

Towards Self-Managing Large Scale Information Systems

Marek Rusinkiewicz

Department of Computer Science
The University of Houston
marek@uh.edu

Abstract. Information systems and the IT infrastructure that is used to support them are becoming too complex to design, maintain, and evolve using traditional approaches. To realize the visions of ubiquitous and adaptive computing, the responsibility for the management of the complexity must be shifted from the designers and operators of the systems, to the systems themselves.

Current work in this area goes in the direction of creating systems that are capable of reacting without human intervention to the new requirements and to the changes in the operating environments. In this talk we will review several examples of technologies that bring us closer to achieving this goal. First, we will review the model of policy-based network management, using as examples automatic network configuration and policy-based network security enforcement. We will show how security objectives, formulated using high-level specifications can be automatically translated into appropriate network settings. The network can be then continuously monitored to detect violations, and appropriate adjustments can be made, as needed to enforce the security policies.

In the second part of the talk, we will discuss information exploitation applications developed as communities of collaborating software agents. We will review agent communication standards, including communication languages and agent conversation templates. We will also discuss various methods of creating agent communities, both facilitated, based on brokering/matchmaking, and self-organizing, based on market concepts. We will illustrate this with examples drawn from the InfoSleuth agent system.