;login:
OCTOBER 2014 VOL. 39, NO. 5
ELECTRONIC SUPPLEMENT

Conference Reports

ICAC '14: 11th International Conference on Autonomic Computing

June 18–20, 2014, Philadelphia, PA

Summarized by Chao Li, Ji Xue, and Feng Yan

Model-Driven Management and Self-Adaptation

Summarized by Feng Yan (fyan@cs.wm.edu)

Storage Workload Isolation via Tier Warming: How Models Can Help

Ji Xue and Feng Yan, College of William and Mary; Alma Riska, EMC Corporation; Evgenia Smirni, College of William and Mary

Ji introduced his team's workload prediction model and scheduling method for the purpose of isolating the performance of user traffic from system background work. He first introduced the tiered storage system and then explained that scheduling background work may flush out the user working set in the fast tier and therefore impact the performance of user traffic. Then he showed how their Markovian-based model can predict the high intensity of future-user workload and proactively warm up the fast-tier with the user working set to prevent performance disaster. Finally, he showed performance comparison results between their proposed method and the state-of-the-art industry solution, the feedback method.

Christopher Stewart (Ohio State University) asked about how the model works with more classes. Ji explained the model can deal with more classes by extending the number of classes in the Markovian model. Michael Crouse (Harvard University) asked how their method compares to the static method, for example, scheduling background work at a pre-defined time. Ji explained they found the feedback method used in the paper for comparison with their method actually outperforms the static method, and that is why the feedback method was chosen as a baseline comparison. In addition, the feedback method is also the state-of-the-art method used in the industry.

Model-driven Elasticity and DoS Attack Mitigation in Cloud Environments

Cornel Barna and Mark Shtern, York University; Michael Smit, Dalhousie University; Hamoun Ghanbari and Marin Litoiu, York University

Mark Shtern introduced how their model-driven approach can make use of the assessment of the business values of the workload to scale Web applications and mitigate DoS attacks. The challenge is to find a proper metric to detect the scale for normal uses or attacks. He demonstrated that the ratio of benefit function to cost function is an effective approach, and used OPERA to illustrate how to do short-term predictions. He also explained in detail about the decision engine, e.g., how to decide to grow or shrink, how to decide if the traffic is normal or abnormal, and if abnormal, how to redirect the traffic, etc. Finally, he presented the experimental results on EC2 to verify the correctness and efficiency of their approach.

Someone asked why use delay instead of, for example, workload injection to do calibration in the model. Mark Shtern explained injecting workload is not predictable: One cloud may result in one behavior while another cloud may produce different results. But using the delay can help accurately calculate the demands. Feng (College of William and Mary) asked how they can combine so many different metrics together to detect the attack. Mark Shtern explained that they first convert each metric to money and then combine them.

Integrating Adaptation Mechanisms Using Control Theory Centric Architecture Models: A Case Study

Filip Křikava, University of Lille 1 and Inria; Philippe Collet, Université Nice Sophia Antipolis; Romain Rouvoy, University of Lille 1 and Inria

Filip Křikava showed how to use control theory-centric architecture models to integrate the classical feedback controllers for self-adaptive properties and improvements. He first introduced the use of control theory-centric architecture models and then discussed their requirements such as generality, visibility, and composition. He introduced feedback control definition language and used Znn.com as a case study. For the case study, the goal was to maintain server load around some pre-set value. He discussed the idea that service time equals fixed overhead plus data-size dependent overhead. The prerequisite is preprocessed content, e.g., different quality and size. Finally, he demonstrated the assessment of their method.

Xiaoyun Zhu (VMware) asked about how much knowledge one needs to implement this approach. Filip Křikava explained that one needs domain-specific knowledge and must be familiar with the control theories that are used.

Cloud Resource Management

Summarized by Ji Xue (jxue02@email.wm.com)

ShuttleDB: Database-Aware Elasticity in the Cloud

Sean Barker, University of Massachusetts Amherst; Yun Chi, Square Inc.; Hakan Hacigumus, NEC Laboratories America; Prashant Shenoy and Emmanuel Cecchet, University of Massachusetts Amherst

Prashant first talked about how the growing popularity of database-as-a-service clouds presents the challenge of how to choose the best elasticity mechanism to allow database clouds to meet requirements. Prashant proposed ShuttleDB, the authors' holistic approach that allows for flexible and automated elasticity of database tenants in the cloud. He illustrated how ShuttleDB works and presented experimental evaluation results.

Someone asked whether the method migrates the whole VM. Prashant answered no, that only the database migrates. Someone else asked whether the migration part includes memory. Prashant said yes, but other things are done as well.

Matrix: Achieving Predictable Virtual Machine Performance in the Clouds

Ron C. Chiang, The George Washington University; Jinho Hwang, IBM T. J. Watson Research Center; H. Howie Huang and Timothy Wood, The George Washington University

Ron first exhibited several problems to the audience, whether a bought VM performs the same as the local machine, and how to get to know the best tradeoff between cost and performance. Based on these questions, Ron proposed Matrix, a performance and resource management system. Matrix makes use of clustering models, then estimates the performance of the new workloads in the virtualized environment in order to choose an appropriate VM type. Finally, the authors talked about the evaluation on EC2 of the effectiveness of Matrix.

Someone asked whether time of day affects the model, and Ron answered that the model experiences an average of 30 runs, including 10 nights, 10 weekdays, and 10 weekends. Another person asked how they had modeled the interference. Ron said that for the background noise, via collecting the workload signature, the system can identify under what specific background environment to do the work modeling.

Adaptive, Model-driven Autoscaling for Cloud Application

Anshul Gandhi, Parijat Dube, Alexei Karve, Andrzej Kochut, and Li Zhang, IBM Research

Anshul raised the challenge of how to set values for users' applications to enable dynamic scaling of the applications, and how to scale an unobservable cloud application to provide high user performance. Anshul then introduced their method, which makes use of some known parameters to guess a black-box value in order to proactively scale the infrastructure to meet user-specified performance requirements.

Someone asked whether the latency in the model training is only for the CPU or also for the network. Anshul replied that it is currently only for the CPU. Someone asked whether the training model always converges to the right-one model, and Anshul answered yes, even though it cannot be proved so far, but due to the correlation between parameters, in the tested cases, their method always converges in a short time.

Exploring Graph Analytics for Cloud Troubleshooting

Chengwei Wang, Karsten Schwan, Brian Laub, Mukil Kesavan, and Ada Gavrilovska, Georgia Institute of Technology

Chengwei first discussed the challenge of troubleshooting in larger-scale datacenters. Chengwei then presented VFocus, a platform that uses streaming graph analytics to narrow down the search space. Finally, Chengwei proposed two cases to show that VFocus guidance operations can troubleshoot VM migration failures.

Someone asked what Chengwei meant by "migration failure." Chengwei replied that it refers to migrations that didn't complete.

ICAC Keynote Address II

Summarized by Chao Li (chaol@ufl.edu)

The Enterprise and Big Data Systems: Yesterday, Today, and Tomorrow

Lucy Cherkasova, HP Labs

Lucy Cherkasova started by introducing today's data explosion challenge. Dr. Cherkasova showed us a very interesting figure from Gartner called the "Hype Cycle for Emerging Technologies, 2013." The figure showed that the concept of Big Data is very close to the peak of the curve, which means it holds a great deal of public interest.

The MapReduce framework can be hard to manage. A simple setup can involve tuning more than 100 system parameters. Dr. Cherkasova pointed out that system designers must carefully review three key design considerations: the order of jobs, the amount of resources that have been used, and the way to allocate spare resources. She demonstrated the importance of job execution order with a set of MapReduce jobs that have no dependencies. Her slides showed that the order in which these jobs were executed can greatly affect the overall job turnaround time and the resource utilization of the server cluster. She said that we must carefully mange the spare system resources. If we used up all the resources, we would not have enough computing capacity for handing any new job requests.

To improve resource allocation effectiveness, Dr. Cherkasova introduced a resource allocation scheme called ARIA (Automatic Resource Inference and Allocation) for MapReduce environments. She also talked about a MapReduce simulation framework called SimMR, which reproduces the original job processing with satisfactory accuracy. In addition, she mentioned an interesting resource-sharing framework called YARN. This framework looked like a virtualized system and it allowed multiple different data-processing tasks to share the computing infrastructure.

Finally, Dr. Cherkasova talked about emerging scale-out workloads. Conventional scale-up system design methodology could greatly limit the performance that future cloud workloads could potentially achieve. She particularly mentioned a paper that appeared on ASLPOS '12: "Clearing the Clouds—A Study of Emerging Scale-out Workloads on Modern Hardware." She said "this paper is fantastic" because it demonstrated very well the needs for scale-out system design. Moreover, she introduced an ongoing project in HP called "The Machine." The project uses a high-throughput rack of systems with small form factors. This machine designed for future scale-out workloads could have over 1 PB installed memory.

Ron C. Chiang (George Washington) asked whether a probability distribution function (PDF) was used in the simulation to determine the MapReduce job execution behaviors. Dr. Cherkasova said that they did not use a probability distribution function. Ron also pointed out that in one of the slides there was a frac-

tion of execution phase (between the map and reduce execution phases) that seemed to be workload-independent. Dr. Cherkasova explained that this execution phase consisted of a number of MapReduce system calls, and therefore the runtime behavior looked workload-independent.