Universidade do Minho
Escola de Engenharia

Semana da Escola de Engenharia
October 24 - 27, 2011
IMPROVING THE SCALABILITY AND DEPENDABILITY OF INFORMATION DISSEMINATION IN WEB SERVICES INTEGRATION

Filipe Campos and José Pereira
Department of Informatics
University of Minho
Campus de Gualtar, 4710-057 Braga, Portugal
E-mail: {fcampos, jop}@di.uminho.pt

ABSTRACT

Information dissemination among a large number of participants is a key feature of many service oriented architectures. More recently, there has been a growing interest in Web Services to address interoperability, composability, and long term maintainability of systems of connected devices in a variety of environments, ranging from manufacturing equipment to home automation. Unfortunately, information dissemination aimed at typical enterprise environments assumes a relatively small number of heavyweight participants, for instance, regarding the state that each participant is expected to hold and the precise knowledge of the participants' list itself. These techniques fall short in face of systems composed of very large numbers of lightweight devices with stringent resource constraints. These shortcomings can be circumvented with the introduction of a service oriented architecture for information dissemination based on existing standards and distributed gossiping, which scales to a large number of loosely coupled participants while providing strong atomicity guarantees.

Although service-orientation is for some time an integral part of enterprise computing, there is a trend of enabling such principles in systems of connected devices in a variety of environments, ranging from industrial manufacturing equipment to home automation. The first enabler for this, is the arrival to market of cheaper and more powerful devices. The second enabler is the observation that service-orientation provides the best solution for interoperability, composability, and long term maintainability challenges in these scenarios. In fact, the current trend in connected devices is expected to accelerate as the Internet of Things becomes a reality.

In this context, being able to expose the joint capabilities of large sets of devices as logical services is enticing, as much as coordinating and composing services in business processes has been important in enterprise service-oriented computing. Again, one is faced with information dissemination as the key technological enabler. Besides the obvious issue of efficiently routing information to all destinations, scalable and dependable dissemination raises a number of traditionally hard problems. For instance, the effort involved in keeping track of destinations themselves in a dynamic environment, namely, how the effort is spread across different components (Stoica, I. et al., 2001). Or how reliable delivery is ensured, namely, how to manage acknowledgement and retransmission without depending on any single system component (Birman, K., 1999). Or even, how to pace transmission in order to avoid overwhelming any particular destination (Piantoni, R. and Stancescu, C., 1997). Such problems should not burden information dissemination at an architectural level.

Two main approaches have been proposed to deal with the need for dependable and efficient information dissemination in service-oriented computing. The first approach is to embody the information dissemination logic in middleware, namely, by using an Enterprise Service Bus or by using JMS or XMPP transports for Web Services. This approach takes advantage of mature technology which solves the aforementioned challenges and is already deployed and well-known in many organizations. On the other hand, it introduces a dependency on a specific software stack and, often, on centralized messaging servers. Moreover, the required middleware is often not available for the entire range of devices that need to be supported, either because it assumes computing resources of typical enterprise systems or simply because the middleware vendor does not target the desired device platform. Finally, by providing messaging as a black box, it supports only a fixed range of information exchange patterns and limits the range of problems and environments that can be addressed. The second approach has been to provide

information dissemination at the service level, by means of specifications that can be combined to expose different message exchange patterns with various qualities of service. An example of this approach is the OASIS WS-Notification family of standards, which besides simple notification and subscription, supports topic based publish-subscribe and brokered dissemination for scalability. An alternative is the simpler WS-Eventing specification, which, although lacking explicit support for brokered dissemination, embodies a flexible filtering mechanism in the base specification, favoring lightweight implementations and the many-to-one dissemination scenario. It has therefore been the preferred choice for standards such as WS-Management or the Devices Profile for Web Services (DPWS). Both specifications can be combined with WS-ReliableMessaging for end-to-end acknowledged message delivery or with WS-AtomicTransactions for multi-party transactional atomicity guarantees.

Note however that both approaches emphasize on information dissemination from one source to several destinations and do it by means of setting up a centralized broker infrastructure. Dependability rests on the central server supporting replication in a cluster, while end-to-end reliability and atomicity depend on participants having sufficient memory for buffering and stable storage for a transactional log. In short, the assumption of a centralized heavy-weight infrastructure permeates most existing solutions for information dissemination in service-oriented computing. We address this need by introducing a service-oriented architecture for information dissemination based on existing standards and distributed gossiping.

Inspired by the form of gossip in social networks, and by the way viruses spread in a biological community, hence also being known as epidemic protocols, gossip protocols enable distributed message dissemination and acknowledgements' processing, in order to avoid network congestion and nodes to be flooded with acknowledgments (Kermarrec and van Steen 2007). These protocols are known to achieve probabilistic atomic delivery and reliable multicast and to consume few resources, which contributes to their inherent scalability. The usage of such protocols better adapts to scenarios where messages are sent from a single source to many recipients, or vice-versa. This mimics quite closely scenarios such as configuration and monitoring or alerting, respectively. Although one could easily use an existing gossip-based messaging middleware as a black-box transport protocol, this would still have many of the shortcomings of the first approach outlined above. Instead, in line with the second approach, we propose the use of service specifications that can be combined to architect a variety of gossip-style interactions, allowing any service-oriented architecture to leverage gossip for multiple information dissemination purposes, by using any of the four provided variants. In comparison to WS-PushGossip (Campos, F. and J. Pereira, 2008) and WS-Membership (W. Wogels and C. Ré, 2003), both WS-Coordination based, the proposed services are better suited for devices and their inherent limitations, as the implementation and use of such a resource exhausting protocol is not necessary or even feasible.

REFERENCES

Birman, K. 1999. “A review of experiences with reliable multicast”. *Software Practice and Experience,* 29(9).

Campos, F. and Pereira, J. 2008. “Gossip-based Service Coordination for Scalability and Resilience”. *Proc. of the 3rd Workshop on Middleware for Service Oriented Computing (MW4SOC'08).*

Kermarrec, A.-M. and van Steen. M. 2007. “Gossiping in distributed systems”. *SIGOPS Oper. Syst. Rev.* 41, Issue 5 (Oct), 2-7.

Piantoni, R. and Stancescu, C. 1997. “Implementing the swiss exchange trading system”. *Fault-Tolerant Computing. FTCS-27. Digest of Papers., Twenty-Seventh Annual International Symposium on*, pp. 309–313.

Stoica, I.; Morris, R.; Karger, D.; Kaashoek, M. 2001. “Chord: A scalable peer-to-peer lookup service for internet applications”. *Proc. of the 2001 Conference on Applications, technologies, architectures, and protocols for computer communications*.

Wogels, W. and Ré, C. 2003. “WS-Membership - Failure management in a web-services world.” *Intl. World Wide Web Conference (WWW).*

AUTHORS’ BIOGRAPHIES

FILIPE CAMPOS was born in Barcelos, Portugal and went to the University of Minho, where he studied systems and informatics engineering and obtained his degree in 2006. He worked for a year at Nokia Siemens Networks before starting his Ph.D. research in 2007 in the University of Minho. His e-mail address is : `fcampos@di.uminho.pt` and his Web-page can be found at `http://gsd.di.uminho.pt/members/fjoc`.