



## Creating a Trustworthy Active Web

Ken Birman

Distributed computing techniques are about to experience a disruptive transformation. Up to the present, distributed computing has been an esoteric technology used mostly inside major technology products from companies like IBM, Oracle and Microsoft, or used by researchers building components of the data centers at Amazon.com and Google. Overnight, distributed systems are posed to become the core technology that will pervade a new kind of Internet that we're calling "the active web".

Cornell's new Live Objects technology enables this active web. The basic idea is that instead of today's rarely-changing static content, an active web would contain active objects that are extremely dynamic: the content may evolve at network speeds, and it will be possible to interact with these objects. Students familiar with the online role-playing game Second Life can think of the Active Web as "Second Life meets the Internet" -- except that unlike Second Life, which needs to be hosted on a data center, the Active Web can run in a purely peer-to-peer configuration, with the players generating and consuming the content, and where servers play minor support roles. If you haven't played Second Life, think about the scene in Minority Report where Tom Cruise interacts with the fancy display and pulls in data from all sorts of sensors and databases -- he was using the active web!

The active web will be a completely transformative environment. We can use it to create a new kind of Second Life game, but we can also use it to create new forms of business tools, collaboration platforms, military systems, and even electronic health systems.

During this tutorial, we'll study the architecture of the active web and will learn about some of the distributed computing technologies that will play critical roles in enabling it. A prototype of this technology exists; we'll see a demo during the tutorial, and students interested in trying it out can download it and extend it. The real emphasis of the class will be on foundational questions involving trustworthiness of the active web. Once we understand the basic ideas, what do we need to do to ensure that it will be secure, scalable, reliable, privacy preserving, self-managing and adequately flexible to accommodate new technologies and new kinds of active content?