

Demo: Dynamic Generation of Adaptive Real-time Dashboards for Continuous Data Stream Processing

Timo Michelsen, Marco Grawunder, Dennis Geesen and H.-Jürgen Appelrath

University of Oldenburg, Germany

{timo.michelsen|marco.grawunder|dennis.geesen|appelrath}@uni-ol.de

Abstract. Conventional database management systems are usually not capable to deal with continuous processing of potentially infinite data streams. Therefore, special data stream management systems and frameworks are developed. They use continuous queries which produce also streams as results, so that static visualization is not feasible, since the results are changing constantly. To handle this, we developed a dashboard concept, which we want to propose in this demonstration. A dashboard can be considered as an control or monitoring panel for real time data stream results. The user is free to define and configure individual dashboard parts. Each part is connected to a (user defined) continuous query, whose results are received and visualized in real-time. In this demonstration, we provide different data stream sources, continuous queries and dashboard parts. With them, the user can compose his own individual presentation of his processing results.

Key words: data stream, continuous queries, continuous visualization, framework, Odysseus

1 Introduction

Active data sources like sensors continuously producing data and send them to processing systems. The resulting data streams are potentially infinite and cannot be persistently stored at once. Conventional database management systems (DBMS) are not capable to deal with these type of sources. It is inevitable that each incoming data element has to be processed immediately. For this purpose, data stream management systems (DSMS) are developed. Instead of queries, which are executed once at a given time, DSMS use continuous queries. Continuous queries are installed into the DSMS and are running infinitely, producing streams of results [4].

Many DSMS are hard-coded, specified for one given task. Others provide entire frameworks with basic components to develop application-specific DSMS at a higher level. Odysseus is such a framework. It uses the OSGi Service Platform to provide basic components, which can be easily replaced or extended to meet specific requirements. For instance, Odysseus provides CQL, a SQL-like query language. A developer can implement an additional, more specific language, if

needed. The graphical front-end of Odysseus uses the Eclipse Rich Client Platform (RCP). RCP is build on top of OSGi, providing the same extensibility as Odysseus itself [2].

2 Dashboards

The continuous processing of data streams prohibits a static visualization of query results. The results have to be presented in real-time. So the visualization component of a data stream management system has to update its visualization each time a new result is calculated. Additionally, the system has to deliver the user the tools to individualize the visualizations. For instance, if data streams come from special sources or the results of multiple queries should be shown in one picture. The most suitable solution allows changes even during the continuous data stream processing.

To meet these challenges, we developed a framework for visualizing multiple results in one huge visual surface and implemented it in Odysseus (see figure 1). Therefore, we split the visualization into smaller parts which can be handled

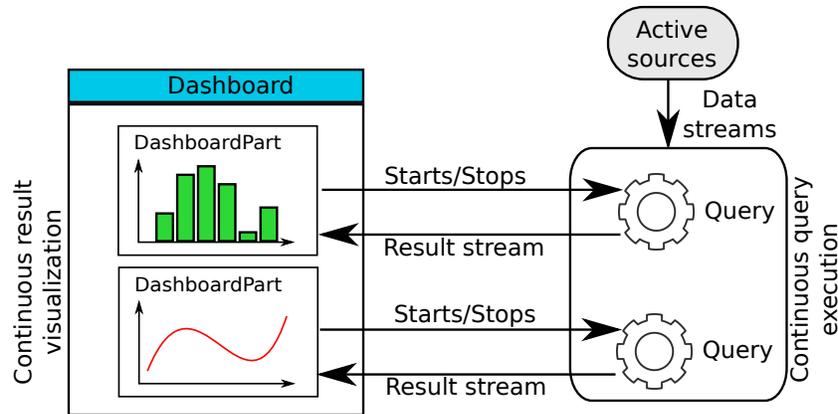


Fig. 1. Correlation between dashboards and continuous query execution in Odysseus.

easier. If a user defines a continuous query, he can connect its (future) results to so-called *dashboard parts*. Dashboard parts are responsible for the execution of the queries: if one part is shown on screen, the underlying query will be installed, immediately started and the results are send to the presenting part. Since the queries are defined independently from the dashboard parts, the queries can be as complex as needed (with complex results and visualizations). Additionally, the dashboard part is highly configurable: First, the user can decide which type of part he wants to show (e. g. classical line diagrams or scatter plots). A developer can easily add additional (application specific) types, if needed. Second, each

part can be individually configured in more detail (e. g. which exact attributes of one data stream should be visualized).

With the underlying RCP-Framework it is possible to save dashboard parts persistently and send them to other users who want to work with them. Next, a set of dashboard parts is composed to one complete visualization. In our concept, we call them *dashboards*, which usually occupy the whole screen. A dashboard is a central control and monitoring panel, which enables the user to see the processing data (and its results) graphically pleasant in real-time. A dashboard is responsible for the positions and sizes of the containing dashboard parts. Beginning with an empty dashboard, the user can include his predefined parts to it (e. g. by drag-and-drop), and move or resize each part inside the dashboard. At the same time, the results of the underlying continuous queries are already shown in real-time. Most likely, all parts are placed in a way, that each part is shown completely. In many situations, the user is only interested in showing his dashboard parts (and does not care about the exact positions and sizes). Therefore, we provide a function to layout all parts automatically (so-called *layouter*). In its simplest form, it positions all parts side-by-side in a grid-pattern. A developer can implement additional layouter for special requirements.

3 Demonstration contents

In this demonstration, we want to show our concept of the dashboards: the dynamic generation and real-time presentation of continuous data streams. Some predefined active data sources continuously send data stream elements to Odysseus. We propose a number of predefined queries which are running in the background. Of course the user can define his own continuous queries if he wants to. For generic data streams, we provide the classical set of diagrams as dashboard parts: line graphs, scatter plots, pie charts, etc. These are suitable for fast and easy visualizations. We also provide some specialized dashboard parts, which are implemented for one specific data stream and their corresponding queries. This shows the high adaptability of our concept. In this demo, we provide two data sources. First, the NREL dataset, containing data from about 32,000 wind power stations. Each station sends a data element each ten minutes which describes the actual amount of generated power, the speed of the wind and several other attributes [3]. Specialized dashboard parts (and dashboards) visualize these data in real-time to give an overview of the current status of the wind power stations. The second data source is a data set provided by the DEBS Grand Challenge 2013, containing position coordinates and moving data from a recorded soccer game [1].

Finally, the user can create an empty dashboard and drag some of our dashboard parts into it. During this process, the continuous queries connected in the background are installed and executed. The user can actually see the results in real-time. If the user removes one part, the corresponding query will be stopped and removed from Odysseus. Figure 2 shows a exemplary picture of a dashboard. It contains two different visualizations/dashboard parts: A specialized

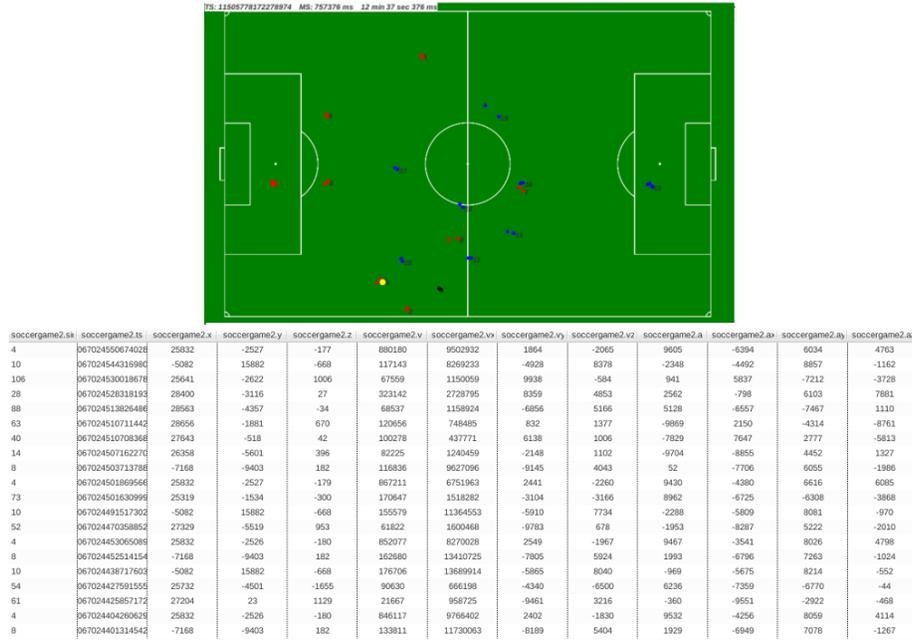


Fig. 2. Screenshot of a dashboard containing two dashboard parts for a soccer game.

part, showing the current position of each player, and a classical table, showing the raw data (e. g. useful for debugging purposes). These parts are rendered during the data stream processing in real-time, which is indicated in a static picture here. It shows, how our concept can be used to build a dynamic, component based real-time visualization of streaming data.

References

1. Debs grand challenge 2013. <http://www.orgs.ttu.edu/debs2013/index.php?goto=cfchallengedetails> (Last visit: 06/12/13), 2013.
2. H.-J. Appelrath, D. Geesen, M. Grawunder, T. Michelsen, and D. Nicklas. Odysseus: a highly customizable framework for creating efficient event stream management systems. In *Proceedings of the 6th ACM International Conference on Distributed Event-Based Systems, DEBS '12*, pages 367–368, New York, NY, USA, 2012. ACM.
3. S. Hammond. Challenges and opportunities in renewable energy and energy efficiency. In *Proceedings of the international conference on Supercomputing, ICS '11*, pages 151–151, New York, NY, USA, 2011. ACM.
4. J. Krämer. Continuous queries over data streams - semantics and implementation. *PhD thesis, University of Marburg*, 2007.