

Enterprise Energy Management using a Linked Dataspace for Energy Intelligence

Edward Curry, Souleiman Hasan, Sean O’Riain
Digital Enterprise Research Institute
National University of Ireland, Galway
{ed.curry, souleiman.hasan, sean.oriain}@deri.org

Abstract—Energy Intelligence platforms can help organizations manage power consumption more efficiently by providing a functional view of the entire organization so that the energy consumption of business activities can be understood, changed, and reinvented to better support sustainable practices. Significant technical challenges exist in terms of information management, cross-domain data integration, leveraging real-time data, and assisting users to interpret the information to optimize energy usage. This paper presents an architectural approach to overcome these challenges using a Dataspace, Linked Data, and Complex Event Processing. The paper describes the fundamentals of the approach and demonstrates it within an Enterprise Energy Observatory.

Keywords: *Energy Intelligence; Energy Management; Dataspace; Linked Data; Semantic Sensor Networks; Complex Event Processing*

I. INTRODUCTION

Society has an energy consumption problem [1], and is increasing the CO₂ level in the atmosphere due to too much energy consumption derived from fossil fuels. According to the International Energy Agency’s statistics in four years (2004–2008) the world population increased 5%, annual CO₂ emissions increased 10%, and gross energy production increased 10% [2]. The research field of Energy Informatics [3] recognizes the role that Information Systems can play in reducing energy consumption, and thus CO₂ emissions. Energy informatics is concerned with analyzing, designing, and implementing systems to increase the efficiency of energy demand and supply. The core idea requires the collection and analysis of energy data to support optimization of energy. Substantial potential exists for Energy Information Systems to bring business processes and resource planning together with direct and indirect energy consumption, to effect positive changes across the activities of governments, organizations and individuals.

While organizations are fighting a data deluge [4], there is a significant lack of data on sustainability concerns. A 2010 survey of more than 600 CIOs and Senior IT Managers highlighted that few organizations are performing well at measuring the effectiveness of their sustainability efforts [5]. The paucity of sustainable information is a significant challenge that needs to be addressed.

Building on the work of Business Intelligence [6], Energy Intelligence can be leveraged to help organizations manage power consumption more efficiently through monitoring and controlling. In order to provide a holistically real-time management of the energy within an Enterprise, Energy Intelligence platforms will need to provide a functional view of the entire organization so that the energy consumption of business activities can be understood, changed, and reinvented to better support sustainable practices. Within this context Energy Intelligence platforms will need to be cross-domain (i.e. finance, manufacturing, facilities, IT, etc.) and multi-level (i.e. Organization, business function, individual). They need to put energy consumption within the context of the associated business activities. Energy Intelligence platforms pose significant technical challenges in terms of information management, cross-domain data integration, and real-time processing of sensor streams and events. They must also assist users to interpret information to understand how it can be used to optimize energy usage.

This paper presents the Linked dataspace for Energy Intelligence (LEI) based on Dataspaces, Linked Data, Complex Event Processing and Situational Awareness. This paper describes the fundamentals of the approach and demonstrates it within an Enterprise Energy Observatory for a Small-to Medium size Enterprise (SME).

II. ENTERPRISE ENERGY MANAGEMENT

Creating a sustainable business requires an enterprise-wide perspective on the use, flows and destiny of energy along with monetary information on environment-related costs, earnings, and savings. This type of information is critical if we are to understand the causal relationships between the various energy saving actions that can be taken, and their impact on energy performance.

Improving energy performance, especially through changing the way an organization operates, requires a number of practical steps which will include the need for a systematic and holistic approach for information-gathering and analysis. Creating a holistic view of energy information for an organization is not a straightforward task, the complexities of real-world organizations mean they use energy in many different ways. Energy Intelligence platforms need to support four key requirements:

- Holistic energy consumption
- Multi-level energy analysis

- Business context Energy Consumption
- Energy Situational awareness

A. Holistic Energy Consumption



Figure 1. Sources of Energy Consumption within the an Enterprise

For an enterprise to manage its energy usage it must consider all the direct and in-direct sources of energy consumption for which it is responsible. Energy consumption takes place across the enterprise in many forms (see Fig. 1), from the direct electricity consumed by Office IT and Data Centers, to the indirect energy from fuel used in business travel and the daily commute of workers (many other form of energy consumption exist including manufacturing, logistics, etc.). Energy Intelligence platforms will need to holistically manage energy in its many forms.

B. Multi-Level Energy Analysis

Energy must be managed at multiple levels within an enterprise. For example, managing energy at the organization-level is useful for macro-level strategic decision making, however it is less useful for operational processes such as optimizing energy usage for a single production line. In the same way, micro-level production line energy consumption is less effective for macro-level corporate decision making when viewed in isolation. Energy intelligence platforms must support multi-level models of enterprise energy usage providing the appropriate level of information.

C. Business Context of Energy Consumption

For an enterprise to manage its energy it is important for it to understand the business context associated with energy consumption. Energy management is a multi-domain problem, and energy data needs to be interpreted within the context of other business information including Enterprise Resource Planning (ERP), finance, facility management, human resources, asset management and code compliance. Little interaction exists between these islands of information and it is therefore difficult to leverage cross-domain data. A well-designed energy intelligence platform will need to provide views of energy together with cross-domain business context for different stakeholders.

D. Energy Situation Awareness

In order to optimize energy usage within the enterprise, energy managers and business stakeholders will need help to interpret energy data and real-time energy events. Situation awareness [7] creates a single holistic real-time picture of entities and their relationships within the enterprise. Energy intelligence platforms need to provide energy situation awareness for the enterprise that helps business users to understand the operational environment and make appropriate decisions. These services need to support both energy performance objectives and business performance objectives from across domains, such as human resources (i.e. occupancy comfort) and enterprise resource planning (i.e. room utilization).

III. BACKGROUND

In order to deliver enterprise energy management systems that support the four requirements a number of technical challenges need to be overcome. The main challenges are 1) a flexible approach to information management, together with 2) an agile approach to data interoperability, and 3) support for real-time processing of dynamic and background data. To meet these challenges we propose the use of three key technologies Dataspace, Linked Data, and Complex Event Processing.

A. Dataspace

A Dataspace is an emerging data management architecture that is very distinct from current approaches. The dataspace approach recognizes that in large-scale integration scenarios, involving thousands of data sources, it is difficult and expensive to obtain an upfront unifying schema across all sources [8]. Dataspace shifts the emphasis to providing support for the co-existence of heterogeneous data that does not require an upfront investment into a unifying schema. Data is integrated on an “as needed” basis with the labor-intensive aspects of data integration postponed until they are required. Dataspaces reduce the initial effort required to setup data integration by relying on existing matching and mapping generation techniques. This can result in a loosely integrated set of data sources, when tighter integration is required it can be achieved in an incremental “pay-as-you-go” fashion by more closely integrating the required data sources. Similarly a dataspace may only provide weak guarantees of consistency and durability. As stronger guarantees are desired, more effort can be put into making agreements among the various systems, see Table I.

TABLE I. DATASPACE COMPARISON

	Model	Formats	Control	Query	Integration
DBMS	Relational	Homog.	Complete	Precise	Explicit
Dataspace	All	Heterog.	Partial	Approx.	Implicit/ Incremental

B. Linked Data

Emerging from research into the Semantic Web, Linked Data proposes an approach for information interoperability based on creation of a global information space [9]. Linked data leverages the existing open protocols and standards of

the World Wide Web architecture for sharing structured data on the web. As part of the Open Data movement many organizations, governments and scientific communities are actively publishing data on the web as Linked Data [10].

1) Linked Data Principles

Linked data technology uses web standards in conjunction with four basic principles for exposing, sharing and connecting data. These principles are:

- *Use Uniform Resource Identifiers (URI) as names for things:* such as a person, a place, a product, an organization, an event or even concepts such as net profit.
- *Use HTTP URIs so that people can look them up:* URIs are used to retrieve data about objects using standard web protocols. For an employee this could be their organization and job classification, for a product this may be its specification, availability, price, etc.
- *When a URI is looked up, provide useful information:* provide data in standardized format such as RDF.
- *Including links to other URIs so that related data can be discovered:* retrieved data may link to other data sources, thus creating a data network e.g., data about a product may link to all the components it is made of, which in turn may link to their supplier.

2) Resource Description Framework

The Resource Description Framework (RDF) is the basic machine-readable representational format used to represent linked data. RDF is a general method for encoding graph-based data that does not follow a predictable structure; it is schema-less and self-describing, meaning that the labels of the graph describe the data. RDF facts are specified as statements and are expressed as atomic constructs of a subject, predicate and object, also known as a triple. The statement “Room 202e is occupied by Edward Curry” is expressed in triple format as:

Subject - “Room 202e”
 Predicate - “is occupied by”
 Object - “Edward Curry”

RDF is designed for use in web-scale decentralized graph data models. For this reason the statement parts need to be identified so that they can be readily and easily reused. RDF uses URIs for identification, expressing the previous statement in RDF then becomes:

<http://lab.linkeddata.deri.ie/2010/deri-rooms#r202e>
<http://vocab.deri.ie/rooms#occupant>
http://www.deri.ie/about/team/member/edward_curry#me

Multiple triples can be joined together to build up a graph of information. Fig. 2 provides an example of a graph segment that models a description of a Room: its name, occupant, containing building, and height dimension.

C. Complex Event Processing

Complex Event Processing can be used to process real-time information flows to detect situations of interest with the use of event patterns. When a pattern is matched, a new higher-level event is generated and can participate in further processing, be visualized in a dashboard or sent to a business process management tool.

The bridge between Complex Event Processing and Linked Data can be achieved with the use of Semantic Sensor networks [11] (SSN). The SSN ontology has been designed to represent common concepts in a sensing environment. It handles metadata information such as the sensor types and platforms as well as classes to describe real-time sensor observations. The use of sensor networks that respect linked data principles [31] when publishing data forms a solid basis for a real-time energy platform.




Figure 2. An RDF graph describing Room 202e in the DERI building.

IV. PROPOSED APPROACH

This paper proposes the use of a dataspace, linked data and complex event processing to meet the challenges for real-time holistic enterprise energy management. The realization of a dataspace using linked data as the enabling technology creates a *Linked Dataspace* that supports the co-existence of heterogeneous data with an incremental approach to interoperability using web-based standards. The inclusion of Complex Event Processing with Semantic Sensor Networks provides support for real-time processing of dynamic and background data for energy situation awareness. These technologies have been combined within a single architecture called the Linked dataspace for Energy Intelligence (LEI) that can be used to build energy management applications.

A. System Architecture

The LEI approach can be accommodated with minimal disruption to existing information infrastructure. The objective is to expose the data within existing systems, but only link the data when it needs to be shared. Representing energy data within the linked data format makes it open allowing it to be easily combined with linked data from other relevant domain silos.

The main components of the architecture, as illustrated in Fig 3, are the linked data wrappers on existing systems, the Linked dataspace consisting of a linked data cloud & support services, and the resulting energy management applications.

1) Linked Data Wrappers

At the bottom of the architecture are the existing operational legacy information systems. Wrappers perform

the “RDFization” process, which transforms multiple formats and legacy data and lifts it to the dataspace.

2) Linked Energy Cloud

The Linked Dataspace links at the information-level (data) not the infrastructure-level (system) by focusing more on the conceptual similarities (shared understanding) between information. This is achieved by following an entity-centric approach that focuses on the concepts that exist within the systems, for example, business entities like employees, products, customers, equipment, assets, buildings, rooms, etc. Entities within the dataspace are enriched with data from multiple systems. This results in a cloud of interlinked resources that reflects virtual or actual entities with links to relevant knowledge and contextual information from across all the information systems that have exposed linked data. Sources may be added in an incremental manner to the cloud where they can be reused. Each entity within the cloud has a dereferenceable URI that returns data in a machine-readable format describing the resource. The resulting Linked Energy Cloud is rich with knowledge and semantics about energy-related performance indicators and forms the basis for a real-time energy analytics and other applications with the help of support services.

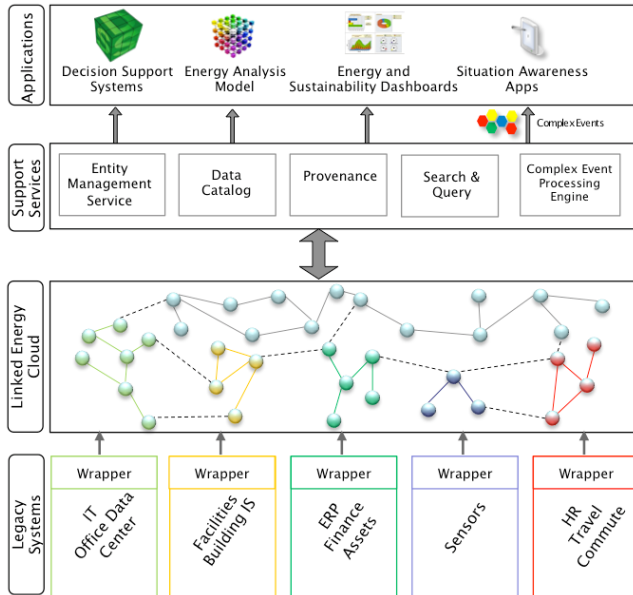


Figure 3. Linked dataspace for Energy Intelligence (LEI)

3) Support Services

Dataspace support services are designed to simplify the consumption of the linked data cloud by encapsulating common services for reuse. Example support services are:

- *Entity Management Service* to improve data quality and inter-linkage between entity data scattered among legacy systems. The EMS can leverage automatic entity consolidation algorithms and crowdsourcing for collaborative data management [12].
- *Complex Event Processing engine* [13] is used to assess situations of interest that are encoded as event/action rules. Real-time information from sensors networks are supported via the SSN Ontology [11].
- *Data Catalogue and Provenance service* [14] track data within the data cloud with specific attributes such as freshness and publisher.
- *Search and Query services* [15] allow users to interact with the dataspace using structured or natural language interfaces.

4) Energy Management Applications

At the top of the architecture are the energy management applications that consume the resulting data and events from the linked data cloud. An example Enterprise Energy Observatory implementation using the LEI approach is now discussed.

V. USE CASE: ENTERPRISE ENERGY OBSERVATORY

To demonstrate a Linked dataspace for Energy Intelligence an enterprise energy observatory has been implemented within the Digital Enterprise Research Institute (DERI). DERI is representative of a typical SME with ~130 staff located in a single dedicated 1675m² building with office space, meeting rooms, café, and a data center. The observatory enables a holistic energy management (facilities, office IT, Data Centre, business travel, commute, etc.), which can be used to manage energy at different-levels (organization, business process, individual) and from different domains perspectives (energy, finance, sustainability, human resources, etc.).

The Observatory utilizes a multi-level model, described in Table II, to scope energy activities. The model contains 3 levels that match the DERI organization, however many other levels are possible within an enterprise [16]. In the model energy can flow between levels, allowing the granularity of energy analysis to be scoped as needed.

TABLE II. MULTI-LEVEL MODEL FOR ENTERPRISE ENERGY OBSERVATORY

Level	Stakeholders	Information Requirements	Energy Metrics
Organization	Executive team, shareholders, regulators, suppliers, consumers, ...	Corporate sustainable reports, sustainability plans, sustainable business objectives, business function performance, ...	- Total energy consumption - Total electricity consumption - % Renewable energy sources
Function	Function management team, organization management, employees, ...	Function sustainability performance, functions sustainability objectives, ...	- Facility electricity consumption - Internal electricity consumption - DC electricity consumption
Individual	Employees	Direct and indirect cost of actions of the individual, private travel, energy use, work commute, impacts of products and services consumed.	- Business travel (flights) - Employee commuting - Individual IT consumption - Laptop electricity consumption - Individual Office energy usage

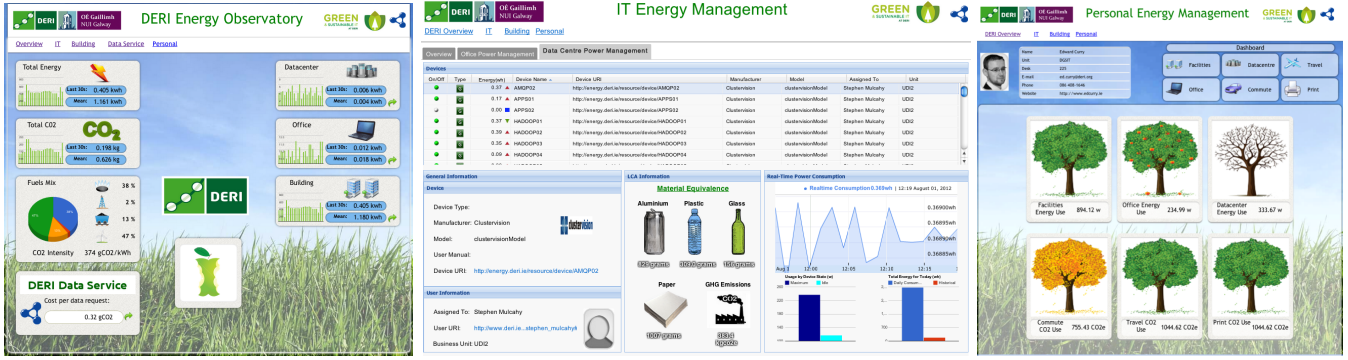


Figure 4. DERI Energy Observatory: (a) left screen is Organisational-Level (b) middle screen is Functional-Level (c) right screen is Individual-Level

The Energy Observatory can be considered as a System of Systems [17] as it consists of over 10 legacy systems, see Table III, with LEI creating the necessary linkages required to integrating the cross-domain legacy information system.

TABLE III. SAMPLE OF LEGACY INFORMATION SYSTEMS

System	Type	Information
Building Resource Allocation	Spreadsheet	Assignment of building resources to people and groups within DERI
Office IT Assets Database	SharePoint	Office IT equipment assignment
Employee Database	SharePoint	Person detail, group affiliations and current work assignments
DC Resource Allocation	Spreadsheet	Assignment of DC resources to services
DC Resource Monitoring	Ganglia Monitoring System	Real-time monitoring of hardware resources within the data center such as CPU load averages or network utilization)
DC IT Assets Database	SharePoint	Repository of equipment within Data Centre
DC Service Management	Zabbix	Real-time monitoring of network services and hardware
Energy Sensor Network	Enistic Sensor Platform	Real-time sensor reading from power meters within building

A. Multi-Level Energy Management

LEI delivers a merged cross-domain multi-level model that provides a holistic view of the energy management within the enterprise. This view can be used to build energy management applications for stakeholders at different levels within DERI. In each these applications the user can see the relative impact of individual KPIs, see trends within and between KPIs over time. These dashboards can assist in decision making via:

- Real-time direct action items that employees can respond to while in the planning stages of emission-producing activities.
- Daily direct action items that employees can respond to after reviewing their previous day's activities.
- Monthly direct action items that employees can respond to by understanding the extrapolated effects of their daily activities.
- Monthly indirect action items that employees can respond to by understanding the effect of their day-to-day activities as averaged over the enterprise

1) Organization-Level

The organization level of the Energy Observatory provides a summary of energy consumption displayed with an executive level dashboard. In the interface, illustrated in Fig 4(a), executives can view the real-time consumption of energy across all enterprises domains, IT, facilities, travel, etc., as well as specific business processes, in this case a cloud data service within the DC. Electricity consumption information is augmented with energy mix (renewable vs. fossil fuels) information to determine the real-time CO₂ footprint of the organization. The energy situation assessment of the organization is conveyed with the “Apple” widget, the more the apple is eaten, the more energy hungry the organization.

2) Function-Level

At the function-level energy management apps have been created for the IT and Facilities functions. The IT Energy Management system, see Fig. 4 (b), is responsible for all of IT related energy consumption, from office IT equipment such as printers, laptops, desktops and monitors, to Data Centers with servers, networking equipment and HVAC. In addition to tracking the physical equipment, the system associated the energy usage of the equipment with the relevant business activities allows for fine-grained understanding of what business activities are responsible for IT energy usage, and can enable IT to bill appropriately.

3) Individual-Level

At the individual level the “myEnergy” interface, see Fig. 4 (c), allows employees to understand the direct and indirect energy consumption. myEnergy gives an employee real-time energy consumption data on their IT, Facilities, Travel, Commute and DC usage. Extensive use of situation awareness widgets allows the employee to quickly understand their performance under each category of energy usage.

B. DERI Energy Management Applications

A number of other energy management applications were developed using LEI in the DERI deployment, these include:

- The Facilities Energy Explorer [18] identifies energy leaks and non-ecological behaviors within the DERI building.
- An environmental chargeback model for cloud computing and data center services [19] provides an

indication of the power and CO₂ costs of services invoked by consumers.

- Sustainable DERI [20] provides a near-real time Corporate Social Responsibility Report on the impacts of DERI business activities

VI. RELATED WORK

The EU GAMES project [21] focuses on the improvement of IT service centers energy performance with respect to quality of service agreements during service composition. The energy efficiency improvement is based on a knowledge base of application level impacts on the IT service center energy efficiency. GAMES focuses mainly on data centers energy efficiency improvement and the problem of multi-domain energy management has not been tackled.

The Los Angeles Smart Grid Project [22] aims at employing semantic technologies and complex event processing to maximize the benefits of a smart grid. While the underlying technologies of the proposed approach and the one suggested in this paper are similar, this paper targets an enterprise scale with more focus on integrating non-energy related data sources for a contextualized understanding of energy usage. In addition the Smart Grid Project does not utilize the concept of dataspace within its architecture.

VII. CONCLUSIONS & FUTURE WORK

Energy Intelligence platforms pose significant technical challenges in terms of information management, cross-domain data integration, and real-time processing of sensor streams and events. This paper presented the Linked dataspace for Energy Intelligence (LEI) based on Dataspaces, Linked Data, Complex Event Processing and Situational Awareness. The approach was demonstrated within an Enterprise Energy Observatory for a Small-to-Medium size Enterprise (SME). Future work will focus on realizing more detailed multi-level enterprise energy models within the LEI, and the need for semantically aware dataspace support services [23].

VIII. ACKNOWLEDGMENTS

This work has been funded by Science Foundation Ireland under Grant No. SFI/08/CE/I1380 (Lion-2).

REFERENCES

- [1] M. I. Hoffert, K. Caldeira, G. Benford, D. R. Criswell, C. Green, H. Herzog, A. K. Jain, H. S. Khesghi, K. S. Lackner, J. S. Lewis, and others, "Advanced technology paths to global climate stability: energy for a greenhouse planet," *science*, vol. 298, no. 5595, p. 981, 2002.
- [2] "International Energy Agency: Key World Energy Statistics 2010," OECD Publishing, 2010.
- [3] R. T. Watson, M. C. Boudreau, and A. J. Chen, "Information systems and environmentally sustainable development: energy informatics and new directions for the IS community," *MIS Quarterly*, vol. 34, no. 1, pp. 23–38, 2010.
- [4] K. Cukier, "A special report on managing information: Data, data everywhere," *The Economist*, vol. 12, no. 8, 2010.
- [5] A. O'Flynn and G. Executive, "Green IT: The Global Benchmark," white paper, Fujitsu, 2010.
- [6] H. P. Luhn, "A business intelligence system," *IBM Journal of Research and Development*, vol. 2, no. 4, pp. 314–319, 1958.
- [7] M. R. Endsley and D. J. Garland, *Situation awareness: analysis and measurement*. CRC, 2000.
- [8] M. Franklin, A. Halevy, and D. Maier, "From databases to dataspace: a new abstraction for information management," *Sigmod Record*, vol. 34, no. 4, pp. 27–33, 2005.
- [9] T. Heath and C. Bizer, "Linked Data: Evolving the Web into a Global Data Space," *Synthesis Lectures on the Semantic Web: Theory and Technology*, vol. 1, no. 1, pp. 1–136, Feb. 2011.
- [10] E. Curry, A. Freitas, and S. O'Riain, "The Role of Community-Driven Data Curation for Enterprises," in *Linking Enterprise Data*, D. Wood, Ed. Boston, MA: Springer US, 2010, pp. 25–47.
- [11] L. Lefort, C. Henson, K. Taylor, and P. Barnaghi, "Semantic Sensor Network XG Final Report, W3C Incubator Group Report (28 June 2011)," 2011.
- [12] U. ul Hassan, S. O'Riain, and E. Curry, "Leveraging Matching Dependencies for Guided User Feedback in Linked Data Applications," in *9th International Workshop on Information Integration on the Web (IIWeb2012)*, 2012.
- [13] S. Hasan, E. Curry, M. Banduk, and S. O'Riain, "Toward Situation Awareness for the Semantic Sensor Web: Complex Event Processing with Dynamic Linked Data Enrichment," in *4th International Workshop on Semantic Sensor Networks 2011 (SSN11)*, 2011, pp. 69–81.
- [14] A. Freitas, T. Knap, S. O'Riain, and E. Curry, "W3P: Building an OPM based provenance model for the Web," *Future Generation Computer Systems*, vol. 27, no. 6, pp. 766–774, Jun. 2011.
- [15] A. Freitas, J. Oliveira, S. O'Riain, E. Curry, and J. Pereira da Silva, "Querying Linked Data using Semantic Relatedness: A Vocabulary Independent Approach," in *Proceedings of the 16th International Conference on Applications of Natural Language to Information Systems, NLDB 2011*, 2011, vol. 6716, pp. 40–51.
- [16] E. Curry, B. Guyon, C. Sheridan, and B. Donnellan, "Developing an Sustainable IT Capability: Lessons From Intel's Journey," *MIS Quarterly Executive*, vol. 11, no. 2, pp. 61–74, 2012.
- [17] E. Curry, "System of Systems Information Interoperability using a Linked Dataspace," in *IEEE 7th International Conference on System of Systems Engineering (SOSE 2012)*, 2012, pp. 101–106.
- [18] E. Curry, J. O'Donnell, and E. Corry, "Building Optimisation using Scenario Modeling and Linked Data," in *First International Workshop on Linked Data in Architecture and Construction (LDAC 2012)*, 2012, pp. 6–8.
- [19] E. Curry, S. Hasan, M. White, and H. Melvin, "An Environmental Chargeback for Data Center and Cloud Computing Consumers," in *First International Workshop on Energy-Efficient Data Centers*, 2012, pp. 117–128.
- [20] E. Curry, S. Hasan, U. ul Hassan, M. Herstand, and S. O'Riain, "An Entity-Centric Approach To Green Information Systems," in *Proceedings of the 19th European Conference on Information Systems (ECIS 2011)*, 2011.
- [21] M. Bertoini, B. Pernici, I. Salomie, and S. Wesner, "GAMES: Green Active Management of Energy in IT Service Centres," *Information Systems Evolution*, pp. 238–252, 2011.
- [22] Y. Simmhan, Q. Zhou, and V. Prasanna, "Semantic Information Integration for Smart Grid Applications," *Green It: Technologies and Applications*, p. 361, 2011.
- [23] S. Hasan, S. O'Riain, and E. Curry, "Approximate Semantic Matching of Heterogeneous Events," in *6th ACM International Conference on Distributed Event-Based Systems (DEBS 2012)*, 2012, pp. 252–263.