CHOReOS: Scaling Choreographies for the Internet of the Future

Hugues Vincent, Valérie Issarny, Nikolaos Georgantas, Emilio Francesquini,
Alfredo Goldman, Fabio Kon

CHOReOS: Scaling Choreographies
for the Internet of the Future

Hugues Vincent
Thales Group
hugues.vincent@thalesgroup.com

Valérie Issarny
Nikolaos Georgantas
ARLES Research Team,
INRIA Paris-Rocquencourt
valerie.issarny@inria.fr
nikolaos.georgantas@inria.fr

Emilio Francesquini
Alfredo Goldman
Fabio Kon
Department of Computer Science
University of São Paulo
{emilio,fabio.kon,gold}@ime.usp.br

http://www.choreos.eu

ABSTRACT

The Internet has been growing at a impressive rate in many aspects such as size, heterogeneity, and usage. This growth forces the continuous improvement of Internet infrastructure technologies. The Future Internet concept magnifies the required shift for Internet technologies, which shall allow supporting the continuously growing scale of the converging networking world together with new generations of services made available to and brought by the broad mass of end users. The CHOReOS project positions itself in this vision of the Future Internet, whilst focusing on the Future Internet of Services. This research project aims at assisting the engineering of software service compositions in this novel networking environment by devising a dynamic development process, and associated methods, tools and middleware, to sustain the composition of services in the form of large-scale choreographies for the Internet of the future.

Categories and Subject Descriptors

D.2.1 [**Software Architectures**]: Service-oriented Architecture

General Terms

Performance, Design, Reliability, Experimentation, Security, Human Factors, Standardization, Languages, Verification.

Keywords

Choreography, Internet of the Future, Web Services

1. INTRODUCTION

The growth in web services adoption during the last years is evident. The usage of this technology has grown to a point that it has become the *de facto* standard for the communication among high-level Internet systems. This popularization, however, highlighted some problems that were not readily apparent in the previous integration efforts as they scarcely achieved the scale of integration the web services systems now boast. One of the problems is related to service composition.

Software developers realized that hard-coding service composition logic into their software was not a good solution, as they would either rewrite routines to control transactions, parallel execution, and flow-control decisions repeatedly or just stick to a explicit programming language in an otherwise language-agnostic system integration. To address these composition problems a few web service composition standards were created until, finally, in 2003, WS-BPEL was consolidated as an open OASIS standard [1]. Currently, there are several WS-BPEL execution engines available as it enjoys a widespread adoption in the market. At the same time that web services orchestrations grew larger and more complex, they became a new source of concern. Their centralized approach to composition, although straightforward and simple, has scalability and single point of failure problems. The next natural step is then to compose the web services in a non-centralized distributed way. Web services compositions organized in this fashion, i.e., in a decentralized, distributed manner, with no single point of failure are called *Choreographies*.

Choreographies are intrinsically more resilient than, although not as easily manageable as, orchestrations. As non-centralized distributed systems they have higher fault tolerance, adaptability, configurability, and freedom to grow. The joint operation and execution of several web services composing a choreography is called "enactment". Even though some choreography standards (e.g., WSCI and WS-CDL) have been defined, to the best of our knowledge, none of them has been completely implemented, there are very few development tools available and there is still little research into the actual usage (deployment and enactment) of choreographies. Moreover as the current ad hoc choreographies get larger and more intricate, they can easily become unmanageable.

The CHOReOS project intends to tackle these issues implementing a middleware that will make possible the actual definition, deployment and enactment of large-scale choreographies, such as those that will be needed for the Internet of the Future [2,3], whilst sustaining adaptation and quality assurance.

2. RESEARCH CHALLENGES

The construction of the target middleware forces us to tackle some problems such as: (i) integration of Grid and Cloud middleware so that service execution may scale up to the required load, (ii) further development of current Distributed

Service Bus (DSB) technology to meet the large-scale challenge, (iii) further support for choreography deployment and execution, (iv) leveraging and evolving service-oriented middleware technology for pervasive networks to face the challenges of the Internet of Things, (v) and implementing a service-oriented middleware for large-scale decentralized choreographies based on the integration of the existing service-based technologies. In particular, the project will investigate the level of scalability that will be required by future Internet applications and the level of scalability that can be provided by alternate architectures. Our initial target is to study choreographies composed of tens to hundreds of services, involving hundreds to thousands of computing nodes and thousands to millions of users.

3. PROPOSED MIDDLEWARE

The execution of large-scale choreographies within the Future Internet heavily relies on adequate middleware support. We will base the implementation of our middleware on three middleware technologies: distributed service bus, pervasive middleware technology and grid and cloud computing technologies. Figure 1 depicts the proposed CHOReOS middleware architecture.

Concerning the Distributed Service Bus, we will rely on the PEtALS ESB-based middleware solution [4]. PEtALS is already further evolving towards a Distributed Service Bus (DSB) to scale to millions of services. The evolution of PEtALS into a DSB relies on a federated architecture based on a P2P overlay network. The desired outcome will then be a highly scalable service bus that allows choreographing heterogeneous services, thanks to the service bus principle and its evolution to cope with the features of the Future Internet.


Figure 1: CHOReOS Middleware Architecture

In the CHOReOS framework, the Internet of Things concept [5] is that of a highly dynamic and heterogeneous networking environment integrating an ultra-large number of devices, including many with limited resources. The CHOReOS middleware needs to meet these requirements accounting for the resource-constraints of devices together with the environment's high dynamics and heterogeneity. This will be done in a way that is compatible with the recently-approved OASIS standard [6] for Web services on resource-constrained devices. We want also to provide CHOReOS with the high-performance computing power available in Grid [7] and Cloud Computing [8] infrastructures. Thus, the computationally intensive processes that will be required to serve millions of users issuing thousands of simultaneous service requests to thousands of services will be able to be processed by Grid and Cloud services. The task will involve efforts in (1) Software Architecture and Engineering to implement the interaction protocols and choreography engines onto the specific context of Grid and Cloud middleware infrastructures and (2) investigation of the CHOReOS methods for creating, managing, and processing choreographies so that its computation can be delegated to a high performance computing engine.

4. CONCLUSION

Existing ESB, Grid/Cloud, and pervasive middleware technologies emerged independently to cope with different scalability issues. Nevertheless, the Future Internet calls for an integrated solution. To this end, CHOReOS (http://www.choreos.eu) will build upon the aforementioned baseline of the individual CHOReOS partners to refine individual solutions so that they meet the challenges of the Future Internet, and moreover to develop a unified middleware infrastructure that enables: (1) service provisioning for the ultra-large number of Future Internet users based on available Grid and Cloud technologies, (2) networking a large number of heterogeneous services via ESB-based middleware, and (3) networking services from the Internet of Things based on middleware for pervasive networks.

5. REFERENCES

[1] OASIS Web Services Business Process Execution Language (WSBPEL), http://www.oasis-open.org/committees/wsbpel, Retrieved 2010-10-07.

[2] ERCIM News. Special Theme: Future Internet Technology. Number 77. April 2009.

[3] DG INFSO Task Force on The Future Internet Content & External technical experts. Draft Report on the Task Force on Interdisciplinary Research Activities applicable to the Future Internet. Version 4.1 - 2009. http://forum.future-internet.eu

[4] PEtALS ESB, http://petals.ow2.org, Retrieved 2010-10-07

[5] Koosha Paridel et al., "Middleware for the Internet of Things, Design Goals and Challenges" *in CAMPUS'2010.*

[6] OASIS Devices Profile for Web Services (DPWS) Version 1.1, http://docs.oasis-open.org/ws-dd/dpws/1.1/os/wsdd-dpws-1.1-spec-os.pdf*, Retrieved 2010-10-07*

[7] Raphael Y. de Camargo, et al. "The Grid Architectural Pattern: Leveraging Distributed Processing Capabilities" *in*

PloPD5. Dragos Manolescu; James Noble; Markus Völter (editors). Addison Wesley, 2006, pages 337-56.

[8] Qi Zhang, Lu Cheng and Raouf Boutaba. "Cloud computing: state-of-the-art and research challenges" *in Journal of Internet Services and Applications*, 1(1):7-18, Springer London, 2010.