# Grid Meets Sensors, Sensors Meet Grid

F. Lelli[1,2], G. Maron[1], S. Orlando[2] and S. Pinter[3]
[1]*Istituto Nazionale di Fisica Nucleare, Laboratori Nazionali di Legnaro, ITALY*
[2]*Dipartimento di Informatica, Università Ca' Foscari di Venezia, ITALY*
[3]*IBM, Haifa Research Lab, Haifa University Campus, ISRAEL*

*Abstract:* - Modern Grid systems use web services to interact with remote and distributed subsystem components to achieve an effective and common connection. Next generation Grid solutions will include instruments and sensors as part of the Grid, in addition to computational units (CE) and disk storage (SE). The Instrument Element (IE) is a new component that provides access to these remotely distributed pieces of hardware. The IE can organize generic devices in a "Grid of Instruments", providing a set of facilities to allow interaction between sensors and the existing classical grid infrastructure. In this paper, we describe the architecture of an Instrument Element system and present a collection of applications that successfully integrate and utilize this framework.

*Key-Words:* - Sensors**,** Grid, instrument, Web Services, Instrument Element, Remote Control

## 1 Introduction

Grid computing refers to the coordinated and secured sharing of computing nodes among dynamic collections of individuals, institutions, and resources. It involves the distribution of computing resources among geographically separate sites (creating a 'grid' of computers), all of which are configured with specialized software for routing jobs, authenticating users, monitoring resources, and so on. Shared, site-based computing resources may include computing and/or storage nodes, software, data, and various scientific instruments. The aim is to provide reliable and secure access to widely scattered resources for authorized users located virtually anywhere in the world. In Grid terminology, the words 'sensor', 'instrument', and 'device' are used to define a piece of equipment that needs to be initialized, configured, operated (start, stop, standby, resume, application specific commands), monitored, or reset. While remote control of, and data collection from, instruments and sensors was part of the initial Grid concept, most recent Grid implementations concentrate on sharing distributed computational and storage resources.

The Grid Enabled Remote Instrumentation with Distributed Control and Computation (GridCC) project [1] provides extensions to include access to and control of distributed instruments and sensors. In other words, the goal of the GridCC project is to build a widely distributed system that is able to remotely control and monitor complex devices. Instruments can range from a set of sensors used by geophysical stations monitoring the state of the earth to a network of small power generators supplying the European power grid [2] and more.

The Instrument Element (IE) component of the GridCC provides a set of services that permit interaction with remote instruments. In this paper, we provide details of the IE architecture and services, and describe several real-life use-cases that are utilizing our first implementation.

The rest of this document is organized as follows. In Section 2, we present and discuss the requirements needed for interaction between computational Grid and sensors. In Section 3, we provide details of the IE component architecture and services. Finally, in Section 4, we present a few use-cases that utilize the IE component.

## 2 Interactions with Sensors in the Grid Framework

Operation and maintenance of different sensors and instruments in the grid share common needs, such as the possibility of being accessed and monitored remotely, which in turn introduce requirements for real-time or timely interactive operation with computing Grid nodes to process the collected information. We believe that creating a coherent collection of services to allow remote configuration, partitioning, control, and tuning a set of physical sensors is necessary for their effective integration in the computational Grid. We call this set of services **Instrument Element (IE).**

The IE component has a set of Web Service [3], [4], [5], [6] interfaces called Virtual Instrument Grid Service (VIGS) that allows the user to remotely access the real instruments, thus plugging the system itself into the Grid. Web Services are an excellent choice when there is a need to provide a common language for cross-domain collaboration and for hiding the internal implementation details of accessing specific sensors. Standards like SWE [11], JMX [13], OGSI [18], WSRF [26], or IVI [12] and P2P systems like JINI [16], Freenet [14], [15], or JAXTA [17] have been analyzed as candidates for the VIGS; specifically, we tested that the VIGS methods provide generic device virtualization.

The IE provides facilities for interactive cooperation between computing Grid nodes to applications that have real-time requirements or need fast interaction with CEs and SEs. In addition, this component can be linked to existing sensor instrumentation to provide Grid interaction and remote control and tuning of standalone resources. Figure 1 provides a high-level view of the relationship between the IE and its users, and between the IE and other Grid components.

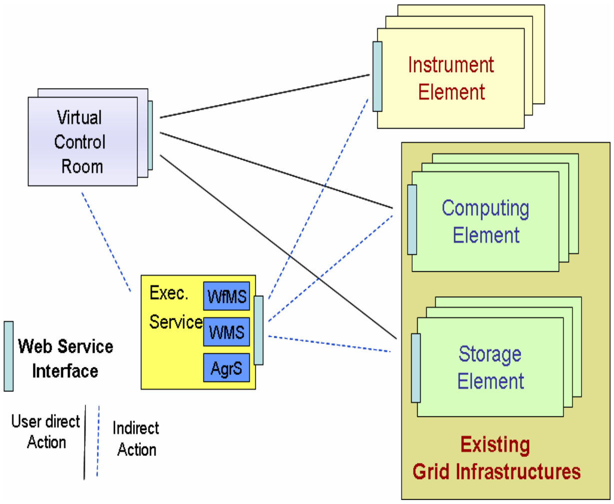

Fig.1: *Interaction between the Instrument Element and other Grid components*

Users can directly access the controlled sensors using a Virtual Control Room [7] to achieve fast, high-level interaction with the IE component and the sensors that it manages. As an additional option, a generic device operation can be part of a complex workflow [8], [9], [10] of a Grid execution service that allows the IE to access and 'talk' with the computational Grid. These ways to control the sensors are independent and can be performed in parallel where needed.

## 3 Instrument Element Architecture

The term Instrument Element describes a set of services that provide the needed interface to enable remote control and monitoring of physical devices. Figure 2 represents the IE architecture. The IE needs to be flexible; in the simplest scenario this component controls a simple geospatial sensor or an FPGA card that performs a specific function, while for a more complex sensor network it can be used as bridge between sensors and computational grid nodes. Finally, the IE can be part of the sensor instrumentation. In this case, the instruments can be organized in a network that enables grid interaction.

We next describe each of the subsystems and discuss their proper use.

**Resource Service (RS):** This service organizes all the resources that belong to the real devices, in sessions, through logical partitions. Sensors can be discovered, allocated, and queried. The complexity of the information managed by this service is use-case dependent and ranges from practically fixing a configuration to configuration and orchestration of thousands of nodes.

**Information and Monitor Service (IMS):** This service uses a publish/subscribe system to disseminate monitoring data to interested partners. More specifically, each instrument publishes monitoring information through the IMSProxy and the IMS arranges them in topics. Subscribers can register for specific messages that are selected based on the source of the information, the error severity level, and the type of messages, or a combination of these tags.

To support the above requirements, we developed RMM-JMS, a JMS implementation atop of a high performance Reliable Multicast Messaging (RMM) layer [27]. This enables the GridCC IE to have high-throughput low-latency transport services designed for one-to-many data delivery or many-to-many data exchange in a message-oriented middleware publish/subscribe fashion, which is also JMS compliant. RMM-JMS supports peer-to-peer communication in both brokered and brokered-less modes and is singled out by its high performance capabilities. To achieve high throughput and low latency, RMM uses information on the network and buffers status to adaptively handle the situation.

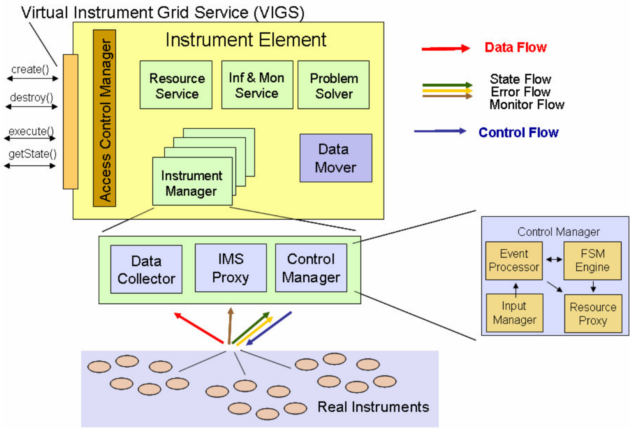

Fig.2: *Instrument Element architecture*

**Data Mover (DM):** This service provides the SRM [19], [20] interface to any external storage (SE) or processing elements (CE). It finds the 'best' mechanism such as GridFTP [21], [22] or other transport protocols to move information from sensors to a storage resource.

**Instrument Managers (IM)**: The IMs are the parts of the IE that perform the actual communication with the devices. They act as protocol adapters that implement the sensor-specific protocols for accessing its functions and reading its status. Since the instruments are heterogeneous in nature, there is a need to support many instrument managers in the same container, one instance for each logical set of sensors.

The main components of each Instrument Manager:

- The **Data Collector** acts as a proxy for the higher-level Data Mover service. This sensor collection service is the heart of the system. It has to withstand a high rate of input data and be available at any time. Losing data items can cause unpredictable decisions; see for example [28].
- The **IMS Proxy** collects information such as logs, errors, states, and critical parameters of the instrument and publishes it via the IMS system.
- The **Control Manager** controls the devices. A typical use of this sub-system is to receive input from the users that control the sensors. Since this component also receives input such as states or errors from the devices, it can also react autonomously to unexpected behaviour of the controlled resources.

**Problem Solver (PS):** This service allows automatic recovery procedures if device malfunctions are detected.

**Access Control Manager (ACM):** This service is responsible for checking user credentials and deciding whether the IE should process an external request.

The use of each IE subcomponent is use-case dependent; in a simple sensor or an FPGA card one IM can be used. The ACM is used mainly in cases where security is mandatory to prevent over-consumption of critical and limited resources in the device. On the other hand, in a more complex scenario, where an entire network of sensors provides the services, the use of components such as the Resource Service and the Problem Solver becomes critical.

## 4 Use Cases

The first prototype is currently used to manage the following cases:

- In Intrusion Detection System, distributed anomaly sensors measure various network elements to detect Distributed Denial of Service (DDoS) attacks or worm propagation [23].
- In electrical utility networks (or power grids [2]): the introduction of a very large number of 'embedded' power generators, often using renewable energy sources, creates a severe challenge for utility companies.
- In the Compact Muon Solenoid (CMS) [24] Data Acquisition (DAQ), the Instrument Element is the master controller when the experiment is collecting data. Approximately $O(2*10^7)$ electronic devices need to be accessed and controlled by $O(10^4)$ different machines scattered over a high-bandwidth network.
- In the Elettra [25] Synchrotron Radiation Storage Ring, the IEs control and monitor many instruments (mainly sensors). The rate of incoming control and monitoring data is different from the previous use-cases; it is

smaller than the data rate in the CMS scenario, yet since it resides at a critical point of alerting against major catastrophe, it imposes high requirements on response times for alerts at the human-machine level interface.

- Operating since 2003, IMAA installed a network of sensors devoted to providing real-time information on several meteorology, physical-chemical parameters of soil and sub-soil useful for describing landslide dynamics, and possibly to timely detect warning conditions [29], [30]. Monitoring is done with a multi-channel data-logger [31] on the monitored side to collect/aggregate information coming from different, both passive and active, sensors [32]. Different sensors are wired to a data-logger while this device is remotely controlled via a wireless network. The instrument element virtualizes the access to the sensors controlling multiple data-acquisition devices allowing grid access to the instruments.

We are currently finishing a complete performance analysis of each IE subcomponent. We plan to include this analysis in the full paper and discuss the scalability and the flexibility of the presented architecture; in addition, we can provide a demo where the IE monitors and controls a set of sensors that reconstruct a Muon trace during a physical experiment.

## Acknowledgment

The GridCC project is supported under EU FP6 contract 511382.

## References

[1] Gaetano Maron, Angelos Lenis, Sakis Moralis, Mary Grammatikou, Theodoros Karounos, Symeon Papavassiliou, Vasilis Maglaris, Paris Sphicas, Symeon Papavassiliou, Tiziana Ferrari, Constantinos A. Kotsokalis, Andrew Stephen McGough, Tatiana Kalganova, Peter Hobson, Roberto Pugliese, Francesco Lelli, and David Colling, The GridCC Architecture (available at www.gridcc.org) May 2005.

[2] M. Irving, G. Taylor, and P. Hobson, *IEEE Power & Energy Magazine March* April 2004, pp 40.

[3] UDDI Specification. Version 2.0, 3.0. http://www.uddi.org/specification.html 2005.

[4] Simple Object Access Protocol (SOAP) 1.1/1.2 W3C http://www.w3.org/TR/SOAP/ 2005.

[5] W3C. Web Services Description language (WSDL) 2.0. Note, http://www.w3.org/TR/wsdl20-primer 2005.

[6] Ian J. Taylor, from P2P to Web Services and Grids, Peers *in a Client/Server World*, Springer, October 2004.

[7] R. Pugliese, F. Asnicar, L. Del Cano, L. Chittaro, R. Ranon, L. De Marco, and A. Senerchia, Collaborative Environments for the GRID: the GRIDCC Multipurpose Collaborative Environment, *Tyrrhenian Workshop*, Sorrento, July 2005.

[8] WS-Agreement specification: http://www.gridforum.org/Meetings/GGF11/Documents/draft-ggf-graap-agreement.pdf.

[9] P. Andreetto et al. Pratical Approaches to Grid Workload and Resource Management in the EGEE Project. *In Proceedings of the International Conference on Computing in High Energy Physics (CHEP2004)*, Interlaken, Switzerland. 27 September-1 October 2004.

[10] T.Ferrari, and E.Ronchieri, gLite Architecture for Resource Reservation and Allocation and Network Element, *2nd EGEE Conference*, Den Haag (NL), Nov 24 2004

[11] OGC Sensor Web Enablement: http://www.opengeospatial.org/functional/?page=swe.

[12] Interchangeable Virtual Instrument Foundation: http://www.ivifoundation.org/.

[13] E. McManus and Sun Microsystems, Inc. JSR 160: JavaTM Management Extensions (JMX) Remote API 1.0 October 2003.

[14]The Freenet Project, http://freenet.sourceforge.net/.

[15] Ian Clarke, S .G. Miller, T. W.Hong, O. Sandberg and B. Wiley Protecting Freedom of Information Online with Freenet, *IEEE Internet Computing*, February 2002.

[16]Jini home page http://www.jini.org/.

[17] S. R. Waterhouse, D. M. Doolin, G. Kan and Y. Faybishenko, JXTA Search: A distributed search framework for peer-to-peer networks, *IEEE Internet Computing,* Vol. 6, pp 68-73, 2002.

[18] Open Grid Services Infrastructure (OGSI) v1.0, http://forge.gridforum.org/projects/ggfeditor/document/draft-ogsi-service1/en/1.

[19] EGEE Middleware Architecture and planning, EGEE Project Deliverable, EGEE-DJRA1.1-594698-v1.0, Chapter 9. Jul 2005 (also available at https://edms.cern.ch/document/594698/).

[20] SRM: Storage Management Working Group http://sdm.lbl.gov/srm-wg/.

[21] W. Allcock, J. Bester, J. Bresnahan, A. Chervenak, L. Liming, and S. Tuecke, GridFTP: Protocol Extensions to FTP for the Grid, http://www-fp.globus.org/datagrid/gridftp.html.

[22] Vladimir Silva, Transferring files with GridFTP Apr 2003: http://www-128.ibm.com/developerworks/grid/library/gr-ftp/?ca=dgr-lnxw03GridFTP.

[23] C. Siaterlis, A. Lenis, A. Moralis, P. Roris, G. Koutepas, G. Androulidakis, V. Chatzigiannakis, M. Grammatikou, D. Kalogeras, S. Papavassiliou & V. Maglaris, Distributed Network Monitoring and Anomaly Detection as a Grid Application, *HP Openview University Association Plenary Workshop (HP-OVUA)*, Porto, Portugal Workshop, July 11 2005.

[24] Sergio Cittolin, Wesley Smith, Joao Varella, Attila Racz, Michel Della Negra, Alain Herve, CMS TDR 6.2, *The TriDAS Data Acquisition project and High-Level Trigger*, CERN/LHCC, December 2002.

[25] Synchrotron Radiation Sotorage Ring Elettra http://www.elettra.trieste.it/index.php.

[26] WSRF specification also availlable at: http://docs.oasis-open.org/wsrf/wsrf-primer-1.2-primer-cd-01.pdf.

[27] B. Carmeli, G. Gershinsky, A. Harpaz, N. Naaman, H. Nelken, J. Satran and P. Vortman, High throughput reliable message dissemination, *Proceedings of the 2004 ACM symposium on Applied computing*, Nicosia, Cyprus, Pages: 322 – 327.

[28] Oren Ben-Zwi and Shlomit S. Pinter, Handling Sensed Data in Hostile Environments, *Proceedings of the International Conference on Mobile Ad-hoc and Sensor Networks (MSN)*, Wuhan, China, Lecture Notes in Computer Science, No. 3794, Springer-Verlag, December 2005, pp. 433--442.

[29] A. Perrone, A. Iannuzzi, V. Lapenna, P. Lorenzo, S. Piscitelli, E. Rizzo, and F. Sdao, (2004). High-resolution electrical imaging of the Varco d'Izzo earthflow (southern Italy), *Journal of Applied Geophysics*, 56 (1), 17-29.

[30] A. Perrone, G. Zeni, S. Piscitelli, A. Pepe, A. Loperte, V. Lapenna, and R. Lanari, (2005). On the joint analysis of SAR interferometry and Electrical Resistivity Tomography surveys for investigating ground deformations: the case-study of Satriano di Lucania (Potenza, Italy). Remote Sensing of Environment.

[31] Data acquisition system (Keithley Instruments model 2701, with plug-in switching modules model 7702).

[32] Non-polarizable Petiau electrodes (SDEC electrodes, Pb/PbCl2-NaCl) Soil probes made by http://www.campbellsci.com/sensors: Time Domain Reflectometer (TDR) and a thermometer. Atmospheric probes made by http://www.campbellsci.com/sensors: Pluviometer and thermometer.