acm's women in computing
acm-w
ACM-W CIS Newsletter:
Celebrating, Informing & Supporting Women in Computing

in this issue

Research Spotlights

Verification of Safe Aircraft Separation

By Sarah Loos, Carnegie Mellon University, Pennsylvania, USA
sloos@cs.cmu.edu

As airspace becomes ever more crowded, air traffic management must reduce both space and time between aircraft to increase throughput, making on-board collision avoidance systems ever more important. These systems are impossible to analyze completely via testing or simulation because of the complexities that arise from non-linear flight trajectories - no amount of simulation could capture all possible aircraft behavior.

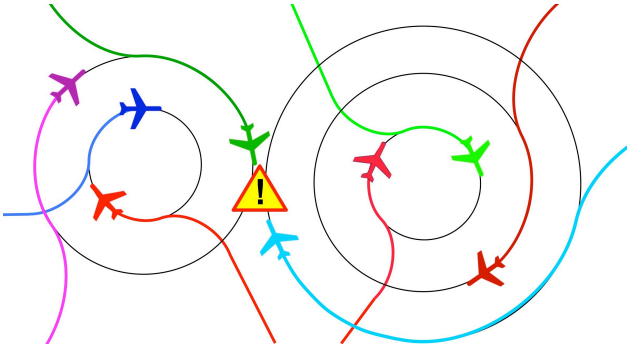

Figure 1: Two safe collision avoidance maneuvers, when executed in close proximity, may actually cause a new, unavoidable collision.

With thousands of aircraft in flight at any given time, it is difficult to predict whether executing a local collision avoidance maneuver may cause unintended behavior elsewhere, thereby doing more harm than good. For example, simple roundabout maneuvers, which are safe for a few aircraft, may be a terrible idea for crowded airspace. As illustrated in Figure 1, two maneuvers executed in close proximity may cause a new, unavoidable collision.

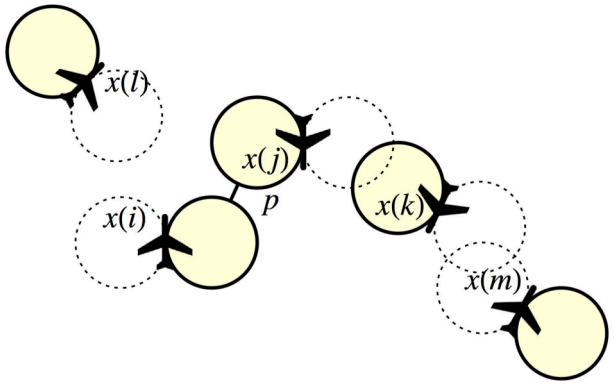

Figure 2 The collision avoidance scheme pictured here utilizes buffer discs associated with each aircraft. Safe separation is probably maintained for an arbitrary number of aircraft using this protocol [2].

This unexpected effect, called an *emergent* behavior, is one of several challenges that Sarah Loos, a Ph.D. student, works on with her advisor, André Platzer at Carnegie Mellon University. They use formal verification methods for cyber-physical systems to ensure the safety of these complex systems. They employ their theorem prover, KeYmaeraD [1], which soundly handles the continuous and non-linear flight trajectories and the infinite behaviors that can result from an arbitrary number of planes entering and exiting maneuvers. Using KeYmaeraD, they proved safe separation for an arbitrary number of aircraft flying under the disc-based collision avoidance scheme pictured in Figure 2 [2], a level of verification that has yet to be accomplished for any other distributed and flyable collision avoidance protocol.

About the Author

Sarah Loos is a Ph.D. student in the computer science department at Carnegie Mellon University. Her research interests include logical analysis and formal verification of distributed hybrid systems. She is a DOE computational science graduate fellow and an NSF graduate research fellow.

References

[1] D. W. Renshaw, S. M. Loos, and A. Platzer, "Distributed Theorem Proving for Distributed Hybrid Systems," in ICFEM, 2011, vol. 6991, pp. 356–371.

[2] S. M. Loos, D. Renshaw, and A. Platzer, "Formal verification of distributed aircraft controllers," in Proceedings of the 16th international conference on Hybrid systems: computation and control, New York, NY, USA, 2013, pp. 125–130.

Fault Tolerant Scheduling & Load Balancing for Cloud Computing – A Survey

by Dr. Jasma Balaangameshwara, Atria Institute of Technology, Bangalore, India
jasma2002@gmail.com

Cloud Computing is a large-scale distributed computing technology, in which a collection of dynamically-scalable and virtualized computing power, storage and services are delivered to customers on demand over the internet.

The focus of the research is on the existing approaches to scheduling, load balancing and fault-tolerance problems of Cloud Computing. Although scheduling, load balancing, and fault tolerance are active research areas in Cloud Computing, these areas have largely been and continue to be developed independent of one another - each focusing on different aspects of computing.

Hence, a survey is being conducted, to show that robust applications can provide efficient results by collectively considering these areas in the design phase. To this end, an introduction to the motivation, scheduling, load balancing and fault tolerance concepts of cloud computing are being discussed along with the works that have provided significant contributions to each of these areas since its inception until 2014. Their advantages, disadvantages and the suitability are being analyzed for usage in a dynamic environment.

This is ongoing research and from the survey conducted so far, it can be concluded that, while important advancements have been made in each of these areas individually, high performance approaches that cumulatively consider these areas

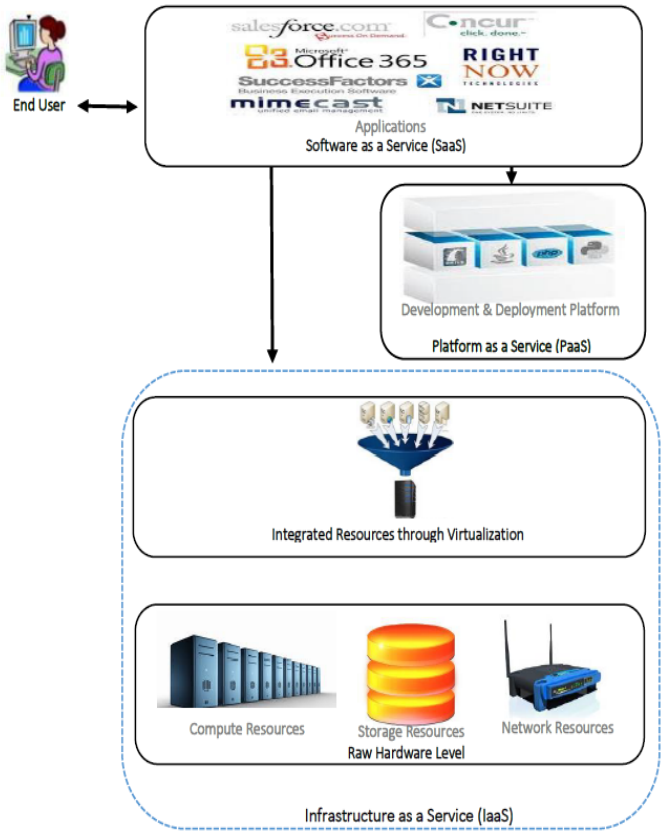

Figure 1: Cloud Computing Architecture

still remain to be explored. The survey concludes with ideas about what the community should be considering for future research pursuits.

About the Author

Dr. Jasma Balasangameshwara is currently working as an Associate Professor in the department of Computer Science and Engineering, Atria Institute of Technology, Bangalore, India. Her research interests include Distributed Computing, Fault Tolerant Scheduling and Data Mining. She is also providing her services as an Information Technology consultant to ANVITA, Bangalore, India, an organization with a mission to develop innovative solutions to diagnose and prevent the spread of Tuberculosis.

References

[1] Foster, Ian, Yong Zhao, Ioan Raicu, and Shiyong Lu. "Cloud computing and grid computing 360-degree compared." In *Grid Computing Environments Workshop, 2008. GCE'08*, pp. 1-10. IEEE, 2008.

Leveraging Thermal Power in Data Processing Frameworks

By Madhavi Vaidya, University of Mumbai, Mumbai, India

vamadhavi03@vesasc.org

A large public power company in India combined sensor data from the smart grid, where sensors are spread across multiple locations in the public power grid for capturing and analyzing data, to generate power by monitoring voltage, current and frequency. Monitoring the health of the entire grid requires the capture and analysis of data from every generator in the grid. Expected load on the grid is calculated one day in advance and by considering load pattern (such as public holidays) of industrial zones in India.

Based on my research, I would recommend that the India power company use the Hadoop cluster to capture and store the sensor data in the network. HDFS (Hadoop Distributed File System) is file system component which is patterned after the UNIX file system. It consists of a single Master Node and many Worker Nodes. The master, called as the Job Tracker, is responsible for accepting jobs from clients, dividing those jobs into tasks, and assigning those tasks to be executed by worker nodes. Each worker node runs a Task Tracker process that manages the execution of tasks currently assigned to that node.

Master nodes can be maintained on the generator in the grid, while the Worker Nodes can control many sensors. The master node would be responsible for connecting the generators to sensors. The MapReduce framework can be used to process the data generated from the sensors to execute forensic analysis. Hadoop can store data from the sensors to keep long-term historical data around for forensic analysis. As a result, the power company can see and react to, long-term trends and problems in the grid.

About the Author

Madhavi Vaidya is an Assistant Professor in Mumbai, India. Her research area is in the Hadoop Architecture. Professor Vaidya is a teacher at heart - she teaches in the areas of databases and operating systems. She published 12 research papers in National and International conferences. Her hobbies are reading, listening to music, meeting and communicating with people. Connect with her via email at vamadhavi@gmail.com