior.* This concerns *how* the setpoint is reached by the system.
- *Robustness.* The controller should be able to cope with inaccurate measurements, delayed data, or other uncertainties not captured in the system model.
- *Disturbance rejection.* The closed-loop system should reach its goal in spite of other disturbing actions happening simultaneously in the system.

These properties are often translated into the corresponding control properties of stability, no overshooting, quick settling time and robustness to model errors and disturbances.

Use Cases: There have been many published studies of developing self-awareness capabilities using feedback control based techniques.

One broad application area has been resource management in computing systems. For example, [36] describes a control-based approach to assign CPU resource shares of virtualized web server applications. This approach emphasizes the CPU allocation around periods of workload changes and uses H^∞ filters to minimize the maximum controller error. The work also extends previous work on optimal configurations for CPU-resource entitlement of virtualized multi-tier web applications. More recently [67], a control-theoretic approach was described to provide performance, dependability and cost guarantees for online cloud services, with time-varying workloads. The approach is validated through case studies and extensive experiments with online services hosted on Amazon EC2. One case study demonstrates SLA guarantees for a cluster-based multi-tier e-commerce service. In [50], application managers automatically learn a quantitative model that correlates app-level performance with resource utilizations and use control theory to derive the optimal resource control settings (limits, reservations) for individual virtual machines such that the multi-tier application can achieve its performance target. Finally, [20] generalizes and automates core allocation based on various metrics. The approach provides formal guarantees on the application behavior in spite of external disturbances like additional load on the machine. This shows the versatility of feedback control as an approach for building self-aware and self-adaptive systems.

Another broad application area for feedback control is cyber-physical systems (see Section 20.3.8). Feedback control is central in smart grids [16], intelligent transportation, modern critical care units, etc. Parameters and offsets of local loops might be obtained from remote repositories. The critical need for these applications typically requires continued, correct operation even in the presence of connectivity problems to the cloud, bad data, and other failures, which often leads to joint investigation of control and safety.

20.5.2 Metric Optimization with Constraints

Description: System design, configuration, and management decisions often come down to a choice between various options. Self-aware systems can reason about the impact of different choices, but that still leaves the question of which choice best achieves the system goal. Optimization techniques require first a formal specification of the system objective and the constraints under which the system operates, and then a solution that attempts to optimize the objective. Solution methods can be exact or approximate, depending upon the situation, but the explicit specification of goals is essential.

Expected Impact: Using optimization techniques for system decisions has several advantages [41]. First, it guarantees clarity: system decisions often come down to choosing between conflicting goals, and specifying an objective forces the designer to explicitly choose how much weight to assign to them. Secondly, when

approximate solution techniques are used. It enables a quantitative evaluation of how far the results are from the optimal. Thirdly, it allows use of the enormous toolbox of optimization techniques that already exists, allowing the system designer to focus on the design aspects.

Details: System designer sometimes use ad hoc heuristics to make system decisions, on the basis that specifying there are conflicting, sometimes intangible requirements that are hard to quantify, and the heuristics produce results that are “good enough”. However, it is then unclear what “good enough” means, or if the heuristics in fact achieve it. Expressing the trade-offs into a common currency – whether execution time, throughput, monetary cost, or a utility function – enables a definition of goodness. A wide variety of optimization techniques can be applied, depending on the formulation and requirements: mathematical optimizations, such as linear programming or mixed integer programming; constraint programming, if only a feasible solution is required; meta-heuristic techniques such as simulated annealing or genetic algorithms, when the optimization problem cannot be solved exactly in the available time; or other approximate optimization techniques that bound the error in the resulting solution.

Use Cases: Examples of the use of optimization in systems include Maestro [54], where online optimization was used with feedback-control to provide performance differentiation between applications in a disk array, and Janus [3], where offline optimization was used to determine allocation of flash resources to workloads. Meta-heuristic techniques have been used for finding the sequence of recovery operations to use after a failure to minimize the cost of the downtime [40] and for designing a redundancy configuration for large enterprise storage systems, in order to minimize the overall cost of the system including operating overheads and potential downtime costs.

Another example from the commercial world is the VMware Distributed Resource Scheduler (DRS) [27], a widely used feature in the VMware vCenter management software. DRS manages a set of virtual machines running on a cluster of physical hosts and performs dynamic load balancing to avoid hot spots and improve application performance, by solving an online optimization problem with hill-climbing heuristics and by taking into account the cost/benefit tradeoff of each move.

Finally, this approach can be applied to the applications described in Section 20.3.4. For example, the SQPR query planner [38] allocates physical resources of heterogeneous clusters to data stream processing queries. SQPR models query admission, allocation and stream reuse as a single constrained optimization problem and solves an approximate version to achieve scalability. The SQPR approach monitors the resource utilization across the cluster and performance progress of running queries to decide on queries’ placement and allocation. SQPR adapts to operating conditions through continuous monitoring and modelling of cluster and queries performance. SQPR uses an off-the-shelf optimization solver for optimal solutions.

20.5.3 Machine Learning-Based Techniques

Description: Machine learning is the science of using data to “uncover an underlying process” [1]. More specifically, it involves designing, implementing, and validating a set of algorithms that can extract insights from data regarding the relationships among objects and events, often captured in the form of statistical models.

Expected Impact: Besides the successful application of machine learning to financial applications, e-commerce, and medical applications, there has been a great deal of research in the past decade on leveraging machine learning based techniques to creating self-awareness for business critical applications (Sec. 20.3.1).

Details: Supervised learning, unsupervised learning, and reinforcement learning are the three main types of machine learning approaches. They are commonly applied to perform clustering, classification, and prediction, using a variety of statistical models, including decision trees, regression models, neural networks, Bayesian networks, or support vector machines.

A system that utilizes machine learning typically has the following components:

- *Sensors.* Software or hardware modules that measure and collect metrics of interest for the target system or process.
- *Preprocessor.* Raw data collected from systems are rarely perfect. It is not uncommon to have missing data from certain components or during certain periods of time, or data corrupted during collection. Such data need to be “cleaned up” or aggregated before being fed into an analysis engine.
- *Analyzer.* This is where statistical learning algorithms are being run, on top of the collected data, to extract relationships, and to build models that can represent the learned behavior in a concise form.
- *Reporter/Predictor.* This is where the insights gained from data are presented to human operators in the forms of alerts, charts, or dashboards. Alternatively, the model learned can be used to generate predictions for the target system or metrics, or to recommend remediation actions.

Use Cases: In [9], Tree-Augmented Bayesian Network (TAN) models are learned on top of instrumentation data collected from a three-tier Internet service to automatically identify the top system-level metrics that have likely contributed to the observed violation of service level objectives (SLOs).

In [57], an autoscaling system employs reinforcement learning to automatically learn the performance behavior of a multi-tier application when instances of an elastic tier are added or removed, and then uses the knowledge to scale the application horizontally as its workload demand varies over time.

20.5.4 Portfolio Scheduling

Description: Traditional scheduling policies are designed a specific workload, and sometimes even for specific applications within a workload; in practice, reuse of old

and adoption of new scheduling policies happens rarely. Instead, portfolio scheduling, which is a self-aware and self-expressive technique, considers a set (portfolio) of scheduling policies, from which it selects at the appropriate moments (e.g., periodically) the policy which promises the best results for the current and expected conditions. In this way, the portfolio scheduler combines the strengths of each individual policy in its portfolio, which also means that designers of scheduling policies can focus on simpler policies that do not need to address every type of workload.

Expected Impact: Portfolio schedulers promise to deliver performance at least as good as any of their constituent policies, without poor performance when the workload changes. In this context, performance includes traditional performance metrics, such as application response time and system utilization, and non-traditional metrics, such as availability, cost-performance efficiency [11], and risk of SLA violations [75], etc.

Details: The concept of portfolio scheduling derives from economics, where stock brokers can use portfolio theory to select policies for managing their stocks, to balance risks and rewards. For cloud applications, portfolio scheduling uses the following four-stage iterative process. In the configuration stage, the portfolio scheduler is equipped with a set of scheduling policies. Then, the portfolio scheduler goes through a selection stage, which results in the selection of an optimal policy; through an application stage, which results in applying the optimal policy to the current scheduling problem (queue) and in monitoring the results; and through a reflection stage, where stale policies are possibly eliminated and the portfolio can compare its operation relative to the goals of the system. The next cycle can be triggered periodically or, if enough resources allow for timely completion of stages, whenever an event can lead to system reconfiguration (e.g., at the arrival of a new request in the system).

Use Cases: Portfolio scheduling has been used for business-critical applications, with a focus on reducing the risk of SLA violations [75]. The selection stage is simulation-based, online, and applied after each arrival of a job and periodically. The optimal policy, from the policies included in the portfolio, is then applied until the next selection occurs. The results obtained by applying this portfolio scheduler on the workloads of a real cloud provider indicate that this portfolio scheduler is better than its constituent policies, but also that the initial configuration of the portfolio is very important.

Portfolio scheduling has been used for online gaming applications, but only tested under laboratory, albeit realistic, conditions [69]. Here, the portfolio is configured with various typical online scheduling policies, but also with an optimal solver of a linear integer problem; in the selection stage, the portfolio is given a limited amount of time to decide.

20.5.5 *Self-Aware Architecture Reconfiguration*

Description: Self-aware architecture reconfiguration is a family of techniques that can re-configure at runtime the architecture of the system or of the application.

Expected Impact: The generality of this approach is an advantage, but also makes the actual impact difficult to predict in advance. By affecting the essence of the entire system, this approach can affect every metric and virtually every application. When applied to applications, this approach can change the operational characteristics of the application. This approach addresses the runtime architectural reconfiguration problem.

Details: By monitoring the environment and the (queued) workload, by predicting the performance of the system or application, and by reflecting on the overall goals of the system or application, self-aware architecture reconfiguration leads to changing of operational characteristics (queue size and policy, structure of application components), but also the way in which the system or application operate (e.g., how queues share load).

Use Cases: Koala-C [17], which services the Dutch research cloud DAS4, creates a system with multiple queues, which can be instructed to shed load to only one other or a group of other queues. Each queue is aimed to run jobs of a specific runtime (job size). If a job scheduled by a queue does not finish in the time allocated for that queue (exceeds the job size for that queue), it is stopped and moved to a queue of larger size. The approach does not require any prior knowledge about the input workload. Instead, each job is submitted to the queue(s) servicing the shortest job; larger jobs traverses progressively the chain of queues, with probabilistic guarantees in terms of performance and wasted resources. (Queues can also be equipped with autoscaling mechanisms.) Another use case [31] is that of the cloud provider Blue Yonder: through a parameterized application performance model, a set of modeled adaptation points, and an adaptation model, the Blue Yonder applications can adjust their resources to workload changes and can share resources between customers.

20.5.6 *Stochastic Performance Models*

Description: Stochastic performance models (e.g., Queuing Networks) enable the prediction of the expected system performance in terms of throughput, response time and utilization for a given workload. In a self-aware system, these models are automatically created and maintained in the learning phase.

Expected Impact: In contrast to machine-learning or feedback control based techniques, they promise the ability to predict also the performance for previously unseen workloads or configurations which are significantly different to the current operating point. Stochastic performance models have been applied to dynamically reconfigure enterprise multi-tier applications for auto-scaling [73] and performance isolation [47].

Details: Classic queueing networks model the computing at processing resources (e.g., CPU, hard disk, network). Furthermore, extended formalisms, such as Layered Queueing Networks or Queueing Petri nets, are able to capture the influence of passive resources (e.g., software or memory resources). In Chapter 16, a more extensive discussion of stochastic performance models can be found.

Use Cases: In [47], the authors describe a request admission controller for multi-tenant applications (e.g., SaaS application) that enables performance isolation and SLA differentiation between tenants. Using a combination of statistical estimation of resource demands and operational analysis from queueing theory, the admission control is able to determine the current resource usage of individual tenants and determine which requests are admitted and which are delayed.

In [73], a controller for vertical CPU scaling of virtualized applications is presented based on a Layered Queueing Network. The model is dynamically parameterized based on monitoring data from the system and also captures contention effects due to hypervisor scheduling. By using the performance model, oscillating reconfigurations are avoided and the parameters can be estimated automatically at runtime.

20.5.7 Other Approaches

A number of other approaches have been developed by related fields, such as design science [59]. Although these approaches do not share the entire spectrum of characteristics of, for example, control-based systems, they can be seen as proto-self-awareness approaches. We enumerate in the following several of these approaches.

Other approaches for auto-scaling: For data-intensive batch processing, which typically has large state, auto-scalers such as Amazon Elastic MapReduce and FAWKES [25] use various types of system feedback to decide on scaling. Amazon's approach is based on the topology-aware S3 storage. Focusing on application semantics, FAWKES considers various types of VMs, including VMs that store data transiently or permanently. These approaches consider the MapReduce programming model, which has not been designed to enable interactive processing (e.g., for fast decision making) or tightly coupled data items (e.g., for graph processing).

In the area of data stream processing, systems such as Apache S4 and Storm exploit intra-query parallelization to scale out operators and eventually handle demanding resource requirements per operator. There are two main challenges in operators' parallelization to support a scalable DSPS, i.e., how to handle operators' state and when to scale out/in. Existing research has shown that simple heuristics to detect violations on resource-based thresholds can be used to horizontally scale out to additional stream processing operators to handle excess load [7, 28]. When handling operator state, most systems focus on the parallelization of stateless operators. In [7, 60] different generic approaches to manage the parallelization of both stateful and stateless operators are introduced. In [18] an advanced approach to handle operators state alongside specialized programming primitives is discussed.

Other approaches for runtime architectural reconfiguration and load balancing: Better load balancing within a cluster of compute nodes [27] or storage nodes [26], or between clusters [17], can lead to avoiding hot spots and to improving the performance of many types of applications, including compute-intensive scientific applications (Section 20.3.2), business-critical applications (Section 20.3.1), and batch-processing applications (Section 20.3.2). Reconfiguring the system partitions, not only in scale but also in the way they interact, can greatly benefit batch workloads where long and short jobs can coexist [17]. A genetic algorithm for load rebalancing (GALR) was proposed to deal with affinity or anti-affinity constraints [74]. In experiments running the PlanetLab dataset in a cluster up to 300 hosts, GALR performs close to optimality in load rebalancing, within 4-5 seconds.

Other approaches for fault-tolerance in distributed systems: Experimental results demonstrated that using brownout [44] and smart load-balancing [13] application can tolerate more replica failures and that the novel load-balancing algorithms improve the number of requests served with optional content, and thus the revenue of the provider by up to 5%, with high statistical significance [45]. Various self-testing and self-diagnostic techniques exist (for example, SMART, which is a commercial technology for hard-disks), but self-repair through self-reconfiguration and self-rebooting is still a largely open field.

20.6 Open Challenges for Self-Aware Cloud Applications

In this section we focus on identifying and analyzing directions for future research in developing self-aware cloud applications. We ask questions that we find challenging, yet promising, that is, whose answers we foresee being put in practice in the next 3-5 years.

To what extent are self-awareness techniques necessary for cloud applications? Currently, non-self-aware techniques are prevalent in the management of increasingly large cloud datacenters and their applications. Uncertainty about the need and possible gains limit adoption of self-awareness techniques for this setting. Anecdotal evidence such as that gathered in this chapter and the existing small controlled experiments conducted by scientists provide some evidence of the benefit of using self-awareness techniques, but on their own still cannot provide the needed strong evidence. Instead, inspired by the evolution of related fields, it would be beneficial to collect and share many operational traces from real-world deployments into an open-access Trace Archive, and to provide a layer of fundamental understanding and knowledge by analyzing and sharing quantitative information such as the frequency of self-aware decisions and metrics about their impact. Large-scale experimental comparisons of self-aware and non-self-aware techniques, published as both technical material and open-access data, could provide a useful complement and further the acquisition of fundamental knowledge.

How can self-aware computing and communications enable or improve upon emerging applications for increasingly capable mobile devices — including smart

phones, wearable devices, and IoT devices? Today's cloud architecture, which relies on the use of a "dumb pipe" sitting between it and a smart edge including increasingly capable mobile devices, is likely to be a productive area for future research [53]. In particular, there is a fundamental limitation in how much power a wearable or hand-held device can consume before it becomes "too hot to wear" or "too hot to handle." Consequently, off-loading to an external computing platform, currently a cloud, becomes attractive. However, given the current Internet and Telecom system architectures, data from apps need to cross a dumb pipe before it reaches the cloud. This can be a severely limiting factor, for example when it comes to supporting emerging applications that need low-latency, high-bandwidth communications. Examples of such emerging applications include rich media (including video), unmanned vehicles, and gaming.

How can self-aware computing learn and maintain knowledge of itself if it is subject to frequent releases? The development of many popular Internet applications (e.g., Facebook, Netflix) is characterized by very short release cycles (e.g., daily). The term DevOps is often used to describe approaches and processes to align development and operation of software systems with the goal of continuous delivery of new versions with new or changed functionality. As a result, such software systems are often in transient phases where different versions of software run in parallel.

How can self-aware computing learn the characteristics of the workloads and update the allocation of their resources in a distributed manner with global coordination for efficient use of data center resources? Current cloud data centers span thousands of physical servers and host tens of thousands of virtual machines running different workloads. Managing such large systems while satisfying individual workload performance requirements and making efficient use of cluster resources is an open challenge. Current scalable approaches use either distributed scheduling (e.g., Mesos) or optimistic resource allocation using shared state (e.g., Omega [66]). However, such approaches cannot guarantee global scheduling optimality with respect to multiple (possibly competing) goals across all resources and workloads in the data center.

Can self-aware computing offer techniques for fully automated root cause analysis (RCA) that can be applied to a variety of management operations in the cloud? Over the past decade, the wide deployment of monitoring solutions in data centers and access to real time telemetry have advanced the art in diagnosis and troubleshooting. At the same time, *automated RCA* still remains an unsolved puzzle for operators. This is mainly due to the large number of hardware (compute, storage, networking) and software (OS, hypervisor, container, middleware, application) components that can potentially contribute to an observed failure or performance drop, and the complex interactions among them. Furthermore, typical statistical analysis and learning approaches discover correlations not causality between different metrics or events, and hence can only provide hints for the real root cause.

20.7 Conclusion

Cloud computing and its applications are already an important branch of ICT, with interesting benefits and challenges. Due to sheer scale, but also to the increasing sophistication of both their stakeholders and their infrastructure, clouds and their applications are increasingly relying on self-aware management techniques. In this chapter we have proposed a systematic framework to explain existing self-awareness approaches, and to facilitate the analysis of self-awareness in the future.

Our framework proposes a structured way to analyze the types of self-awareness approaches used in practice, in cloud computing and its applications. The framework focuses on the types of applications, of problems, and of approaches relevant to self-awareness. The framework also structures the discussion and the analysis of open challenges.

In this chapter, we have used the framework to analyze seven types of self-awareness approaches used in practice: feedback control based techniques, metric optimization with constraints, machine learning based techniques, portfolio scheduling, self-aware architecture reconfiguration, and stochastic performance models. We conduct a systematic survey of self-awareness techniques. We focus on eight types of applications, among which most are already established, whereas applications such as online gaming, partial processing, and cyber-physical applications are still emerging. We analyze ten types of traditional and novel problems. Novel, we focus on problems that have developed beyond their traditional scope or even emerged altogether in the space of cloud computing, such as autoscaling, energy-proportionality, performance isolation, and intrusion detection and prevention. We also identify four open challenges for self-awareness in cloud computing and its applications.

The future of this work is in facilitating, for the authors and for the self-awareness community at large, work addressing the open challenges. We also hope the framework will be used for analyzing other self-awareness approaches, new problems, and new applications.

Acknowledgements This work is partially supported by the Dutch STW/NWO Veni personal grant @large and Vidi personal grant MagnaData, by the Dutch national program COMMIT and COMMISSIONER sub-project, by the Dutch KIEM project KIESA, by a generous ERO gift from Oracle, by the European FP7 research project AMADEOS Grant Agreement 610535 on Systems of Systems, by the Swedish Research Council (VR) for the projects “Cloud Control” and “Power and temperature control for large-scale computing infrastructures”, and through the LCCC Linnaeus and ELLIIT Excellence Centers.

References

1. Yaser S. Abu-Mostafa, Malik Magdon-Ismail, and Hsuan-Tien Lin. *Learning From Data*. AMLBook, 2012.

2. Orna Agmon Ben-Yehuda, Assaf Schuster, Artyom Sharov, Mark Silberstein, and Alexandru Iosup. Expert: Pareto-efficient task replication on grids and a cloud. In *IPDPS*, 2012.
3. Christoph Albrecht, Arif Merchant, Murray Stokely, Muhammad Waliji, François Labelle, et al. Janus: Optimal flash provisioning for cloud storage workloads. In *USENIX ATC*, 2013.
4. Luiz Andre Barroso and Urs Hölzle. *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*. Morgan & Claypool, 2009.
5. Jean Arnaud and Sara Bouchenak. *Performance and Dependability in Service Computing*, chapter Performance, Availability and Cost of Self-Adaptive Internet Services. IGI, 2011.
6. Aniruddha Basak, Irina Brinster, and Ole J. Mengshoel. MapReduce for Bayesian network parameter learning using the EM algorithm. In *Proc. of Big Learning: Algorithms, Systems and Tools*, 2012.
7. Raul Castro Fernandez, Matteo Migliavacca, Evangelia Kalyvianaki, and Peter Pietzuch. Integrating scale out and fault tolerance in stream processing using operator state management. In *SIGMOD*, 2013.
8. Arthur Choi, Adnan Darwiche, Lu Zheng, and Ole J. Mengshoel. A tutorial on Bayesian networks for system health management. In A. Srivastava and J. Han, editors, *Data Mining in Systems Health Management: Detection, Diagnostics, and Prognostics*. Chapman and Hall/CRC Press, 2011.
9. Ira Cohen, Moises Goldszmidt, Terence Kelly, Julie Simons, and Jeff Chase. Correlating instrumentation data to system states: A building block for automated diagnosis and control. In *OSDI*, 2004.
10. Louis Columbus. Roundup of cloud computing forecasts and market estimates, 2015. Forbes Tech Report, 2015.
11. Kefeng Deng, Junqiang Song, Kaijun Ren, and Alexandru Iosup. Exploring portfolio scheduling for long-term execution of scientific workloads in iaas clouds. In *SC*, 2013.
12. Yixin Diao, Joseph L. Hellerstein, Sujay Parekh, Rean Griffith, Gail E. Kaiser, and Dan Phung. A control theory foundation for self-managing computing systems. *IEEE J. on Selected Areas in Communications*, 23(12):2213–2222, 2006.
13. Jonas Dürango, Manfred Dellkrantz, Martina Maggio, Cristian Klein, Alessandro Vittorio Papadopoulos, et al. Control-theoretical load-balancing for cloud applications with brownout. In *CDC*, 2014.
14. E.N. Elnohazy et al. System resilience at extreme scale. White paper. Defense Advanced Research Project Agency (DARPA) report, 2009.
15. European Commission. Uptake of cloud in europe. Final Report. Digital Agenda for Europe report. Publications Office of the European Union, Luxembourg, 2014.
16. Xi Fang, Satyajayant Misra, Guoliang Xue, and Dejun Yang. Smart grid; the new and improved power grid: A survey. *IEEE Communications Surveys Tutorials*, 14(4):944–980, 2012.
17. Lipu Fei, Bogdan Ghit, Alexandru Iosup, and Dick H. J. Epema. KOALA-C: A task allocator for integrated multicloud and multicluster environments. In *CLUSTER*, 2014.
18. Raul Castro Fernandez, Matteo Migliavacca, Evangelia Kalyvianaki, and Peter Pietzuch. Making state explicit for imperative big data processing. In *USENIX ATC*, 2014.
19. Antonio Filieri, Carlo Ghezzi, Alberto Leva, and Martina Maggio. Self-adaptive software meets control theory: A preliminary approach supporting reliability requirements. In *ASE*, 2011.
20. Antonio Filieri, Henry Hoffmann, and Martina Maggio. Automated design of self-adaptive software with control-theoretical formal guarantees. In *ICSE*, 2014.
21. Andrew Forward and Timothy C. Lethbridge. A taxonomy of software types to facilitate search and evidence-based software engineering. In *Conference of the Centre for Advanced Studies on Collaborative Research*, page 14, 2008.
22. Matthieu Gallet, Nezih Yigitbasi, Bahman Javadi, Derrick Kondo, Alexandru Iosup, and Dick H. J. Epema. A model for space-correlated failures in large-scale distributed systems. In *Euro-Par*, 2010.
23. Anshul Gandhi, Mor Harchol-Balter, Ram Raghunathan, and Michael A. Kozuch. Autoscale: Dynamic, robust capacity management for multi-tier data centers. *ACM Trans. Comput. Syst.*, 30(4):14, 2012.

24. Shravan Gaonkar, Kimberly Keeton, Arif Merchant, and William H. Sanders. Designing dependable storage solutions for shared application environments. *IEEE Trans. Dependable Secur. Comput.*, 7(4):366–380, 2010.
25. Bogdan Ghit, Nezh Yigitbasi, Alexandru Iosup, and Dick H. J. Epema. Balanced resource allocations across multiple dynamic mapreduce clusters. In *SIGMETRICS*, 2014.
26. Ajay Gulati, Chethan Kumar, Irfan Ahmad, and Karan Kumar. BASIL: Automated IO load balancing across storage devices. In *FAST*, 2010.
27. Ajay Gulati, Ganesha Shanmuganathan, Anne Holler, Carl Waldspurger, Minwen Ji, and Xiayun Zhu. VMware Distributed Resource Management: Design, implementation, and lessons learned. *VMware Technical Journal*, 1(1), 2012.
28. Vincenzo Gulisano, Ricardo Jimenez-Peris, Marta Patino-Martinez, Claudio Soriente, and Patrick Valduriez. Streamcloud: An elastic and scalable data streaming system. *IEEE Trans. Parallel Distrib. Syst.*, 23(12):2351–2365, 2012.
29. Joseph L. Hellerstein, Yixin Diao, Sujay Parekh, and Dawn M. Tilbury. *Feedback Control of Computing Systems*. John Wiley & Sons, 2004.
30. Jin Heo and Tarek Abdelzaher. Adaptguard: Guarding adaptive systems from instability. In *ICAC*, pages 77–86, 2009.
31. Nikolaus Huber, Jürgen Walter, Manuel Bähr, and Samuel Kounev. Model-based Autonomous and Performance-aware System Adaptation in Heterogeneous Resource Environments: A Case Study. In *ICCAC*, 2015.
32. IDC. Worldwide and regional public it cloud services: 2013-2017 forecast. IDC Tech Report. [Online] Available: www.idc.com/getdoc.jsp?containerId=251730, 2013.
33. Alexandru Iosup and Dick H. J. Epema. Grid computing workloads. *IEEE Internet Computing*, 15(2):19–26, 2011.
34. Alexandru Iosup, Nezh Yigitbasi, and Dick H. J. Epema. On the performance variability of production cloud services. In *CCGrid*, 2011.
35. Gideon Juve, Ann L. Chervenak, Ewa Deelman, Shishir Bharathi, Gaurang Mehta, and Karan Vahi. Characterizing and profiling scientific workflows. *Future Generation Comp. Syst.*, 29(3):682–692, 2013.
36. Evangelia Kalyvianaki and Themistoklis Charalambous. A Min-Max framework for CPU resource provisioning in virtualized servers using H-infinity Filters. In *CDC*, 2010.
37. Evangelia Kalyvianaki, Themistoklis Charalambous, Marco Fiscato, and Peter Pietzuch. Overload Management in Data Stream Processing Systems with Latency Guarantees. In *Feedback Computing*, 2012.
38. Evangelia Kalyvianaki, Wolfram Wiesemann, Quang Hieu Vu, Daniel Kuhn, and Peter Pietzuch. Sqpr: Stream query planning with reuse. In *ICDE*, 2011.
39. Rini T. Kaushik, Tarek Abdelzaher, Ryota Egashira, and Klara Nahrstedt. Predictive data and energy management in greenhdfs. In *IGCC*, pages 1–9, 2011.
40. Kimberly Keeton, Dirk Beyer, Ernesto Brau, Arif Merchant, Cipriano Santos, and Alex Zhang. On the road to recovery: Restoring data after disasters. In *EuroSys*, 2006.
41. Kimberly Keeton, Terence Kelly, Arif Merchant, Cipriano A. Santos, Janet L. Wiener, Xiayun Zhu, and Dirk Beyer. Don't settle for less than the best: Use optimization to make decisions. In *HotOS*, 2007.
42. Kimberly Keeton and Arif Merchant. A framework for evaluating storage system dependability. In *DSN*, 2004.
43. Mohammad M. H. Khan, Hieu Khac Le, Hossein Ahmadi, Tarek F. Abdelzaher, and Jiawei Han. Troubleshooting interactive complexity bugs in wireless sensor networks using data mining techniques. *ACM Trans. Sen. Netw.*, 10(2):31:1–31:35, 2014.
44. Cristian Klein, Martina Maggio, Karl-Erik Årzén, and Francisco Hernández-Rodríguez. Brownout: Building more robust cloud applications. In *ICSE*, 2014.
45. Cristian Klein, Alessandro V. Papadopoulos, Manfred Dellkrantz, Jonas Durango, Martina Maggio, et al. Improving cloud service resilience using brownout-aware load-balancing. In *SDRS*, 2014.
46. Joseph A. Konstan and John Riedl. Recommended to you. *IEEE Spectrum*, 2012.

47. Rouven Krebs, Simon Spinner, Nadia Ahmed, and Samuel Kounev. Resource usage control in multi-tenant applications. In *CCGrid*, 2014.
48. William S. Levine. *The control handbook*. The electrical engineering handbook series. CRC Press New York, 1996.
49. Shen Li, Shiguang Wang, Fan Yang, Shaohan Hu, Fatemeh Saremi, and Tarek Abdelzaher. Proteus: Power proportional memory cache cluster in data centers. In *ICDCS*, pages 73–82, 2013.
50. Lei Lu, Xiaoyun Zhu, Rean Griffith, Pradeep Padala, Aashish Parikh, Parth Shar, and Evgenia Smirni. Application-driven dynamic vertical scaling of virtual machines in resource pools. In *NOMS*, 2014.
51. Luo Mai, Evangelia Kalyvianaki, and Paolo Costa. Exploiting time-malleability in cloud-based batch processing systems. In *LADIS*, 2013.
52. Ole J. Mengshoel, Mark Chavira, Keith Cascio, Scott Poll, Adnan Darwiche, et al. Probabilistic model-based diagnosis: An electrical power system case study. *IEEE Trans. on Systems, Man and Cybernetics, Part A: Systems and Humans*, 40(5):874–885, 2010.
53. Ole J. Mengshoel, Bob Iannucci, and Abe Ishihara. Mobile computing: Challenges and opportunities for autonomy and feedback. In *Feedback Computing '13*, 2013.
54. Arif Merchant, Mustafa Uysal, Pradeep Padala, Xiaoyun Zhu, Sharad Singhal, and Kang G. Shin. Maestro: quality-of-service in large disk arrays. In *ICAC*, 2011.
55. Alexandru-Corneliu Olteanu, Alexandru Iosup, and Nicolae Tapus. Towards a workload model for online social applications. In *ICPE*, 2013.
56. Alexandru-Corneliu Olteanu, Nicolae Tapus, and Alexandru Iosup. Extending the capabilities of mobile devices for online social applications through cloud offloading. In *CCGrid*, 2013.
57. Pradeep Padala, Anne Holler, Lei Lu, Xiaoyun Zhu, Aashish Parikh, and Madhuri Yechuri. Scaling of cloud applications using machine learning. *VMware Technical Journal*, 2014.
58. Pradeep Padala, Kai-Yuan Hou, Kang Shin, Xiaoyun Zhu, Mustafa Uysal, Zhijui Wang, Sharad Singhal, and Arif Merchant. Automated control of multiple virtualized resources. In *EuroSys*, 2009.
59. Ken Peffers, Tuure Tuunanen, Marcus A. Rothenberger, and Samir Chatterjee. A design science research methodology for information systems research. *J. of Management Information Systems*, 24(3):45–77, 2008.
60. Zhengping Qian, Yong He, Chunzhi Su, Zhuojie Wu, Hongyu Zhu, et al. Timestream: Reliable stream computation in the cloud. In *EuroSys*, 2013.
61. Ragnathan Rajkumar, Insup Lee, Lui Sha, and John Stankovic. Cyber-physical systems: The next computing revolution. In *DAC*, 2010.
62. Charles Reiss, Alexey Tumanov, Gregory R. Ganger, Randy H. Katz, and Michael A. Kozuch. Heterogeneity and dynamicity of clouds at scale: Google trace analysis. In *SOCC*, 2012.
63. Brian Ricks and Ole J. Mengshoel. Diagnosis for uncertain, dynamic, and hybrid domains using bayesian networks and arithmetic circuits. *International Journal on Approximate Reasoning*, 55(5):1207–1234, 2014.
64. Matthew A. Russell. *Mining the Social Web: Analyzing Data from Facebook, Twitter, LinkedIn, and Other Social Media Sites*. O'Reilly Media, Inc., 1st edition, 2011.
65. Johann Schumann, K. Y. Rozier, T. Reinbacher, Ole J. Mengshoel, T. Mbaya, and C. Ippolito. Real-time, on-board, hardware-supported sensor and software health management for unmanned aerial systems. In *Annual Conf. Prognostics and Health Mgmt. Soc.*, 2013.
66. Malte Schwarzkopf, Andy Konwinski, Michael Abd-El-Malek, and John Wilkes. Omega: Flexible, scalable schedulers for large compute clusters. In *EuroSys*, 2013.
67. Damián Serrano, Sara Bouchenak, Yousn Kouki, Thomas Ledoux, Jonathan Lejeune, et al. Towards QoS-oriented SLA guarantees for online cloud services. In *CCGrid*, 2013.
68. Damián Serrano, Sara Bouchenak, Yousri Kouki, Frederico Alvares de Oliveira Jr, Thomas Ledoux, et al. SLA guarantees for cloud services. *Future Generation Comp. Sys.*, 2015.
69. Siqi Shen, Kefeng Deng, Alexandru Iosup, and Dick H. J. Epema. Scheduling jobs in the cloud using on-demand and reserved instances. In *Euro-Par*, 2013.
70. Siqi Shen, Alexandru Iosup, Assaf Israel, Walfredo Cirne, Danny Raz, and Dick H. J. Epema. An availability-on-demand mechanism for datacenters. In *CCGRID*, 2015.

71. Siqi Shen, Vincent van Beek, and Alexandru Iosup. Statistical characterization of business-critical workloads hosted in cloud datacenters. In *CCGRID*, 2015.
72. Snir et al. Addressing failures in exascale computing. *IJHPCA*, 28(2):129–173, 2014.
73. Simon Spinner, Samuel Kounev, Xiaoyun Zhu, Lei Lu, Mustafa Uysal, Anne Holler, and Rean Griffith. Runtime vertical scaling of virtualized applications via online model estimation. In *SASO*, 2014.
74. Priya K. Sundararajan, Eugen Feller, Julien Forgeat, and Ole J. Mengshoel. A constrained genetic algorithm for rebalancing of services in cloud data centers. In *CLOUD*, 2015.
75. Vincent van Beek, Jesse Donkervliet, Tim Hegeman, Stefan Hugtenburg, and Alexandru Iosup. Mnemos: Self-expressive management of business-critical workloads in virtualized datacenters. *IEEE Computer*, 48(7):46–54, 2015.
76. Qiushi Wang and Katinka Wolter. Reducing task completion time in mobile offloading systems through online adaptive local restart. In *ICPE*, 2015.
77. Pengcheng Xiong, Calton Pu, Xiaoyun Zhu, and Rean Griffith. vPerfGuard: An automated model-driven framework for application performance diagnosis in consolidated cloud environments. In *ICPE*, 2013.
78. Yong Yang, Lu Su, Mohammad Khan, Michael Lemay, Tarek Abdelzaher, and Jiawei Han. Power-based diagnosis of node silence in remote high-end sensing systems. *ACM Trans. Sen. Netw.*, 11(2):33:1–33:33, 2014.
79. Nezhir Yigitbasi, Matthieu Gallet, Derrick Kondo, Alexandru Iosup, and Dick H. J. Epema. Analysis and modeling of time-correlated failures in large-scale distributed systems. In *GRID*, 2010.