

Electronic Communications of the EASST  
Volume X (2011)



Proceedings of the  
Workshops der wissenschaftlichen Konferenz  
Kommunikation in verteilten Systemen 2011  
(WowKiVS 2011)

An Approach for Urban Sensing with Quality-Aware Situation Detection  
and Efficient Communication  
*Position Paper*

Gerald G. Koch, Andreas Benzing, and Christoph P. Mayer

10 pages

# An Approach for Urban Sensing with Quality-Aware Situation Detection and Efficient Communication Position Paper

Gerald G. Koch<sup>1</sup>, Andreas Benzing<sup>2</sup>, and Christoph P. Mayer<sup>3</sup>

<sup>1</sup>[gerald.koch@ipvs.uni-stuttgart.de](mailto:gerald.koch@ipvs.uni-stuttgart.de), <sup>2</sup>[andreas.benzing@ipvs.uni-stuttgart.de](mailto:andreas.benzing@ipvs.uni-stuttgart.de)

Institute for Parallel and Distributed Systems  
University of Stuttgart, Germany

<sup>3</sup>[mayer@kit.edu](mailto:mayer@kit.edu)

Institute of Telematics  
Karlsruhe Institute of Technology (KIT), Germany

**Abstract:** Urban Sensing employs physical-world mining to create a digital model of the physical world using a large number of sensors. Handling the large amount of data generated by sensors is costly and requires energy-saving measures for sensing and sensor data transmission. Such schemes often affect data quality and message delay. However, the detection of real-world situations using *Complex Event Processing* on sensor data has to be dependable and timely and requires precise data. In this position paper, we propose an approach to integrate the contradicting optimization goals of energy-efficient wireless sensor networks and dependable situation detection. It separates the system into the following tiers: First, to support energy-efficiency and allow sparse, unconnected sensor networks, we exploit the mobility of people through *Delay Tolerant Networking* for collecting sensor data. This frees sensor nodes from energy-expensive routing. Second, we employ *Diagnostic Simulation* which provides data that is complete, precise and in time and therefore supports quality-aware situation detection.

**Keywords:** Urban Sensing, Internet of Things, Complex Event Processing, Delay Tolerant Network, Distributed Diagnostic Simulation.

## 1 Introduction

With the pervasiveness of sensing devices in the upcoming Internet of Things, a large amount of sensor data from real-world environments becomes available in the digital domain. Innovative applications use this data to detect real-world phenomena and react on them. For instance, the growing awareness of human impact on ecology leads to large-scale distributed applications in fields like dynamic traffic guidance, environmental observation, and decentralized management of power consumers and power stations [CHPP09]. These applications require a large amount of sensor data to be collected, evaluated and transmitted. They are therefore often supported by *Complex Event Processing* (CEP) [Luc01] systems which create high-level abstractions out of low-level sensor data. For instance, applications that reduce atmospheric pollution in a given

area receive information about detected situations—e. g. the threat of exceeding limit values or a looming traffic gridlock—from DCEP, which observed the situation by analyzing and correlating a large amount of data on exhaust fumes and fine dust, traffic situations, wind-fields, and atmospheric behavior over multiple large areas.

The success of such systems heavily relies on the combination of robust and energy-efficient data collection mechanisms through sensors, together with powerful and dependable situation detection. The authors of [CHPP09] propose an architecture divides information processing in the Internet of Things into three *domains*: i) an *edge domain* in the physical world where sensors are located, ii) a *processing domain* in an information system where sensor events are processed, and iii) a *business domain* that is driven by processed events. While we agree that these concerns need to be separated, we observe that particular attention must be paid to the amount, quality and temporal and geographical relation of data that is processed at the domains' interfaces.

In the edge domain, we expect sensor networks that are heavily energy-constrained. A promising way to save energy is to reduce the communication tasks of sensors, because communicating data through the sensor network consumes orders of magnitude more energy than other tasks [DGMS07]. This holds especially when sparsely deployed sensor networks provide sufficient coverage but require high power levels for communication which either increases energy consumption, or prevents communication altogether. It is therefore not always desirable for a sensor network to form a connected communication network and autonomously route data to sinks. Equipping sensor nodes for direct communication with the infrastructure (e. g. with chips for cellular networks) is not desirable because it is costly—due to high cost for hardware, energy, and data plan—and easily exceeds the monetary value of data acquired through sensing. Therefore, sensor networks' interfaces provide a large amount of—possibly incomplete and delayed—data that needs to be collected close to the place where it occurred.

Innovative applications in the business domain are faced with strict user requirements. For instance, the actuation of a traffic guidance system (redirecting or constraining traffic, adjusting electricity pricing and pollution rights) is only acceptable to users if the system's decisions are timely and dependable. Therefore, applications require a specific, complete and fresh view on the observed physical world. To provide this specific view, auxiliary systems in the processing domain reduce the amount of sensor data to observed situations. To assure the completeness and freshness of the view, event processing systems need to be *quality-aware*, i. e. to make sure that the sensor data they process is complete, precise, and up-to-date. These characteristics, however, are not provided at the interface between the edge domain and the event processing domain.

Clearly, an approach is required that integrates the energy efficiency in the edge domain with quality-awareness in the processing and business domain. In this position paper, we propose an approach which, instead of searching for a trade-off, divides the system into tiers where the sensor network tier optimizes the system's energy efficiency, while the situation detection tier considers data access, data quality and delay. To bridge the gap between the domains we employ additional techniques, as shown in Figure 1: Between sensor network routing and situation detection, distributed instances of *Diagnostic Simulation* (DS) provide data that is complete and fresh regardless of sensor data incompleteness or message delay. Location-independent access to this data is granted by a network of distributed data streams [BKVR10]. In the edge domain, we exploit the mobility of human-carried devices for energy-efficient sensor data collection using *Delay Tolerant Networks* (DTN) [Zha07].

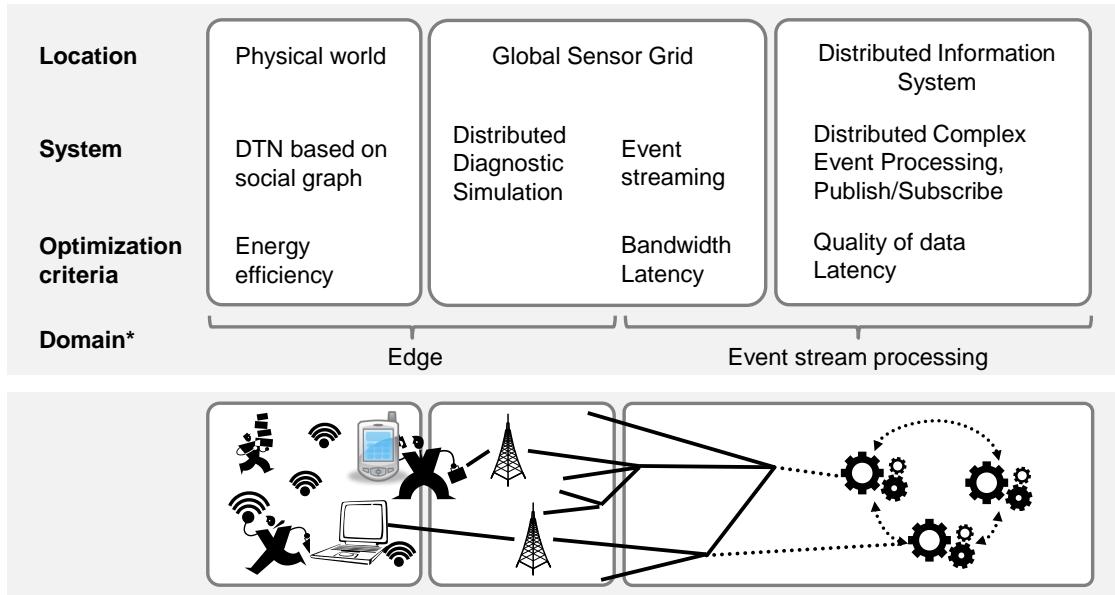


Figure 1: Functionality and optimization criteria (\* in accordance to [CHPP09]).

This position paper is structured as follows: In Section 2 we discuss related approaches for urban sensing with focus on general approaches and architectures. In Section 3 we propose our approach for urban sensing, present the tiers of our envisioned solution and their interaction in detail. Section 4 concludes this position paper.

## 2 Related Approaches

The first projects dealing with environmental sensing evolved in the late 1990s with *Smartdust* [KKP99] and NASA's *Sensor Webs* [BA07] projects. Their goal was to develop autonomous sensing and communication systems based on miniaturized sensor nodes. If sensor nodes form a connected graph through hop-by-hop links, protocols in the scope of *Wireless Sensor Networks* (WSNs) can be employed [AY05] to enable end-to-end routing. On the link layer, hop-by-hop communication based on sleep cycles lowers the wake time and, consequently, energy consumption. On the routing layer, aggregation mechanisms [RV06] save energy by reducing the data volume. For monitoring large geographic areas, however, the WSN approach requires too large a number of sensor nodes to form a connected graph. If, though, the density of sensor nodes is low, unreasonably high energy power-levels are required for the wireless hop-by-hop channel. We therefore argue that a general architecture for large-scale urban sensing based on large WSN is not feasible. As an alternative, geographically local WSNs could be deployed, but those require interconnection with remote systems through other communication means.

An approach to such interconnected sensor networks is the *Global Sensor Network* (GSN) project [AHS06]. It enables a static configuration of sensors at gateway PCs, as well as interconnection of such gateway PCs through the Internet. Provided sensor streams are defined in

configuration files and can be queried by remote systems through the static PC gateways. While this approach provides a good basis for the interconnection of sensor networks, it requires sensors connected to gateways in a static way and predefined configurations for sensor streams.

*IrisNet* [GKK<sup>03</sup>] takes an approach based on distributed databases for sensing with focus on stream querying and processing. While IrisNet provides a general architecture for global sensing, the authors use webcams as exemplary sensing devices that have static Internet access through a PC. IrisNet manages a distributed database defined by XML specifications which are written by service providers that describe their services. The focus of IrisNet are not communication aspects, it rather focuses on finding, collecting, and distributed processing of sensor feeds.

*Hourglass* [SPL<sup>04</sup>] provides a communication infrastructure based upon a *Peer-to-Peer* (P2P) network. The P2P network is used as distributed infrastructure for service registration, discovery, and routing of data streams from sensors to registered clients. As P2P machines are defined to be well-connected and stable, the approach is similar to a publish/subscribe system for sensor data registration and delivery. Similar, *HiFi* [CEF<sup>05</sup>] provide a general architectures to uncouple data providers from data consumers. The focus of Hifi is the functional abstraction into 5 layers. At the lowest layer sensors provide input data that is cleaned and given to gateways. Those Internet-connected gateways smooth the data and forward it to proxies that remove duplicates and perform correlation. Regional systems validate the data and correlate it for business data. Finally, a headquarter root node performs overall analysis and actuation decisions.

Work that come closest to the approach presented in this position paper are *Data Mules* [SRJB03] and *Urban Sensing* [CEL<sup>06</sup>]. They shift the complexity and energy demand of sensor network routing from sensor nodes to human-carried mobile devices. Such devices, like mobile phones, have rich resources and a simple charging-cycle. The architecture, however, can introduce long delays and low delivery probabilities which collides with the goal of quality-awareness for sensor data. In this paper we show that such an approach is still beneficial in terms of energy consumption, and we present a way to handle quality of sensed data.

### 3 Integrating Energy Efficiency and Quality Awareness

In the following we detail on our proposed approach to integrate widely distributed heterogeneous sensor networks with quality-aware Distributed CEP (DCEP). We first show the benefits of DTNs for carrying sensor data in Section 3.1—which corresponds to the left tier in Figure 1 labeled ‘Physical world’. The DTN routing provides energy-efficiency but inherently induces delay and low delivery probabilities into the data supply. In Section 3.2—corresponding to the mid tier in Figure 1 labeled ‘Global Sensor Grid’—we show how Diagnostic Simulation and efficient querying help to bridge the data quality gap. This allows for high-quality data delivery, as required by CEP (right tier in Figure 1) which we describe in Section 3.3.

#### 3.1 DTN and Sensor Data

The approach for data gathering used in this work is based upon *Delay Tolerant Networks* (DTNs), through the use of opportunistic routing [Zha07]. We exploit the mobility of human-carried devices to collect data from sparsely distributed sensor nodes, similar to [SRJB03, CEL<sup>06</sup>]: Mo-

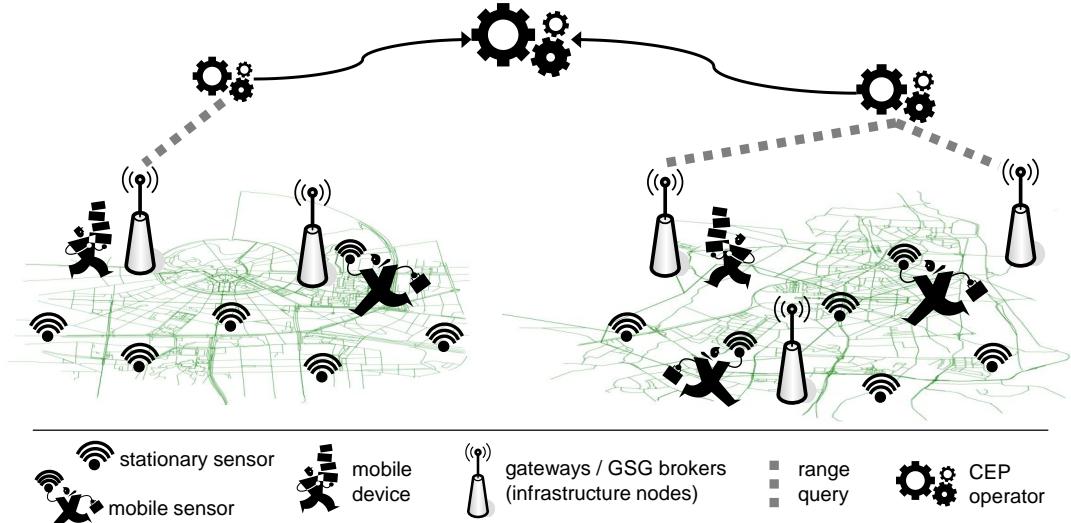


Figure 2: System components and their interaction in the proposed approach

mobile devices query data from stationary sensors when opportunistically within communication range and cooperatively route sensed data to sinks. Figure 2 depicts such a scenario. For routing collected sensor data cooperatively to sinks, mobile devices perform DTN-based store-carry-forward routing based on opportunistic encounters between mobile devices. As mobile devices are exposed to human mobility, they therefore allow for collecting data from stationary sensors in large geographic areas like cities. In addition, sensors can be attached to the mobile devices, which increases geographic coverage of sensor values, as they too are exposed to mobility.

Our intent for using DTN routing is to relieve sensors from sensor network routing which is both complex, and energy-consuming. Shifting the routing task away from sensors furthermore reduces the complexity and cost of sensors. This way, we enable sensors to be simple and energy-efficient. At the same time, DTNs can carry high capacity of data in a cheap fashion [GT02]. The human-carried devices—like mobile phones, or PDAs—which perform the DTN routing of sensor data are, in contrast to sensor nodes, computationally strong and have a simple charging cycle through its user. As an incentive for owners of mobile devices to participate in sensor data routing we imagine, for example, a discount on the cell-phone plan. Clearly, security mechanisms are required for end-to-end encryption between sensor nodes and the system operator to prevent rogue DTN devices to inject corrupt sensor data. As drawback, however, DTNs provide no quality in communication. They feature merely probabilistic end-to-end delay and message delivery probabilities. This poses further challenges as human mobility is non-deterministic and sensor data may not be available for the system. We focus on this challenge in the next Section.

### 3.2 Distributed Diagnostic Simulation

The interface between the edge and the processing domain consists of *gateways* where the DTN routing system presented above delivers sensor data from nearby areas which may be incomplete and considerably delayed. However, dependable situation detection using DCEP requires fresh

and complete sensor data from local and remote areas. We solve this contradiction using two additional auxiliary system at the gateway nodes of the DTN system. The first is distributed Diagnostic Simulation which, for phenomena that follow a known model, generates a complete and delay-free local view of the real world based on incomplete and delayed data received from the DTN system. The second system is the Global Sensor Grid which provides efficient event streaming for location-independent data access by multiple distributed applications. We describe the Diagnostic Simulation system first.

There are different approaches for enriching data gathered from sensors. Intermediate, missing values could be completed using interpolation [Mei02]. However, the information content actually stays constant, i. e., gaps in incomplete data sets persist as no additional knowledge is integrated. Such additional information can only be given by mathematical models that describe the correlations between measurements for Model-driven Data Acquisition [DGM<sup>+</sup>04]. This way, incomplete sensor data measurements provided by the DTN can be employed to generate information that is complete in time and space.

The approach of *Diagnostic Simulations* takes this even further: it employs a model of the complete underlying physical process instead of just modeling correlations between measurements. Information about the measurements is no longer computed using tools like regression or interpolation. Unlike a predictive simulation, however, Diagnostic Simulations run in real-time and maintain an internal state as a current representation of the physical world. Several such models, for example for wind fields, have been proposed and evaluated over the last decades. In this case the simulation also considers turbulence which cannot be covered by statistical analysis at all. Although the first wind field models have been proposed over 40 years ago, they are still subject to current research [FJNS09]. Many kinds of simulations have been developed for different physical phenomena. Some of them couple the simulation of multiple physical processes to model complex procedures. Particle simulations for the distribution of exhaust fumes and other pollutants for example rely on the wind fields as well as precipitation and chemical reactions. However, there still remain phenomena that up to now cannot be adequately simulated.

As soon as the DTN routing provides data at gateways, the Diagnostic Simulation uses the data for two purposes: first, for verification of the simulation, and second, for adaptation of the simulation state. Specific parts of the simulation state are archived and compared to measurements when those become available, so that the system can automatically verify whether the simulation model represents the real world correctly. If a measurement deviates significantly from the simulation state, the reading is directly forwarded to the CEP layer. The CEP system clearly needs to be able to cope with out-of-order events in this case. This behavior ensures that unexpected situations that may indicate an arising emergency are not filtered out as outliers.

Approaches like *Data Management in the Worldwide Sensor Web* [BDF<sup>+</sup>07] adapt the simulation state. A major advantage of Diagnostic Simulations over other approaches is that even outdated measurements can be used to update the simulation. Despite the higher age of measurements due to the DTN approach for gathering sensor data in an energy efficient manner, their values still present additional information. Such outdated measurements are only considered with reduced weight since their impact on the current state decays with their age. For now, the weight is only determined from the age using a simple mapping function. In the future, more advanced determination of weights based on results from information theory might be included to fully exploit the information contained in the measurements.

Additionally, the simulation includes knowledge about the physical phenomenon observed in the form of mathematical descriptions—for example, the differential equations governing the behavior of air currents. The simulation therefore converges towards the physical world even in areas where no measurements can be acquired as additional measurements are continuously used to update the simulation state over time.

As a result, Diagnostic Simulation at gateways provides a complete set of fresh data to the event processing domain. However, this data still has a strong geographic relation to the location where it occurred, so that a high amount of data has to be streamed over the network to satisfy remote applications. Approaches for managing high bandwidth data streams of sensor data by a *Global Sensor Grid* (GSG) have already been proposed in the literature [BKVR10]. However, to provide the quality properties required by DCEP, it is not sufficient to reduce the network load. Therefore, we first describe the details of quality aware situation detection and its prerequisites. Afterward, Section 3.3 describes how the GSG can be extended and bi-directional communication between the GSG and DCEP can be used to provide the required quality.

### 3.3 Quality-aware Situation Detection

*Complex Event Processing* (CEP) is a popular method to analyze event streams and detect situations on event collections, which is beneficial to use if the amount of data and the complexity of the observed patterns would overcharge the application itself. Furthermore, *Distributed Complex Event Processing* (DCEP) systems with expressive event algebras [KKR10] provide flexible situation detection efficiently in terms of overall computational effort and message overhead in a setting where event sources (gateways) and subscribers to situation detection (applications) are widely distributed (see the example in the Introduction). Distributed operators of the DCEP system can also consecutively detect increasingly complex situations. They communicate using a self-organizing publish/subscribe system [TKK<sup>+</sup>10]. Publish/subscribe provides efficient communication in our setting where heterogeneous information sources (sensor networks, databases, other applications) should be decoupled from data sinks and do not maintain explicit end-to-end connections.

DCEP is *quality-aware* only if it takes into account that the real world's representation in sensor readings is imperfect. Sensor data might be imprecise, delayed, or missing. Without restrictions or measures on these sources uncertainty it is impossible to provide dependable situation detection. Our proposed architecture supports quality-aware DCEP in the following manner.

Before all, data quality is not impaired by data aggregation because with DTN routing, no information is lost on its path through the sensor network. DTN routing achieves energy efficiency not by aggregating data to reduce the message load, but by not maintaining a static routing structure of its own.

The uncontrollable amount of latency and decreased delivery probability that is introduced by DTN routing is in turn handled by employing Diagnostic Simulation at gateways as described above. Diagnostic Simulation provides a stream of data which supports quality-aware DCEP in three aspects. First, the stream is temporally and spatially *complete*, because readings are published at the simulation rate for each data point in the simulation state. Second, the stream is *fresh* since the message delay induced by DTN inside the sensor network becomes invisible at the sensor network gateways that run Diagnostic Simulation. Similarly, between the DCEP instances,

the latency-aware publish/subscribe system limits message latency at user request [TKK<sup>+</sup>10]. Third, sensor data that stems from simulation follows a model which can be equipped with meta-data about its preciseness. Quality-aware situation detection can take this information into account and measure the dependability of each detected situation individually [KKR08]. The only remaining problem is that the data streams are only locally available at gateways.

This is fixed by the GSG as it accepts range queries on sensor data at any GSG broker, for arbitrary areas. The GSG autonomously optimizes routing from gateways to accessed brokers and reuses simulation results for all users simultaneously by exploiting intersections of queries. The benefits from this flexibility are twofold. First, DCEP is free to migrate situation detection functionality to appropriate nodes in the network, considering its own and the application's latency requirements rather than the location of sensor network gateways. Second, DCEP can influence the latency and quality of data provided by the GSG in a certain way. Querying large areas at one access broker allows the GSG to optimize the routing of data streams, which supports data aggregation. However, due to the hierarchical structure of the system, large queries result in longer paths inside of the GSG and therefore increase message latency. Therefore, alternatively, DCEP can split a query into a number of small areas. The corresponding data streams will be provided close to the gateways in the network topology and therefore at a smaller latency. In this case, however, DCEP and the publish/subscribe system need to cope with a higher number of incoming data connections, which leads to considerable control overhead.

In any case, the GSG allows DCEP operators to process complete and fresh sensor data from local and remote areas independently on the location of gateways and thus to provide dependable situations in-time at the interface to the business domain.

## 4 Conclusion

In this work we propose an approach to detect high-level situations from a large amount of sensor data generated in the Internet of Things. We address the energy efficiency requirements in sensor networks by exploiting the mobility of humans that collect data from sensors and forward collected data through collaborative DTN routing to data sinks. This enables low-cost data collection—with minimal spectrum load in comparison to sensor network routing—and transfer to dedicated sinks at the cost of low delivery probability and high delay. Sensor data is handled by hierarchically organized Distributed Complex Event Processing (DCEP) which detects high-level situations through event pattern detection. As situation detection needs to be dependable and quality-aware DCEP is impossible with missing or delayed data, we propose the use of Diagnostic Simulation to bridge the gap between DTN and DCEP. Through simulation of the real-world behavior, Diagnostic Simulation provides a complete, precise, and fresh stream of event data which alleviates the disadvantages introduced by the DTN transport.

## Acknowledgment

This work was funded by the Landesstiftung Baden-Württemberg as part of the “Spontaneous Virtual Networks” (SpoVNet) project and by the German Research Foundation (DFG) in the Cluster of Excellence “Simulation Technology” (SimTech).

## Bibliography

[AHS06] K. Aberer, M. Hauswirth, A. Salehi. A Middleware for Fast and Flexible Sensor Network Deployment. In *Proceedings of Very Large Data Bases (VLDB)*. Pp. 1199–1202. Seoul, Korea, Sept. 2006.

[AY05] K. Akkaya, M. Younis. A Survey on Routing Protocols for Wireless Sensor Networks. *Ad Hoc Networks* 3(3):325–349, 2005.

[BA07] M. Botts, R. Alex. Bringing the Sensor Web Together. *Geosciences* 6:46–53, Oct. 2007.

[BDF<sup>+</sup>07] M. Balazinska, A. Deshpande, M. Franklin, P. Gibbons, J. Gray, S. Nath, M. Hansen, M. Liebhold, A. Szalay, V. Tao. Data Management in the Worldwide Sensor Web. *Pervasive Computing, IEEE* 6(2):30–40, April–June 2007.

[BKVR10] A. Benzing, B. Koldehofe, M. Völz, K. Rothermel. Multilevel Predictions for the Aggregation of Data in Global Sensor Networks. In *DS-RT '10: Proceedings of the 14th International Symposium on Distributed Simulation and Real Time Applications*. Pp. 1–10. IEEE, October 2010.

[CEF<sup>+</sup>05] O. Cooper, A. Edakkunni, M. J. Franklin, W. Hong, S. R. Jeffery, S. Krishnamurthy, F. Reiss, S. Rizvi, E. Wu. Design Considerations for High Fan-in Systems: The HiFi Approach. In *Proceedings of Conference on Innovative Data Systems Research (CIDR)*. Pp. 290–304. Asilomar, CA, USA, Jan. 2005.

[CEL<sup>+</sup>06] A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, R. A. Peterson. People-centric Urban Sensing. In *Proceedings of International Conference on Wireless Internet (WICON)*. Boston, MA, USA, Aug. 2006. (digital).

[CHPP09] C.-H. Chen-Ritzo, C. Harrison, J. Paraszczak, F. Parr. Instrumenting the Planet. *IBM Journal of Research and Development* 53(3):1–16, 2009.

[DGM<sup>+</sup>04] A. Deshpande, C. Guestrin, S. R. Madden, J. M. Hellerstein, W. Hong. Model-Driven Data Acquisition in Sensor Networks. In *Proceedings of International Conference on Very Large Data Bases (VLDB)*. Pp. 588–599. VLDB Endowment, Toronto, Canada, Aug. 2004.

[DGMS07] Y. Diao, D. Ganesan, G. Mathur, P. Shenoy. Rethinking Data Management for Storage-centric Sensor Networks. In *CIDR '07: Proceedings of the 3rd Biennial Conference on Innovative Data Systems Research*. Asilomar, California, January 2007.

[FJNS09] C. Flores, H. Juárez, M. A. Núñez, M. L. Sandoval. Algorithms for Vector Field Generation in Mass Consistent Models. *Numerical Methods for Partial Differential Equations* 26(4):826–842, April 2009.

[GKK<sup>+</sup>03] P. B. Gibbons, B. Karp, Y. Ke, S. Nath, S. Seshan. IrisNet: An Architecture for a World-Wide Sensor Web. *IEEE Pervasive Computing* 2(4):22–33, Oct. 2003.

[GT02] M. Grossglauser, D. N. C. Tse. Mobility Increases the Capacity of Ad Hoc Wireless Networks. *IEEE/ACM Transactions on Networking* 10(4):477–486, Aug. 2002.

[KKP99] J. M. Kahn, R. H. Katz, K. S. J. Pister. Next Century Challenges: Mobile Networking for Smart Dust. In *Proceedings ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM)*. Pp. 271–278. Seattle, Washington, USA, Aug. 1999.

[KKR08] G. G. Koch, B. Koldehofe, K. Rothermel. Higher Confidence In Event Correlation Using Uncertainty Restrictions. In *Proc. 28th IEEE Int. Conf. on Distributed Computing Systems Workshops (ICDCSW '08)*. Pp. 417–422. IEEE Computer Society, Jun 2008.

[KKR10] G. G. Koch, B. Koldehofe, K. Rothermel. Cordies: Expressive event correlation in distributed systems. In *DEBS '10: Proceedings of the 4th International Conference on Distributed Event-Based Systems*. ACM New York, NY, USA, 2010.

[Luc01] D. C. Luckham. *The Power of Events: An Introduction to Complex Event Processing in Distributed Enterprise Systems*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2001.

[Mei02] E. Meijering. A Chronology of Interpolation: From Ancient Astronomy to Modern Signal and Image Processing. *Proceedings of the IEEE* 90(3):319–342, Mar. 2002.

[RV06] R. Rajagopalan, P. K. Varshney. Data Aggregation Techniques in Sensor Networks: A Survey. *IEEE Communications Surveys and Tutorials* 8(4):48–63, Dec. 2006.

[SPL<sup>+</sup>04] J. Shneidman, P. Pietzuch, J. Ledlie, M. Roussopoulos, M. Seltzer, M. Welsh. Hourglass: An Infrastructure for Connecting Sensor Networks and Applications. Harvard technical report TR-21-04, Harvard, 2004.

[SRJB03] R. C. Shah, S. Roy, S. Jain, W. Brunette. Data MULEs: Modeling a Three-tier Architecture for Sparse Sensor Networks. In *IEEE SNPA Workshop*. Pp. 30–41. Anchorage, Alaska, USA, May 2003.

[TKK<sup>+</sup>10] M. A. Tariq, G. G. Koch, B. Koldehofe, I. Khan, K. Rothermel. Dynamic publish/subscribe to meet subscriber-defined delay and bandwidth constraints. In *The Sixteenth International Conference on Parallel Computing (Euro-Par 2010)*. Pp. 1–13. Springer, August 2010.

[Zha07] Z. Zhang. Routing in Intermittently Connected Mobile Ad Hoc Networks and Delay Tolerant Networks: Overview and Challenges. *IEEE Communications Surveys & Tutorials* 8(1):24–37, Mar. 2007.