Project N°:	**FP7-318338**
Project Acronym:	**Optique**
Project Title:	**Scalable End-user Access to Big Data**
Instrument:	**Integrated Project**
Scheme:	**Information & Communication Technologies**

Deliverable D8.2
Siemens Use Case Report 2

Due date of deliverable:	(T0+24)
Actual submission date:	November 21, 2014

SEVENTH FRAMEWORK PROGRAMME

Start date of the project: **1st November 2012** Duration: **48 months**

Lead contractor for this deliverable: **SIEMENS**

Dissemination level: **CO** – Confidential, only for members of the consortium (including Commission Services)

Draft Revision November 17, 2014

Executive Summary:
Siemens Use Case Report 2

This document summarises deliverable D8.2 of project FP7-318338 (Optique), an Integrated Project supported by the 7th Framework Programme of the EC. Full information on this project, including the contents of this deliverable, is available online at `http://www.optique-project.eu/`.

The main objective of Work Package 8 is the specification and realization of a real-world industrial use case for the Optique Platform. The work package will provide requirements for the Optique technical work packages, set-up an evaluation infrastructure and provide realistic test datasets.

Deliverable D8.2 discribes the results of Work Package 8 for Year 2. The following list gives an overview of the highlghts of what has been achieved until the end of Year Two:

- *Year-Two Optique Platform successfully installed at Siemens*
  The single most significant achievement for Year Two from a Siemens perspective has been the deployment of the current Optique platform at Siemens and its connection to actual turbine data from the Remote Diagnostic Centres. This has, for the first time in the Optique project, allowed Siemens to test the Optique Platform on realistic data. This achievement is reported in detail in Chapter 4. It should be stresses that access to realistic data could not only be achieved in the sense of querying for historic data but also in the sense of accessing live data streams using STARQL. This is discussed in more detail in Section 4.2.4.

- *Use-case specific visualisation widgets*
  In collaboration with Fluid Operations, Siemens has enhanced the data visualisation widgets used by the Optique platform at Siemens. The main challenges here were related to the large amount of data generated by turbine control systems that had to be visualised meaningfully. This enhancement includes the addition of R-plugins capable of providing sophisticated statistical analysis for signal data. The user-interface of the platform at Siemens is described in detail in Section 4.2 with the R-plugin being introduced in Subsection 4.2.4.

- *End-user workshop*
  Siemens has conducted the end-user workshop at Siemens AS in Oslo, Norway, on 30 October 2014. For the first time, service engineers from the Remote Diagnostic Centres both in Lincoln, UK, and Finspang, Sweden, have evaluated the Year-2 demonstrator. This evaluation involved performing realistic diagnostic tasks on actual trubine data provided by Siemens. The tasks for the end-user workshop are described in Section 2.4.

- *Further analysis of use-case "predictive analysis"*
  Siemens has conducted additional interviews with engineers from the Remote Diagnostic Centres to ensure that the use-case developed within the Optique project is aligned with the actual work of the service engineers. The key task of their predictive work is to anticipate major malfunctions before they happen. This involves analysing both short-term (24 hours) historical data as well as live data streams. The use-case is described in detail in Section 2.2.

- *Definition of Failure Mode Scenarios*
  The insight gained from direct collaboration with service engineers has led to the definition of some

core scenarios that reflect common cases of the engineer's daily diagnostic work. Examples of such failure modes are: discrepancies between pairs or groups of signals, unexpectedly high rates of change of signals, or certain patterns in the event message queue. The failure modes are described in more detail in Section 2.3.

Altogether, especially with the Optique Platform being installed at Siemens with access to real turbine data both in the sense of querying historic data as well as accessing live data streams using STARQL, it is fair to conclude that great progress has been made at Siemens with Optique. This is complemented by an improved understanding of the requirements of the end-users. This understanding will provide the basis for an even more focused further development of the Optique Platform at Siemens in Year Three.

List of Authors

Sebastian-Philipp Brandt (SIEMENS)
Steffen Lamparter (SIEMENS)
Alexey Fishkin (SIEMENS)
Mikhail Roshchin (SIEMENS)
Nina Solomakhina (SIEMENS)
Stuart Watson (SIEMENS)
Richard Arnatt (SIEMENS)
Gary Over (SIEMENS)

List of Contributors

Martin Giese (UiO)
Evgeny Kharlamov (UOXF)
Rudolf Schlatte (UiO)
Ahmet Soylu (UiO)
Artem Kozlov (FOP)
Johannes Trame (FOP)

Contents

Chapter 1

Introduction

The present deliverable details the main results of Year Two of WP8, the Siemens Use Case. The aim of the Optique project is to provide access to big data resources for users not specialising in IT but rather in a domain of discourse underlying the relevant data sources. Within Optique the key to reaching this aim is to devise an ontology representing the domain knowledge necessary for translating data between two divergent perspectives: the perspective of the engineer interpreting the data and the perspective of the database storing it. Optique enables the domain expert to query the knowledge resource using concepts familiar to her or his daily routine. The evaluation of such queries is done by translation into a potentially large number of data sources for which the structure and organisation principle may be totally different.

For Siemens, on an industrial scale, this kind of capability has the potential to greatly increase the effectiveness and cost-efficiency of the management of gas and steam turbines. In particular, it would provide an attractive solution for the already existing problem of vastly heterogeneous data sources. Moreover, the uniformity of the access layer is a key to protecting Siemens' investment into highly sophisticated data anlytics architecture and solutions that, currently, need to replicate the heteronegeity of the data sources they are accessing. The latter aspect is of particular interest for Siemens as the demands of its turbomachinery equipment are such as cannot be met by fixing faults as they occur. Rather, Siemens needs to employ the most sophisticated analytics techniques in order to anticipate problems with steam or gas turbines before they occur and take the relevant countermeasures.

The scope of the content of this deliverable, as defined in the Description of Work (DoW, see [3]) and in the Project Proposal [4] is given in the following section.

1.1 Scope

As laid down in the DOW [3], WP8 has the goal to utilise the capabilities provided by Optique for the work-flow of the Siemens Energy Service environment. The present document summarises the achievements at the end of the second year. These refer to the following tasks of the DOW:

- Infrastructure (Task T8.1)
- Siemens Use Case Requirements (Task T8.2)
- Data Models (Task T8.3)
- Installation preparation and tool development (Task 8.4)
- Installation (Task T8.5)
- End-user evaluation

The present deliverable focuses on further refinements made in the Siemens use-case (Task T8.2), and the successful development and deployment of the Optique platform within the IT infrastructure at Siemens. Altogether, the main achievements of Year Two at Siemens can be summarised as follows:

1.1.1 Infrastructure

In the first year the Siemens Energy Services infrastructure has been surveyed and options for the deployment of Optique have been derived. In close collaboration the main stakeholders including diagnostics engineers, Energy Services IT experts as well as management, Siemens has found a viable option for the installation of the Optique platform within the highly secured de-militarised zone. This location enables Optique to access actual turbine data as provided by the Remote Diagnostic Centres rather than simulated or anonymised data.

Note that the topic of IT infrastructure will be covered together with the installation of the Optique Platform at Siemens in Chapter 4 rather than in a chapter on its own.

1.1.2 Siemens Use Case Requirements

Due to the increased collaboration between stakeholders at Siemens the second year has given Siemens a better working knowledge of the daily routine of engineers working in the Remote Diagnostic Centres. In particular, there is a better underastanding of the different diagnostic modes provided by the diagnostic engineers and the role of retrospective and prospective diagnostics. The main outcome of this better understanding is that retrospective diagnostics, i.e., analysis of historic data, and predictive analytics, i.e., analysis of live data, are not as strictly separated as previously thought. For the Optique project and its application within the diagnostic workflow at Siemens this translates into even more opportunities to benefit from the Optique platform.

In addition to the better general understanding of the daily routine of the diagnostic engineers some considerable time has been invested into refining specific use-cases for Optique. The foundation for this work has been laid in the first year of the project. Now, at the end of year two, our use cases reflect the actual working scenario of the diagnostic engineers in more detail.

1.1.3 Data models

In order to provide access to actual turbine data collected in the Remote Diagnostic Centres the Optique Platform at Siemens had to be provided with a new database connection and, subsequently, with new mappings. These mappings connect the largely unchanged doman model (i.e., domain ontologies) to the new database schema and tables.

1.1.4 Installation

The most significant achievement for Year 2 from a Siemens perspective has been the deployment of the current Optique platform at Siemens and its connection to actual turbine data from the Remote Diagnostic Centres. Siemens is now, for the first time in the Optique project, in the position to test the Optique Platform on real data from Siemens turbines. Acces could not only be provided in the form of retrospective querying of historic data but also in the form of STARQL-based processing of live data streams. This fully meets the development stage planned for the end of Year Two and, with respect to stream processing, arguably exceeds it.

1.1.5 End-user evaluation

The Optique Platform as installed within Siemens has been evaluated by Siemens service engineers accessing original Siemens turbine data. This has happened within an end-user evaluation workshop at Siemens AS in Oslo, Norway, on 30 October. The fact that such an evaluation could be succesfully conducted at all shows that significant progress has been made within Year Two at Siemens. The result of the user survey accompanying the end-user workshop will be the subject of Deliverable D1.2 [9].

Chapter 2

Siemens Use Case

Within Year Two the collaboration between the various stakeholders at Siemens was successfully deepened and extended. The result of that is a more in-depth understanding of the daily routine and requirements of the service engineers who embody the experts behind the use-case at Siemens. In addition to more in-depth knowledge of diagnostic work driven by so-called "failure modes" the general relationship between retrospective analysis of historical data and real-time analysis could be clarified as well. This aspect is discussed in more detail in Sections 2.2.1 and 2.3.5.

In this chapter we begin by recalling some basic facts about Siemens gas turbines in general and the way these are maintained. In this sense Section 2.1 and the beginning of Section 2.2 should provide the reader with all the necessary knowledge to follow the discussion of the use-cases in Sections 2.3.

2.1 Siemens gas turbines

In order to make the remainder of this chapter more easy to follow we begin by providing a quick overview over the basic of Siemens gas and steam turbines and their maintenance. This topic has been discussed in Deliverable D8.1 [8] already. Some of the illustrations and basic explanations shown here are re-called from [8].

Siemens Energy Services maintains thousands of devices related to power generation, including gas and steam turbines – generally called *appliances*, *units* or *rotating equipment*. Operational support for these appliances is provided through a global network of more than 50 Remote Diagnostic Centres. These centres are linked to a common database center, where the data generated by each appliance data is stored in a large number of databases. In the following we describe the structure and monitoring facilities of an appliance. Figure 2.1 presents a high-level overview of the data flow from appliances, i.e., gas and steam turbines for our use-case, all the way through to the Remote Diagnostic Centres. The latter are called "service center" in Figure 2.1.

Each turbine comprises several industrial computers that operate based upon information from sensors. These computers serve the functions of (i) a *control unit* and (ii) a *data collector*. Overall, approximately 2000 hardware sensors and measuring devices are used to monitor the correct operation of a single turbine. The tasks of the control unit and data collector can be described as follows [8]:

Control Unit. Each control unit serves the following functions: receiving and processing sensor measurements, real-time monitoring of the turbine, and communication of all information to a data collector.

Receiving measurement data from sensors triggers analysis on two levels in the control unit:

1. So-called *soft sensors* apply predefined simple rules and triggers, to identify, e.g., threshold exceedance and trends, and provide this information as derived sensor data.
2. Additional simple analysis and information preprocessing is conducted on information from physical and soft sensors. One typical example is the computation of Fast Fourier Transformations to facilitate

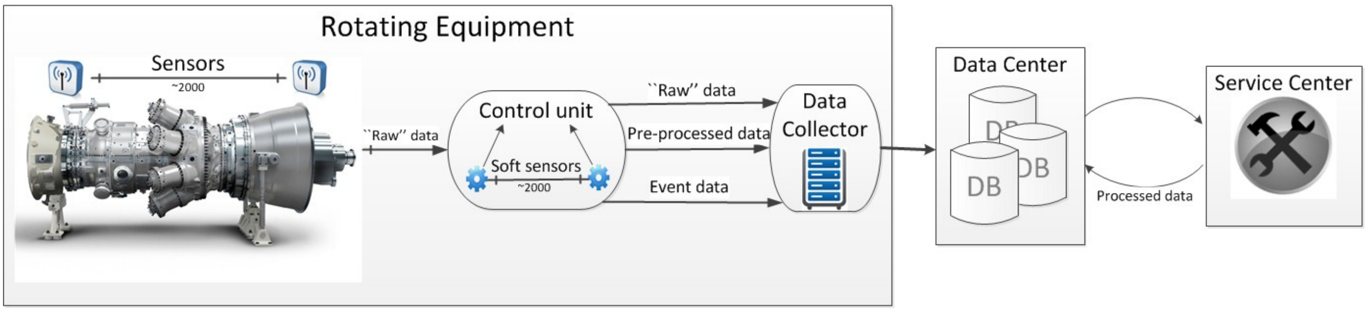

Figure 2.1: Appliance structure and data flow.

the analysis of time series data.

Based on the results from these processing steps, the control unit generates short messages called *events* that describe the status of a unit. Important events are immediately displayed on a small monitor accessible to the operator of the turbine. Some turbine types additionally store a local history of events for a limited period of time (e.g. one hour) to permit local analysis by the operator. In order to persist the operational information, the control unit passes the information on to the data collector.

Data Collector. The main function of the data collector is to accumulate the information forwarded by the control unit and send it to the central database in a regular interval (e.g. daily). Note, however, that due to unstable and slow connections between plants and data center (typically caused by geographic distance), it is not possible to transfer the complete raw data. Instead, the data collector transmits events and soft sensor data plus a selected subset of raw data.

Data Center. Typically one common data center corresponds to several Remote Diagnostic Centres. At the data center, information sent from the data collectors is stored in a service-center specific database. Within each of these databases there exists a large and heterogeneous landscape of different schemas. The reason for this is two-fold: Firstly, different makes and generations of turbines have been designed historically to report different sets of data in different formats. Secondly, sub-components of an appliance might be from different vendors, introducing still other data formats.

For more information on how a Siemens gas turbine actually looks like and what its performance characteristics are, please find public data on-line[1] or in the relevant trade publications [2].

2.2 Gas turbine diagnostics at Siemens

For the purpose of describing the use-case at Siemens in more detail, we begin by discussing the workflow for the analysis of turbine data in more detail. In comparison to Figure 2.1 the information flow from Siemens equipment deployed globally to their relevant diagnostic centres is sketched in more detail in Figure 2.2 [5].

The illustration shows that the first database to actual store large amounts of signal and event message data transmitted from Siemens appliances is a dedicated global data centre for Siemens equipment. In addition to just storing raw data, especially signal data is enhanced here and made accessible to the Remote Diagnostic Centres subscribing to the data. Hence, the retrospective analysis of historic data in the centres is based on data available through the global data centre.

Due to increased intraction with service engineers from Remote Diagnostic Centres the actual workflow relevant to our use-case is better understood at the end of Year Two in comparison to the first year. In particular, there are different modes of operation present in the Remote Diagnostic Centres that are highly relevant to the use-case of Optique.

[1]See, e.g., `http://www.energy.siemens.com/hq/pool/hq/power-generation/gas-turbines/sgt-750/SGT-750_brochure.pdf`

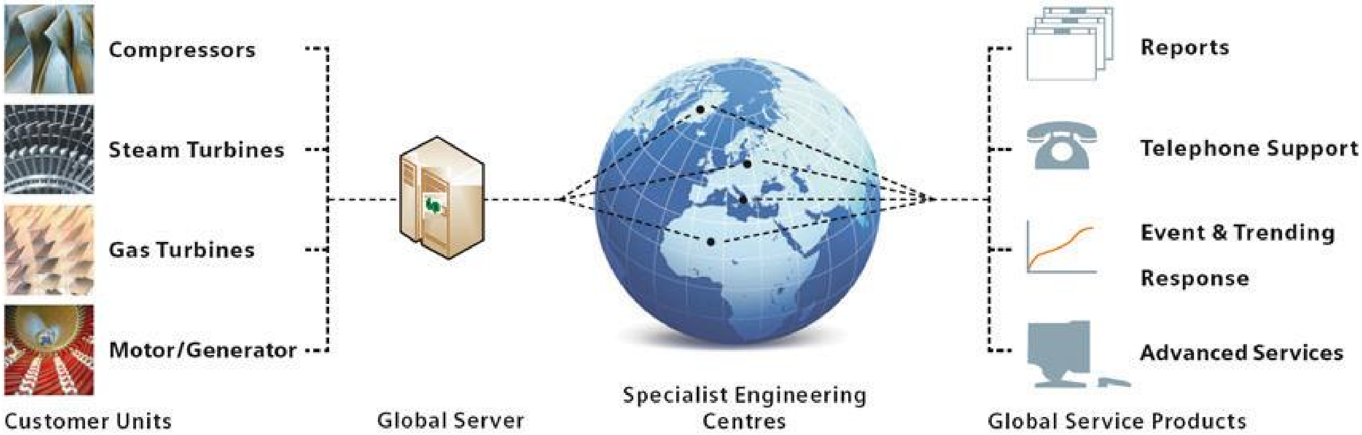

Figure 2.2: Remote Diagnostic Services

2.2.1 Daily monitoring vs help desks

Looking at the Situation today there are two major modes in which the Remote Diagnostic Centres of Siemens provide support to the customers who run Siemens appliances. This depends on the kind of service program the customers have signed up to. The more basic one is the Help Desk and the more comprehensive one the Daily Monitoring Programme. The differences between the two can be summarised as follows:

- *Help Desk*
  The work of the Help Desk is driven by the customer's encounter of problems with their appliances in general or turbines in particular. If the customer notices anything unusual they contact the Help Desk and create a ticket. A service engineer of the Help Desk then analyses the live data stream generated by the turbine of interest in order to assess the situation. Should there be a problem or fault it may be necessary that historic data of the relevant turbine needs to be analysed as well.

- *Daily Monitoring*
  Daily Monitoring, in contrast to the work of the Help Desk, happens pro-actively by the Remote Diagnostic Centers without individual request by the customer. Every service engineer for the Daily Monitoring progam has a certain set of turbines under her or his supervision. For these the engineer analyses the data in batch mode on an hourly or daily basis, depending on the type of the contract. This work is usually accomplished with the help of sophisticated signal processing and statistics tools.

In Deliverable D8.1 [8] the scenario of the Help Desk has been discussed in dept while the role of the Daily Monitoring program was not fully understood. The general perspective is that today the work of the Help Desk looks more into the present while the Daily Monitoring programme looks into the recent past.

While this difference in perspective as sketched above is true in general the crucial difference lies in the fact that the main purpose of the Daily Monitoring programme is to prevent larger-scale faults of turbines by taking corrective actions as soon as small-scale problems become apparent. Hence, the perspective of the Daily Monitoring programme is inherently predictive and hence predictive analytics of data are highly relevant. But with growing importance and sophistication of predictive analytics real-tme analysis of data streams becomes inevitable. And as the Daily Monitorng programme runs continually without explicit request by the customer this constitutes a much stronger use-case for real-time analytics of data.

In the following section we will go into more detail about the kind of analysis carried out in the Remote Diagnostic Centres. The relationship between reactive and predictive analytics will be re-visited in Section 2.3.5.

2.3 Failure modes

One of the main tasks of the service engineer in one of the Remote Diagnostic Centres at Siemens is to detect certain configurations of signal data (and, in many cases, event messages) and take the appropriate countermeasures before larger-scale deviations from the operating envelope and hence potential outages occur. These configurations can be categorised by their root causes and analysed further. [7]

In the context of faults of turbines, these categories are commonly called failure modes. In the following section we will introduce the most common failure modes relevant to the service engineers at the Remote Diagnostic Centres.

2.3.1 Message-signal dependencies

Although this may be the most generic category of faults, the relationship between event messages sent by the control system and the underlying signal data yields valuable links for the re-construction of the root cause of a given problem. The task is easily illustrated from the perspective of a service engineer working for the Daily Monitoring program of turbines.

Consider the case that the engineer finds that some turbine (T) has had an unplanned stop during some given 24-hour period. The abnormal stop, usually called "Trip", will be visible in the event message queue of the turbine in the form of some message like "Timestamp X: Running trip (Shutdown)", where "Shutdown" is the event category. The usual first diagnostic step of the engineer now is to inspect the event message queue prior to time point X in search for some other event that potentially caused the turbine to shut down.

If the engineer finds that some message M was sent at time point Y prior to X where M indicates a serious deviation from the usual operating envelope then the engineer has already accomplished a significant step towards reconstructing the cause of the shut down. The key step now is to infer from message M which kinds of signal data needs to be inspected around time point Y (or possibly even earlier) in order to check why message M has been generated.

This knowledge about the link between event messages and signals is currently not accessible by automated analytics but rather part of the experience that service engineers in the Remote Diagnostic Centres need to have to carry out their work. Depending on the message of interest this link can be trivial or at least appear trivial. For instance, if the message says "Oil temperature high at journal bearing" then the signal of interest is obviously the temperature sensor for the oil system responsible for lubricating the relevant bearing. Nevertheless, lots of other signals need to be inspected in order to assess the relevance of the high temperature reading: oil pressure, vibration at the relevant bearing, load of the turbine, and even ambient temperature.

Again, mastering the associations between the individual signals is currently a skill of the service engineer and not a service provided automatically. This means, however, that the service engineer spends valuable time selecting signals for inspection manually instead of being presented the entire picture.

2.3.2 Signal discrepancies

In contrast to the previous case, signal discrepancies represent a very basic failure mode that can be discussed in the absence of event messages. Like other complex machines, turbines rely on negative feed-back loops in which certain operational parameters are set and the enforced by the control system. The turbine can be run safely only if all the relevant feed-back loops work, i.e., the control system manages to ensure that certain parameters remain set to the value given by the control system.

A simple example of this is the task of controlling a valve of which turbines have several. The valve is equipped with (1) a servo motor capable of setting the valve to an arbitrary angle, from completely shut to completely open, and (2) a sensor that actually measures the opening angle of the valve. If the turbine runs normally then the control system keeps setting a certain angle for the valve, with the valve responding by actuating the servo motor accordingly. The control system then uses the sensor to check whether the demand value is actually set or not.

Figure 2.3: Blow-off valve discrepancy 1/2

A typical example of such a feed-back loop malfunctioning is shown in Figure 2.3. Here the demand is represented by the red line maintaining the value 0 throughout. The actual angle of the valve measured by the relevant sensor is shown by the blue line fluctuating from between 1 and 2 to around zero to higher values again. The threshold to which discrepancies between demand and actual value or valve setings are tolerated vary between machines but is 3 degrees in the graph shown in Figure 2.3. Consequently, the control system has correctly raised an alarm message between 15:27 and 15:44 o'clock when the threshold was exceeded.

The simple example discussed above seems to suggest that no specific intelligence is necessary to handle deviations of demand values such as the blow-off valve angle. The following instance in Figure 2.4 shows that the complete picture is slightly more involved. The figure again shows demand (red) and actual value (blue) of one of the valves of a turbine. This time the control system has again raised an alarm, however for a only minimal violation of the 3-degree threshold.

If a deviation such as shown in the first case in Figure 2.3 happens during start-up of a turbine then there is no other choice than to abort the start. In the case shown in in Figure 2.4, however, there may be several options to tolerate the deviation for the time being and try to take corrective measures. This case suggests that real-time monitoring of the valve situation together with a sufficient domain model of the turbine may offer a chance to counteract the problem before a more serious situation occurs.

2.3.3 Rates of change

Another common failure mode that deals with sensors in isolation concerns the rate at which a signal changes. During normal operation of a turbine most of the operating parameters remain in a relatively steady state. That means that extreme changes of signal values usually do not occur. The one reservation to this rule regards noisy signals where adjacent data points can be quite different from each other despite an almost constant moving average of the signal.

An example of a rapidly changing signal is shown in Figure 2.5. The graph shows the readings of the

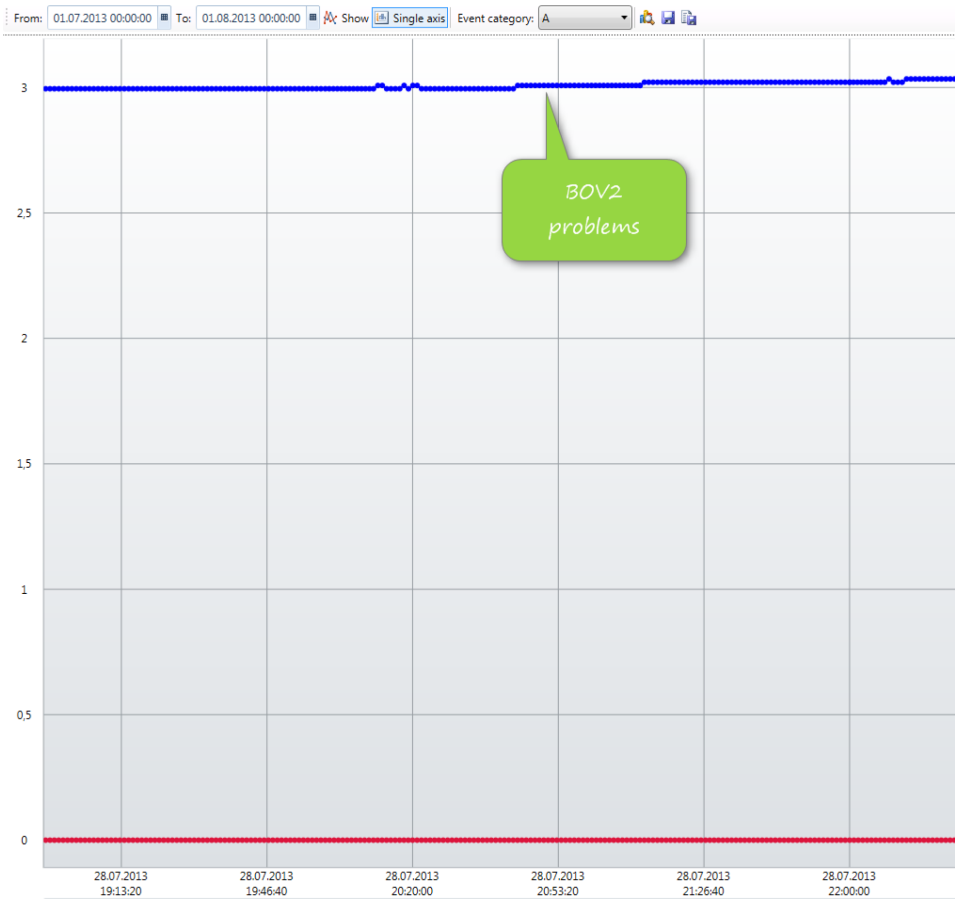

Figure 2.4: Blow-off valve discrepancy 2/2

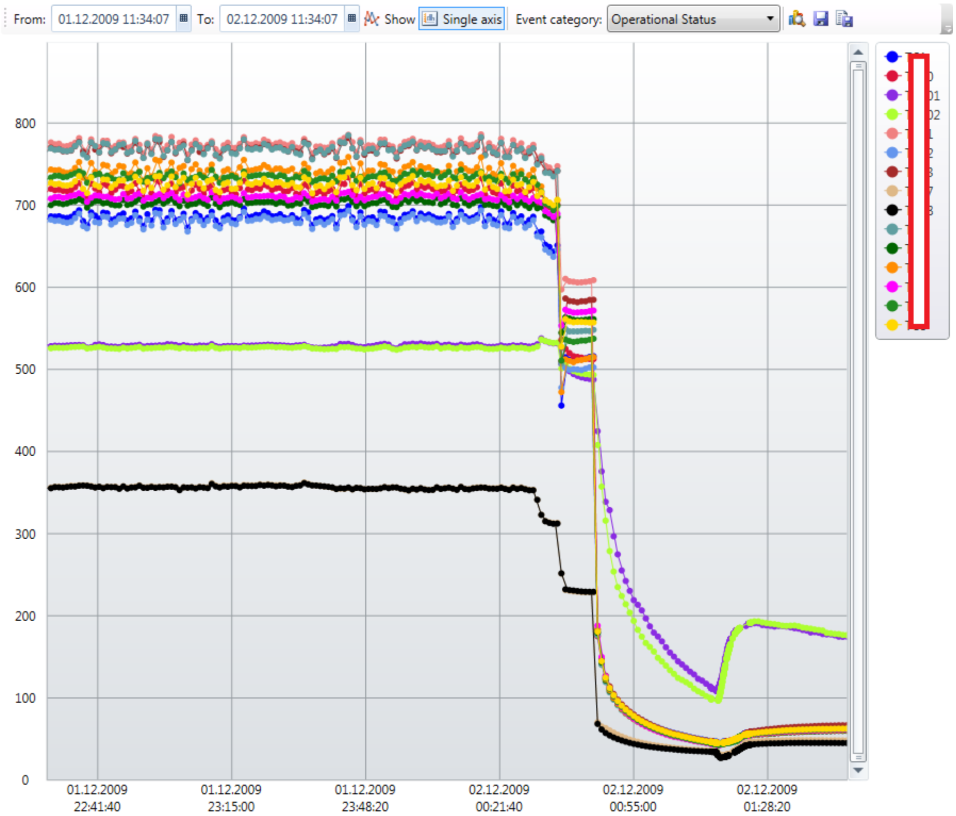

Figure 2.5: Thermocouples

main temperature sensors in the gas path of the gas turbine. The Y-axis is labeled in degrees celsius. Hence, the temperature readings range from between 700 and 800 degrees at the hottest part to around 350 degrees in the exhaust.

The graph shows that the temperatures dropped rapidly at some point in time between 0:21 and 0:55 o'clock. Such termperature drops cannot be tolerated if they happen too fast because too quick temperature changes translate to mechanical stress on the rotor of the turbine. Hence, detecting abnormal rates of change is an important capability both in terms of retrospective analysis as well as analysis of real-time data. In Section 2.3.4 an example will be shown where signals did not change fast enough in order to trigger the emergency routines of the control system but still did cause serious damage after some delay.

What makes the rates of change failure mode interesting for Optique is that by only looking at the signal itself it is usually not easy to discriminate the serious deviations from operating parameters from cases where appropriate countermeasures can still be taken. Hence, if Optique provides both a domain model and analytics capabilities for streaming data this would be an attractive basis for taking appropriate countermeasures before large-scale problems occur.

Incidentally, Figure 2.5 also illustrates the problem of noisy sensors. Computing the rate of change solely by the difference of two adjacent measurement points would lead to ountless false alarms. Here the promise of Optique is that adding a layer of statistics processing is much easier to do than altering the core algorithms of the sensor or the data collector attached to the turbine.

2.3.4 Complex example: water stroke

The previous examples have focused on anomalies of event message and sensor readings that can be understood in relative isolation. For instance, observing the behaviour of the demand and actual value for the blow-off valve is enough to ensure safe operation of the valve group of the turbine. If we focus our attention to a single turbine then an engineer may still have the opinion that dealing with the situations described above should be done without complex knowledge models and semantic technologies but rather with slightly improved turbine control systems.

The following example is intended to illustrate that this view does not hold even if we limit the perspective to a single machine: we are going to illustrate a case in which only a substantial amount of domain knowledge on top of the control system would have been able to avoid serious damage to one of Siemes' steam turbines. We consider a single gas-steam power unit installed on one of the Siemens locations. It is designed to support the following working ranges:

- Active power 0–20 MW
- Revolution speed 0–7000 rpm
- Turbine flange temperature 180–500 ^{0}C
- Fresh steam temperature 180–500 ^{0}C
- Wheel chamber pressure 0–20 bar

The present use case represents the data set collected in the period from 2004 to 2014. Besides a collection of all "events", i.e., information, alert and warning massages, reported by the control system, it also consists of online measurements coming from 2000 sensors every minute.

For simplicity, here we only describe a so-called "water-stroke" scenario. In this day the operator used an "unusual" fossil source for producing "dry-steam". Due to specific thermodynamic properties the steam has started to condense in the working chamber. Since there were only tiny deviations from the "usual" operating mode, the control system was not registering any significant deviations in the critical control values. However, accumulating over time the condensed water started to cause very specific "functional" problems in different power unit subsystems which were recognized by the control system. However, the operator was taking them as "usual" warnings without reacting on them properly. So, the turbine was running until an " emergency stop" has not been remitted by the control system due to the "extreme" control parameter values. The resulted "trip" stop of the turbine from 7000rpm to 0rpm within several seconds caused significant mechanical forces in the flange and vibrations.

Figure 5.1 depicts the situation on 12 September at 15:16. One can see that the active power dropped from 19.9MW to 4.9MW within 2:25 minutes. The fresh steam temperature was showing 465^{0}C. Next, at 15:19 a warning message "Fresh steam pressure $<$ SSV min" and then one more time at 15:30.

The situation starting from 15:29 is shown in Fugure 2.7. The fresh steam temperature SSV fell down from 419^{0}C to 285^{0}C within 2 minutes. In the same time the inside turbine flange temperature has fallen from 358^{0}C to 194^{0}C within 2 minutes. The generator active power has dropped from 4.63MW to 0MW within 16 seconds.

At 15:30 a red alarm "Generator Power Protection" was emitted. This tripped the turbine. At 15:31 a boiler "trammel water level" alarm was emitted. The turbine speed has been dropping from 6842 rpm to 46 rpm within 2 minutes. The normal turbine cast-down takes around 17 minutes.

At 15:32 there was a "red alarm". At 15:32 there were several warnings indicating vibrations in the turbine and gear units. The first "stand-still" was 18 minutes. Then, the turbine was rotating at 116 rpm for 13 minutes. The second "stand-still" was for 9 minutes. The warning messages regarding "vibrations" were coming at 15:51.

The main cause of the dangerous vibrations was a significant drop in the "fresh steam" temperature followed by a significant drop in the "turbine flange" temperature. In the same time, the coast down of turbine took just 3 minutes instead of required 17 minutes. This form shows that there were significant axial rubbing forces which caused vibrations and possible turbine damages.

The above example is intended to illustrate that the control system of the turbine is not able to predict the effect of "water-stroke" in advance because this scenario requires too deep knowledge of critical longer-term consequences of relatively non-critical sensor readings. The control system reacted to the present situation by just reactively stopping the turbine as soon as the temperature, pressure and vibration values

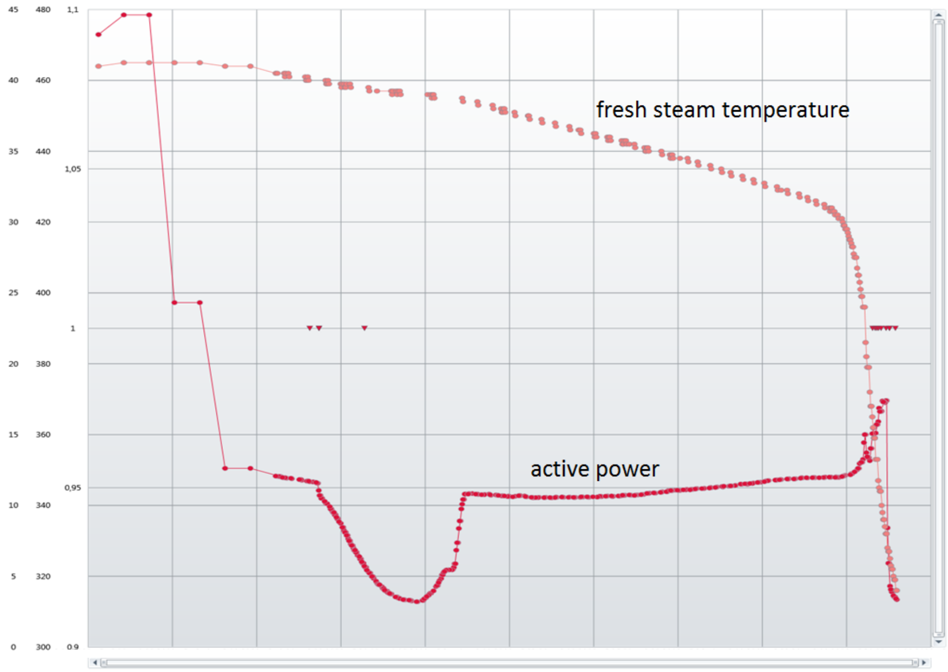

Figure 2.6: Water stroke scenario 1/3

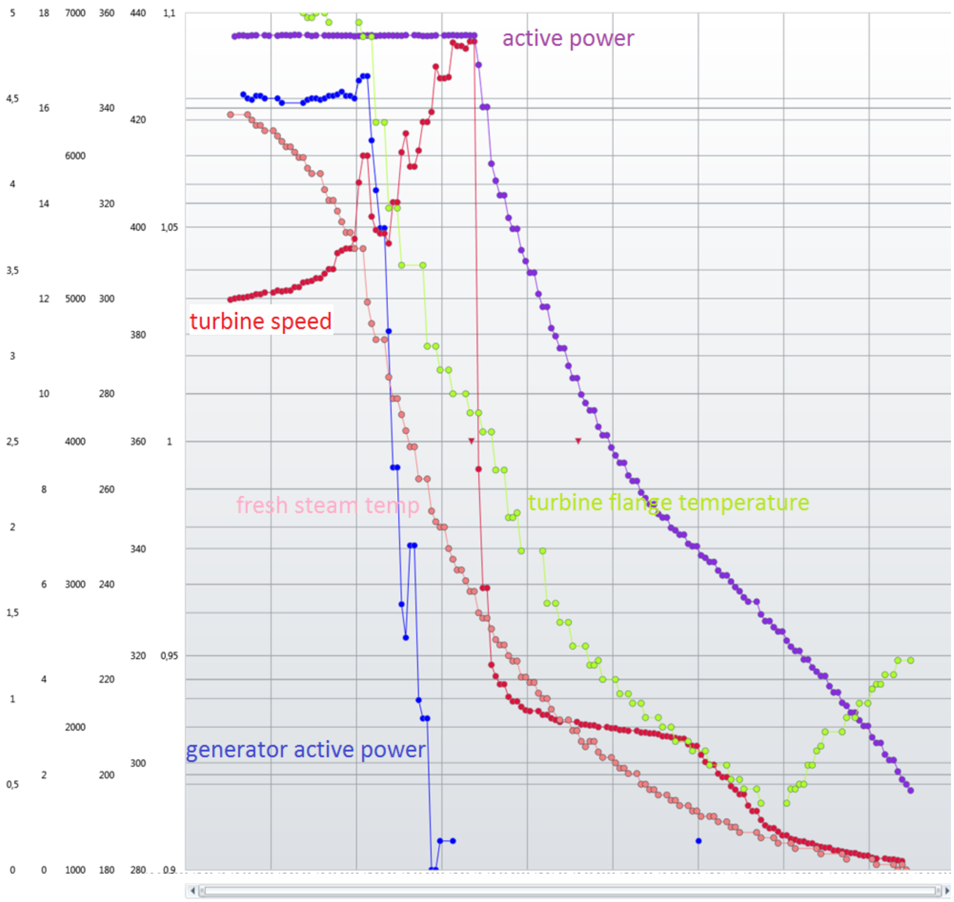

Figure 2.7: Water stroke scenario 2/3

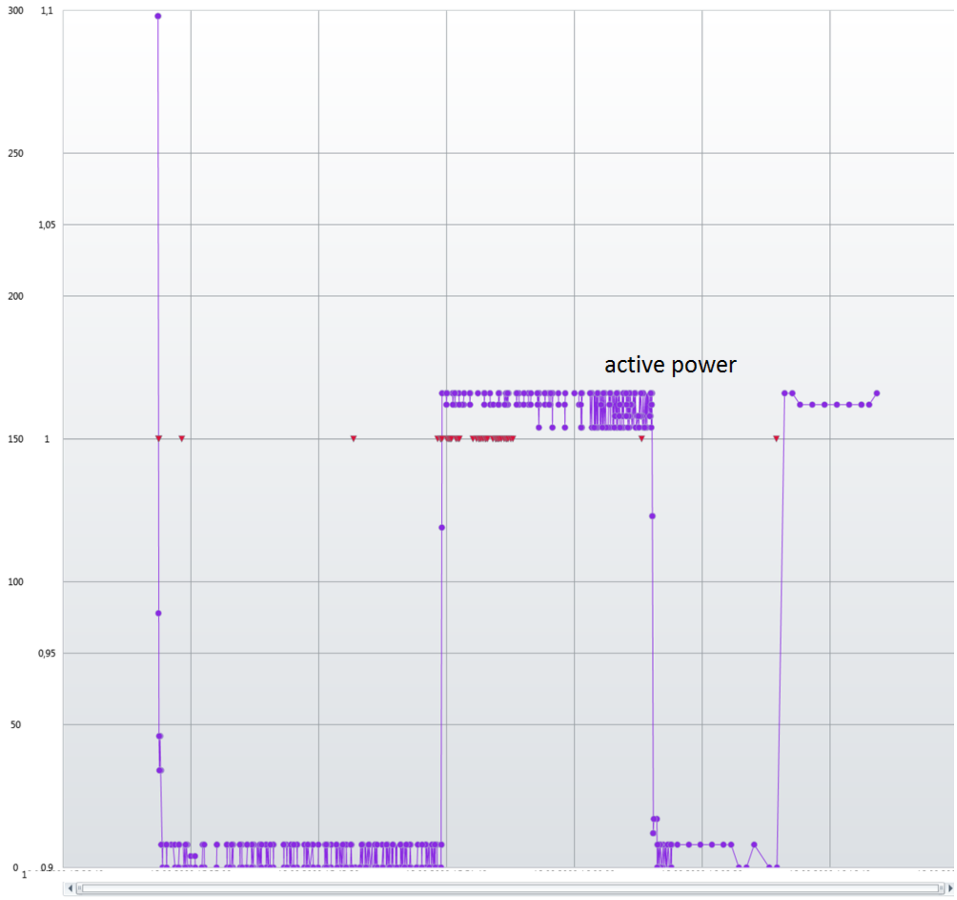

Figure 2.8: Water stroke scenario 3/3

had exceeded the pre-defined extreme control values. However, we can see that the control system was sending “warning” messages from different subsystems. Furthermore, we can also see that there was enough information available from the measurements to indicate critical mechanical and thermodynamic states.

We expect that Optique will provide the capabilities to prevent situations such as the above. The key difference that Optique makes is that it enables engineers to represent higher-level domain knowledge that enable an automated system on the level above the actual control system to interpret the data from the control system in such as way as to predict dangerous constellations of signal data in time to take the appropriate countermeasures.

2.3.5 Reactive vs predictive diagnostics

In Section 2.2.1 the roles of the Help Desk and the Remote Diagnostic Centres and their apparent respective focus on real-time analysis and retrospective analysis have been discussed. In the respective sections above we have seen that all failure modes have obvious applications for retrospective analysis: if the running parameters of the turbine are too far beyond the envelope permitted by the control system then the control system will simply shut down the turbine, giving the service engineer in the Remote Diagnostic Centre the task to re-construct the course of events prior to the shut-down. An important part of this re-construction is to go through similar steps as the ones shown above.

We have seen, however, that all failure modes have a real-time dimension as well. If the relationship between alarms and the underlying signals are clear enough and if deviations from the norm are small then it becomes highly likely that appropriate countermeasures will avoid the necessity to shut down the turbine or other large-scale effects. Hence, if smaller-scale deviations with respect to any of failure modes above can be detected automatically by processing the relevant live data streams and the deviation brought to the attention of either the service engineer or rather the local operators on-site then countermeasures can be taken to avoid shutting down the machine in the first place. The last example in Section 2.3.4 reenforces the view that such functionality is unlikely to be implemented on the level of the individual sensors or on the level of the control system simply because the relationship between various different signals is too complex.

In terms of requirements for Optique this means that the capability of processing live data streams seems even more important at the end of Year Two than with respect to the use-case as understood at the end of Year One.

The understanding gained in the area of failure modes was used to re-visit the query catalog built up as a repository of standard tools for diagnostic work with Optique. We have been able to verify that the query catalog as developed in Year One was sufficient to answer most of the questions related to the failure modes described above, if not in a fully-automated way. The only query that was added has been one designed to return all turbines that had sent a message of clategory “Shutdown” in the previous 24 hours.

Another key insight gained in connection with the failure modes, however, refers to the topic of migration of analytics solutions from one machine to a new one and, similarly, to the topic of maintaining and enhancing analytics solutions. Without an abstract domain model describing the structure of Siemens equipment on a generic level all failure modes inevitably have to be implemented in a way that directly refers to the specific sensor tags of the machines and, similarly, to all the peculiarities of the control system, data collector, and turbine database. This makes transferring and upgrading solutions quite hard and, additionally, has the disasvantage of changes to the database structure invalidating analytics solutions. With an abstract model of the mechines to be monitored Optique offers the unique alternative to implement the analysis of failure modes and other analytics solution in general on an abstract level that (1) does not dirctly reference machine particulars and (2) is not rigidly tied to any specific database infrastructure. This means that an analytics solution based on Optique can be migrated to new machines easily and is robust w.r.t. changes to the database structures. Considering the complete picture of protecting the investment in sophisticated analytics solutions this aspect of Optique may be the most important one.

2.4 End-user evaluation

In order to evaluate the Optique Platform as installed at Siemens at the end of Year Two, a workshop has been organised on the premises of Siemens AS in Oslo, Norway, on 30 October. The participants of this workshop included Optique partners, 10 representatives of Siemens AS, and, most importantly, four engineers from two Siemens Remote Diagnostic Centres for steam and gas turbines. The latter had in-depth experience managing and maintaining Siemens gas and steam turbines as well as the entire IT infrastructure around this process. Hence, the four engineers represented the perfect audience to provide feed-back on the usability and maturity of the Optique Platform at the end of Year Two.

After a general introduction into the aims of Optique the four engineers were given an initial survey with the following questions:

1. What is your age?

2. What is your occupation?

3. What is your level of education?

4. I have technical skills (i.e., computer) such as programing and query languages (e.g., SQL, Java, PHP, SPARQL etc.)
(Agreement to be expressed on a five-point scale.)

5. I am familiar with tools similar to Optique.
(Agreement to be expressed on a five-point scale.)

The main part of the evaluation consisted of a numbr of tasks that have to be considered very simple for the level of expertise of the perticipating engineers. The aim of the tasks was to encourage the engineers to explore the capabilities of the Optique Platform.

It should be emphasized that the Optique Demonstrator used for the end-user workshop had access to original Siemens data. Hence, the engineers were not evaluating Optique on the basis of toz data but rather on the basis of realistic turbine data. The tasks given to the engineers were as follows:

- *Task 1*
Using the Visual Query Formulation Interface, find all assemblies that exist in Optique.

- *Task 2*
Using the Visual Query Formulation Interface, show all Messages that Turbine T generated from 01.12.2009 to 02.12.2009.

- *Task 3*
Using the Dashboard for Turbine T, show all Messages that this machine generated from 01.12.2009 to 02.12.2009.

- *Task 4*
Using the Dashboard for Turbine T, find all messages between 01.12.2009 and 02.12.2009 that contain the text “Trip”.

- *Task 5*
Using the Visual Query Formulation Interface, show all turbines that sent a message containing the text “Trip” between 1.12.2009 and 2.12.2009.

- *Task 6*
Using the Visual Query Formulation Interface, show all event categories known to Optique.

- *Task 7*
  Using the Visual Query Formulation Interface, show all turbines that sent a message of category "Shutdown" between 1.12.2009 and 2.12.2009.

- *Task 8*
  Try to find out whether Turbine T is currently running.

At the end of the actual diagnostic tasks, the participants were given the main part of the questionnaire. The main part of the questions related to the topics of usability and fitness for purpose:

1. I think that I would like to use this system frequently.
   (Agreement to be expressed on a five-point scale.)

2. I found the system unnecessarily complex.
   (Agreement to be expressed on a five-point scale.)

3. I thought the system was easy to use.
   (Agreement to be expressed on a five-point scale.)

4. I think that I would need the support of a technical person to be able to use this system.
   (Agreement to be expressed on a five-point scale.)

5. I found the various functions in this system were well integrated.
   (Agreement to be expressed on a five-point scale.)

6. I thought there was too much inconsistency in this system.
   (Agreement to be expressed on a five-point scale.)

7. I would imagine that most people would learn to use this system very quickly.
   (Agreement to be expressed on a five-point scale.)

8. I found the system very cumbersome to use.
   (Agreement to be expressed on a five-point scale.)

9. I felt very confident using the system.
   (Agreement to be expressed on a five-point scale.)

10. I needed to learn a lot of things before I could get going with this system.
    (Agreement to be expressed on a five-point scale.)

11. What did you like about the tool?

12. What didn't you like about the tool?

13. Do you have any other comments?

The outcome of the end-user evaluatin is going to be reported in Deliverable D1.2 [9]. Independently of this it should be noted that the fact that four engineers from Remote Diagnostic Centres were willing to join the Optique partners in Oslo for an evaluation session is indicative of a considerable amount of interest in Optique within Siemens.

Chapter 3

Siemens Data Models

As we will discuss in detail in Chapter 4, one of the main achievements of Year Two has been the successful implementation of an Optique Platform at Siemens and its connection to the actual Siemens turbine data sources. These data sources represent turbine data, both event messages and signals, from the Remote Diagnostic Centres. The underlying schema of the relevant database is shown in Figure 3.1.

3.1 Database schema

As the diagram shows, the schema is straightforward: turbines, more generically called 'assemblies' here, contain sensors that in turn generate measurements. Each mesurement has a timestamp, a sensor it belongs to, and a value. In addition to the measurements, the assembly (turbine) also generates event messages each of which belongs to some category. In addition to that, messages contain administrative information not immediately relevant to the use-case.

It should be noted that some of the data hosted by the Remote Diagnostic Centres is represented w.r.t. a more complex schema than shown in Figure 3.1. While this more complex schema could be obtained and is currently being processed for further use within the Optique Project, access to real turbine data could already be established on the basis of the simpler schema. Recall that the general infrastructure of the Remote Diagnostic Centres has been introduced in Sectiona 2.1 and 2.2.

3.2 Domain ontology

In comparison to the state at the end of Year One, the domain ontologies used for the Optique Platform at Siemens have not seen any major changes. Note that such changes have not been in the focus of the development of Year Two at Siemens. We therefore refer to Deliverable D8.1 [8] for a more compehensive

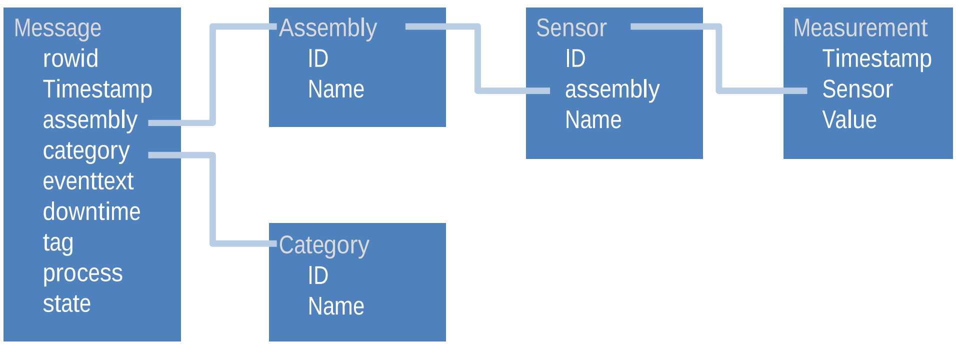

Figure 3.1: Database schema

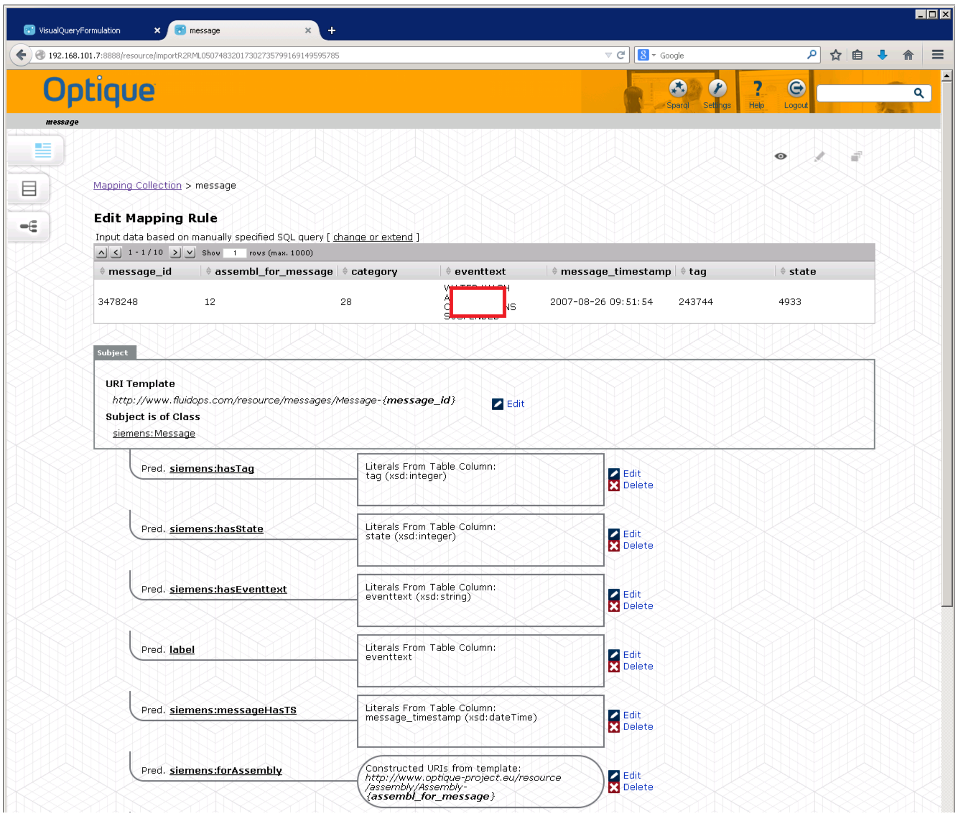

Figure 3.2: Optique mapping rule editor

discussion on the structure of the domain ontologies. The core of the ontology relevant to driving the front-end of the Optique Platform is listed in Appendix A.

3.3 Mappings

The main prerequisite for connecting the Optique Platform to the data sources at Siemens has been the provision of appropriate semantic mappings. These have been developed together with Fluid Operations and were deployed at Siemens. The full mappings are listed in Appendix B.

Apart from modifications necessary to accommodate the new data source the general structure of the mappings has not changed significantly. It should be stressed, however, that the changes in the mappings have not been made by hard-coding the sructure of the new data source into the Optique Platform at Siemens. Rather, Siemens has used the capabilities of the Optique front-end to simply edit the underlying mapping rules. Figure 3.2 shows the mapping rule editor for the mapping rule of event messages.

As an interesting detail, the lowest (rounded) box on the screen shows how messages are connected to assemblies via an object property (`siemens:forAssembly`) provided by the domain model while the table on top of the screen shows the database table from which event messages will be generated in the language of the domain model.

3.4 Anonymised data

Siemens turbine data such as the data accessed by the Optique Platform at Siemens is highly sensitive and classified as strictly confidential by Siemens security policies. That means that this data has to be hosted on secured servers without direct intranet or internet connection and hence cannot be shared with the Optique consortium. For the purpose of tool development and testing this poses a major challenge as the Optique partners need test input to help improve their algorithms and tools. In order to meet this need, Siemens has agreed to create an additional anonymised versions of their data sets exceeding the sizes of the anonymised data sets already provided within Year One.

Within Year Two Siemens has succeeded to generate anonymised data sets of substantially larger size than has been possible in Year One. Altougth the relevant data sets could not yet be sent to our project partners, Siemens plans to distribute about 1 Terabyte of test data to the Optique Consortium.

Chapter 4

Infrastructure and Installation

One of the main achievements of Year Two in the Optique Project has been the successful installation of the Optique Platform at Siemens. Just as the preliminary versions of Optique, the platform installed at Siemens is built up on the Information Workbench (IWB) [6], a software platform developed and maintained by fluid Operations.[1] The IWB is an open, data-centric development platform specifically designed to support the whole lifecyle of interacting with semantic data – from integration to access, visualization, exploration, and data interaction. Based on the architecture and language standards agreed in Deliverable D2.1, the Optique platform is designed to provide interfaces for data access. It provides a plugin-mechanisms for query language extensions as well as the actual front-end components. The Information Workbench already provides a number of extension points in order to automatically plug-in new modules.

4.1 The Optique Platform at Siemens

One of the main challenges for setting up Optique at Siemens has been posed by the highly secured IT infrastructure. All data generated by Siemens equipment represents highly confidential material. Accessing such material without authorisation is highly likely to harm both Siemens as a service provider and the customer using Siemens equipment. Therefore, data security has the highest priority in the Siemens Energy Services business field in general and in the Remote Diagnostic Centres where data is being held in particular.

Siemens has succeeded in allocating space on one of its servers in the de-militarised zone where Optique has been installed. More precisely, the Optique Platform is running on the host system the shape of a virtual Linux machine. The main advantages of virtualising the Optique Platform in this way are:

- *Independence*
  Setting up the Information Workbench driving the Optique Platform on a virtual machine requires no knowledge of the host system at Siemens. Hence the software developers at Fluid Operations had complete freedom to set up their system in the optimal way without being limited by restrictions of the operating system environment on the host system.
- *Maintainability*
  As the server hosting the Optique Platform is situated in a highly secured environment at Siemens, it has been impossible to grant Fluid Operations even limited access to the machine for maintenance purposes. With a virtual machine, however, Fluid Operations was in a position to carry out further developments of the platform on a copy of the virtual machine on their own premises. After finalisation of their build, the virtual machine could be simply be sent to Siemens and be deployed by a Siemens administrator with sufficient access privileges.
- *Security*
  Permitting an external software system into a highly secured environment necessitates tight control

[1] `http://www.fluidops.com/information-workbench/`

on the access rights of the relevant software system. This kind of access control is much easier to enforce and guarantee if the external system is wrapped up in a virtual machine. Hence, the solution via a virtual machine proved much easier to find an agreement on with respect to existing Siemens IT policies.

- *Development support*
  With the Optique Platform running at Siemens in the shape of a virtual machine it has been easy to store milestones of the platform while experimentally changing the setup. This means Siemens has been able to keep track not only of the incremental changes to the Information Workbench but of the setup of the entire solution at Siemens.

The machine chosen for hosting the Optique Platform at Siemens is a Windows Server with parameters as shown in Tables 4.1, 4.2, 4.3. Hence, the Optique Platform still uses a machine of the same computing power as used for the Year One prototype [8]. For more Details see the Primergy Data Sheet[1].

Table 4.1: Hardware of the server.

System Manufacturer	FUJITSU
System Model	PRIMERY RX600 S6
System Type	x64-based PC
Network Card(s)	5 NIC(s) installed

Table 4.2: Computing power of the server.

Processors:	4 Intel64 Processors installed Each: 8 Core $\sim$ 2000 MHz 18MB Cache

Table 4.3: Storage capabilities of the server.

RAM:	8 x 64 GB
Internal:	2 x 600 GB Disk 6 x 1000 GB Disk
External:	8 x 3000 GB Disk

It should be noted, however, that the server is not running Optique exclusively and hence it has not been possible to derive solid performance figures at Siemens. Nevertheless, the server has access to the highly secured turbine data sources at Siemens and is capable of providing data to the Optique Platform at sufficient speeds.

The main advances at the end of Year Two over the situation in Year One are:

- Access to real Siemens turbine data, both by retrospective querying and by real-time stream processing.
- Improved UI for proper visualisation of signal and event message data.
- R-plugin for sophisticated analysis of signals by means of R scripts.

In the following section details of the actual functionality of the Optique Platform will be described in more detail.

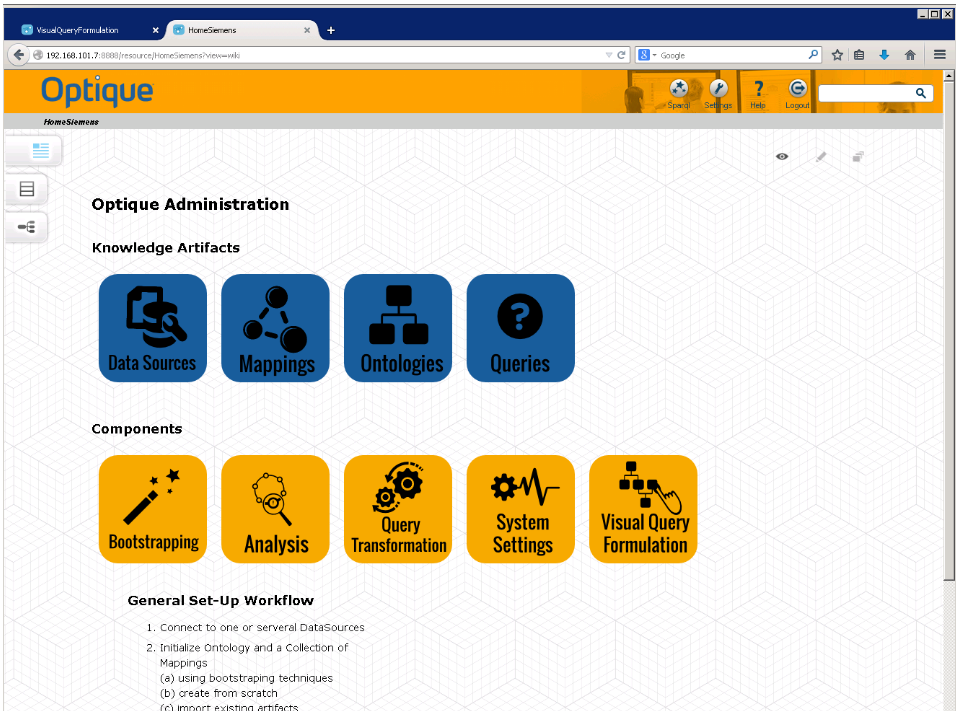

Figure 4.1: Optique Platform main screen

4.2 Optique's user front-end at Siemens

In Year Two of the Optique Project the Optique Platform has been successfully installed at Siemens, giving Optique access to real-world turbine data generated by actual Siemens gas and steam turbines. The provision of real-world data, however, was the starting point of further development of the platform. In contrast to experimental installations on test data, the Optique Platform at Siemens was intended to show data in such a way as to be meaningful to actual service engineers monitoring turbines. In order to achieve this goal, several UI elements had to be overhauled – especially in order to be more aligned with the way Siemens engineers are used to looking at the data. This will be discussed in more detail in the context of the machine dashboard below. We begin by taking a tour through the Optique installation starting from its main screen.

4.2.1 Main screen

The starting point of the Optique Platform at Siemens is the home screen shown in Figure 4.1. The main screen serves as entry point both for maintaining the data sources and mappings that supply Optique with information as well as actually accessing data by means of the visual query interface or a bespoke turbine dashboard.

The main funcionality provided by the home screen has not changed significantly. Under 'Data Sources' the user can add and maintain Optique's connection to the actual data. In the case of the Siemens instance of Optique, two data sources have been configured:

- A 'historic' data source proving the link to (one of) the Siemens turbine database(s)
- A 'real-time' data source capable of consuming live data streams

Under 'Mappings', the connection of the data sources to Optique's domain model are maintained. The domain model is provided in the form of an OWL ontology that can be registered to the platform under 'Ontologies'. Finally, the item 'Queries' leads to the query catalog in which archetypical queries are stored.

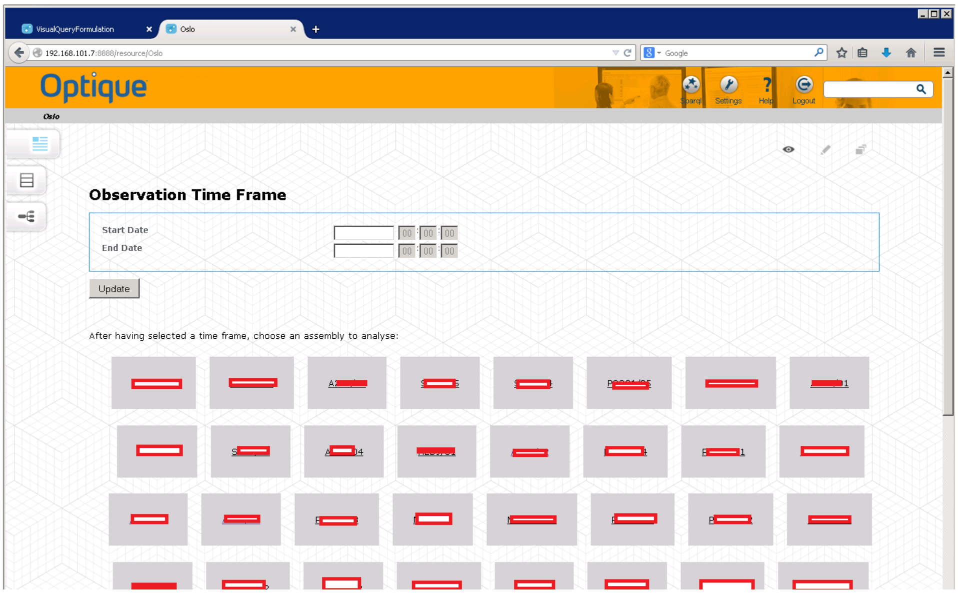

Figure 4.2: Optique Platform turbine overview

4.2.2 Turbine overview

In order to meet the needs of the service engineers at Siemens, the Year Two front-end has been extended by an additional page providing an overview over all the steam and gas turbines in scope, i.e., under the care of an individual engineer in terms of our use case. This view is shown in Figure 4.2. In this view, each box represents one steam or gas turbine by its ID[2] familiar to the engineer.

Above the turbine grid there are entry fields for the time period in which the relevant turbines are supposed to be analysed. We plan to use this information to provide a quick overview over the machine state by colour coding the boxes depending on the number of serious problems found in the relevant time period.

4.2.3 Turbine dashboard

Once the engineer clicks on one of the turbines in the overview she or he is taken to the turbine dashboard of the relevant machine. This is the central view in which all information relevant to the machine as a whole is being show. Figure 4.3 shows the turbine dashboard for one particular turbine.[3]

The dashboard for one machine starts by showing the reference time period for the data visualised further down. The main vier of the dashboard is a chart with sensor signals in the reference time period. In Figure 4.3, data is shown for the time from 1 to 2 December 2009. The majority of the signals depicted are temperature readings from various positions in the gas path of a gas turbine. The temperatures are mesured in ^{0}C. The chart shows that during the 24 hour period under consideration the machine has been running at about 700^{0}C to 800^{0}C for a short time, has then been shut down and been re-started about three hours later, only to continue running for the rest of the time.

The signal view does not only show numerical signal data but additionally indicates time points at which the control system of the turbine has sent event messages. These are depicted as boxes with "E" on the X-axis. As discussed in Section 2.3.1 the relationship between event messages and signal data is very important for the service engineer's ability to effectively find the root cause of a problem. By moving the mouse over one of the event message boxes in the chart the actual message text is displayed.

Below the signal view the turbine dashboard shows two views on event message data sent by the control

[2]The actual IDs of the turbines on display in Figure 4.2 had to be anonymised to meet the security restrictions imposed by Siemens Energy Services.

[3]The actual ID of the turbine, the IDs of the sensors and the contents of the message queue had to be anonymised in order to meet the security restrictions imposed by Siemens Energy Services.

system during the time period of interest: a chronological list on the left and a frequency analysis on the right. The frequency analysis shows which messages have been repeated most often in the time period of interest. Firstly, this quickly gives the engineer an impression of what kind of program the machine has been put on – i.e., have there been many start attempts, are their many mainenance messages. Secondly, if error or warning messages, even non-critical ones, occur very often during a short time period then this indicates that the machine needs to be investigated further for some persistent problems that have been registered by the control system to some extent without forcing the machine to stop operating.

The chronological list of event messages on the left gives the service engineer a finer grained account of what happened. In particular, it provides a starting point for re-constructing the course of events prior to emergency messages indicative of forced shut-downs or other major deviations from the planned load profile of the machine.

4.2.4 Sensor dashboard

In order to provide the service engineers with an even more fine-grained analysis tool beyond what is shown in the turbine dashboard, the Optique Platform has additional dashboards for each sensor of the turbine under consideration. For the example of a temperature sensor, such a single-sensor dashboard is depicted in Figure 4.4.[4]

The upper part of the single-sensor dashboard is structured similarly to the turbine dashboard: it shows the time period for which data is being displayed as well as a graph depicting the relevant sensor readings during that time period. Again, the occurrence of event messages is depicted by icons in the chart, giving the service engineer a feeling for the interaction between messages and signals. The actual event message can be displayed by moving the mouse over the relevant message indicator in the chart.

The remainder of the page provides deeper analytics into the sensor signal during the reference time period: below the raw signal plot an R-analytics[5] plugin shows the result of an outlier detection algorithm. Although the choice of actual R-script here should be viewed as preliminary the fact that arbitrary R analytics is possible within the Optique Platform at Siemens is highly significant for two reasons:

- For the service engineers in the Remote Diagnostic Centres R-scripts range among the standard tools for in-depth analysis. Some considerable time, resources, and training has been invested into the creation and use of R libraries that support engineers in their daily work. By accommodating R, the Optique Platform preserves the investment that has been made in R analytics and simplifies the transition to Optique from the perspective of the end-user of the system.

- The seamless integration of R into the Optique Platform shows that the plugin system of Optique is actually working as planned. Hence, by the end of Year Two Optique is alreay delivering on the promise that the platform can be integrated with pre-existing tools in daily use at Siemens. While the complete set of analytics tools to integrate into Optique is not finalised by the end of Year Two the fact a fully integrated plugin can already be demonstrated is a great achievement.

Below the R-plugin, under "Live Data", the single-sensor dashboard shows a live stream of signal data from the sensor under consideration. This is shown at the bottom of Figure 4.4: the live view just shows a live chart of the raw data stream, updated every second.

Smilarly to the situation with the R-plugin, the live data view is likely to change over the course of the next two years. Nevertheless, the view shows that Optique shows live data streams of actual Siemens turbine data. This highlights two major achievements at the end of Year Two:

- Access to data streams of real Siemens turbine data is provided within the IT infrastructure set up for the Optique Platform at Siemens.

[4] The actual ID of the sensor had to be anonymised in order to meet the security restrictions imposed by Siemens Energy Services.

[5] See http://www.r-project.org/

- Optique is capable of processing live data streams using STARQL and driving UI views with it.

As discussed above in the context of noisy data it is likely that more sophisticated processing of live data will be necessary in ordet to create tools that support service engineers in a meaningful way. Still, at the end of Year Two the fact that real data streams can already be processed and visualised using SPARQL is a major achievement arguably ahead of the project's planned timeline.

4.2.5 Visual query formulation

The visual query formulation tool of the Optique Platform has been discussed in Deliverable D8.1[8] already and the tool itself has remained largely unchanged in the Year Two prototype. Nevertheless, due to the successful connection of Optique to real Siemens turbine data the visual query formulation tool is now, for the first time, capable of supporting queries over real data.

In Figure 4.5 on the left an example of a visual query assembled in the tool is shown. The query asks for all assemblies that have generated a message of category "Shutdown". The respective result is shown on the right of the same Figure.[6]

Although other approaches [10] to user-friendly query formulation are being evaluated at Siemens we believe that the visual query formulation tool will give service engineers the ability to quickly change or add views onto their data without being distracted by the syntax of the query language generating the views. The visual query formulation tool has been a major part of the end-user workshop and will be evaluated for further development in WP1. See Deliverable D1.2 [9] for details.

[6] The actual ID of the turbine and the actual message texts had to be anonymised in order to meet the security restrictions imposed by Siemens Energy Services.

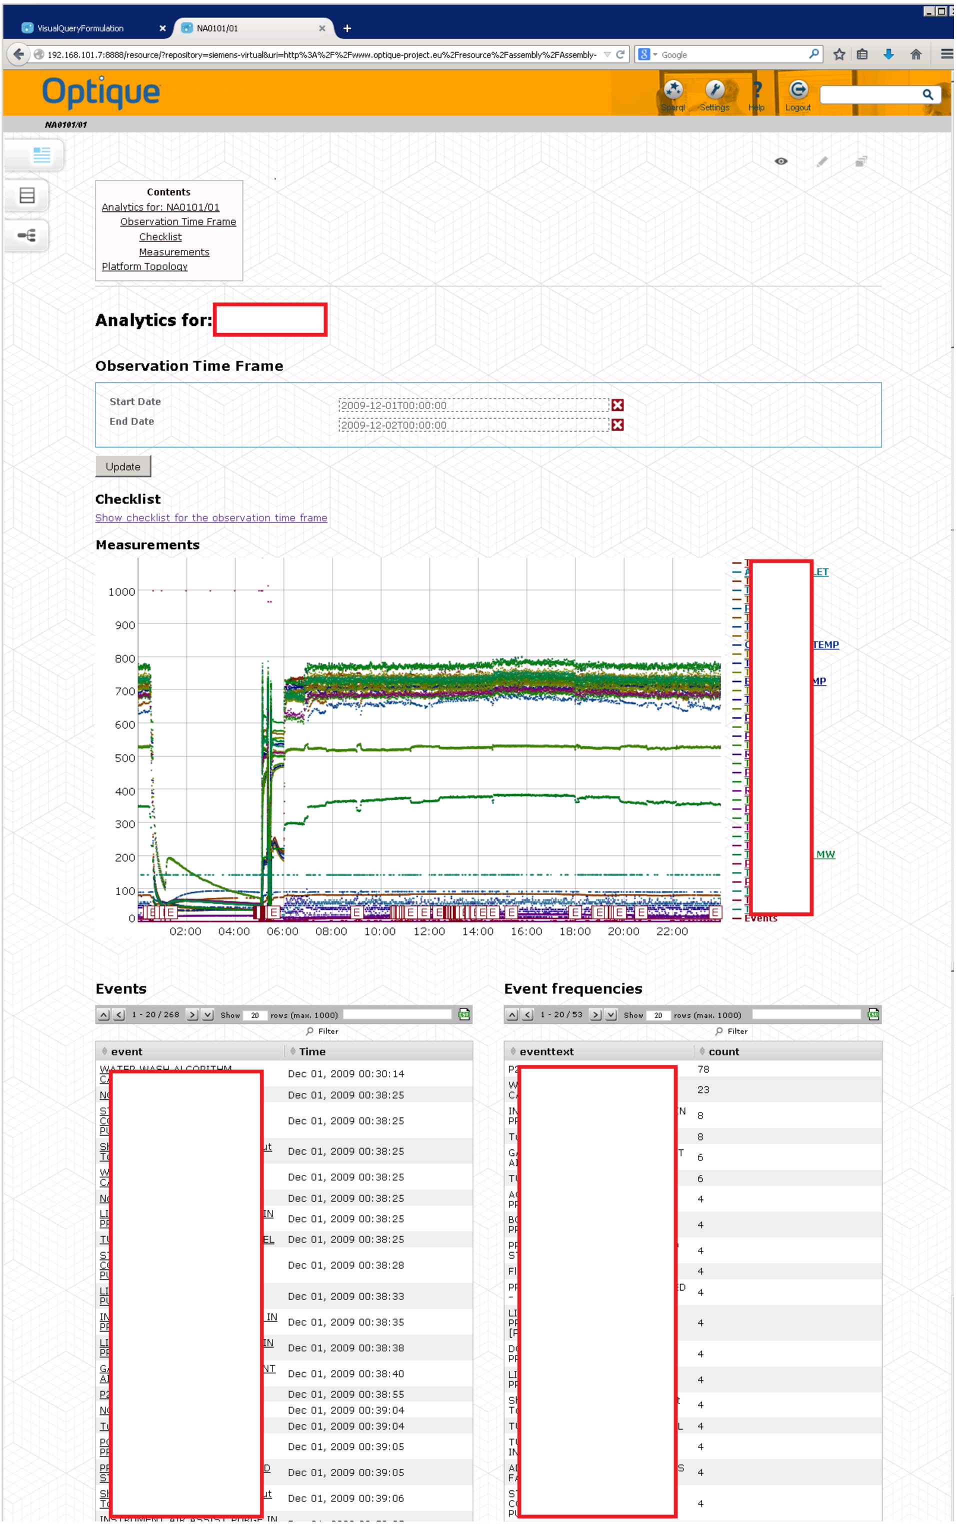

Figure 4.3: Optique Platform single-turbine dashboard

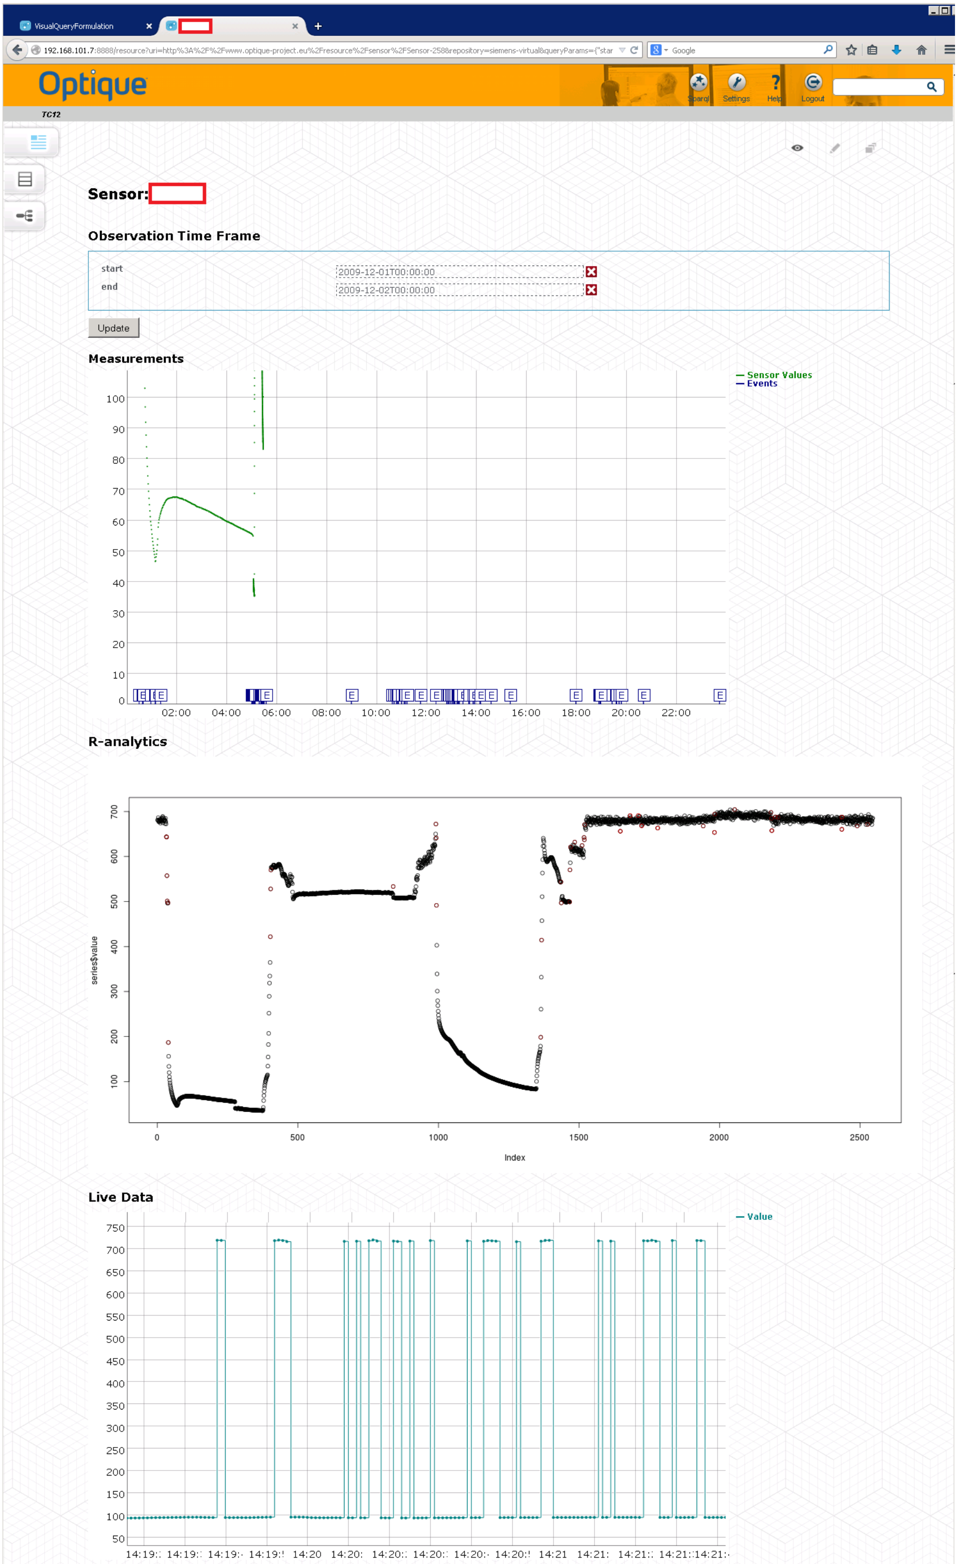

Figure 4.4: Optique Platform single-sensor dashboard

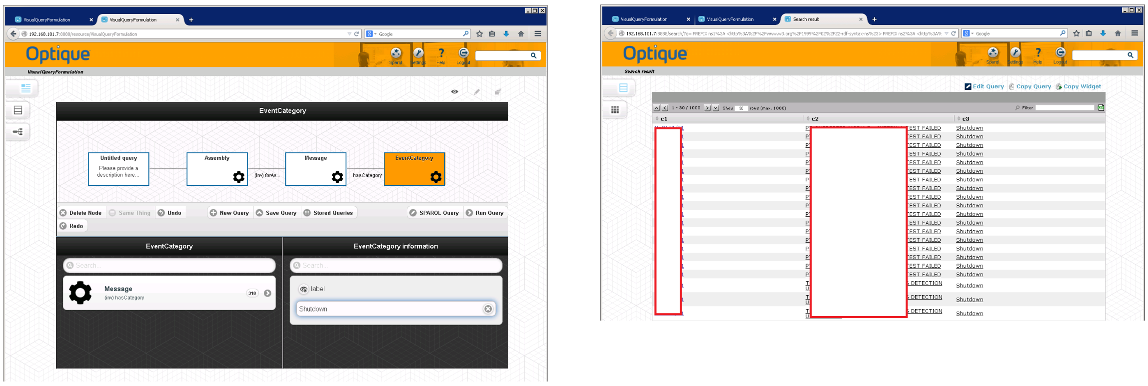

Figure 4.5: Optique Platform visual query formulation tool and results

Chapter 5

Conclusion

In the present document we have described the work carried out in WP8 of the Optique Project in Year Two. According to the DOW [3], the major milestones we set out to reach in the second year have been as follows:

1. Creating an improved instance of the Optique Platform inside of the IT infrastructure of Siemens;
2. Enabling the Optique Platform at Siemens to access to real turbine data from the Siemens Remote Diagnostic Centres; and
3. Having the Optique Platform evaluated by actual Siemens service engineers accessing real turbine data familiar to them.

At the end of Year Two these three major goals have been accomplished, as described in the chapters of the present report.

Several other achievements have been made with respect to data models and functionality of the Optique Platform at Siemens. The platform is able to consume live data streams of real turbine data using STARQL and process the data to drive views on the user-interface. Moreover, the plugin architecture of Optique has been utilised properly for the first time, implementing a module that carries out signal analytics by means of arbitrary R scripts. Just as the rest of the platform, the R analytics is carried out on real turbine data.

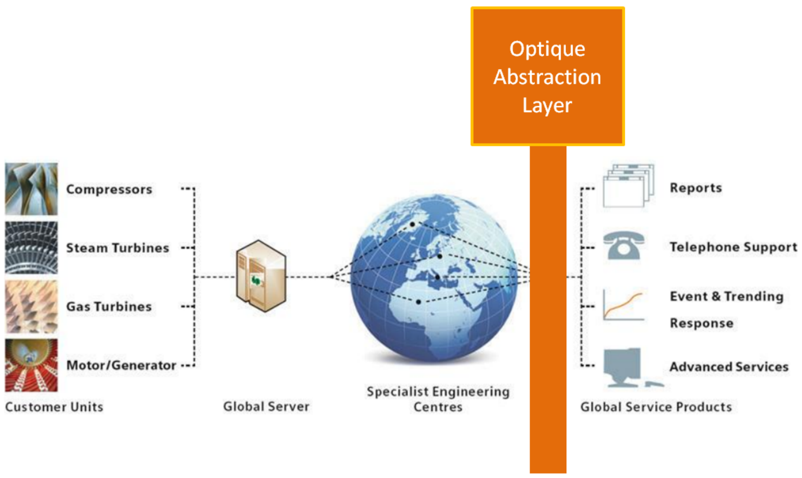

Figure 5.1: Remote Dianostic Services and Optique

Returning to the discussion in Section 2.3.5 Siemens plans to utilise Optique as an abstraction layer between the data sources feeding into the analytics workflow of the Remote Diagnostic Centres and the top-level analytics tools providing the relevant analytics solutions. In this way Siemens' expectation is that Optique will be the key to maintaining flexibility w.r.t. the underlying data sources as well as protecting the investment into the analytics tool set.

List of Figures

Bibliography

[1] Fujitsu primergy data sheet. Last visited October 18th, 2013: url = "http://www.fujitsu.com/fts/products/computing/servers/primergy/rack/rx600/".

[2] Meet the SGT-750. proven technology. perfected results, 2010.

[3] Optique - scalable end-user access to big data, grant agreement no. 318338, description of work (dow), June 4th 2012.

[4] Optique - scalable end-user access to big data, proposal, January 17th 2012.

[5] Remote diagnostic services for industrial gas turbines up to 15MW, 2013.

[6] fluidOps. Information workbench. Last visited October 18th, 2013: `"http://www.fluidops.com/information-workbench/"`.

[7] W. E. Forsthoffer. Forsthoffer's best practice handbook for rotating machinery, 2011.

[8] Thomas Hubauer, Steffen Lamparter, Christian Neuenstadt, Mikhail Roshchin, and Gerd Völksen. D8.1: Deliverable d8.1 interim siemens use case report 1. Deliverable, The Optique Project, 2013.

[9] Akrivi Katifori and Rudolf Schlatte. D1.2: Requirement analysis and evaluation framework. Deliverable, The Optique Project, 2014. To appear.

[10] Malte Sander, Ulli Waltinger, Michail Roshchin, and Thomas Runkler. Ontology-based translation of natural language queries to SPARQL. *In AAAI 2014 Fall Symposium*, 2014.

Appendix A

Domain ontology

```
<!DOCTYPE Ontology [
    <!ENTITY xsd "http://www.w3.org/2001/XMLSchema#" >
    <!ENTITY xml "http://www.w3.org/XML/1998/namespace" >
    <!ENTITY rdfs "http://www.w3.org/2000/01/rdf-schema#" >
    <!ENTITY rdf "http://www.w3.org/1999/02/22-rdf-syntax-ns#" >
]>

<Ontology xmlns="http://www.w3.org/2002/07/owl#"
     xml:base="http://www.siemens.com/demo/VisualQueryFormulation"
     xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
     xmlns:xsd="http://www.w3.org/2001/XMLSchema#"
     xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
     xmlns:xml="http://www.w3.org/XML/1998/namespace"
     ontologyIRI="http://www.siemens.com/demo/VisualQueryFormulation">
    <Prefix name="" IRI="http://www.semanticweb.org/owl/owlapi/turtle#"/>
    <Prefix name="owl" IRI="http://www.w3.org/2002/07/owl#"/>
    <Prefix name="rdf" IRI="http://www.w3.org/1999/02/22-rdf-syntax-ns#"/>
    <Prefix name="xsd" IRI="http://www.w3.org/2001/XMLSchema#"/>
    <Prefix name="rdfs" IRI="http://www.w3.org/2000/01/rdf-schema#"/>
    <Prefix name="annotations" IRI="http://eu.optique.ontology/annotations#"/>
    <Annotation>
        <AnnotationProperty abbreviatedIRI="rdfs:label"/>
        <Literal datatypeIRI="&rdf;PlainLiteral">Siemens Turbine Ontology
  (Visual Query Formulation)
</Literal>
    </Annotation>

    <Declaration>
        <Class IRI="http://www.siemens.com/demo#Assembly"/>
    </Declaration>

    <Declaration>
        <Class IRI="http://www.siemens.com/demo#EventCategory"/>
    </Declaration>

    <Declaration>
        <Class IRI="http://www.siemens.com/demo#Measurement"/>
```

```
</Declaration>

<Declaration>
    <Class IRI="http://www.siemens.com/demo#Message"/>
</Declaration>

<Declaration>
    <Class IRI="http://www.siemens.com/demo#Sensor"/>
</Declaration>

<Declaration>
    <ObjectProperty IRI="http://www.siemens.com/demo#forAssembly"/>
</Declaration>

<Declaration>
    <ObjectProperty IRI="http://www.siemens.com/demo#hasCategory"/>
</Declaration>

<Declaration>
    <ObjectProperty IRI="http://www.siemens.com/demo#inAssembly"/>
</Declaration>

<Declaration>
    <ObjectProperty IRI="http://www.siemens.com/demo#isRelatedTo"/>
</Declaration>

<Declaration>
    <ObjectProperty IRI="http://www.siemens.com/demo#measuredBy"/>
</Declaration>

<Declaration>
    <DataProperty IRI="http://www.siemens.com/demo#eventHasDescription"/>
</Declaration>

<Declaration>
    <DataProperty IRI="http://www.siemens.com/demo#hasEventtext"/>
</Declaration>

<Declaration>
    <DataProperty IRI="http://www.siemens.com/demo#hasState"/>
</Declaration>

<Declaration>
    <DataProperty IRI="http://www.siemens.com/demo#hasTag"/>
</Declaration>

<Declaration>
    <DataProperty IRI="http://www.siemens.com/demo#hasValue"/>
</Declaration>

<Declaration>
```

```
    <DataProperty IRI="http://www.siemens.com/demo#isProcess"/>
</Declaration>

<Declaration>
    <DataProperty IRI="http://www.siemens.com/demo#measurementHasTS"/>
</Declaration>

<Declaration>
    <DataProperty IRI="http://www.siemens.com/demo#messageHasTS"/>
</Declaration>

<Declaration>
    <DataProperty abbreviatedIRI="rdfs:label"/>
</Declaration>

<Declaration>
    <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
</Declaration>

<Declaration>
    <AnnotationProperty abbreviatedIRI="annotations:range_class"/>
</Declaration>

<ObjectPropertyDomain>
    <ObjectProperty IRI="http://www.siemens.com/demo#forAssembly"/>
    <Class IRI="http://www.siemens.com/demo#Message"/>
</ObjectPropertyDomain>

<ObjectPropertyDomain>
    <ObjectProperty IRI="http://www.siemens.com/demo#hasCategory"/>
    <Class IRI="http://www.siemens.com/demo#Message"/>
</ObjectPropertyDomain>

<ObjectPropertyDomain>
    <ObjectProperty IRI="http://www.siemens.com/demo#inAssembly"/>
    <Class IRI="http://www.siemens.com/demo#Sensor"/>
</ObjectPropertyDomain>

<ObjectPropertyDomain>
    <ObjectProperty IRI="http://www.siemens.com/demo#measuredBy"/>
    <Class IRI="http://www.siemens.com/demo#Measurement"/>
</ObjectPropertyDomain>

<ObjectPropertyRange>
    <ObjectProperty IRI="http://www.siemens.com/demo#forAssembly"/>
    <Class IRI="http://www.siemens.com/demo#Assembly"/>
</ObjectPropertyRange>

<ObjectPropertyRange>
    <ObjectProperty IRI="http://www.siemens.com/demo#hasCategory"/>
    <Class IRI="http://www.siemens.com/demo#EventCategory"/>
```

```
</ObjectPropertyRange>

<ObjectPropertyRange>
    <ObjectProperty IRI="http://www.siemens.com/demo#inAssembly"/>
    <Class IRI="http://www.siemens.com/demo#Assembly"/>
</ObjectPropertyRange>

<ObjectPropertyRange>
    <ObjectProperty IRI="http://www.siemens.com/demo#measuredBy"/>
    <Class IRI="http://www.siemens.com/demo#Sensor"/>
</ObjectPropertyRange>

<DataPropertyDomain>
    <DataProperty IRI="http://www.siemens.com/demo#eventHasDescription"/>
    <Class IRI="http://www.siemens.com/demo#Measurement"/>
</DataPropertyDomain>

<DataPropertyDomain>
    <DataProperty IRI="http://www.siemens.com/demo#hasEventtext"/>
    <Class IRI="http://www.siemens.com/demo#Measurement"/>
</DataPropertyDomain>

<DataPropertyDomain>
    <DataProperty IRI="http://www.siemens.com/demo#hasState"/>
    <Class IRI="http://www.siemens.com/demo#Measurement"/>
</DataPropertyDomain>

<DataPropertyDomain>
    <DataProperty IRI="http://www.siemens.com/demo#hasTag"/>
    <Class IRI="http://www.siemens.com/demo#Measurement"/>
</DataPropertyDomain>

<DataPropertyDomain>
    <DataProperty IRI="http://www.siemens.com/demo#hasValue"/>
    <Class IRI="http://www.siemens.com/demo#Measurement"/>
</DataPropertyDomain>

<DataPropertyDomain>
    <DataProperty IRI="http://www.siemens.com/demo#measurementHasTS"/>
    <Class IRI="http://www.siemens.com/demo#Measurement"/>
</DataPropertyDomain>

<DataPropertyDomain>
    <DataProperty IRI="http://www.siemens.com/demo#messageHasTS"/>
    <Class IRI="http://www.siemens.com/demo#Message"/>
</DataPropertyDomain>

<DataPropertyRange>
    <DataProperty IRI="http://www.siemens.com/demo#eventHasDescription"/>
    <Datatype abbreviatedIRI="xsd:string"/>
</DataPropertyRange>
```

```
<DataPropertyRange>
    <DataProperty IRI="http://www.siemens.com/demo#hasEventtext"/>
    <Datatype abbreviatedIRI="xsd:string"/>
</DataPropertyRange>

<DataPropertyRange>
    <DataProperty IRI="http://www.siemens.com/demo#hasState"/>
    <Datatype abbreviatedIRI="xsd:integer"/>
</DataPropertyRange>

<DataPropertyRange>
    <DataProperty IRI="http://www.siemens.com/demo#hasTag"/>
    <Datatype abbreviatedIRI="xsd:int"/>
</DataPropertyRange>

<DataPropertyRange>
    <DataProperty IRI="http://www.siemens.com/demo#hasValue"/>
    <Datatype abbreviatedIRI="xsd:double"/>
</DataPropertyRange>

<DataPropertyRange>
    <DataProperty IRI="http://www.siemens.com/demo#measurementHasTS"/>
    <Datatype abbreviatedIRI="xsd:dateTime"/>
</DataPropertyRange>

<DataPropertyRange>
    <DataProperty IRI="http://www.siemens.com/demo#messageHasTS"/>
    <Datatype abbreviatedIRI="xsd:dateTime"/>
</DataPropertyRange>

<DataPropertyRange>
    <DataProperty abbreviatedIRI="rdfs:label"/>
    <Datatype abbreviatedIRI="xsd:string"/>
</DataPropertyRange>

<AnnotationAssertion>
    <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
    <IRI>http://www.siemens.com/demo#eventHasDescription</IRI>
    <IRI>http://www.siemens.com/demo#Measurement</IRI>
</AnnotationAssertion>

<AnnotationAssertion>
    <AnnotationProperty abbreviatedIRI="annotations:range_class"/>
    <IRI>http://www.siemens.com/demo#eventHasDescription</IRI>
    <AbbreviatedIRI>xsd:string</AbbreviatedIRI>
</AnnotationAssertion>

<AnnotationAssertion>
    <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
    <IRI>http://www.siemens.com/demo#forAssembly</IRI>
```

```
        <IRI>http://www.siemens.com/demo#Message</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:range_class"/>
        <IRI>http://www.siemens.com/demo#forAssembly</IRI>
        <IRI>http://www.siemens.com/demo#Assembly</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="rdfs:comment"/>
        <IRI>http://www.siemens.com/demo#forAssembly</IRI>
        <Literal datatypeIRI="&rdf;PlainLiteral">
   Relation between message and assembly.
</Literal>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
        <IRI>http://www.siemens.com/demo#hasCategory</IRI>
        <IRI>http://www.siemens.com/demo#Message</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:range_class"/>
        <IRI>http://www.siemens.com/demo#hasCategory</IRI>
        <IRI>http://www.siemens.com/demo#EventCategory</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="rdfs:comment"/>
        <IRI>http://www.siemens.com/demo#hasCategory</IRI>
        <Literal datatypeIRI="&rdf;PlainLiteral">
   Relation between message and event category.
</Literal>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
        <IRI>http://www.siemens.com/demo#hasEventtext</IRI>
        <IRI>http://www.siemens.com/demo#Measurement</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:range_class"/>
        <IRI>http://www.siemens.com/demo#hasEventtext</IRI>
        <AbbreviatedIRI>xsd:string</AbbreviatedIRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
```

```
        <IRI>http://www.siemens.com/demo#hasState</IRI>
        <IRI>http://www.siemens.com/demo#Measurement</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:range_class"/>
        <IRI>http://www.siemens.com/demo#hasState</IRI>
        <AbbreviatedIRI>xsd:integer</AbbreviatedIRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
        <IRI>http://www.siemens.com/demo#hasTag</IRI>
        <IRI>http://www.siemens.com/demo#Measurement</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:range_class"/>
        <IRI>http://www.siemens.com/demo#hasTag</IRI>
        <AbbreviatedIRI>xsd:int</AbbreviatedIRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
        <IRI>http://www.siemens.com/demo#hasValue</IRI>
        <IRI>http://www.siemens.com/demo#Measurement</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:range_class"/>
        <IRI>http://www.siemens.com/demo#hasValue</IRI>
        <AbbreviatedIRI>xsd:double</AbbreviatedIRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
        <IRI>http://www.siemens.com/demo#inAssembly</IRI>
        <IRI>http://www.siemens.com/demo#Sensor</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:range_class"/>
        <IRI>http://www.siemens.com/demo#inAssembly</IRI>
        <IRI>http://www.siemens.com/demo#Assembly</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="rdfs:comment"/>
        <IRI>http://www.siemens.com/demo#inAssembly</IRI>
        <Literal datatypeIRI="&rdf;PlainLiteral">
  Relation between sensor and assembly.
```

```
</Literal>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
        <IRI>http://www.siemens.com/demo#measuredBy</IRI>
        <IRI>http://www.siemens.com/demo#Measurement</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:range_class"/>
        <IRI>http://www.siemens.com/demo#measuredBy</IRI>
        <IRI>http://www.siemens.com/demo#Sensor</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="rdfs:comment"/>
        <IRI>http://www.siemens.com/demo#measuredBy</IRI>
        <Literal datatypeIRI="&rdf;PlainLiteral">
   Relation between measurement and sensor.
</Literal>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
        <IRI>http://www.siemens.com/demo#measurementHasTS</IRI>
        <IRI>http://www.siemens.com/demo#Measurement</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:range_class"/>
        <IRI>http://www.siemens.com/demo#measurementHasTS</IRI>
        <AbbreviatedIRI>xsd:dateTime</AbbreviatedIRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
        <IRI>http://www.siemens.com/demo#messageHasTS</IRI>
        <IRI>http://www.siemens.com/demo#Message</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:range_class"/>
        <IRI>http://www.siemens.com/demo#messageHasTS</IRI>
        <AbbreviatedIRI>xsd:dateTime</AbbreviatedIRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
        <AbbreviatedIRI>rdfs:label</AbbreviatedIRI>
        <IRI>http://www.siemens.com/demo#Assembly</IRI>
```

```
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
        <AbbreviatedIRI>rdfs:label</AbbreviatedIRI>
        <IRI>http://www.siemens.com/demo#EventCategory</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
        <AbbreviatedIRI>rdfs:label</AbbreviatedIRI>
        <IRI>http://www.siemens.com/demo#Measurement</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
        <AbbreviatedIRI>rdfs:label</AbbreviatedIRI>
        <IRI>http://www.siemens.com/demo#Message</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:domain_class"/>
        <AbbreviatedIRI>rdfs:label</AbbreviatedIRI>
        <IRI>http://www.siemens.com/demo#Sensor</IRI>
    </AnnotationAssertion>

    <AnnotationAssertion>
        <AnnotationProperty abbreviatedIRI="annotations:range_class"/>
        <AbbreviatedIRI>rdfs:label</AbbreviatedIRI>
        <AbbreviatedIRI>xsd:string</AbbreviatedIRI>
    </AnnotationAssertion>

</Ontology>
```

Appendix B

Mappings to Siemens Database Schema

B.1 Database mapping

```
f@prefix optique: <http://www.optique-project.eu/resource/> .
@prefix r2rml: <http://www.w3.org/ns/r2rml#> .
@prefix siemens: <http://www.siemens.com/Optique/OptiquePattern#> .

optique:importR2RML04903171478346635248131217121904 a r2rml:PredicateMap ;
    r2rml:constant <http://www.siemens.com/demo#measurementHasTS> .

optique:importR2RML05074832017302735799169149595785 a r2rml:TriplesMap ;
    rdfs:label "message" ;
    r2rml:logicalTable optique:logicalTable64458945986892970289404351385074 ;
    r2rml:predicateObjectMap optique:importR2RML67142100850026223055186190059032 , optique:importR
    r2rml:subjectMap optique:subjectMap94817192609138823992111066019080 .

optique:importR2RML10485079562877503657201870964972 a r2rml:ObjectMap ;
    r2rml:column "eventtext" ;
    r2rml:datatype xsd:string .

optique:importR2RML17706402239273390100661426096200 a r2rml:TriplesMap ;
    rdfs:label "measurement" ;
    r2rml:logicalTable optique:logicalTable64457776764972067866884407984625 ;
    r2rml:predicateObjectMap optique:importR2RML33335013478798341412147703532695 , optique:importR
    r2rml:subjectMap optique:subjectMap51385798095651157762782890726693 .

optique:importR2RML30694980870683629511519676100800 a r2rml:PredicateMap ;
    r2rml:constant <http://www.siemens.com/demo#hasTag> .

optique:importR2RML33335013478798341412147703532695 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:importR2RML95638677821349335037686026554578 ;
    r2rml:predicateMap optique:importR2RML77751073146786491942287527471268 .

optique:importR2RML49442555401925497376934801730702 a r2rml:PredicateMap ;
    r2rml:constant <http://www.siemens.com/demo#hasState> .

optique:importR2RML53661422957751077440772451409673 a r2rml:TriplesMap ;
    rdfs:label "eventCategory" ;
```

```
    r2rml:logicalTable optique:logicalTable1944139948654983569291919l627146 ;
    r2rml:predicateObjectMap optique:predicateObjectMap6049981325761343227326713l736379 ;
    r2rml:subjectMap optique:subjectMap5489612977556840035377l447451358 .

optique:importR2RML6708505047495978l645188176815090 a r2rml:ObjectMap ;
    r2rml:column "state" ;
    r2rml:datatype xsd:integer .

optique:importR2RML67142100850026223055186190059032 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:importR2RML70783571822261235781103279277377 ;
    r2rml:predicateMap optique:importR2RML30694980870683629511519676100800 .

optique:importR2RML70783571822261235781103279277377 a r2rml:ObjectMap ;
    r2rml:column "tag" ;
    r2rml:datatype xsd:integer .

optique:importR2RML75255124316443167413378634326653 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:importR2RML67085050474959781645188176815090 ;
    r2rml:predicateMap optique:importR2RML49442555401925497376934801730702 .

optique:importR2RML77751073146786491942287527471268 a r2rml:PredicateMap ;
    r2rml:constant <http://www.siemens.com/demo#hasValue> .

optique:importR2RML78703408784487444181498169631378 a r2rml:PredicateMap ;
    r2rml:constant <http://www.siemens.com/demo#hasEventtext> .

optique:importR2RML81068643265870846875604137380300 a r2rml:TriplesMap ;
    rdfs:label "sensor" ;
    r2rml:logicalTable optique:logicalTable20356428278882997466118707281926 ;
    r2rml:predicateObjectMap optique:predicateObjectMap50231424501398728316229291465162 , optique:
    r2rml:subjectMap optique:subjectMap12239008720520167692604304143485 .

optique:importR2RML85176285385530499735680629570739 a r2rml:TriplesMap ;
    rdfs:label "assembly" ;
    r2rml:logicalTable optique:logicalTable20281130462567802259621964533198 ;
    r2rml:predicateObjectMap optique:predicateObjectMap44771278385369536572689423566624 ;
    r2rml:subjectMap optique:subjectMap85402635829079525126330582779882 .

optique:importR2RML86037211015979910522847644222085 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:importR2RML93342043559907900931588431330303 ;
    r2rml:predicateMap optique:importR2RML04903171478346635248131217121904 .

optique:importR2RML86096462369920387213827052832382 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:importR2RML10485079562877503657201870964972 ;
    r2rml:predicateMap optique:importR2RML78703408784487444181498169631378 .

optique:importR2RML93342043559907900931588431330303 a r2rml:ObjectMap ;
    r2rml:column "timestamp" ;
    r2rml:datatype xsd:dateTime .

optique:importR2RML95638677821349335037686026554578 a r2rml:ObjectMap ;
```

```
    r2rml:column "value" ;
    r2rml:datatype xsd:double .

optique:logicalTable1944139948654983569291919l627146 a r2rml:LogicalTable , r2rml:R2RMLView ;
    r2rml:sqlQuery "SELECT category.id as category_id, category.name as category_name FROM categor

optique:logicalTable2028113046256780225962l964533198 a r2rml:LogicalTable , r2rml:R2RMLView ;
    r2rml:sqlQuery "SELECT assembly.id as assembly_id, assembly.name FROM assembly" .

optique:logicalTable2035642827888299746611870728l926 a r2rml:LogicalTable , r2rml:R2RMLView ;
    r2rml:sqlQuery "SELECT sensor.id as sensor_id, sensor.name as sensor_name, sensor.assembly as

optique:logicalTable6445777676497206786688440798462 5 a r2rml:LogicalTable , r2rml:R2RMLView ;
    r2rml:sqlQuery "SELECT measurement.\"Timestamp\" as timestamp, measurement.sensor, measurement

optique:logicalTable6445894598689297028940435138507 4 a r2rml:LogicalTable , r2rml:R2RMLView ;
    r2rml:sqlQuery "SELECT message.rowid as message_id, message.assembly as assembl_for_message, m

optique:objectMap04179249561677750368993325062653 a r2rml:ObjectMap ;
    r2rml:column "category_description" ;
    r2rml:datatype xsd:string .

optique:objectMap10838686251506810791145385634210 a r2rml:ObjectMap ;
    r2rml:column "name" ;
    r2rml:datatype xsd:string .

optique:objectMap11216121966928262775793743694590 a r2rml:ObjectMap ;
    r2rml:column "category_name" ;
    r2rml:datatype xsd:string .

optique:objectMap11649717299821629295634278350586 a r2rml:ObjectMap ;
    r2rml:template "http://www.optique-project.eu/resource/assembly/Assembly-{assembl_for_message}

optique:objectMap13929097951947055459564363368895 a r2rml:ObjectMap ;
    r2rml:template "http://www.optique-project.eu/resource/sensor/Sensor-{sensor}" .

optique:objectMap31244770072524483855964834755217 a r2rml:ObjectMap ;
    r2rml:column "sensor_name" ;
    r2rml:datatype xsd:string .

optique:objectMap38681559675275180176298138987383 a r2rml:ObjectMap ;
    r2rml:column "eventtext" .

optique:objectMap63515287759293689673922993696610 a r2rml:ObjectMap ;
    r2rml:column "message_timestamp" ;
    r2rml:datatype xsd:dateTime .

optique:objectMap76539671892003926933942949451270 a r2rml:ObjectMap ;
    r2rml:template "http://www.optique-project.eu/resource/assembly/Assembly-{sensor_in_assembly}"

optique:objectMap89407633019517956477381274582512 a r2rml:ObjectMap ;
```

```
    r2rml:template "http://www.optique-project.eu/resource/event-category/EventCategory-{category}

optique:predicateMap00587622061915359082894761291530 a r2rml:PredicateMap ;
    r2rml:constant rdfs:comment .

optique:predicateMap01398155043244776277673188150034 a r2rml:PredicateMap ;
    r2rml:constant rdfs:label .

optique:predicateMap03143796755943562713247923771494 a r2rml:PredicateMap ;
    r2rml:constant <http://www.siemens.com/demo#forAssembly> .

optique:predicateMap07991435530056532772681473670364 a r2rml:PredicateMap ;
    r2rml:constant rdfs:label .

optique:predicateMap23292667840648295048665397431255 a r2rml:PredicateMap ;
    r2rml:constant <http://www.siemens.com/demo#eventHasDescription> .

optique:predicateMap36359402384916955151547031430255 a r2rml:PredicateMap ;
    r2rml:constant <http://www.siemens.com/demo#hasCategory> .

optique:predicateMap58128559008240293679312840123623 a r2rml:PredicateMap ;
    r2rml:constant rdfs:label .

optique:predicateMap58416353754643802313182050377264 a r2rml:PredicateMap ;
    r2rml:constant <http://www.siemens.com/demo#messageHasTS> .

optique:predicateMap62601684015963192473193950724981 a r2rml:PredicateMap ;
    r2rml:constant rdfs:label .

optique:predicateMap71209030473479238047570576874631 a r2rml:PredicateMap ;
    r2rml:constant <http://www.siemens.com/demo#measuredBy> .

optique:predicateMap93437950130131464529973218081766 a r2rml:PredicateMap ;
    r2rml:constant <http://www.siemens.com/demo#inAssembly> .

optique:predicateObjectMap05115047701679010179567157120265 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:objectMap04179249561677750368993325062653 ;
    r2rml:predicateMap optique:predicateMap23292667840648295048665397431255 .

optique:predicateObjectMap14829871717521904903630719737367 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:objectMap81914496837550063583680737006619 ;
    r2rml:predicateMap optique:predicateMap00587622061915359082894761291530 .

optique:predicateObjectMap15752172423069957911233762954556 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:objectMap38681559675275180176298138987383 ;
    r2rml:predicateMap optique:predicateMap62601684015963192473193950724981 .

optique:predicateObjectMap33663009323695180813290466055863 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:objectMap63515287759293689673922993696610 ;
    r2rml:predicateMap optique:predicateMap58416353754643802313182050377264 .
```

```
optique:predicateObjectMap35774058489099100530251854040244 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:objectMap11649717299821629295634278350586 ;
    r2rml:predicateMap optique:predicateMap03143796755943562713247923771494 .

optique:predicateObjectMap44771278385369536572689423566624 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:objectMap10838686251506810791145385634210 ;
    r2rml:predicateMap optique:predicateMap01398155043244776277673188150034 .

optique:predicateObjectMap46160380175336377940066105954890 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:objectMap13929097951947055459564363368895 ;
    r2rml:predicateMap optique:predicateMap71209030473479238047570576874631 .

optique:predicateObjectMap50231424501398728316229291465162 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:objectMap76539671892003926933942949451270 ;
    r2rml:predicateMap optique:predicateMap93437950130131464529973218081766 .

optique:predicateObjectMap53966808549313303369022103294534 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:objectMap31244770072524483855964834755217 ;
    r2rml:predicateMap optique:predicateMap58128559008240293679312840123623 .

optique:predicateObjectMap59534346283284707771215352003463 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:objectMap89407633019517956477381274582512 ;
    r2rml:predicateMap optique:predicateMap36359402384916955151547031430255 .

optique:predicateObjectMap60499813257613432273267131736379 a r2rml:PredicateObjectMap ;
    r2rml:objectMap optique:objectMap11216121966928262775793743694590 ;
    r2rml:predicateMap optique:predicateMap07991435530056532772681473670364 .

optique:subjectMap12239008720520167692604304143485 a r2rml:SubjectMap ;
    r2rml:template "http://www.optique-project.eu/resource/sensor/Sensor-{sensor_id}" ;
    r2rml:class <http://www.siemens.com/demo#Sensor> .

optique:subjectMap51385798095651157762782890726693 a r2rml:SubjectMap ;
    r2rml:template "http://www.optique-project.eu/resource/measurement
/Measurement-{timestamp}-{sensor}" ;
    r2rml:class <http://www.siemens.com/demo#Measurement> .

optique:subjectMap54896129775568400353771447451358 a r2rml:SubjectMap ;
    r2rml:template "http://www.optique-project.eu/resource/event-category
/EventCategory-{category_id}" ;
    r2rml:class <http://www.siemens.com/demo#EventCategory> .

optique:subjectMap85402635829079525126330582779882 a r2rml:SubjectMap ;
    r2rml:template "http://www.optique-project.eu/resource/assembly
/Assembly-{assembly_id}" ;
    r2rml:class <http://www.siemens.com/demo#Assembly> .

optique:subjectMap94817192609138823992111066019080 a r2rml:SubjectMap ;
    r2rml:template "http://www.optique-project.eu/resource/messages
/Message-{message_id}" ;
    r2rml:class <http://www.siemens.com/demo#Message> .
```

B.2 Stream mapping

```
@prefix r2rml: <http://www.w3.org/ns/r2rml#> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix siemens: <http://www.siemens.com/Optique/OptiquePattern#> .

<http://www.optique-project.eu/resource/hasVal> a r2rml:TriplesMap ;
    r2rml:logicalTable [ a r2rml:R2RMLView ;
        r2rml:sqlQuery "SELECT \"value\", sensor FROM Measurements"
    ] .

<http://www.optique-project.eu/resource/hasVal> r2rml:predicateObjectMap [
        r2rml:predicate siemens:hasVal ;
        r2rml:objectMap [
            r2rml:template "{value}"
        ]
    ] .

<http://www.optique-project.eu/resource/hasVal> r2rml:subjectMap [ a r2rml:TermMap ;
        r2rml:template "http://www.siemens.com/Optique/OptiquePattern#{sensor}"
    ] .

<http://www.optique-project.eu/resource/hasAssembly> a r2rml:TriplesMap ;
    r2rml:logicalTable [ a r2rml:R2RMLView ;
        r2rml:sqlQuery "SELECT assembly, sensor FROM Measurements"
    ] .

<http://www.optique-project.eu/resource/hasAssembly> r2rml:predicateObjectMap [
        r2rml:predicate siemens:hasAss ;
        r2rml:objectMap [
            r2rml:template "http://www.siemens.com/Optique/OptiquePattern#{assembly}"
        ]
    ] .

<http://www.optique-project.eu/resource/hasAssembly> r2rml:subjectMap [ a r2rml:TermMap ;
        r2rml:template "http://www.siemens.com/Optique/OptiquePattern#{sensor}"
    ] .

<http://www.optique-project.eu/resource/hasEvent> a r2rml:TriplesMap ;
    r2rml:logicalTable [ a r2rml:R2RMLView ;
        r2rml:sqlQuery "SELECT assembly, eventtext FROM Messages"
    ] .

<http://www.optique-project.eu/resource/hasEvent> r2rml:predicateObjectMap [
        r2rml:predicate siemens:hasEvent ;
        r2rml:objectMap [
            r2rml:template "{eventtext}"
        ]
    ] .

<http://www.optique-project.eu/resource/hasEvent> r2rml:subjectMap [ a r2rml:TermMap ;
```

```
        r2rml:template "http://www.siemens.com/Optique/OptiquePattern#{assembly}"
    ] .

<http://www.optique-project.eu/resource/partOf> a r2rml:TriplesMap ;
    r2rml:logicalTable [ a r2rml:R2RMLView ;
        r2rml:sqlQuery "SELECT sensor, assembly FROM Measurements"
    ] .

<http://www.optique-project.eu/resource/partOf> r2rml:predicateObjectMap [
        r2rml:predicate siemens:partOf ;
        r2rml:objectMap [
            r2rml:template "http://www.siemens.com/Optique/OptiquePattern#{assembly}"
        ]
    ] .

<http://www.optique-project.eu/resource/partOf> r2rml:subjectMap [ a r2rml:TermMap ;
        r2rml:template "http://www.siemens.com/Optique/OptiquePattern#{sensor}"
    ] .

<http://www.optique-project.eu/resource/removeDueToSensor> a r2rml:TriplesMap ;
    r2rml:logicalTable [ a r2rml:R2RMLView ;
        r2rml:sqlQuery "SELECT sensor, assembly FROM Measurements"
    ] .

<http://www.optique-project.eu/resource/removeDueToSensor> r2rml:predicateObjectMap [
        r2rml:predicate siemens:removeDueToSensor ;
        r2rml:objectMap [
            r2rml:template "http://www.siemens.com/Optique/OptiquePattern#{sensor}"
        ]
    ] .

<http://www.optique-project.eu/resource/removeDueToSensor> r2rml:subjectMap [ a r2rml:TermMap ;
        r2rml:template "http://www.siemens.com/Optique/OptiquePattern#{assembly}"
    ] .

<http://www.optique-project.eu/resource/Sensor> a r2rml:TriplesMap ;
    r2rml:logicalTable [ a r2rml:R2RMLView ;
        r2rml:sqlQuery "SELECT DISTINCT sensor FROM Measurements"
    ] .

<http://www.optique-project.eu/resource/Sensor> r2rml:subjectMap [ a r2rml:TermMap ;
        r2rml:template "http://www.siemens.com/Optique/OptiquePattern#{sensor}" ;
        r2rml:class siemens:Sensor
    ] .

<http://www.optique-project.eu/resource/Critical> a r2rml:TriplesMap ;
    r2rml:logicalTable [ a r2rml:R2RMLView ;
        r2rml:sqlQuery "SELECT \"value\", sensor FROM Measurements WHERE \"value\" > 80"
    ] .

<http://www.optique-project.eu/resource/Critical> r2rml:predicateObjectMap [
```

```
        r2rml:predicate siemens:hasVal ;
        r2rml:objectMap [
            r2rml:template "{value}"
        ]
    ] .

<http://www.optique-project.eu/resource/Critical> r2rml:subjectMap [ a r2rml:TermMap ;
        r2rml:template "http://www.siemens.com/Optique/OptiquePattern#{sensor}" ;
        r2rml:class siemens:Critical
    ] .

<http://www.optique-project.eu/resource/Assembly> a r2rml:TriplesMap ;
    r2rml:logicalTable [ a r2rml:R2RMLView ;
        r2rml:sqlQuery "SELECT DISTINCT assembly FROM Measurements"
    ] .

<http://www.optique-project.eu/resource/Assembly> r2rml:subjectMap [ a r2rml:TermMap ;
        r2rml:template "http://www.siemens.com/Optique/OptiquePattern#{assembly}" ;
        r2rml:class siemens:Assembly
    ] .

<http://www.optique-project.eu/resource/TempSensor> a r2rml:TriplesMap ;
    r2rml:logicalTable [ a r2rml:R2RMLView ;
        r2rml:sqlQuery "SELECT domain FROM a_static WHERE range = 'TempSensor'"
    ] .

<http://www.optique-project.eu/resource/TempSensor> r2rml:subjectMap [ a r2rml:TermMap ;
        r2rml:template "http://www.siemens.com/Optique/OptiquePattern#{domain}" ;
        r2rml:class siemens:TempSensor
    ] .
```