

The Semantic Web in Large Scale Distributed Systems

David De Roure

I see a future with a digital world behind-the-scenes — a world with masses of computation and data — which intersects with the physical world through a large number and variety of pervasive computing devices. Today this is perhaps best exemplified through applications of the Grid, which intersects with the physical world through many instruments, sensors and devices (see *Where the Grid meets the Physical World*). Where does the Semantic Web fit this vision?

The Semantic Grid experience

Grid computing is (quoting Foster) ‘distinguished from conventional distributed computing by its focus on large-scale resource sharing, innovative applications, and, in some cases, high-performance orientation.’ The Grid vision — and the e-Science vision — promises a high degree of easy-to-use and seamless automation, enabling flexible collaborations and computations on a global scale. I believe that techniques for machine-processable knowledge are essential to achieve this vision — where the knowledge is not just in the application but also describes the components of the system and its behaviour. We call this the Semantic Grid and since 2001 we have been applying Semantic Web technologies to various aspects of the Grid problem. In a project with chemists, for example, we have over 80 million RDF triples online — a significant contribution to the Semantic Web (see *What the Grid can do for the Semantic Web*). The knowledge is collected from instruments in the smart laboratory and chemists see their experimental results on their phones and PDAs in the bar — a case study in future computing.

Things on my mind

So we see increasing ease of SW rollout at "datagrid" level, but the description and automatic composition of resources is not a solved problem and it is attracting considerable effort, for example in e-Science workflow and in Semantic Web Services. We are drawing on the techniques of multi-agent systems, particularly with respect to negotiation (see *Agents and the Grid*). In fact we pursue a kind of reflection — you could say we are trying to use the Semantic Web to deliver the infrastructure of the Semantic Web.

But is this really the Semantic Web *per se*? The real magic of the Semantic Web comes from the fact that there is one Semantic Web and from the network effects of sharing. So currently I'm trying to figure out how to achieve these benefits when working down ‘underneath the applications’. And is the global sharing magic precluded in the pervasive computing context when there is limited availability of information?

In summary, I anticipate a decentralised and self-organising future infrastructure, where the self-organisation is enabled by the explicit machine-processable descriptions using Semantic Web technologies. The Semantic Web is really good at joining things up. The decentralised self-organising vision involves letting go of things, and I'm hoping that the Semantic Web will help us bring them back together!

David De Roure, Salzburg, September 2005