

Odysseus - An Extensible Research Platform for Streaming Applications

Marco Grawunder

University of Oldenburg, Department of Computing Science
Marco.Grawunder@uni-oldenburg.de,
<http://odysseus.uni-oldenburg.de>

Abstract. Stream processing has received much attention in academia and industry in the recent years. Especially, in the research area it is hard to develop and evaluate new concepts in this area because of the large overhead necessary to get a streaming system running. There is a bunch of streaming systems around. Probably, one of the most well know is Apache Storm. While Storm is very good in distributed processing of streaming data, Storm burdens the user even with simple filtering or correlation of stream. The user has to write a Java program that represents the streaming application.

Odysseus is our streaming engine developed since 2007 at the University of Oldenburg. The main goal was the development of a software platform to evaluate stream processing methods. From the beginning an important aspect was easy extensibility and maintenance. Odysseus is written in Java and uses a plugin mechanism based on OSGi. Odysseus provides about 3 million lines of code and about 400 plugins.

In this talk I will present the conceptual operator based model behind Odysseus that allows for many cases a deterministic processing. A very important part here is the processing of windows, which is done in a different way to many other systems by annotating stream elements with processing metadata. Odysseus provides different declarative query languages (e.g., CQL or StreamingSPARQL). One main problem with an extendable system is how to use the new functions provided by some plugin without touching the main system. For this, we developed an operator based language called PQL (Procedural Query Language) which can be extended easily and allows the integration of new functionality, even at runtime.