

When Things Matter: A Data-Centric View of the Internet of Things

YONGRUI QIN, The University of Adelaide

QUAN Z. SHENG, The University of Adelaide

NICKOLAS J.G. FALKNER, The University of Adelaide

SCHAHRAM DUSTDAR, Vienna University of Technology

HUA WANG, Victoria University

ATHANASIOS V. VASILAKOS, National Technical University of Athens

With the recent advances in radio-frequency identification (RFID), low-cost wireless sensor devices, and Web technologies, the Internet of Things (IoT) approach has gained momentum in connecting everyday objects to the Internet and facilitating machine-to-human and machine-to-machine communication with the physical world. While IoT offers the capability to connect and integrate both digital and physical entities, enabling a whole new class of applications and services, several significant challenges need to be addressed before these applications and services can be fully realized. A fundamental challenge centers around managing IoT data, typically produced in dynamic and volatile environments, which is not only extremely large in scale and volume, but also noisy, and continuous. This article surveys the main techniques and state-of-the-art research efforts in IoT from data-centric perspectives, including data stream processing, data storage models, complex event processing, and searching in IoT. Open research issues for IoT data management are also discussed.

Categories and Subject Descriptors: C.2.4 [Distributed Systems]: Distributed Databases; H.3.5 [Information Storage and Retrieval]: On-line Information Services—*Web-based services*

General Terms: Algorithms, Performance

Additional Key Words and Phrases: Internet of Things, data management, applications

1. INTRODUCTION

The Internet is a global system of networks interconnecting computers using the standard Internet protocol suite. It has significant impact on the world as it can serve billions of users worldwide. Millions of private, public, academic, business, and government networks, of local to global scope, all contribute to the formation of

Authors' addresses: Y. Qin, Q. Z. Sheng, N. Falkner, School of Computer Science, The University of Adelaide, Adelaide SA 5005, Australia; emails: {yongrui, qsheng, nick}@cs.adelaide.edu.au. S. Dustdar, Institute of Information Systems, Vienna University of Technology; email: dustdar@infosys.tuwien.ac.at. H. Wang, Center for Applied Informatics, Victoria University, Melbourne, Victoria 3122, Australia; email: hua.wang@vu.edu.au; A. Vasilakos, Department of Electrical and Computer Engineering, National Technical University of Athens, Greece; email: vasilako@ath.forthnet.gr.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2014 ACM 0000-0000/2014/0000-0001 \$5.00

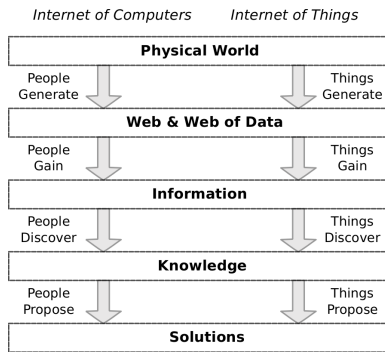


Fig. 1. Internet of Computers v.s. Internet of Things

the Internet. It is a network of networks and each network connects various numbers of computers. Hence, the traditional Internet has a focus on computers and can be called the Internet of Computers. In contrast, evolving from the Internet of Computers, the Internet of Things (IoT) emphasizes things rather than computers [Ashton 2009]. It aims to connect everyday objects, such as coats, shoes, watches, ovens, washing machines, bikes, cars, even humans, plants, animals, and changing environments, to the Internet to enable communication/interactions between these objects. The ultimate goal of IoT is to enable computers to see, hear and sense the real world. It is predicted by Ericsson that the number of Internet-connected things will reach 50 billion by 2020. Electronic devices and systems exist around us providing different services to the people in different situations: at home, at work, in their office, or driving a car on the street [James et al. 2009].

“Changes brought about by the Internet will be dwarfed by those prompted by the networking of everyday objects”, says a report by a United Nation (UN) [Biddlecombe 2005]. IoT is widely regarded as the number one of top 10 technologies that will change the world in the next 10 years [Bort 2011]. National Intelligence Council [Anonymous 2008] foresees that “by 2025, Internet nodes may reside in everyday things – food packages, furniture, paper documents, and more. Widespread diffusion of an Internet of Things (IoT) could contribute invaluable to economic development.”

There are several definitions or visions of IoT from different perspectives. From the viewpoint of services provided by things, IoT means “a world where things can automatically communicate to computers and each other providing services to the benefit of the human kind” [CASAGRAS 2000]. From the viewpoint of connectivity, IoT means “from anytime, anyplace connectivity for anyone, we will now have connectivity for anything” [ITU 2005]. From the viewpoint of communication, IoT refers to “a world-wide network of interconnected objects uniquely addressable, based on standard communication protocols” [INFSO 2008]. Finally, from the viewpoint of networking, IoT is the Internet evolved “from a network of interconnected computers to a network of interconnected objects” [Commission 2009].

We focus on our study of the Internet of Things from a *data perspective*. As shown in Fig. 1, data is processed differently in the Internet of Things and traditional




Fig. 2. Roadmap of this survey

Internet environments (i.e., Internet of Computers). In the Internet of Computers, both main data producers and consumers are human beings. However, in the Internet of Things, the main actors become *things*, which means things are the majority of data producers and consumers. Therefore, we give our definition of the Internet of Things as follows:

“In the context of the Internet, addressable and interconnected things, instead of humans, act as the main data producers, as well as the main data consumers. Computers will be able to learn and gain information and knowledge to solve real world problems directly with the data fed from things. As an ultimate goal, computers enabled by the Internet of Things technologies will be able to sense and react to the real world for humans.”

As of 2012, 2.5 quintillion (2.5×10^{18}) bytes of data are created daily¹. In IoT, connecting all of the things that people care about in the world becomes possible. All these things would be able to produce much more data than nowadays [Barnaghi et al. 2013]. The volumes of data are vast, the generation speed of data is fast and the data/information space is global [James et al. 2009]. Indeed, IoT is one of the major driving forces for *big data analytics*. Given the scale of IoT, topics such as storage, distributed processing, real-time data stream analytics, and event processing are all critical, and we may need to revisit these areas to improve upon existing technologies for applications of this scale [Barnaghi et al. 2013; James et al. 2009].

In this survey, we systematically investigate the key technologies related to the development of IoT and its applications, particularly from a data-centric perspective. The aim of this work is to provide a better understanding of the current research activities and issues. Fig. 2 shows the roadmap of this survey. As can be seen from the figure, we review and compare technologies including data streams, data storage models, searching, and event processing technologies, which play a vital role in enabling the vision of IoT. We also describe some relevant applications from several representative areas. Although some surveys about IoT have been conducted recently (e.g., [Atzori et al. 2010; Zeng et al. 2011; Perera et al. 2013; Li et al. 2014]), they focus on high level general issues and are mostly fragmented. In addition, these surveys do not specifically cover techniques on data processing and management, which is fundamentally critical to fully embrace IoT. To the best of our knowledge, this is the first survey that studies and discusses state-of-the-art techniques of IoT from the data-centric perspective.

The remainder of the article is organized as follows. Section 2 surveys the data

¹<http://www-01.ibm.com/software/data/bigdata/>

streaming techniques and Section 3 focuses on the data models and storage technologies for IoT. Search and event processing technologies are discussed in Sections 4 and 5, respectively. In Section 6, some typical ongoing and/or visionary IoT applications where data techniques for IoT can bring significant changes are described. Finally, Section 7 highlights some research open issues on IoT from the data perspective and Section 8 offers some concluding remarks.

2. DATA STREAMS

A data stream is a sequence of data objects, of which the number is potentially *unbounded*. A data stream may be continuously generated at a rapid rate. In the data stream, each data object can be described by a multidimensional attribute vector within a continuous, categorical, or mixed attribute space [de Andrade Silva et al. 2013]. There are some typical characteristics of data streams:

- Continuous arrival of data objects
- Disordered arrival of data objects
- Potentially unbounded size of a stream
- Normally no persistence of data objects after being processed
- Changing probability distributions of the unknown data generation process

Due to the excessive amount of data produced by all kinds of things in the era of IoT, data streams play an important role in data processing and analysis. This section will focus on related data stream research efforts that can help handle IoT data. Our discussions include general data stream processing, RFID data stream processing, and RDF triple stream processing.

2.1 General Data Stream Processing

Data streams can be generated in various scenarios, including a network of sensor nodes, a stock market or a network monitoring system and so on. In many scenarios such as the sensor network scenario, sensor nodes are normally powered by batteries or solar panels. Therefore, in typical a sensor data processing system, one of the challenging issues is power constraints. In most applications, communication across sensor networks or with a centralized server requires the largest amount of energy as sensing consumes less energy [Subramaniam and Gunopulos 2007]. If sensor nodes send their raw sensing data to a server without consideration of the amount of energy needed to communicate, the battery life of the sensor nodes could be drastically reduced. Consequently, sensor data processing techniques, including data aggregation, data compression, modeling and online querying, should be performed *on-site* or *in-network* to reduce communication cost [Subramaniam and Gunopulos 2007]. Furthermore, numerous demands on efficient data processing algorithms for sensor systems arise due to the limitations of computational power of sensor nodes as well as the existence of inaccuracy and bias in the sensor readings. In other scenarios, such as stock market and network monitoring systems, there also exist challenges in processing high-rate data streams.

2.1.1 Query Processing. There are several important queries to be considered [Subramaniam and Gunopulos 2007]:

- Aggregate Queries.* Aggregate Queries is an important class of queries in sensor systems, including MIN, COUNT and AVG operators. Various techniques have been proposed to efficiently process these aggregate operators in sensor systems, which can help to effectively reduce power consumption. Considering the properties of the aggregate functions, the in-network partial data could be preprocessed first, which can then be utilized to produce the final results for the issued queries.
- Join Queries.* An example of join queries is “Return the objects that were detected in both regions R1 and R2” [Subramaniam and Gunopulos 2007]. To evaluate the query, stream readings from the sensors in regions R1 and R2 should be joined first before we can determine whether an object was detected in the two designated regions. Join queries are useful in many applications, such as monitoring an environment where multiple sensors are deployed, tracking moving objects that are monitored by several types of sensors, etc.
- Top-k Monitoring.* Babcock and Olston [2003] investigated the general problem of monitoring top-k values from distributed data streams. A technique is proposed to ensure the validity of the most recently communicated top-k answers by maintaining some specified arithmetic constraints at the stream sources. User specified error tolerance is also considered in order to provide high-quality answers. This technique can help reduce the overall communication cost between different sources.
- Continuous Queries.* To monitor designated changes in an environment, sensors are typically required to answer queries in continuous manner. For instance, motion or sound sensors might be used to evaluate some continuous queries, such as “Turn lights off if no motion is detected in area A in the past 10 minutes”. When the query constraints are satisfied, the action of turning lights off could be automatically triggered by these sensors. If there are more than one continuous query evaluated over the same sensor readings, the storage and computation can be optimized by exploiting the fact that the sources of the queries and their partial results could overlap [Subramaniam and Gunopulos 2007].

2.1.2 *Stream Mining.* Stream mining can extract useful rules/information from data streams. The following lists some typical tasks for stream mining:

- Clustering.* Clustering is the task of grouping a set of objects in such a way that objects in the same group (called a cluster) are more similar to each other than to those in other groups (clusters). Clustering techniques for data streams typically continuously cluster objects on memory constrained devices with some time limitations. Due to these restrictions, there are some requirements to consider when designing algorithms for clustering data streams [Gama 2010]: (i) providing clustering results via fast and incremental processing of data objects; (ii) rapidly detecting new clusters or changes of existing clusters; (iii) scaling to the potentially unbounded number of objects in data streams; (iv) providing a model representation that is consistently compact regardless the number of data objects; (v) rapidly detecting the presence of outliers and acting accordingly; and (vi) dealing with different data types, such as XML trees, DNA sequences, GPS temporal and spatial information.
- Classification.* Classification [Wang and Liu 2011] uses prior knowledge to guide

the partitioning process to construct a set of classifiers to represent the possible distribution of patterns. Basically, compared with clustering, classification is a supervised learning process whereas clustering is an unsupervised learning process. More formally, a typical classification algorithm can be defined as follows [Wang and Liu 2011]: given a predefined classifier and two sets of data, labeled data and unlabeled data, the labeled data is used to train the classifier and the unlabeled data can then be classified by the trained classifier.

- *Outlier and Anomaly Detection.* In outlier and anomaly detection, the main task is to find data points that are most different from the remaining points in a given data set. Most existing outlier detection algorithms are based on the distance between every pair of points. The points that are most distant from all other points will be marked as outliers [Knorr and Ng 1998]. This kind of algorithms suffer from the same performance issue as they all run in $O(n^2)$ time. Hence, it is difficult to extend such approaches to distributed streaming data sets because points in those data sets normally arrive at multiple distributed end-points and must be processed incrementally.
- *Frequent Itemset Mining.* Frequent itemset mining is to find sets of items or values that co-occur frequently, or in other words, to find co-occurrence relationships in a transactional data set. Here a transactional data set refers to a data set where a set of items appear together in some specified context. Given a predefined support s , the goal in frequent itemset mining is to find all subsets of items that occur at least s number of times, or in other words, that appear in at least s transactional data sets at hand. Frequent itemset mining is both CPU and I/O intensive. Therefore, it is costly to completely re-mine a dynamic data set, which will be a typical case in IoT.

2.1.3 Data Stream Processing in IoT. In IoT, multiple data streams processing [Guo et al. 2011; Xie et al. 2012] would be more preferable as data streams can be generated at anywhere around the world and can be accessed globally via the Internet if being made public. For example, SmartSantander² proposes a city-scale experimental research facility in support of typical applications and services for a smart city. Around 20,000 sensors have been deployed to provide a variety of services, such as static environmental monitoring, mobile environmental monitoring, parks and gardens irrigation, outdoor parking area management, guidance to free parking lots and traffic intensity monitoring. A large number of data streams have to be processed efficiently to provide real-time monitoring of a smart city.

2.2 RFID Data Stream Processing

In 2003, a nonprofit open forum called the Ubiquitous ID Center³ was established. So far, more than 500 companies and organizations worldwide have contributed to it, publishing uID standards and industrial open standard specifications. uID standards are based on the uID architecture [Koshizuka and Sakamura 2010], which identifies real-world entities via Radio-Frequency Identification (RFID) tags or barcodes, determines contextual information such as environmental parameters from

²<http://www.smartsantander.eu/>

³uID Center: www.uidcenter.org

networked sensors, and adapts information services according to the data it obtains.

RFID systems consist of radio frequency (RF) tags (also called transponders) and RF tag readers (also called transceivers). Readers may be able to both read data from and write data to a transponder. RFID is a promising electronic identification technology that enables real-time monitoring and tracking applications in a variety of domains. Streams of RFID data, whose basic form is a triplet $\langle tagid; readerid; timestamp \rangle$, raise new challenges since the data may be insufficient, incomplete, and voluminous [Sheng et al. 2008].

2.2.1 RFID Data Cleaning (Uncertainty and Unreliability). SMURF (Statistical sMoothing for Unreliable RFid data) [Jeffery et al. 2006] is the first declarative, adaptive smoothing filter for cleaning raw RFID data streams. Unlike conventional techniques which expose the smoothing window parameter to the application, SMURF adapts the window size automatically and continuously over the lifetime of the system based on observed readings.

Periods of dropped readings and periods when a tag has moved are difficult to distinguish, which poses some challenges for the design of SMURF. To overcome such difficulty, a statistical sampling-based approach is put forward in SMURF. The main motivation is that RFID data streams can be modeled as a random sample of the tags in a reader's detection range. This sample-based view of observed RFID readings enables SMURF to develop algorithms based on statistical sampling theory to adapt the window size effectively. Basically, the false reads in RFID streams can be classified into two categories [Liao et al. 2011]:

- Missing-Reads.* Though an RFID tag indeed locates in the range of a reader, it might not be read at all, thereby leading to a false prediction that the tag is not present. This may be caused by the weakness of RF signal, shortage of power, shield of signal between the tag and the reader, and the collision between tags. This type of errors is also referred as false negatives.
- Cross-Reads.* When an RFID tag locates outside the range of a reader, but it might be captured by this reader which leads another false prediction that the object is present in the scope of this reader (sometimes called *ghost reading*). Cross-reads may be arisen by the reflection of metal items, the abrupt strengthen of RF, and the change of antenna directions. This type of errors is also called false positives.

SMURF cannot eliminate the cross-reads generated by physical factors. A kernel density-based probability cleaning method, called KLEAP, can be used to filter the cross-reads in RFID data streams [Liao et al. 2011]. KLEAP considers cross-reads as outliers, thus, the determination of cross-reads is transformed into the issue of detecting outliers on data streams. The density-based methods often perform well than the distance-based one, so KLEAP applies the density-based methods to detect cross-reads. It detects the exact positions of tags over the RFID data streams through examining the kernel densities of each tag captured by multiple readers.

Fazzinga et al. [2014] exploit the knowledge on the map of the real world and on the motility characteristics (such as the maximum speed) of the monitored objects, even if the users who analyze the data are typically acquainted with these aspects.

From this knowledge of the domain, constraints can be naturally derived on the connectivity between pairs of locations (direct unreachability constraints) and/or on the time needed for reaching a location starting from another one (traveling-time constraints). These constraints can be used to discard interpretations of the data corresponding to inconsistent trajectories. Then a graph is built in the following way: its nodes correspond to pairs $\langle location, timestamp \rangle$ and inside the graph, paths from source to target nodes one-to-one correspond to the valid trajectories in real world. Each node or edge is assigned a probability obtained by revising the a priori probability of the corresponding pair $\langle location, timestamp \rangle$, so that the overall probability of a source-to-target path is the conditioned probability of the corresponding trajectory. In this way, trajectories of RFID-monitored objects can be cleaned.

2.2.2 RFID Data Inference and Compression. RFID data inference techniques are closely related to RFID data cleaning techniques because inference techniques will need to clean RFID data first and then they can infer to the high level information about the tagged objects, i.e., location and containment relationships. Since raw RFID data contains a large amount of redundancies, RFID data compression is also applied to reduce space requirements after inference results have been obtained. RFID data compression is a further step beyond inference, where compression is performed based on the results of inference to remove the redundant data.

Tran et al. [2009] consider noisy, raw data streams from mobile RFID readers and employ a probabilistic approach to translate these streams into clean, rich event streams with location information. Their probabilistic model is built based on the mobility of the reader, object dynamics, and noisy readings. *Particle filtering* is used to infer clean information about object locations from raw streams captured from mobile RFID readers.

The aforementioned data cleaning and inference techniques focus on smoothing over time, where containment relationships are not considered. Containment refers to *inter-object* relationships, e.g., containment between objects, cases, and pallets. Containment queries can be useful for enforcing packaging and shipping regulations. Cao et al. [2011b] provide some examples of containment queries, such as “raise an alert if a flammable item is not packed in a fireproof case” or “verify that the food containing peanuts is never exposed to other food cases for more than an hour”. They also observe that some known containment relations can be used to determine object locations by smoothing over these facts. For example, suppose that we can infer that a specific set of objects have been packed in the same container. According to such knowledge, if one object in the container is read, all of the other objects must be in the same place. However, the fact is that the containment relationships are not known in advance. Therefore, a graphic model is proposed to infer containment relationships and to detect changes in containment relationships [Cao et al. 2011b; Nie et al. 2012].

2.3 RDF Triple Stream Processing

Linked Data is a method for publishing structured data and interlink such data to make it more useful⁴. It builds upon standard Web technologies such as HTTP, RDF and URIs and extends these technologies to share information. Linked Data can be understandable by computers. Data from different sources can be connected and queried in the form of Linked Data. Basically, Linked Data refers to a set of best practices to be followed in order to publish and link data on the Web, using the following basic principles⁵:

- Use URIs as names for things.
- Use HTTP URIs so that people can look up those names.
- When someone looks up a URI, provide useful information, using appropriate standards (RDF, SPARQL).
- Include links to other URIs, so that more things can be discovered.

The concept of *Linked Stream Data* applies the Linked Data principles to streaming data, so that data streams can be published as part of the Web of Linked Data. Stream reasoning can provide the abstractions, foundations, methods and tools required to integrate data streams, the Semantic Web and reasoning systems. Substantial research efforts have been put forward, focusing on how to apply reasoning on streaming data, how to publish raw streaming data and connect them to the existing data sets on the Semantic Web, and how to extend the SPARQL query language to process streaming data [Zhang et al. 2012]. These research efforts lay some foundations of semantic IoT technologies, facilitating machine-to-machine communication in IoT.

2.3.1 *Linked Stream Processing and Reasoning.* Sequeda and Corcho [2009] initiate efforts to apply the linked data principles to stream (sensor) data, so that this wealth of information could be easily included in the Linked Data cloud⁶.

There are three typical streaming RDF/SPARQLS engines, including Streaming SPARQL [Bolles et al. 2008], SPARQLStream [Calbimonte et al. 2010], C-SPARQL [Barbieri et al. 2010], and EP-SPARQL [Anicic et al. 2011]. Each of these systems also proposes its own SPARQL extension for streaming data processing. In these studies, SPARQL has been extended to have sliding window operators for RDF stream processing.

For example, Streaming SPARQL extends SPARQL to support window operators. But it does not consider performance issues, specially when designing the data structures. Further, it does not consider the sharing of computing states for continuous execution. Another example is SPARQLStream, which aims at enabling ontology-based access to streaming data. It defines a SPARQLStream language, which can be translated into another relational stream language based on mapping rules.

C-SPARQL (Continuous SPARQL) [Barbieri et al. 2010] attempts to facilitate reasoning upon rapidly changing information. In C-SPARQL, continuous queries

⁴en.wikipedia.org

⁵<http://www.w3.org/DesignIssues/LinkedData.html>

⁶<http://linkeddata.org/>

Table I. Comparisons of Linked Stream Processing and Reasoning.

Approach	Native	Aggregation Support	Reasoning Support	SPARQL 1.1 Support
Streaming SPARQL [Bolles et al. 2008]	No	Limited	Limited	Limited
SPARQLStream [Calbimonte et al. 2010]	No	Limited	Limited	Limited
C-SPARQL [Barbieri et al. 2010]	No	Rich	Limited	Limited
EP-SPARQL [Anicic et al. 2011]	No	Limited	Rich	Limited
CQELS [Phuoc et al. 2011]	Yes	Limited	Limited	Limited

are divided into static and dynamic parts and streaming data is transformed into non-streaming data within a specified window in order to apply standard algebraic operations, such as aggregate functions like `COUNT`, `COUNT DISTINCT`, `MAX`, `MIN` and `AVG`. The static parts will be loaded into relations, and the continuous queries are executed by processing the stream data against these relations. Event Processing SPARQL (EP-SPARQL), a language to describe event processing and stream reasoning, can be translated to ETALIS [Anicic et al. 2011], a Prolog-based complex event processing framework. First, RDF-based data elements are transformed into logic facts, and then EP-SPARQL queries are translated into Prolog rules.

Different from the above approaches, CQELS [Phuoc et al. 2011] is a native streaming RDF/SPARQL system built from scratch. CQELS defines and implements a native processing model in the query engine. Its query execution framework can also dynamically adapt the query processor to changes in the input data. By using data encoding and caching of intermediate query results, CQELS reduces external disk access on large Linked Data collections. Some indexing techniques are also adopted to enable faster data access. Table I compares all these systems from various aspects.

2.3.2 Extracting RDF Triples from Unstructured Data Streams. Gerber et al. [2013] point out that although the current Linked Open Data (LOD) cloud has tremendously grown over the last few years, it delivers mostly encyclopedic information (such as albums, places, kings, etc.) and fails to provide up-to-date information. Based on such observation, they develop RdfLiveNews, an approach that allows extracting RDF from unstructured (i.e., textual) data streams in a fashion similar to the live versions of the DBpedia⁷ and LinkedGeoData⁸ datasets. RdfLiveNews takes unstructured data streams as its input. It firstly removes duplicates in the streams. Then it uses the cleaned streams as a basis to extract patterns for relations between known resources. Next, the patterns will be clustered to labeled relations and finally will be used as a basis for generating RDF triples.

⁷<http://live.dbpedia.org/sparql>

⁸<http://live.linkedgeo.org/sparql>

3. DATA STORAGE MODELS

The nature of data produced by the Internet of Things calls for a revisit of data storage techniques, which will be further discussed in this section.

3.1 New Architecture

Traditional Database Management Systems (DBMSs) employ record-oriented (i.e., a record is represented by a row in a relational table) storage systems. With this row store architecture, a single disk write is able to store a single record with multiple attributes to disk. Records writes and updates are normally of high performance in these systems. Therefore, a DBMS with a row store architecture can be called a *write-optimized* system. In contrast, some systems need to deal with ad-hoc querying of large amounts of data, where read performance is of more importance. For such systems, *read-optimized* is the major design factor. Take data warehouses as an example. They represent one class of read-optimized system. In these read-optimized systems, a column-store architecture is a better choice. This is because in a column-store system, the values for each single column (or attribute) are stored contiguously, which can be easily optimized for high-performance querying.

Stonebraker et al. [2005] designed C-Store, a column-store architecture that supports the standard relational logical data model. Compared with the traditional DMBS architecture, the major differences are: (i) data in C-Store is not physically stored using its related relational logical data model; and (ii) whereas most row stores implement physical tables directly and then add various indexes to speed access, C-Store implements only projections. Here, projections are sorted subsets of the attributes of a table. Furthermore, Stonebraker et al. [2007] show superior performance of column store based systems over the major RDBMS (relational DBMS) system. It is experimentally demonstrated that specialized engines in the data warehouse, stream processing, text, and scientific database markets can speed up the querying performance by 1-2 orders of magnitude using the column-store architecture. They also suggest that the DBMS vendors (and the research community) should start from scratch and design novel systems for requirements to be fulfilled in the near future, rather than just adapting current systems for those new requirements.

3.2 Large-Scale Storage in Distributed Environments

Storage issues in large scale systems have arisen due to the arrival of the big data era. For example, users of websites such as Facebook, Ebay and Yahoo! usually demand fast response times. One solution for this is to replicate data across globally distributed datacenters. However, Kadambi et al. [2011] discover that to replicate all data to all locations may waste huge amounts of resources since users from different locations may have different data consumption needs. For example, an European server may not need to maintain a replica of some rare accessed records in an Asian server. By exploiting such observations and selectively replicating large-scale web databases on a record-by-record basis, bandwidth and disk costs can be saved.

To meet the exceptional demands of data storage in IoT, developments of large-scale, distributed storage systems are of essential. There are three factors or re-

Table II. Comparisons of three types of distributed storage systems.

Type	Pros	Cons	Representatives
CA	<ul style="list-style-type: none"> —Single copy of data —Consistency is easily ensured —Availability is assured by the excellent design of databases 	<ul style="list-style-type: none"> —Could not handle network failures 	<ul style="list-style-type: none"> —Traditional small-scale relational databases
CP	<ul style="list-style-type: none"> —Maintain several copies of the same data —A certain level of fault tolerance is ensured —Consistency is ensured by guaranteeing multiple copies of data to be identical 	<ul style="list-style-type: none"> —Could not ensure sound availability due to the high cost for consistency assurance 	<ul style="list-style-type: none"> —BigTable [Chang et al. 2008] —Hbase [Apache 2014]
AP	<ul style="list-style-type: none"> —Maintain several copies of the same data —A certain level of fault tolerance is ensured —Availability is assured by the design of distributed storage systems 	<ul style="list-style-type: none"> —Strong consistency is not ensured —May cause a certain amount of data errors 	<ul style="list-style-type: none"> —Dynamo [DeCandia et al. 2007] —Cassandra [Lakshman and Malik 2010]

requirements to be considered when designing a distributed storage system [Chen et al. 2014]:

- Consistency*: Consistency means to ensure that multiple copies of the same data are identical since server failures and parallel storage may cause inconsistency.
- Availability*: Availability refers to the requirement that the entire distributed storage system (which contains multiple servers) should not be seriously affected by some extent of server failures and should be able to provide satisfactory reading and writing performance.
- Partition Tolerance*: Since multiple servers are interconnected by a network and the data is partitioned across the network, the distributed storage system should have a certain level of tolerance to problems caused by network failures. This refers to partition tolerance requirement.

Interestingly, it has been proved by Gilbert and Lynch [2002] that a distributed storage system could not simultaneously meet the requirements on consistency, availability, and partition tolerance, and at most *two of the three requirements* can be satisfied at the same time. On top of this theory, there are three types of distributed storage systems: (1) a CA system, which ignores partition tolerance; (2) a CP system, which ignores availability; and (3) an AP system, which ignores consistency. The comparisons of these systems and some of their representative works are summarized in Table II.

3.3 Storage on Resource-Constrained Devices

Storage issues also arise in resource-constrained scenarios in IoT. For example, in sensor networks, communication activity normally plays a more important role than storage. But Mottola [2010] argues that for batch data collection, delay-tolerant mobile applications, and disconnected operations in static networks, the storage-centric paradigm becomes more critical. It is favored by decreasing costs and increasing capacity of storage hardware. SQUIRREL is also proposed in the same work, which is a lightweight run-time layer allocating data to different storage areas, based on data size versus energy trade-offs.

Yang et al. [2009] developed SolarStore, a power storage service for solar-powered storage-centric sensor networks. The main goal of SolarStore is to improve the total amount of data that can be eventually retrieved from the network. It adaptively balances data reliability against data sensing since solar energy is renewable and dynamic. For example, it chooses to replicate data in the network until the next opportunity to upload data to the server. The degree of data replication also varies dynamically depending on the availability of solar energy and sensor storage.

Early database systems for sensor networks such as TinyDB and Cougar only act as filters for data collection networks and not as databases, i.e., no data is stored in, or retrieved from, any database. Tsiftes and Dunkels [2011] presented a database management system for resource-constrained sensors named Antelope. Antelope supports run-time creation and deletion of databases and indexes and hence is a dynamic database system. It is the first DBMS for resource-constrained sensor devices, which enables a class of sensor network systems where every sensor holds a database. Tsiftes and Dunkels [2011] also envisioned that database techniques would become increasingly important in the progress of sensor network applications and energy-efficient storage. Further, indexing and querying would play important roles in emerging storage-centric applications.

Besides, Nath [2009] proposes to use flash storage for logging data on a sensor node, called *amnesic storage* systems. An amnesic storage system archives streaming data using two key techniques: (i) data is compressed (usually with lossy compression methods) in an online fashion before being archived; and (ii) an amnesic storage system uses aging archived data by reducing the fidelity of older data to make space for newer data.

4. SEARCH TECHNIQUES

Searching and finding relevant objects from billions of things is one of the major challenges for the future Internet of Things and can bring about huge potential impact to humans [Barnaghi et al. 2013; Christophe et al. 2011; Wang et al. 2010; Ostermaier et al. 2010]. Supporting technologies for searching things in the IoT are very different from those used in searching Web documents because things are tightly bound to contextual information (e.g., location) and have no easily indexable properties (e.g., human readable text in the case of Web documents). In addition, the state information of things is dynamic and rapidly changing. Things discovery calls for innovative ways of managing and searching from dynamic data, which makes it different from traditional Web searching. This section overviews the relevant areas such as the Deep Web, Semantic Web and then discusses state-

of-the-art techniques in searching things in the IoT environments.

4.1 Deep Web and Semantic Web

Deep Web refers to data stored outside Web pages and accessible from the Web, typically through HTML forms. The size of the Deep Web is estimated to be several orders of magnitude larger than that of the so-called *Surface Web* (the Web that is accessible and indexable by text search engines) [Cali and Martinenghi 2010].

Executing structured, high-level queries on deep web data sources involves a number of challenges because query execution engines have a very limited access to data [Wang and Agrawal 2011]. Besides, hidden data on the Deep Web may not provide the domain information for an attribute. Hence, Jin et al. [2011] argue that *domain discovery* becomes a critical challenge as a broad range of existing techniques on third-party analytical and mash-up applications⁹ are being applied over hidden databases. Furthermore, the traditional way to access hidden data on the Deep Web (by manually filling-up HTML forms on search interfaces) is not scalable given the growing size of the Deep Web. Therefore, Khare et al. [2010] argue that automatic access to such hidden data requires an automatic understanding of search interfaces by computers, which would be challenging, as the interfaces are originally designed for human access.

The Deep Web provides a wealth of hidden data in semi-structured form, accessible through Web forms and Web services. Since the data is hidden, to reach the whole content of the World Wide Web by just following hyperlinks is impossible. Regarding such issues, on top of XML, the Semantic Web grows as a common structured data source [Suchanek et al. 2011]. With the W3C standards Resource Description Framework (RDF) and Web Ontology Language (OWL), the Semantic Web aims to unify the way semantic information is stored and exchanged. The Semantic Web makes it possible for machines themselves to not just read, but also “understand” the data from data sources, which enables machine to machine communication. In particular, languages such as as Microformats¹⁰ and schema.org, can be used to add semantics to the descriptions of Web resources (including things).

4.2 Web Search

The frequent changes and the unprecedented scale of the Web together pose enormous challenges to Web search engines, making it challenging to provide the most up-to-date and highly relevant information to its users [Cho and Garcia-Molina 2010]. In IoT, this may become even more challenging as things would scale up the Web further and make the Web change more rapidly. For example, Tsubuyaku Sensor¹¹ is a new wireless device from Japanese Ubiquitous Computing Technology. It can monitor conditions such as temperature, humidity and radiation levels. It then automatically tweet the resulting data via Twitter. In this way, a sensor becomes a virtual Twitter user, which can actively post tweets on the Web.

⁹Here, a mash-up is a Web page or application that integrates complementary elements from two or more sources.

¹⁰<http://www.microformats.org>

¹¹<http://ts.uctec.com/tsensor/index-e.php>

4.2.1 Information Extraction. Information extraction from the Web is of growing importance. For example, objects on the Web are often associated with many attributes that describe the objects, making it essential to extract such attributes and map them to their corresponding objects. However, much attribute information about an object is hidden in the dynamic user interaction and is not on the Web page that describes the object. Huang et al. [2011] build a search model for exploratory Web sites, and algorithms are also proposed for identifying, clustering, and relationship mining of related Web pages based on the model. Besides, Satpal et al. [2011] study the problem of extracting structured data from web pages taking into account both the content of individual attributes as well as the structure of pages and sites. Markov Logic Networks (MLNs) are adopted to capture both content and structural features in a single unified framework, which is able to introduce more accurate inference.

The corpus of a search engine forms a rich source of information of analytical interest to third parties, but the only available access is by issuing search queries through its interface. Zhang et al. [2011] claim that, in order to support data analytics over a search engine's corpus, it is necessary to address two main problems: (i) the sampling of documents (for offline analytics); and (ii) the direct (online) estimation of aggregates. Meanwhile, in order to complete data analytics tasks, only a small number of queries would be issued to a search engine through the keyword-search interface.

4.2.2 Real-time Web Search. Real-time web search refers to the retrieval of very latest content which is in high demand. It is reported that Twitter handled more than 50 million tweets per day. Providing real-time search service is indeed very challenging in such large-scale microblogging systems because thousands of new updates need to be processed per second [Chen et al. 2011].

Sakaki et al. [2010] observe that Twitter is real-time micro-blogging and investigate the real-time interaction of events such as earthquakes in Twitter. They consider each Twitter user as a virtual sensor and apply Kalman filtering and particle filtering for estimating the centers of earthquakes and the trajectories of typhoons. Similarly, Dong et al. [2010] identify two challenges not encountered in non-real-time web search when supporting real-time web search, which are (i) quickly crawling relevant content and (ii) ranking documents with link and click information. Then they propose to use the micro-blogging data stream to detect fresh URLs and to compute novel and effective features for ranking fresh URLs based on micro-blogging data.

4.2.3 Searching information over RDF Data. Searching information from RDF data is important as more and more information is published in the form of RDF (e.g., via Linked Open Data Cloud). Efficient management of RDF data is also an important factor in realizing the Semantic Web vision [Abadi et al. 2007]. Performance and scalability issues need to be addressed as the Semantic Web technology is applied to real-world applications. Unlike the relational database community, the Semantic Web community uses a very different data model, which is RDF.

Tsatsanifos et al. [2011] present MIDAS-RDF, a distributed P2P RDF/S repository that is built on top of a distributed multi-dimensional index structure. It

features fast retrieval of RDF triples satisfying various pattern queries by translating them into multi-dimensional range queries, which can be processed by the underlying index in hops logarithmic to the number of peers. Further, since in IoT, data uncertainty is critical, Lian and Chen [2011] address the problem of efficiently answering queries on probabilistic RDF data graphs. They model RDF data by probabilistic graphs, and an RDF query is equivalent to a search over subgraphs of probabilistic graphs that have high probabilities to match with a given query graph.

4.2.4 Collaborative Web Search. Web search engines often answer user queries based on data and information in relevant structured databases, which will be searched in isolation. Since a single database may not contain sufficient information to answer the query, the search often produces empty or incomplete results. Motivated by this observation, Agrawal et al. [2009] exploit web search results and the items in structured databases together to produce more complete answers to a wide range of queries that traditional web search cannot support well. Take query “light-weight gaming laptop” as an example. Dell XPS M1330 should be considered a match to such query as it is a light-weight laptop and suitable for gaming. But if searching only the query keywords {**light-weight**, **gaming**} on the Web, Dell XPS M1330 may not appear in the search results. Therefore, Agrawal et al. [2009] propose to utilize the web search results (e.g. a set of relevant web documents) to help identify relevant information in some structured databases. Then the user queries could be better answered.

Similarly, Chaudhuri et al. [2009] exploit web search engines in order to define new similarity functions for recognizing named entities such as products, people names, or locations from documents, such as “X61” and the entity “Lenovo ThinkPad X61 Notebook”. The proposed new similarity functions are more accurate than existing string-based similarity functions because they aggregate evidence from multiple documents, and exploit web search engines to measure similarity.

4.3 Search of Things in IoT

In IoT, connecting things enabled by RFID, embedded sensors and sensor networks to the Internet and publishing their output on the Web would become a reality. Real-world objects would have their own Web presence. Considering the potential and profound impact of IoT technologies, search of things in IoT will become as important as today’s document search on the Web [Ostermaier et al. 2010].

4.3.1 Key Words based Search of Things. Unlike search engines such as Google, searching for information in the physical world is more difficult because the physical objects do not have (reliable) connections to virtual space. For example, online books can be easily discovered by searching but physical books at home may be more difficult to find. This observation motivates Wang et al. [2010] to propose Snoogle, a search engine for the physical world. The basic idea behind Snoogle is that sensor nodes carry a textual description of the object they will be attached to. Such description forms the keywords for search of things. Then the key words information of the whole sensor network is indexed using a two-tiered hierarchy. The lower tier contains many mediators, which are also called *index points*. Each index point maintains an aggregate view of all sensors in a local area (e.g. a room)

and every sensor in the same area will be assigned to the same index point. In the top tier, there is a single mediator called the *key index point*. The key index point will maintain an aggregate view of the whole network.

MAX, a system that users can easily locate objects, is also designed [Yap et al. 2008]. The main assumption is that tags are attached to everyday objects and each tag stores a descriptor of the object it is attached to (e.g., the book of **Harry Potter**). Multiple descriptor words are allowed in each tag, enabling users to label the object with richer information, so that others can locate the object based on the label. A three-tiered hierarchy of mediators is used. In the lowest tier, substations represent immobile objects such as tables or shelves, on which mobile tagged objects can be placed. In the middle tier, base stations represent a geographical space such as a room containing multiple substations. In the top tier, the MAX server represents the entire space covered by the system. When searching for an object left behind, it is easy to locate where the object has been left by exploiting the knowledge of which substation and base station it belongs to.

Microsearch is a system that runs on resource constrained small devices capable of being embedded into everyday objects [Tan et al. 2010]. It allows users to do textual search in the local storage of a stand-alone small device, without support from a backend server. The challenge is that Microsearch runs in a resource constrained platform, where conventional search engine design and algorithms cannot be used. Tan et al. [2010] also propose information retrieval (IR) techniques for query resolution which can answer top-k queries in a space-efficient manner.

Another search engine mainly designed for searching things, called Dyser, is proposed in [Elahi et al. 2009; Ostermaier et al. 2010]. Dyser allows users to search for real-world entities with a given state, such as “hot” or “cold”. However, this approach imposes two strong conditions: (i) to perform a query, end users have to know the vocabulary used by sensors (how states are named); and (ii) an entity must be represented by all the sensors that compose it. In order to estimate the probability of a sensor matching a query with sufficient accuracy and to rank sensor matching results, prediction models are adopted. The key idea of sensor ranking is to exploit the periodic nature of people-centric sensors by using appropriate prediction models.

4.3.2 Collaborative Search of Things. Frank et al. [2008] present a comprehensive system for managing and finding everyday objects relying on the collaboration of mobile phones in an urban area as object-sensing devices. For such tasks, the authors argue that the necessary infrastructures for such system include a sensing infrastructure, a communication infrastructure and a commercial infrastructure. Because of these requirements, the modern mobile phone system, which contains mobile sensors, provides a unique opportunity to realize collaborative search of everyday objects. The sensing model of the proposed system associates a probability with locations, meaning that the object currently has a certain probability of being at a certain location, thereby accelerating the search speed and reducing communication cost. Mobility provided by mobile sensors increases spatial coverage and hence the probability of finding a sought object.

5. COMPLEX EVENT PROCESSING

Data streaming techniques typically process incoming data through a sequence of transformations based on common SQL operators, like **selection**, **aggregate**, **join**, and these operators are defined in general by relational algebra. By contrast, the *complex event processing* (CEP) model views the information in the streams as events in the physical world. These events must be filtered, combined and transformed into higher-level events for better understanding by computers and humans [Cugola and Margara 2012]. Similar to traditional publish-subscribe systems, CEP systems allow subscribers to express their interest in composite events. The focus of CEP model is on detecting occurrences of particular patterns of (low-level) events indicating some higher-level events, which may be interesting to some particular event subscribers. In the era of IoT, CEP techniques lay part of the foundation of supporting computers to sense and react to events in the physical world.

5.1 Complex Event Processing

Systems for event processing and in particular event recognition (*event pattern matching*) accept a stream of time-stamped, simple or low-level events as input. A low-level event is the result of applying a computational derivation process to some other event, such as an event coming from a sensor. Using low-level events as input, (complex) event processing systems identify composite or high-level events of interest [Artikis and Paliouras 2014]. They are also collections of events that satisfy certain patterns.

SASE is a complex event processing system designed for monitoring queries over streams of RFID readings [Wu et al. 2006]. The SASE defines its own declarative event language that combines filtering, correlation, and transformation of events. The overall structure of the SASE language contains the **EVENT** clause specifying event patterns, the **WHERE** clause specifying qualifications and the **WITHIN** clause specifying window sizes. To meet the needs of RFID-enabled monitoring applications, several operators are also defined, including the **ANY** operator, the **SEQ_** operator, the **SEQ_WITHOUT** operator, the **Selection** operator and the **WITHIN_** operator. In order to process SASE queries, a query plan in SASE adopts a subset of six operators: **sequence scan**, **sequence construction**, **selection**, **window**, **negation**, and **transformation**. Pipelined execution of the above operators is used. More specifically, if a query matches a current event and some previous events, these events will be emitted from sequence scan and sequence construction immediately and form an event sequence. This event sequence is then pipelined through the subsequence operators, and added to the final output. To realize sequence scan, the basis of the whole process, Non-deterministic Finite Automata (NFA) are used.

Pattern matching over streams has been studied by Agrawal et al. [2008]. It presents two new challenges: (i) compared with languages for regular expression matching, languages for pattern matching over streams are significantly richer; and (ii) the conventional techniques for stream query processing are inadequate for efficient evaluation of pattern queries over streams. In order to represent each pattern query, a new query evaluation model is designed for processing pattern matching over RFID streams, employing a new type of automaton that comprises a non-

deterministic finite automaton (NFA) and a match buffer, named NFA^b [Agrawal et al. 2008]. Because of the powerful expressiveness of NFA, the semantics for the complete set of event pattern queries can be captured by the NFA^b model. Optimizations and query evaluation plans can also be produced and applied based on this model over event streams.

Liu et al. [2011] propose nested CEP language called NEEL to support the flexible nesting of **AND**, **OR**, **Negation** and **SEQ** operators at any level. One NEEL query example is given in Fig. 3, which expresses “a critical condition that after being recycled and washed, a surgery tool is being put back into use without first being sharpened, disinfected and then checked for quality assurance” [Liu et al. 2011]. Several techniques are also proposed to accelerate the evaluation of nested queries. Firstly, nested event expressions will be converted into normal forms by a normalization procedure. Secondly, a group of similar sub-expressions will be processed using prefix caching, suffix clustering methods and a customized physical execution strategy. Thirdly, an optimizer for optimal shared execution method is also designed based on the idea of iterative improvement. Compared with the traditional iterative nested execution, the optimized NEEL execution is up to two orders of magnitude faster. Ray et al. [2013] optimize the processing of Nested Complex Event Processing queries by designing the Continuous Sliding View structures for inner sub-queries.




Fig. 3. Nested CEP Query Example
(Adapted from [Liu et al. 2011])

Recent efforts have also been put on other aspects of complex event processing. For example, Heinze et al. [2013] study complex event processing in a distributed environment and propose FUGU – an elastic allocator for Complex Event Processing systems. FUGU can dynamically allocate and de-allocate both stateless and stateful queries in order to meet the utilization goals. To that end, FUGU relies on bin packing to allocate queries to hosts. Very recently, He et al. [2014] investigate load shedding techniques for complex event processing under various resource constraints. Like other stream systems, CEP systems often face bursty input data. Since over-provisioning the system to the point where it can handle any such burst may be uneconomical or impossible, during peak loads a CEP system may need to “shed” portions of the load. The key technical challenge is to selectively shed work in order to eliminate the less important query results, thereby preserving the more useful query results defined by some utility function. Motivated by this, several load shedding algorithms are designed, including CPU-bound load shedding,

memory-bound load shedding, and dual-bound load shedding (with both CPU- and memory-bound), depending on which resource is constrained.

5.2 Semantic Complex Event Processing

The combination of event processing and knowledge representation can lead to novel semantic-rich event processing engines [Teymourian and Paschke 2009; Zhou et al. 2011; Teymourian et al. 2012a]. These intelligent event processing engines can (i) help to understand what is happening in terms of events, (ii) state and know what reactions and processes it can invoke, and furthermore (iii) decide what new events it can signal. The identification of critical events and situations requires processing vast amounts of data and meta-data within and outside the systems.

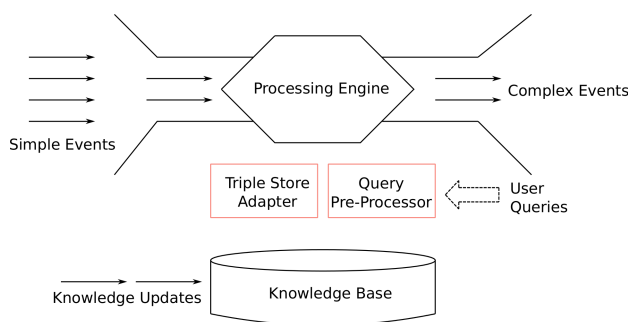


Fig. 4. Semantic Complex Event Processing System Overview
(Adapted from [Teymourian et al. 2012a])

5.2.1 Semantic CEP System. A semantic CEP system is shown in Fig. 4. Semantic models of events can improve event processing quality by using event meta-data in combination with ontologies and rules (i.e., *knowledge bases*). The fusion of background knowledge with data from an event stream can help the event processing engine to know more about incoming events and their relationships to other related concepts. A Knowledge Base (KB) can be used to provide background knowledge about the events and other non-event resources [Teymourian et al. 2012a]. This means that events can be detected based on reasoning on their type hierarchy, temporal/spatial relationships, or their relationship to other objects in the application domain.

The benefits of using background knowledge in complex event processing can be seen as two major advantages over state-of-the-art CEP systems. The first benefit is its higher expressiveness and the second one its flexibility. Expressiveness means that an event processing system can precisely express complex event patterns and reactions to events which can be directly translated into business operations. Flexibility means that a CEP system is able to integrate new business changes into the systems in a fraction of time rather than changing the whole event processing rules. Complex event patterns are independent of current businesses and are defined in a higher level of abstraction based on business strategies. When something is changed in the business environment, it can be considered simply as an update

in the background knowledge and the complex event detection patterns which are defined based on the business plans should not be changed.

5.2.2 Semantic Event Enrichment. The usage of background knowledge about events and their relations to other concepts in the application domain can improve the expressiveness and flexibility of CEP systems. Huge amounts of domain background knowledge stored in external knowledge bases can be used in combination with event processing in order to achieve more knowledgeable complex event processing. Teymourian et al. [2012b] discuss the benefits of background knowledge for event processing and describe different categories of event query rules.

Hasan et al. [2013] identify an information completeness problem in semantic event processing contexts from a different angle. For example, while the basic information item in an event-based system is an event, normal users often require the system to handle information that is not encoded in the event. Such information typically comes from legacy databases or web data sources. This requires some degrees of information completeness or incompleteness for events to be sufficient for tasks such as subscription matching. The process of reducing information incompleteness is called *event enrichment*. Several challenges are identified for event enrichment, including determination of the enrichment source, retrieval of information items from the enrichment source, finding complementary information for an event in the enrichment source and fusion of complementary information with the event. To address these challenges, a model based on unifying enrichment within the event consumer logic and a native enricher that tackles incompleteness before matching are proposed [Hasan et al. 2013].

5.2.3 Approximate Semantic Matching. Approximate semantic matching is first studied by Zhou et al. [2011]. To achieve approximate matching, semantic selection and inexact selection are used. More specifically, the semantic selection evaluates pattern constraints based on the semantic equivalence of attribute meanings captured by the event ontology instead of syntactic identical attribute values, while the inexact selection selects events and allows a limited number of mismatches to detect relevant patterns. A similarity function is associated with the inexact selection to evaluate relevance between matching patterns and target patterns.

Hasan et al. [2012] study approximate semantic matching of heterogeneous events. The motivation is that heterogeneous events are difficult to match in a distributed computing environment as similar or closely related events may not be described using the same words but in a semantically related form. To match all interesting events, users may have to write many slightly different subscriptions and have to know the exact format of all the heterogeneous events. Based such observation, semantic decoupling of events and user's subscriptions becomes necessary. However, after such decoupling, the subscriptions would hardly exactly match the descriptions of events. This indicates that approximate matching and processing of events are inevitable. A model for approximate semantic matching that addresses event semantic decoupling is proposed. The model is evaluated using a hybrid matching approach based on both thesauri, semantic similarity and relatedness measures. After adopting this technique, the number of event subscriptions to achieve sufficiently precise matching results can be greatly reduced because of the decoupling

between events and user subscriptions.

6. POTENTIAL IOT APPLICATIONS

As Ashton [2009] points out that IoT “has the potential to change the world, just as the Internet did”. The ongoing and/or visionary IoT applications show that IoT can bring significant changes in many domains, i.e. cities and homes, environment monitoring, health, energy, and business, etc. Mattern and Floerkemeier [2010] also argue that IoT can bring the ability to react to events in the physical world in an automatic, rapid and informed manner. This also opens up new opportunities for dealing with complex or critical situations and enables a wide variety of business processes to be optimized. In this section, we overview several representative domains where IoT can make some profound changes.

6.1 Smart Cities and Homes

IoT can connect billions of smart things and can help capture information in cities. Based on IoT, cities would become smarter and more efficient. Below are some examples of promising IoT applications in future smart cities. In a modern city, lots of digital data traces are generated there every second via cameras and sensors of all kinds [Guinard 2010]. All this data represents a goldmine for everyone, if people in the city would be able take advantage of it in an efficient and effective way. For example, IoT can facilitate resources management issues for modern cities. Specifically, static resources (e.g. fire stations, parking spots) and mobile resources (e.g. police cars, fire trucks) in a city can be managed effectively using IoT technologies [Gao et al. 2009]. Whenever events (fires, crime reports, cars looking for parking) arise, IoT technologies would be able to quickly match resources with events in an optimal way based on the information captured by smart things, thereby reducing cost and saving time. Taxi drivers in the city would also be able to better serve prospective passengers by learning passenger’s mobility patterns and other taxi drivers’ serving behaviors through the help of IoT technologies [Yuan et al. 2011]. Besides, one study estimated a loss of \$78 billion in 2007 in the form of 4.2 billion lost hours and 2.9 billion gallons of wasted gasoline in the United States alone [Mathur et al. 2010]. IoT could bring fundamental changes in urban street-parking management, which would greatly benefit the whole society by reducing traffic congestion and fuel consumption [Mathew et al. 2014].

Security in a city is of great concerns, which can benefit a lot from the development of IoT technologies. Losses resulted from property crimes were estimated to be \$17.2 billion in the U.S. in 2008 [Guha et al. 2010]. Current security cameras, motion detectors, and alarm systems are not able to help track or recover stolen property. IoT technologies can help to deter, detect, and track personal property theft since things are interconnected and can interact with each other. IoT technologies can also help improve stolen property recovery rates, and disrupt stolen property distribution networks. Similarly, a network of static and mobile sensors can be used to detect threats on city streets and in open areas such as parks [Liu et al. 2011].

With IoT technologies, people can browse and manage their homes via the Web. For example, they would be able to check whether the light in their bedrooms is on and could turn it off by simply clicking a button on a Web page. Similar

operations and management could be done in office environments. Plumbing is ranked as one of the ten most frequently found problems in homes [Lai et al. 2010]. It is important to determine the spatial topology of hidden water pipelines behind walls and underground. In IoT, smart things in homes would be able to report plumbing problems automatically and report to owners and/or plumbers for efficient maintenance and repair.

6.2 Environment Monitoring

IoT technologies can also help to monitor and protect environments thereby improving human’s knowledge about environments. Take water as an example. Dettweiler et al. [2010] propose that understanding the dynamics of bodies of water and their impact on the global environment requires sensing information over the full volume of water. In such context, IoT technologies would be able to provide effective approaches to study water. Also IoT could improve water management in a city. Drinking water is becoming a scarce resource around the world. In big cities, efficiently distributing water is one of the major issues [Guinard 2010]. Various reports show that on average 30% of drinkable water is lost during transmission due to the aging infrastructure and pipe failures. Further, water can be contaminated biologically or chemically due to inefficient operation and management. In order to effectively manage and efficiently transport water, IoT technologies would be of great importance.

Soil contains vast ecosystems that play a key role in the Earth’s water and nutrient cycles, but scientists cannot currently collect the high-resolution data required to fully understand them. Many soil sensors are inherently fragile and often produce invalid or uncalibrated data [Ramanathan et al. 2009]. IoT technologies would help to validate, calibrate, repair, or replace sensors, allowing to use available sensors without sacrificing data integrity and meanwhile minimizing the human resources required.

Sound is another example where IoT technologies can help. Sound is multidimensional, varying in intensity and spectra. So it is difficult to quantify, e.g., it is difficult to determine what kind of sound is noise. Further, the definitions and feelings of noise are quite subjective. For example, some noises could be pleasant, like flowing water, while others can be annoying, such as car alarms, screeching breaks and people arguing. Zimmerman and Robson [2011] design and build a device to monitor residential noise pollution to address the above problems. Firstly, noise samples from three representative houses are used, which span the spectrum of quiet to noisy neighborhoods. Secondly, a noise model is developed to characterize residential noise. Thirdly, noise events of an entire day (24 hours) are compressed into a one minute auditory summary. Data collection, transmission and storage requirements can be minimized in order to utilize low-cost and low-power components, while sufficient measurement accuracy is still maintained.

Intel has developed smart sensors that can warn people about running outside when the air is polluted¹². For example, if someone is preparing to take a jog along his/her regular route, an application on his/her smartphone pushes out a message:

¹²<http://www.fastcoexist.com/1680111/intels-sensors-will-warn-you-about-running-outside-when-the-air-is-polluted>

air pollution levels are high in the park where he/she usually runs. Then he/she could try a recommended route that is cleaner. Currently, many cities already have pollution and weather sensors. They are usually located on top of buildings, far from daily human activities.

6.3 Health

In future IoT environments, an RFID-enabled information infrastructure would be likely to revolutionize areas such as healthcare, and pharmaceutical [Garfinkel and Rosenberg 2005]. For example, a healthcare environment such as a large hospital or aged care could tag all pieces of medical equipment (e.g., scalpels, thermometers) and drug products for inventory management. Each storage area or patient room would be equipped with RFID readers that could scan medical devices, drug products, and their associated cases. Such an RFID-based infrastructure could offer a hospital unprecedented near real-time ability to track and monitor objects and detect anomalies (e.g., misplaced objects) as they occur.

As personal health sensors become ubiquitous, they are expected to become interoperable. This means standardized sensors can wirelessly communicate their data to a device many people already carry today (e.g., mobile phones) [Cornelius and Kotz 2011]. Lester et al. [2009] argue that one challenge in weight control is the difficulty of tracking food calories consumed and calories expended by activity. Then they present a system for automatic monitoring of calories consumed using a single body-worn accelerometer. To be fully benefited from such data for a large body of people, applying IoT technologies in such area would be an promising direction.

Mobile technology and sensors are creating ways to inexpensively and continuously monitor people's health. Doctors may call their clients to schedule an appointment, rather than vice-versa, because the doctors could know their clients' health conditions in real-time. Some projects for such purpose have been initiated. For example, EveryHeartBeat¹³ is a project for Body Computing to "connect the more than 5 billion mobile phones in the world to the health ecosystem". In the initial stage, heart rate monitoring is investigated. Consumers would be able to self track their pulse and studies show heart rate monitoring could be useful in detecting heart conditions and enabling early diagnosis. The future goal is to include data on blood sugar levels, and other biometrics collected via mobile devices.

6.4 Energy

Home heating is a major factor in worldwide energy use [Scott et al. 2011]. In IoT, home energy management applications could be built upon embedded Web servers [Priyantha et al. 2008]. Through such online web services, people can track and manage their home energy consumption. Gupta et al. [2009] present a system for augmenting these thermostats using just-in-time heating and cooling based on travel-to-home distance obtained from location-aware mobile phones. The system makes use of a GPS-enabled thermostat which could lead to savings of as much as 7%. In IoT, as things in homes would become smart and connected to the Internet, similar energy savings could be more effective. For example, by automatically sensing occupancy and sleep patterns in a home, it would be possible to save en-

¹³<http://join.everyheartbeat.org/>

ergy by automatically turning off the home’s HVAC (heating, ventilation, and air conditioning) system [Lu et al. 2010].

Besides home heating, fuel consumption is also an important issue. Ganti et al. [2010] develop GreenGPS, a navigation service that uses participatory sensing data to map fuel consumption on city streets. GreenGPS would allow drivers to find the most fuel-efficient routes for their vehicles between arbitrary end-points. In IoT, fuel consumption would be further reduced by enabling cars and passengers to communicate with each other for ride sharing [Yuan et al. 2011].

6.5 Business

IoT technologies would be able to help to improve efficiency in business and bring other impacts on business [Mattern and Floerkemeier 2010]:

- From a commercial point of view, IoT can help increase the efficiency of business processes and reduce costs in warehouse logistics and in service industries. This is because more complete and necessary information can be collected by interconnected things. owing to its huge and profound impact on the society, IoT research and applications can also trigger new business models involving smart things and associated services.
- From a social and political point of view, IoT technologies can provide a general increase in the quality of life for the following reasons. Firstly, consumers and citizens will be able to obtain more comprehensive information. Secondly, care for aged and/or disabled people can be improved with smarter assistance systems. Thirdly, safety can be increased. For example, road safety can be improved by receiving more complete and real-time traffic and road condition information.
- From a personal point of view, new services enabled by IoT technologies can make life more pleasant, entertaining, independent and also safer. For example, business taking advantages of technologies of search of things in IoT can help locate lost things quickly, such as personal belongs, pets or even other people.

Besides, take improving information handover efficiency in a global supply chain as an example. Stephan et al. [2010] propose the *digital object memories* (DOM), which can store order-related data via smart labels on the item. Based on DOM, relevant life cycle information could be attached to the product itself. Considering the potential different stakeholders including manufacturer, distributor, retailer, and end customer along the supply/value chain, this approach facilitates information handover.

Further, there are many important bits of information in an IoT-based supply chain, such as the 5W (what, when, where, who, which). It is also necessary to integrate them efficiently and in real-time in other operations. The EPCIS (Electronic Product Code Information System) network is a set of tools and standards for tracking and sharing RFID-tagged products in IoT. However, much of this data remains in closed networks and is hard to integrate [Wu et al. 2013]. IoT technologies could be used to make it easier to use all this data, to integrate it into various applications, and to build more flexible, scalable, global application for better (even real-time) logistics.

7. OPEN ISSUES

The development of IoT technologies and applications is merely starting off. Many new challenges and issues have not been addressed, which require substantial efforts from both academia and industry. In this section, we identify some key directions for future research and development from a data-centric perspective.

- Data Quality and Uncertainty*: In IoT, as data volume increases, inconsistency and redundancy within data would become paramount issues. Fan et al. [2010] indicate that one of the central problems for data quality is *inconsistency detection* and when data is distributed, the detection would be far more challenging. This is because inconsistency detection often requires shipping data from one site to another. Meanwhile, inherited from RFID data [Cao et al. 2011a] and sensor data [Prabhakar and Cheng 2009], IoT data would be of great uncertainty, which also presents significant challenges.
- Co-Space Data*: In an IoT environment, the physical space and the virtual (data) space co-exist, and interact simultaneously. Novel technologies must be developed to allow data to be processed and manipulated seamlessly between the real and digital spaces [Ooi et al. 2009]. To synchronize data in both real and virtual worlds, large amount of data and information will flow between co-spaces, which pose new challenges. For example, it would be challenging to process heterogeneous data streams in order to model and simulate real world events in the virtual world. Besides, more intelligent processing is needed to identify and send interesting events in the co-space to objects in the physical world.
- Transaction Handling*: When the data being updated is spread across hundreds or thousands of networked computers/smart things with differing update policies, it would be difficult to define what the transaction is. In addition, most of things are resource-constrained, which are typically connected to the Internet using light-weight, *stateless* protocols such as CoAP (Constraint Application Protocol)¹⁴ and 6LoWPAN (IPv6 over Low Power Wireless Personal Area Networks)¹⁵ and accessed using RESTful Web services. This makes transaction handling in IoT a great challenge. As James et al. [2009] point out that the problem is that the world is changing fast, the data representing the world is on multiple networked computers/smart things and existing database technologies cannot manage. Techniques developed for streamed and real-time data may provide some hints.
- Frequently Updated Timestamped Structured (FUTS) Data*: The Internet, and hence IoT, contains potentially billions of Frequently Updated Timestamped Structured (FUTS) data sources, such as real-time traffic reports, air pollution detection, temperature monitoring, crops monitoring, etc. FUTS data sources contain states and updates of physical world things. Current technologies are not capable in dealing with FUTS data sources [James et al. 2009] because: (i) no data management system can easily display FUTS past data; (ii) no efficient crawler or storage engine is able to collect and store FUTS data; and (iii) querying

¹⁴<http://tools.ietf.org/html/draft-ietf-core-coap-18>

¹⁵<http://tools.ietf.org/wg/6lowpan>

and delivering FUTURE data is hardly supported. All these pose great challenges for the design of novel data management systems for FUTURE data.

- Distributed and Mobile Data*: In IoT, data will be increasingly distributed and mobile [James et al. 2009]. Different from traditional mobile data, distributed and mobile data in IoT would be much more highly distributed and data intensive. In the context of interconnecting huge numbers of mobile and smart objects, centralized data stores would not be a suitable tool to manage all the dynamics of mobile data produced in IoT. Thus there is a need for novel ways to manage distributed and mobile data efficiently and effectively in IoT.
- Semantic Enrichment and Semantic Event Processing*: The full potentials of IoT would heavily rely on the progress of semantic Web. This is because things and machines should play a much more important role than humans in IoT to process and understand data. This calls for new research in Semantic technologies. For example, there are increasing efforts in building public knowledge bases (such as DBpedia, FreeBase, Linked Open Data Cloud, etc.). But how these knowledge bases can be effectively used to add to the understanding of raw data coming from sensor data streams and other types of data streams? To resolve this challenge, semantic enrichment of sensing data is a promising research direction. Further, consider the potential excessively large amount of subscriptions of IoT data. To produce proper semantic enrichment to meet different enrichment needs from different subscribers poses great challenges. Finally, how to effectively incorporate semantic enrichment techniques with semantic event processing to provide much better expressiveness in event processing is still at its initial stage. This will also demand a large amount of research efforts.
- Mining*: Data mining aims to facilitate the exploration and analysis of large amounts of data, which can help to extract useful information for huge volume of IoT data. Data mining challenges may include extraction of temporal characteristics from sensor data streams, event detection from multiple data streams, data stream classification, activity discovery and recognition from sensor data streams. Besides, clustering and table summarization in large data sets, mining large (data, information or social) networks, sampling, and information extraction from the Web are also great challenges in IoT.
- Knowledge Discovery*: Knowledge discovery is the process of extracting useful knowledge from data. This is essential especially when connected things populate their data to the Web. Weikum [2011] identify the following issues related to knowledge discovery in IoT: (i) automatic extraction of relational facts from natural-language text and multi-modal contexts; (ii) large-scale gathering of factual-knowledge candidates and their reconciliation into comprehensive knowledge bases; (iii) reasoning on uncertain hypotheses, for knowledge discovery and semantic search; and (iv) deep and real-time question answering, e.g., to enable computers to win quiz game shows.
- Security*: Due to the proliferation of embedded devices in IoT, effective device security mechanisms are essential to the development of IoT technologies and applications. National Intelligence Council [Anonymous 2008] argues that, to the extent that everyday objects become information security risks, the IoT could distribute those risks far more widely than the Internet has to date. For exam-

ple, RFID security presents many challenges. Potential solutions should consider aspects from hardware and wireless protocol security to the management, regulation and sharing of collected RFID data [Welbourne et al. 2009]. Lin et al. [2009] point out that establishing trust between a group of individuals remains a difficult problem. Besides, Lagesse et al. [2009] argue that there is still no generic framework for deploying and extending traditional security mechanisms over a variety of pervasive systems. Regarding security concerns of the network layer, Kounavis et al. [2010] suggest that the Internet can be gradually encrypted and authenticated based on the observations that the recent advances in implementation of cryptographic algorithms have made general purpose processors capable of encrypting packets at high rates. But how to generalize such algorithms to IoT would be challenging as things in IoT normally only maintain low transmission rates and connections are usually intermittent.

- Privacy*: Privacy protection is a serious challenge in IoT. One of the fundamental problems is the lack of a mechanism to help people expose appropriate amounts of their identity information [Zhu and Zhu 2009]. Embedded sensing is becoming more and more prevalent on personal devices such as mobile phones and multimedia players. Since people are typically wearing and carrying devices capable of sensing, details such as activity, location, and environment could become available to other people. Hence, personal sensing can be used to detect their physical activities and bring about privacy concerns [Klasnja et al. 2009].
- Social Concerns*: Since IoT connects everyday objects to the Internet, social concerns would become a hot topic in the development of IoT. For example, home is a private and intimate place. It may have multiple stakeholders competing priorities and tolerances for what is acceptable and useful [Choe et al. 2011]. To build smart homes in IoT, similar social concerns should be considered. Further, online social networks with personal things information may incur social concerns as well, such as disclosures of personal activities and hobbies, etc. Appropriate economic and legal conditions and a social consensus on how the new technical opportunities in IoT should be used also represents a substantial task for the future [Mattern and Floerkemeier 2010].

8. SUMMARY

It is widely predicted that the next generation of the Internet will be comprised of trillions of connected computing nodes at a global scale. Through these nodes, everyday objects in the world can be identified, connected to the Internet and take decisions independently. In this context, Internet of Things (IoT) is considered a new revolution of the Internet. In IoT, the possibility of seamlessly merging the real and the virtual worlds, through the massive deployment of embedded devices, opens up many new and exciting directions for both research and development. In this article, we have provided an overview of some key research areas of IoT, specifically from a data-centric perspective. It also presents a number of fundamental issues to be resolved before we can fully realize the promise of IoT applications. This article covers investigations on data models, data storage, stream processing, search and event processing. The most relevant application fields have also been reviewed.

Over the last few years, the Internet of Things has gained momentum and is

becoming a rapidly expanding area of research and business. Many efforts from researchers, vendors and governments have been devoted to creating and developing novel IoT applications. Along with the current research efforts, we encourage more insights into the problems of this promising technology, and more efforts in addressing the open research issues identified in this article.

REFERENCES

- ABADI, D. J., MARCUS, A., MADDEN, S., AND HOLLENBACH, K. J. 2007. Scalable Semantic Web Data Management Using Vertical Partitioning. In *Proceedings of the 33rd International Conference on Very Large Data Bases (VLDB)*. ACM, University of Vienna, Austria, 411–422.
- AGRAWAL, J., DIAO, Y., GYLLSTROM, D., AND IMMERMANN, N. 2008. Efficient pattern matching over event streams. In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD’08)*. 147–160.
- AGRAWAL, S., CHAKRABARTI, K., CHAUDHURI, S., GANTI, V., KÖNIG, A. C., AND XIN, D. 2009. Exploiting Web Search Engines to Search Structured Databases. In *Proceedings of the 18th International Conference on World Wide Web (WWW)*. ACM, Madrid, Spain, 501–510.
- ANICIC, D., FODOR, P., RUDOLPH, S., AND STOJANOVIC, N. 2011. EP-SPARQL: a unified language for event processing and stream reasoning. In *Proceedings of the 20th International Conference on World Wide Web (WWW)*. 635–644.
- ANONYMOUS. 2008. National Intelligence Council (NIC), Disruptive Civil Technologies: Six Technologies with Potential Impacts on US Interests Out to 2025, Conference Report CR 2008-07, April 2008, http://www.dni.gov/nic/NIC_home.html.
- APACHE. 2014. Apache HBase Project, <https://hbase.apache.org/>.
- ARTIKIS, A. AND PALIOURAS, G. 2014. Tutorial: Formal Methods for Event Processing. In *Proc. 17th International Conference on Extending Database Technology (EDBT)*. 675.
- ASHTON, K. 2009. That ‘Internet of Things’ Thing. <http://www.rfidjournal.com/article/view/4986>.
- ATZORI, L., IERA, A., AND MORABITO, G. 2010. The Internet of Things: A survey. *Computer Networks* 54, 15, 2787–2805.
- BABCOCK, B. AND OLSTON, C. 2003. Distributed Top-K Monitoring. In *Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data (SIGMOD Conference)*. 28–39.
- BARBIERI, D. F., BRAGA, D., CERI, S., AND GROSSNIKLAUS, M. 2010. An execution environment for C-SPARQL queries. In *Proceedings of the 13th International Conference on Extending Database Technology (EDBT)*. 441–452.
- BARNAGHI, P. M., SHETH, A. P., AND HENSON, C. A. 2013. From Data to Actionable Knowledge: Big Data Challenges in the Web of Things. *IEEE Intelligent Systems* 28, 6, 6–11.
- BIDDLECOMBE, E. 2005. UN predicts ‘Internet of Things’. <http://news.bbc.co.uk/2/hi/technology/4440334.stm>.
- BOLLES, A., GRAWUNDER, M., AND JACOBI, J. 2008. Streaming SPARQL - Extending SPARQL to Process Data Streams. In *Proceedings of the 5th European Semantic Web Conference on The Semantic Web: Research and Applications (ESWC)*. 448–462.
- BORT, J. 2011. 10 technologies that will change the world in the next 10 years. <http://www.networkworld.com/news/2011/071511-cisco-futurist.html>.
- CALBIMONTE, J.-P., CORCHO, Ó., AND GRAY, A. J. G. 2010. Enabling Ontology-Based Access to Streaming Data Sources. In *Proceedings of the 9th International Semantic Web Conference (ISWC)*. 96–111.
- CALÌ, A. AND MARTINENGI, D. 2010. Querying the Deep Web. In *Proceedings of the 13th International Conference on Extending Database Technology (EDBT)*. ACM, Lausanne, Switzerland, 724–727.
- CAO, Z., SUTTON, C., DIAO, Y., AND SHENOY, P. J. 2011a. Distributed Inference and Query Processing for RFID Tracking and Monitoring. *Proceedings of the VLDB Endowment* 4, 5, 326–337.

- CAO, Z., SUTTON, C. A., DIAO, Y., AND SHENOY, P. J. 2011b. Distributed inference and query processing for RFID tracking and monitoring. *Proceedings of the VLDB Endowment* 4, 5, 326–337.
- CASAGRAS. 2000. CASAGRAS (Coordination And Support Action for Global RFID-related Activities and Standardisation).
- CHANG, F., DEAN, J., GHEMAWAT, S., HSIEH, W. C., WALLACH, D. A., BURROWS, M., CHANDRA, T., FIKES, A., AND GRUBER, R. E. 2008. Bigtable: A Distributed Storage System for Structured Data. *ACM Trans. Comput. Syst.* 26, 2.
- CHAUDHURI, S., GANTI, V., AND XIN, D. 2009. Exploiting Web Search to Generate Synonyms for Entities. In *Proceedings of the 18th International Conference on World Wide Web (WWW)*. ACM, Madrid, Spain, 151–160.
- CHEN, C., LI, F., OOI, B. C., AND WU, S. 2011. TI: An Efficient Indexing Mechanism for Real-Time Search on Tweets. In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*. ACM, Athens, Greece, 649–660.
- CHEN, M., MAO, S., AND LIU, Y. 2014. Big Data: A Survey. *MONET* 19, 2, 171–209.
- CHO, J. AND GARCIA-MOLINA, H. 2010. Dealing with Web Data: History and Look ahead. *Proceedings of the VLDB Endowment* 3, 1, 4.
- CHOE, E. K., CONSOLVO, S., JUNG, J., HARRISON, B. L., AND KIENZT, J. A. 2011. Living in a Glass House: A Survey of Private Moments in the Home. In *Proceedings of the 13th International Conference on Ubiquitous Computing (UbiComp)*. ACM, Beijing, China, 41–44.
- CHRISTOPHE, B., VERDOT, V., AND TOUBIANA, V. 2011. Searching the 'Web of Things'. In *Proceedings of the 5th IEEE International Conference on Semantic Computing (ICSC)*. IEEE, Palo Alto, CA, USA, 308–315.
- COMMISSION, E. 2009. European Commission: Internet of Things, An action plan for Europe. http://eur-lex.europa.eu/LexUriServ/site/en/com/2009/com2009_0278en01.pdf.
- CORNELIUS, C. AND KOTZ, D. 2011. Recognizing Whether Sensors Are on the Same Body. In *Proceedings of the 9th International Conference on Pervasive Computing (Pervasive)*. Springer, San Francisco, CA, USA, 332–349.
- CUGOLA, G. AND MARGARA, A. 2012. Processing flows of information: From data stream to complex event processing. *ACM Comput. Surv.* 44, 3, 15.
- DE ANDRADE SILVA, J., FARIA, E. R., BARROS, R. C., HRUSCHKA, E. R., DE CARVALHO, A. C. P. L. F., AND GAMA, J. 2013. Data stream clustering: A survey. *ACM Comput. Surv.* 46, 1, 13.
- DECANDIA, G., HASTORUN, D., JAMPANI, M., KAKULAPATI, G., LAKSHMAN, A., PILCHIN, A., SIVASUBRAMANIAN, S., VOSSHALL, P., AND VOGELS, W. 2007. Dynamo: amazon's highly available key-value store. In *Proceedings of the 21st ACM Symposium on Operating Systems Principles (SOSP)*. 205–220.
- DETWELER, C., DONIEC, M., JIANG, M., SCHWAGER, M., CHEN, R., AND RUS, D. 2010. Adaptive Decentralized Control of Underwater Sensor Networks for Modeling Underwater Phenomena. In *Proceedings of the 8th International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, Zurich, Switzerland, 253–266.
- DONG, A., ZHANG, R., KOLARI, P., BAI, J., DIAZ, F., CHANG, Y., ZHENG, Z., AND ZHA, H. 2010. Time is of the Essence: Improving Recency Ranking Using Twitter Data. In *Proceedings of the 19th International Conference on World Wide Web (WWW)*. ACM, Raleigh, North Carolina, USA, 331–340.
- ELAHI, B. M., RÖMER, K., OSTERMAIER, B., FAHRMAIR, M., AND KELLERER, W. 2009. Sensor Ranking: A Primitive for Efficient Content-based Sensor Search. In *Proceedings of the 8th International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, San Francisco, California, USA, 217–228.
- FAN, W., GEERTS, F., MA, S., AND MÜLLER, H. 2010. Detecting Inconsistencies in Distributed Data. In *Proceedings of the 26th International Conference on Data Engineering (ICDE)*. IEEE, Long Beach, California, USA, 64–75.
- FAZZINGA, B., FLESCA, S., FURFARO, F., AND PARISI, F. 2014. Cleaning trajectory data of RFID-monitored objects through conditioning under integrity constraints. In *Proc. 17th International Conference on Extending Database Technology (EDBT)*. 379–390.

- FRANK, C., BOLLIGER, P., MATTERN, F., AND KELLERER, W. 2008. The Sensor Internet at Work: Locating Everyday Items Using Mobile Phones. *Pervasive and Mobile Computing* 4, 3, 421–447.
- GAMA, J. 2010. *Knowledge Discovery from Data Streams*. Chapman and Hall / CRC Data Mining and Knowledge Discovery Series. CRC Press.
- GANTI, R. K., PHAM, N., AHMADI, H., NANGIA, S., AND ABDELZAHER, T. F. 2010. GreenGPS: A Participatory Sensing Fuel-Efficient Maps Application. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services (MobiSys)*. ACM, San Francisco, California, USA, 151–164.
- GAO, J., GUIBAS, L. J., MILOSAVLJEVIC, N., AND ZHOU, D. 2009. Distributed Resource Management and Matching in Sensor Networks. In *Proceedings of the 8th International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, San Francisco, California, USA, 97–108.
- GARFINKEL, S. AND ROSENBERG, B. 2005. *RFID: Applications, Security, and Privacy*. Addison-Wesley, Boston, USA.
- GERBER, D., HELLMANN, S., BÜHMANN, L., SORU, T., USBECK, R., AND NGOMO, A.-C. N. 2013. Real-Time RDF Extraction from Unstructured Data Streams. In *Proceedings of the 12th International Semantic Web Conference (ISWC)*. 135–150.
- GILBERT, S. AND LYNCH, N. A. 2002. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News* 33, 2, 51–59.
- GUHA, S., PLARRE, K., LISSNER, D., MITRA, S., KRISHNA, B., DUTTA, P., AND KUMAR, S. 2010. AutoWitness: Locating and Tracking Stolen Property While Tolerating GPS and Radio Outages. In *Proceedings of the 8th International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, Zurich, Switzerland, 29–42.
- GUINARD, D. 2010. A Web of Things for Smarter Cities. In *Technical Talk*. 1–8.
- GUO, J., ZHANG, P., TAN, J., AND GUO, L. 2011. Mining frequent patterns across multiple data streams. In *Proceedings of the 20th ACM Conference on Information and Knowledge Management (CIKM)*. 2325–2328.
- GUPTA, M., INTILLE, S. S., AND LARSON, K. 2009. Adding GPS-Control to Traditional Thermostats: An Exploration of Potential Energy Savings and Design Challenges. In *Proceedings of the 7th International Conference on Pervasive Computing (Pervasive)*. Springer, Nara, Japan, 95–114.
- HASAN, S., O’RIAIN, S., AND CURRY, E. 2012. Approximate semantic matching of heterogeneous events. In *Proceedings of the Sixth ACM International Conference on Distributed Event-Based Systems (DEBS)*. 252–263.
- HASAN, S., O’RIAIN, S., AND CURRY, E. 2013. Towards unified and native enrichment in event processing systems. In *Proceedings of the 7th ACM International Conference on Distributed Event-Based Systems (DEBS)*. 171–182.
- HE, Y., BARMAN, S., AND NAUGHTON, J. F. 2014. On Load Shedding in Complex Event Processing. In *Proc. 17th International Conference on Database Theory (ICDT)*. 213–224.
- HEINZE, T., JI, Y., PAN, Y., GRÜNEBERGER, F. J., JERZAK, Z., AND FETZER, C. 2013. Elastic complex event processing under varying query load. In *Proceedings of the First International Workshop on Big Dynamic Distributed Data*. 25–30.
- HUANG, J., WANG, H., JIA, Y., AND FUXMAN, A. 2011. Link-based Hidden Attribute Discovery for Objects on Web. In *Proceedings of the 14th International Conference on Extending Database Technology (EDBT)*. ACM, Uppsala, Sweden, 473–484.
- INFSO. 2008. INFSO D.4 Networked Enterprise & RFID INFSO G.2 Micro & Nanosystems, in: Co-operation with the Working Group RFID of the ETP EPOSS, Internet of Things in 2020, Roadmap for the Future, Version 1.1, 27 May 2008.
- ITU. 2005. International Telecommunication Union (ITU) Internet Reports, The Internet of Things, November 2005.
- JAMES, A. E., COOPER, J., JEFFERY, K. G., AND SAAKE, G. 2009. Research Directions in Database Architectures for the Internet of Things: A Communication of the First International Workshop on Database Architectures for the Internet of Things (DAIT 2009). In *Proceedings of the 26th British National Conference on Databases (BNCOD)*. Springer, Birmingham, UK, 225–233.

- JEFFERY, S. R., GAROFALAKIS, M. N., AND FRANKLIN, M. J. 2006. Adaptive Cleaning for RFID Data Streams. In *Proceedings of the 32nd International Conference on Very Large Data Bases (VLDB)*. 163–174.
- JIN, X., ZHANG, N., AND DAS, G. 2011. Attribute Domain Discovery for Hidden Web Databases. In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*. ACM, Athens, Greece, 553–564.
- KADAMBI, S., CHEN, J., COOPER, B. F., LOMAX, D., RAMAKRISHNAN, R., SILBERSTEIN, A., TAM, E., AND GARCIA-MOLINA, H. 2011. Where in the World is My Data? *Proceedings of the VLDB Endowment* 4, 11, 1040–1050.
- KHARE, R., AN, Y., AND SONG, I.-Y. 2010. Understanding Deep Web Search Interfaces: A Survey. *SIGMOD Record* 39, 1, 33–40.
- KLASNJA, P. V., CONSOLVO, S., CHOUDHURY, T., BECKWITH, R., AND HIGHTOWER, J. 2009. Exploring Privacy Concerns about Personal Sensing. In *Proceedings of the 7th International Conference on Pervasive Computing (Pervasive)*. Springer, Nara, Japan, 176–183.
- KNORR, E. M. AND NG, R. T. 1998. Algorithms for Mining Distance-Based Outliers in Large Datasets. In *Proceedings of 24rd International Conference on Very Large Data Bases (VLDB)*. 392–403.
- KOSHIZUKA, N. AND SAKAMURA, K. 2010. Ubiquitous ID: Standards for Ubiquitous Computing and the Internet of Things. *IEEE Pervasive Computing* 9, 4, 98–101.
- KOUNAVIS, M. E., KANG, X., GREWAL, K., ESZENYI, M., GUERON, S., AND DURHAM, D. 2010. Encrypting the Internet. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*. ACM, New Delhi, India, 135–146.
- LAGESSE, B., KUMAR, M., PALUSKA, J. M., AND WRIGHT, M. 2009. DTT: A Distributed Trust Toolkit for Pervasive Systems. In *Proceedings of the 7th Annual IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, Galveston, TX, USA, 1–8.
- LAI, T.-T., CHEN, Y.-H., HUANG, P., AND CHU, H.-H. 2010. PipeProbe: A Mobile Sensor Droplet for Mapping Hidden Pipeline. In *Proceedings of the 8th International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, Zurich, Switzerland, 113–126.
- LAKSHMAN, A. AND MALIK, P. 2010. Cassandra: a decentralized structured storage system. *Operating Systems Review* 44, 2, 35–40.
- LESTER, J., HARTUNG, C., PINA, L., LIBBY, R., BORRIELLO, G., AND DUNCAN, G. 2009. Validated Caloric Expenditure Estimation Using a Single Body-Worn Sensor. In *Proceedings of the 11th International Conference on Ubiquitous Computing (Ubicomp)*. ACM, Orlando, Florida, USA, 225–234.
- LI, S., XU, L. D., AND ZHAO, S. 2014. The Internet of Things: A Survey. *Information Systems Frontiers*, (to appear).
- LIAN, X. AND CHEN, L. 2011. Efficient Query Answering in Probabilistic RDF Graphs. In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*. ACM, Athens, Greece, 157–168.
- LIAO, G., LI, J., CHEN, L., AND WAN, C. 2011. KLEAP: an efficient cleaning method to remove cross-reads in RFID streams. In *Proceedings of the 20th ACM Conference on Information and Knowledge Management (CIKM)*. 2209–2212.
- LIN, Y.-H., STUDER, A., HSIAO, H.-C., MCCUNE, J. M., WANG, K.-H., KROHN, M. N., LIN, P.-L., PERRIG, A., SUN, H.-M., AND YANG, B.-Y. 2009. SPATE: Small-group PKI-less Authenticated Trust Establishment. In *Proceedings of the 7th International Conference on Mobile Systems, Applications, and Services (MobiSys)*. ACM, Kraków, Poland, 1–14.
- LIU, A. H., BUNN, J. J., AND CHANDY, K. M. 2011. Sensor Networks for the Detection and Tracking of Radiation and Other Threats in Cities. In *Proceedings of the 10th International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, Chicago, IL, USA, 1–12.
- LIU, M., RUNDENSTEINER, E. A., DOUGHERTY, D. J., GUPTA, C., WANG, S., ARI, I., AND MEHTA, A. 2011. High-performance nested CEP query processing over event streams. In *Proceedings of the 27th International Conference on Data Engineering (ICDE)*. 123–134.

- LU, J., SOOKOOR, T. I., SRINIVASAN, V., GAO, G., HOLBEN, B., STANKOVIC, J. A., FIELD, E., AND WHITEHOUSE, K. 2010. The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes. In *Proceedings of the 8th International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, Zurich, Switzerland, 211–224.
- MATHEW, S. S., ATIF, Y., SHENG, Q. Z., AND MAAMAR, Z. 2014. Building Sustainable Parking Lots with the Web of Things. *Personal and Ubiquitous Computing* 18, 4, 895–907.
- MATHUR, S., JIN, T., KASTURIRANGAN, N., CHANDRASEKARAN, J., XUE, W., GRUTESER, M., AND TRAPPE, W. 2010. ParkNet: Drive-by Sensing of Road-Side Parking Statistics. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services (MobiSys)*. ACM, San Francisco, California, USA, 123–136.
- MATTERN, F. AND FLOERKEMEIER, C. 2010. From the Internet of Computers to the Internet of Things. In *From Active Data Management to Event-Based Systems and More*. Springer, 242–259.
- MOTTOLA, L. 2010. Programming Storage-Centric Sensor Networks with Squirrel. In *Proceedings of the 9th International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, Stockholm, Sweden, 1–12.
- NATH, S. 2009. Energy Efficient Sensor Data Logging with Amnesic Flash Storage. In *Proceedings of the 8th International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, San Francisco, California, USA, 157–168.
- NIE, Y., COCCI, R., CAO, Z., DIAO, Y., AND SHENOY, P. J. 2012. SPIRE: Efficient Data Inference and Compression over RFID Streams. *IEEE Trans. Knowl. Data Eng.* 24, 1, 141–155.
- OOI, B. C., TAN, K.-L., AND TUNG, A. K. H. 2009. Sense The Physical, Walkthrough The Virtual, Manage The Co (existing) Spaces: A Database Perspective. *SIGMOD Record* 38, 3, 5–10.
- OSTERMAIER, B., RÖMER, K., MATTERN, F., FAHRMAIR, M., AND KELLERER, W. 2010. A Real-Time Search Engine for the Web of Things. In *Proceedings of the 2010 Internet of Things (IOT)*. IEEE, Tokyo, Japan.
- PERERA, C., ZASLAVSKY, A. B., CHRISTEN, P., AND GEORGAKOPOULOS, D. 2013. Context Aware Computing for The Internet of Things: A Survey. *CoRR abs/1305.0982*.
- PHUOC, D. L., DAO-TRAN, M., PARREIRA, J. X., AND HAUSWIRTH, M. 2011. A Native and Adaptive Approach for Unified Processing of Linked Streams and Linked Data. In *Proceedings of the 10th International Semantic Web Conference (ISWC)*. 370–388.
- PRABHAKAR, S. AND CHENG, R. 2009. Data Uncertainty Management in Sensor Networks. In *Encyclopedia of Database Systems*. Springer US, USA, 647–651.
- PRIYANTHA, N. B., KANSAL, A., GORACZKO, M., AND ZHAO, F. 2008. Tiny Web Services: Design and Implementation of Interoperable and Evolvable Sensor Networks. In *Proceedings of the 6th International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, Raleigh, NC, USA, 253–266.
- RAMANATHAN, N., SCHOELLHAMMER, T., KOHLER, E., WHITEHOUSE, K., HARMON, T., AND ESTRIN, D. 2009. Suelo: Human-Assisted Sensing for Exploratory Soil Monitoring Studies. In *Proceedings of the 7th International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, Berkeley, California, USA, 197–210.
- RAY, M., RUNDENSTEINER, E. A., LIU, M., GUPTA, C., WANG, S., AND ARI, I. 2013. High-performance complex event processing using continuous sliding views. In *Proceedings of Joint 2013 EDBT/ICDT Conferences (EDBT)*. 525–536.
- SAKAKI, T., OKAZAKI, M., AND MATSUO, Y. 2010. Earthquake Shakes Twitter Users: Real-time Event Detection by Social Sensors. In *Proceedings of the 19th International Conference on World Wide Web (WWW)*. ACM, Raleigh, North Carolina, USA, 851–860.
- SATPAL, S., BHADRA, S., SELLMANICKAM, S., RASTOGI, R., AND SEN, P. 2011. Web Information Extraction Using Markov Logic Networks. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*. ACM, San Diego, CA, USA, 1406–1414.
- SCOTT, J., BRUSH, A. J. B., KRUMM, J., MEYERS, B., HAZAS, M., HODGES, S., AND VILLAR, N. 2011. PreHeat: Controlling Home Heating Using Occupancy Prediction. In *Proceedings of*

- the 13th International Conference on Ubiquitous Computing (*Ubicomp*). ACM, Beijing, China, 281–290.
- SEQUEDA, J. AND CORCHO, Ó. 2009. Linked Stream Data: A Position Paper. In *Proceedings of the 2nd International Workshop on Semantic Sensor Networks (SSN09)*. 148–157.
- SHENG, Q. Z., LI, X., AND ZEADALLY, S. 2008. Enabling Next-Generation RFID Applications: Solutions and Challenges. *IEEE Computer* 41, 9 (September), 21–28.
- STEPHAN, P., MEIXNER, G., KOESSLING, H., FLOERCHINGER, F., AND OLLINGER, L. 2010. Product-Mediated Communication through Digital Object Memories in Heterogeneous Value Chains. In *Proceedings of the 8th Annual IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, Mannheim, Germany, 199–207.
- STONEBRAKER, M., ABADI, D. J., BATKIN, A., CHEN, X., CHERNIACK, M., FERREIRA, M., LAU, E., LIN, A., MADDEN, S., O’NEIL, E. J., O’NEIL, P. E., RASIN, A., TRAN, N., AND ZDONIK, S. B. 2005. C-Store: A Column-oriented DBMS. In *Proceedings of the 31st International Conference on Very Large Data Bases (VLDB)*. ACM, Trondheim, Norway, 553–564.
- STONEBRAKER, M., MADDEN, S., ABADI, D. J., HARIZOPOULOS, S., HACHEM, N., AND HELLAND, P. 2007. The End of an Architectural Era (It’s Time for a Complete Rewrite). In *Proceedings of the 33rd International Conference on Very Large Data Bases (VLDB)*. ACM, University of Vienna, Austria, 1150–1160.
- SUBRAMANIAM, S. AND GUNOPULOS, D. 2007. A Survey of Stream Processing Problems and Techniques in Sensor Networks. In *Data Streams - Models and Algorithms*. 333–352.
- SUCHANEK, F. M., VARDE, A. S., NAYAK, R., AND SENELLART, P. 2011. The Hidden Web, XML and the Semantic Web: Scientific Data Management Perspectives. In *Proceedings of the 14th International Conference on Extending Database Technology (EDBT)*. ACM, Uppsala, Sweden, 534–537.
- TAN, C. C., SHENG, B., WANG, H., AND LI, Q. 2010. Microsearch: A Search Engine for Embedded Devices Used in Pervasive Computing. *ACM Trans. Embedded Comput. Syst.* 9, 4, 29.
- TEYMOURIAN, K. AND PASCHKE, A. 2009. Towards semantic event processing. In *Proceedings of the Third ACM International Conference on Distributed Event-Based Systems (DEBS)*.
- TEYMOURIAN, K., ROHDE, M., AND PASCHKE, A. 2012a. Fusion of background knowledge and streams of events. In *DEBS*. 302–313.
- TEYMOURIAN, K., ROHDE, M., AND PASCHKE, A. 2012b. Knowledge-based processing of complex stock market events. In *Proceedings of the 15th International Conference on Extending Database Technology (EDBT)*. 594–597.
- TRAN, T. T. L., SUTTON, C. A., COCCI, R., NIE, Y., DIAO, Y., AND SHENOY, P. J. 2009. Probabilistic Inference over RFID Streams in Mobile Environments. In *Proceedings of the 25th International Conference on Data Engineering (ICDE)*. 1096–1107.
- TSATSANIFOS, G., SACHARIDIS, D., AND SELLIS, T. K. 2011. On Enhancing Scalability for Distributed RDF/S Stores. In *Proceedings of the 14th International Conference on Extending Database Technology (EDBT)*. ACM, Uppsala, Sweden, 141–152.
- TSIFTES, N. AND DUNKELS, A. 2011. A Database in Every Sensor. In *Proceedings of the 9th International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, Seattle, WA, USA, 316–332.
- WANG, F. AND AGRAWAL, G. 2011. Effective and Efficient Sampling Methods for Deep Web Aggregation Queries. In *Proceedings of the 14th International Conference on Extending Database Technology (EDBT)*. ACM, Uppsala, Sweden, 425–436.
- WANG, F. AND LIU, J. 2011. Networked Wireless Sensor Data Collection: Issues, Challenges, and Approaches. *IEEE Communications Surveys and Tutorials* 13, 4, 673–687.
- WANG, H., TAN, C. C., AND LI, Q. 2010. Snoogle: A Search Engine for Pervasive Environments. *IEEE Trans. Parallel Distrib. Syst.* 21, 8, 1188–1202.
- WEIKUM, G. 2011. Database Researchers: Plumbers or Thinkers? In *Proceedings of the 14th International Conference on Extending Database Technology (EDBT)*. ACM, Uppsala, Sweden, 9–10.
- ACM Journal Name, Vol. V, No. N, July 2014.

- WELBOURNE, E., BATTLE, L., COLE, G., GOULD, K., RECTOR, K., RAYMER, S., BALAZINSKA, M., AND BORRIELLO, G. 2009. Building the Internet of Things Using RFID: The RFID Ecosystem Experience. *IEEE Internet Computing* 13, 3, 48–55.
- WU, E., DIAO, Y., AND RIZVI, S. 2006. High-performance complex event processing over streams. In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD'06)*. 407–418.
- WU, Y., SHENG, Q. Z., SHEN, H., AND ZEDADALLY, S. 2013. Modeling Object Flows from Distributed and Federated RFID Data Streams for Efficient Tracking and Tracing. *IEEE Transactions on Parallel and Distributed Systems* 24, 10, 2036–2045.
- XIE, Q., ZHU, J., SHARAF, M. A., ZHOU, X., AND PANG, C. 2012. Efficient buffer management for piecewise linear representation of multiple data streams. In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management (CIKM)*. 2114–2118.
- YANG, Y., WANG, L., NOH, D. K., LE, H. K., AND ABDELZAHER, T. F. 2009. SolarStore: Enhancing Data Reliability in Solar-Powered Storage-Centric Sensor Networks. In *Proceedings of the 7th International Conference on Mobile Systems, Applications, and Services (MobiSys)*. ACM, Kraków, Poland, 333–346.
- YAP, K.-K., SRINIVASAN, V., AND MOTANI, M. 2008. MAX: Wide Area Human-Centric Search of the Physical World. *ACM Trans. on Sensor Networks* 4, 4, 34.
- YUAN, J., ZHENG, Y., ZHANG, L., XIE, X., AND SUN, G. 2011. Where to Find My Next Passenger? In *Proceedings of the 13th International Conference on Ubiquitous Computing (Ubicomp)*. ACM, Beijing, China, 109–118.
- ZENG, D., GUO, S., AND CHENG, Z. 2011. The Web of Things: A Survey (Invited Paper). *Journal of Communications* 6, 6, 424–438.
- ZHANG, M., ZHANG, N., AND DAS, G. 2011. Mining a search engine's corpus: Efficient yet unbiased sampling and aggregate estimation. In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD)*. ACM, Indianapolis, Indiana, USA, 793–804.
- ZHANG, Y., PHAM, M.-D., CORCHO, Ó., AND CALBIMONTE, J.-P. 2012. SRBench: A Streaming RDF/SPARQL Benchmark. In *Proceedings of the 11th International Semantic Web Conference (ISWC)*. 641–657.
- ZHOU, Q., SIMMHAN, Y., AND PRASANNA, V. K. 2011. Towards an inexact semantic complex event processing framework. In *Proceedings of the Fifth ACM International Conference on Distributed Event-Based Systems (DEBS)*. 401–402.
- ZHU, F. AND ZHU, W. 2009. RationalExposure: a Game Theoretic Approach to Optimize Identity Exposure in Pervasive Computing Environments. In *Proceedings of the 7th Annual IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, Galveston, TX, USA, 1–8.
- ZIMMERMAN, T. AND ROBSON, C. 2011. Monitoring Residential Noise for Prospective Home Owners and Renters. In *Proceedings of the 9th International Conference on Pervasive Computing (Pervasive)*. Springer, San Francisco, CA, USA, 34–49.

Received July 2014