

Travel Report from University of Virginia

Dan Henriksson
Lund Institute of Technology

October 13, 2003

1. Introduction

This report will describe my research visit at the Department of Computer Science at the University of Virginia. My stay extended over three months from June to September 2003. During my stay I worked with Tarek Abdelzaher on research involving control-theoretic methods for QoS and real-time scheduling. The trip was partially funded by ARTES travel support.

2. University of Virginia and the Surroundings

University of Virginia (UVa) is located in Charlottesville, about 200 kilometers southwest of Washington D.C. The city of Charlottesville has a population of around 45000, and is a very nice and friendly city at the base of the Blue Ridge Mountains. The surroundings are very beautiful, with lots of hiking areas. The weather, however, was a bit too tropical and unpredictable for my liking.

UVa has about 20000 students and is ranked number one (tied with UCB) in the latest ranking of public universities in USA. The university was founded in 1819 by Thomas Jefferson, the 3rd president of the United States and author of the declaration of American independence. Monticello, the home of Thomas Jefferson, is located just a few miles from Charlottesville.

Virginia as a state is very interesting from many aspects, mainly because of its history. Virginia was the largest of the original British colonies and the state in which most of the battles during the Civil War were fought. Virginia also is the birthplace of many influential Americans, including George Washington and Thomas Jefferson.

3. Department of Computer Science – Real-time Systems Group

The Department of Computer Science at UVa has about 120 graduate students and 35 senior faculty working in most areas of computer science, including graphics, AI, architecture, networking, and real-time systems. The department

is headed by Professor John F. Stankovic.

I was working within the real-time systems group together with Tarek Abdelzaher. The real-time systems group has about 30 graduate students and are doing research in areas such as wireless networks, quality-of-service in software services, real-time scheduling, and real-time databases.

The largest research project is a DARPA-funded project on wireless sensor networks. These types of physically distributed nodes, also referred to as "smart dust", can be used for tracking and surveillance in large areas. The research deals with development of new specially tailored scheduling mechanisms and network protocols. Many practical issues such as power control, stealthiness, and perseverance need to be dealt with. The use of feedback control for resource management in these networks is also of great interest.

Another project is on the use of control-theoretic methods to provide QoS guarantees in modern software services such as, e.g., web servers. The research involves both theoretic issues such as design of controllers and mapping of QoS specifications into feedback control problems, and implementation issues such as middle-ware development and kernel support.

Tarek Abdelzaher also works on schedulability bounds for aperiodic tasks, and has recently come up with many interesting results in the area. The results can be applied in, e.g., web server applications to provide guidance for admission control schemes to avoid server overload.

The research I performed during my stay in a sense connected the QoS control for web servers with the schedulability analysis for aperiodic tasks. I gave a well-received seminar on flexible real-time scheduling of control tasks, and presented the simulator TrueTime, which is a result of our research in Lund. The simulator can be used to simulate real-time control systems, and I used it during my stay to simulate a web server system and to evaluate feedback scheduling strategies for aperiodic tasks.

Feedback scheduling strategies based on schedulability bounds for aperiodic tasks were used in a power-management application. The processing speeds of the aperiodically arriving web requests were scaled to minimize power consumption while still meeting timing requirements. The timing requirements were specified as deadlines for the incoming requests. This corresponds to different QoS specification for different classes of requests. The timing requirements were met by controlling the processing times around the proper schedulability bounds.

4. Conclusions

This has been a very rewarding research visit both from a scientific and social point of view. It was nice to spend these months in an up-front research environment, and to be able to find new friends in the area of integrated control and scheduling. Hopefully, the collaboration between our departments will continue in the future.