Linked Data-Fu: From Data Integration to System Interoperation

Andreas Harth
Karlsruhe Institute of Technology (KIT), Germany
Research Contribution

Abstract

We present an architecture for applications that require data integration and system interoperation in distributed, heterogeneous environments, for example virtual reality applications in the manufacturing industry. The interface to components is based on (semantic) web technologies. We specify the integration of data and the interaction between components in a rule-based language. We describe a data processor, called Linked Data-Fu (ldfu), that is able to continuously execute such rule-based specifications in small intervals, and present a prototype demonstrating the assembly of a product in virtual reality.

1 Introduction

Increasingly, applications need access to data and functionality from multiple sources and systems. Consider, for example, applications in the manufacturing industry. Companies want to test products and production processes in virtual reality before actually manufacturing the product or reconfiguring a manufacturing plant.

Such scenarios require the access to and interplay between different components: sensors (e.g., motion tracking), actuators (e.g., head-mounted displays), and external data (e.g., 3D scenes of the product, parts information). Further, such scenarios often include workflows, i.e., the description of behaviour of a user interacting with the product or a production worker assembling an artefact.

Current monolithic applications hard-wire access to components and hard-code the description of behaviour of the application as well as the behaviour of users and workers. Future systems provide flexible access to distributed components based on standardised interfaces and declarative descriptions of behaviour.

The poster will cover a description of such applications in the context of the i-VISION[1] and ARVIDA[2] projects. We introduce a generic architecture blueprint illustrating the various components involved in industrial virtual reality applications. We further sketch a rule language to express application logic.

2 Architecture Blueprint

In the following, we describe an architectural blueprint for interactive applications, including the encoding of application logic using rules. The various components provide an interface based on web technologies (URIs for identification, HTTP for create-read-update-delete (CRUD) operations). The components expose data and current internal state via resources identified using URIs. CRUD operations use RDF to encode state, but other formats (such as a binary format for point cloud data) are supported as well via content types. Figure 1 shows the overall setting.

We use the rule language to encode:

[1]Immersive Semantics-based Virtual Environments for the Design and Validation of Human-centred Aircraft Cockpits, a European FP7 project. Web site at `http://www.ivision-project.eu/`.

[2]Angewandte Referenzarchitektur fü virtuelle Dienste und Anwendungen (applied reference architecture for virtual services and applications), a German BMBF project. Web site at `http://www.arvida.de/`.

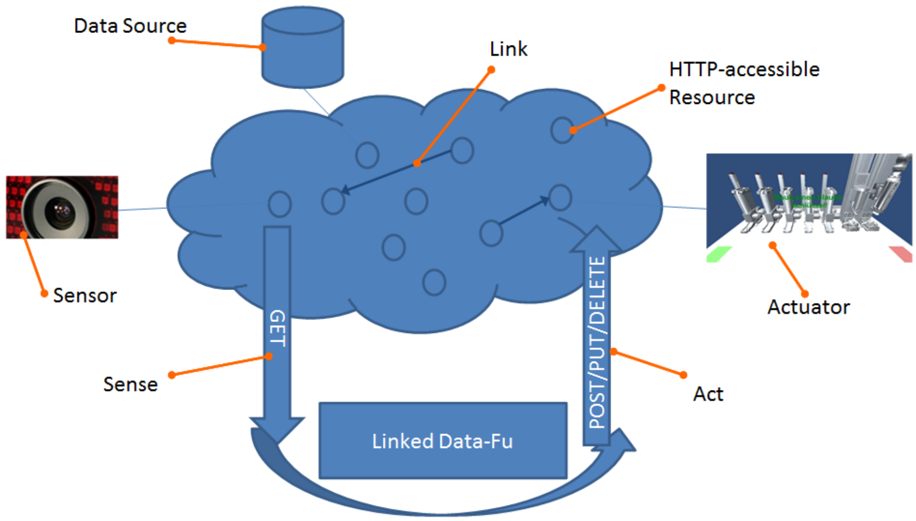

Figure 1: Components (sensors, actuators and data sources) provide access to data and functionality via a resource-based interface. Resources are identified via URIs, and link to each other. Linked Data-Fu coordinates the interaction between components as specified in a rule-based program.

- Deduction rules, that is, a specification of how to derive new triples based on existing triples.
- Request rules, that is, a specification of which requests to issue based on existing triples.

The rules operate on the current state of the world and the application. The current world state is gathered at short intervals (with update rates up to 60 Hertz).

The rule engine (called Linked Data-Fu, or ldfu for short) executes the programs at defined intervals. A basic program just reads the current state of the sensors and updates the states of the actuators accordingly. For example, a program could read the skeleton of a tracked person, and update the pose of a stick figure in virtual reality accordingly. For interactive applications we require a update frequency of at least 30 Hertz (30 program runs per second). More elaborate programs would also keep track of internal state, and include external sources in the application.

With a uniform CRUD interface across components, the same knowledge representation and rule language can be used to integrate data and interoperate between components.

3 Conclusion

We have presented an architecture blueprint for interactive applications, accessing sensors, actuators, and data from external sources using a uniform interface based on URIs and HTTP. The system uses RDF as knowledge representation language for encoding data and state of components. The rule-based Linked Data-Fu programs encode specifications of behaviour, based on data and state of components expressed in RDF. The proposed architecture seamlessly integrates data from very different components, from sensors with high update rates to web-accessible sources that are infrequently updated. The architecture also supports writes, that is, changing of data or effecting change using actuators. We have developed several prototypes implementing the proposed architecture, and continue to apply the architecture to projects involving the automobile and aerospace industries.

Curriculum Vitae

Dr. Andreas Harth is a post-doctoral researcher at Institute AIFB at the Karlsruhe Institute of Technology. His research interests are large-scale data interoperation on the Semantic Web, Linked Data, knowledge representation, computational logic and user interaction on web data. Andreas has published several dozen papers in these areas, and is author of a number of open source software systems. Recently, he has become interested in web-based architectures for cyber-physical systems (e.g., smart energy grids, surgery systems, mixed-reality systems).

Andreas was awarded his Ph.D. by the Digital Enterprise Research Institute (DERI) at the National University of Ireland, Galway. Andreas worked as intern at Fraunhofer Gesellschaft in Würzburg and at IBM's Silicon Valley Lab in San Jose, CA. His Diplom thesis was carried out in collaboration with Centro Politcnico Superior at Universidad de Zaragoza, Spain. He visited USC's Information Sciences Institute in Marina del Rey, CA as a research assistant.

Andreas has participated in numerous EU and national projects and was active in W3C working groups. In addition, he served as program committee member of numerous conferences and was one of the co-organisers of the Consuming Linked Data (COLD) workshop series and of the Semantic Web Challenge.