Decentralized Processing Strategies for Participatory Sensing Data

Heitor Ferreira, Sérgio Duarte and Nuno Preguiça

CITI / DI-FCT-UNL
Quinta da Torre, 2829-516 Caparica, Portugal

Abstract. Participatory Sensing is a new computing paradigm that aims to turn personal mobile devices into advanced mobile sensing networks. In this paper, we describe and evaluate two distributed strategies for processing participatory sensing data. These strategies are built on top of a decentralized architecture composed on user-contributed machines. The evaluation shows that these strategies spread the load among system nodes, thus providing a scalable solution for supporting participatory sensing applications that requires no central node.

Key words: participatory sensing, distributed processing, mobile computing

1 Introduction

Participatory Sensing [1, 4] is an emergent pervasive computing discipline that aims to leverage the growing ubiquity of sensor capable mobile phones as the basis for performing wide-area sensing tasks. Taking advantage of people's movements in their daily routines, Participatory Sensing promises an enormous potential for gathering valuable data at a fraction of the cost associated with the deployment of dedicated sensor infrastructures. Examples of this potential are well illustrated in [7, 12, 15], which concern road conservation and traffic monitoring, based on accelerometer and GPS readings collected by mobile devices.

Realizing the potential of Participatory Sensing applications poses several challenges, in particular, that of large-scale support, since the appeal of the paradigm could make its applications popular, but also because high data redundancy may be required to address accuracy concerns. Many of the application examples found in the literature [7, 12, 15] use centralized architectures. These are appropriate for small-scale, proof-of-concept experiments, or if the backing of a large institution can be secured. Decentralized processing, however, offers an economically viable alternative to build a community sensing platform supported mainly by its users. At the same time, it can also address privacy concerns by foregoing the control of user data to a single entity.

In this paper, we present two distribution strategies for fully decentralized processing of participatory data. We describe them as an integral part of the 4Sensing platform - a cooperative system for supporting participatory sensing applications, and evaluate both in a simulator environment. A case study traffic monitoring application is used to to compare and discuss their expected performance in a large-scale setting.

Fig. 1: 4Sensing architecture

Fig. 2: Snapshot of the traffic simulation application

The rest of the paper is structured as follows. Section 2 provides an overview of the 4Sensing system. Section 3 details the two distribution strategies proposed in this paper, followed by the results of their experimental evaluation in Section 4. The paper concludes with an overview of related work and some conclusions, respectively, in Sections 5 and 6.

2 System Overview

2.1 System Architecture

The 4Sensing system architecture, depicted in figure 1, comprises two kinds of nodes. A large number of mobile nodes, equipped with sensors, are the main source of data for the system, whereas a fixed node infrastructure provides computing and storage resources. Mobile nodes, typically smartphones, are expected to have limited computational power, battery life and communications capabilities. As such, mobile nodes will mainly support user interaction and run data acquisition services on behalf of the applications that run mostly on the fixed infrastructure. The fixed infrastructure can be composed only of personal computers contributed by the users themselves, allowing a self-sufficient system to be assembled and maintained by its own user community. The fixed infrastructure can include also virtual machines and servers hosted by independent entities.

Fixed nodes are organized as a two-tier overlay network of loosely-coupled geographic partitions or clusters. Each fixed node is given a geographic position. Each geographic partition, individually, forms an one-hop DHT [11], as described in [5], meaning that every node strives to maintain complete knowledge about its cluster's node composition (including their coordinates). These partitions are organized according to geography and are centered on major points of interest, such as a large metropolitan area. This arrangement leverages geography to optimize communications latency and throughput, exploiting the strong correlation expected regarding where sensory data is produced and consumed. Inter-cluster connectivity is provided by a top-tier. This tier is also organized as an one-hop DHT and is populated by a few selected nodes (of each partition) with the added task of maintaining routing information to other clusters.

Mobile nodes do not interact directly, and connect to the fixed infrastructure to upload sensory data and request processed information. A privileged link to a *homebase* fixed node is used to facilitate this and serve as a personal gateway. The fixed homebase is useful in that it minimizes the routing logic and network involvement on the mobile side. For instance, the homebase node can be used to automatically refresh data request leases on behalf of its mobile node.

2.2 Data Model

Applications access participatory sensing data by issuing queries over *virtual tables* - the main high-level data abstraction provided by 4Sensing for naming a collection of data with a particular schema and semantics. Virtual tables have a global scope within a particular instance of the 4Sensing platform. Issuing a query over a virtual table serves two purposes, it allows applications to request data and, at the same time, restrict its scope to a subset of the whole, in the form of spatial and temporal constraints.

While flexible regarding the actual schema, virtual tables handle geo-referenced and time-stamped data. As such, virtual table data is represented as a set of named attributes and, as a common denominator, each data tuple must include a temporal attribute, i.e., a timestamp, and a spatial component, in the form of physical coordinates or a geographic extent. A query over a virtual table will then result in a sequence of data tuples to be forwarded to the requesting application, confined to the requested geographic area and time period.

Virtual tables are purely logical entities, in the sense they do not necessarily need to refer to data already present in some physical storage medium. Interfacing applications with persisted (stored) data is just one of the facets of this construct. Virtual tables can also target data that does not yet exist and that may never be stored, including data intended to be produced and consumed with (near) realtime requirements. Moreover, another central concept to virtual tables is that they refer to data produced as a result of applying a set of transformations to raw sensor inputs and other virtual tables. As such, virtual tables embody a generic inference mechanism that creates high-order data from simpler constituents. This design is intended to facilitate and promote data sharing and cooperation among otherwise unrelated applications.

Data transformation Data transformation is achieved by specifying the execution of a *pipeline* of successive operations. Starting with raw sensor samples, or data from other virtual tables, a derived class of data can be produced, based on the output of a combination of pipeline operators. Data transformations typically involve *mapping*, *aggregation* and *condition detection* operations. Mapping adds new attributes to individual data elements, such as tagging a GPS reading with a street segment identifier. Aggregation combines several samples into one and often involves statistical operators, such as *count*, *average*, etc., over temporal data snapshots. Finally, condition detection generates new data indicative of a particular noteworthy event, such as a congested street. Due to space limitations, we only illustrate them with an example, the same used as the basis for the experimental evaluation discussed in Section 4.

Definition 1 TrafficHotspots virtual table specification

```
sensorInput( GPSReading )
dataSource {
    process{ GPSReading r ->
        r.derive( MappedSpeed, [boundingBox: model.getSegmentExtent(r.segmentId) ])
    }
    timeWindow(mode: periodic, size:10, slide:10)
    groupBy(['segmentId']){
        aggregate( AggregateSpeed ) { MappedSpeed m ->
            sum(m, 'speed', 'sumSpeed')
            count(m, 'count')
} } }
globalAggregation {
    timeWindow(mode: periodic, size:10, slide:10)
    groupBy(['segmentId']){
        aggregate( AggregateSpeed ) { AggregateSpeed a ->
            avg(a, 'sumSpeed', 'count', 'avgSpeed')
    } }
    classify( AggregateSpeed ) { AggregateSpeed a ->
        if(a.count > COUNT_THRESHOLD && a.avgSpeed <= SPEED_THRESHOLD * model.maxSpeed(a.segmentId))
            a.derive( Hotspot, [confidence: Math.min(1, a.count/COUNT_THRESHOLD*0.5) ])
} }
```

The *TrafficHotspots* virtual table, c.f. Definition 1, exemplifies the inference logic for the application devoted to realtime traffic monitoring used in the evaluation section - see figure 2 for a screenshot of the simulation. Starting with discrete GPS readings collected by mobile nodes, it outputs a stream of *Hotspot* tuples, representing real-time detections of congested road segments, when certain thresholds are exceeded. In between, data is progressively transformed into intermediate forms, such as *MappedSpeed* and *AggregateSpeed* representing for each road segment, respectively, an individual speed measurement and the average car speed for the last 10 seconds. An in-depth description of virtual tables and the semantics of these pipeline operators is presented elsewhere [8].

A key aspect about data processing in virtual tables is that it is split between two separate transformation pipeline stages. One, *dataSource*, processes data locally available to the node, i.e., data already stored locally or data that is being acquired and uploaded by mobile nodes. The *globalAggregation* merges the contributions of several nodes. Together, the two pipeline stages abstract the 4Sensing distributed data processing facet. Consequently, their operation is very closely tied to the distribution strategy employed, which among other things determines the shape and properties of the infrastructure aggregation tree used, as discussed further on.

Data Dissemination When a node issues a query, the system instantiates data processing pipelines in certain fixed infrastructure nodes, depending on the geographic coverage of the standing queries. For returning results to the client, 4Sensing adopts a push-based model that leverages a publish/subscribe, content-based routing substrate. As data becomes available in the distributed processing pipelines, it is published and routed by the fixed infrastructure to the interested nodes. When results are to be returned to mobile nodes, they are propagated via their respective *homebase* nodes.

An in-depth discussion of this content-based routing substrate is out of scope of this paper and can be found in [5]. However, the following key character-

istics are important to highlight. The substrate spans all fixed infrastructure nodes. Mobile nodes do not participate directly and interface via their respective homebase nodes. It leverages the one-hop DHTs to assemble, on-the-fly, a random dissemination tree for each message separately. The result is trees that span solely the message's publisher node and the actual recipient nodes (i.e., produces no false positives). Successive messages produce different dissemination trees, which improves load-balancing. Fault-tolerance measures also ensure that no false negatives are introduced in the presence of network churn, so that messages will eventually reach all known recipients at the time of publishing, with high probability.

3 Distribution Strategies

The abstract model for distributed processing outlined above can be materialized by different distribution strategies, which determine how the overall infrastructure is organized to support multiple participatory sensing applications. In particular, each strategy defines the role of the homebase, how data is acquired by the system, how queries are distributed to relevant nodes, and how data is aggregated. Different strategies strive for different strengths. We next describe two strategies that we have implemented in the 4Sensing platform.

3.1 RTree

The main rational behind RTree (Random Tree) is to have the homebase as a personal data repository and use it for privacy protection mechanisms - such as the privacy firewall advocated in [9]. Personal stores can be private - an incentive for users to contribute resources to the system infrastructure - or shared by a group of users. All data acquired by a mobile node is uploaded to the user's homebase, which is expected to cover a potentially large geographic area (e.g., a city-wide area) around the user's home, workplace and in-between.

Query Dissemination and Processing The determining aspect of RTree is that every homebase potentially hosts data relevant to any particular query. Targeting every homebase is too inefficient, as a tradeoff, RTree restricts query distribution by having each homebase define the geographical area it is willing to serve - its *query filter*. In general, this area will be restricted to the regions its owner visits more often, thus guaranteeing that all (or most of) collected data is shared with the community.

For the purpose of distributed processing, RTree treats overlapping queries as a single one. Therefore, for each set of overlapping queries, RTree constructs a random tree spanning all homebases whose query filters intercept the union of the query areas. The algorithm that assembles these trees is fully distributed and leverages the same publish/subscribe substrate used to disseminate query results. When a mobile node issues a new query, its homebase publishes the query (characterized by a geographic extent). All homebases with intercepting query filters (and only those) will receive the query and will try to merge it with existing queries. If a node determines that the new query overlaps the boundary of already existing queries, the resulting compound (wider) query is published, repeating the process. Eventually, no more overlaps are produced and the process stops.

Since the publish/subscribe substrate already produces a random tree for each published message, each node only has to record the path each query has taken to determine its place in the random aggregation tree. When individual queries are not renewed and expire, it can cause larger compound queries to shrink or split into disjoint pieces. This is also handled in a distributed fashion. When a node determines a individual query has expired, it eliminates it and re-merges the remaining ones. Again, if there is a change in the resulting set, it publishes the new merged queries, mimicking the initial process. Logical timestamps and other optimizations are used to limit the number of messages exchanged and handle concurrent re-publish operations from different nodes. To improve on load-balancing, merged queries can be re-published periodically, producing a new tree in the process. RTree generates aggregation trees with a default fanout degree of 6, but it can instruct the publish/subscribe substrate to use other values, for instance to control tree depth.

In a given RTree instance, every node (homebase) has the same role and will instantiate both the data source and global aggregation pipeline stages. Since a node cannot determine if it has complete information for a particular spatial extent, data aggregation and processing may have to proceed until the root is reached. This is particularly problematic for virtual tables whose inference logic requires detecting the absence of certain data patterns. Early detections can be published as soon as they are produced in any point of the aggregation tree. As discussed earlier, data dissemination uses dedicated trees and is delivered exclusively to nodes with matching (individual) queries.

3.2 QTree

In QTree (Quad Tree), the idea is to do an *a priori* partition of the data by performing a regular subdivision of the space into quadrants, while the number of infrastructure nodes located in each area exceeds a *minimum occupancy* threshold. Each infrastructure node belongs simultaneously to all the quadrants that contain it, down to the smallest - called its *maximum division quadrant*. Under this scheme, acquired sensor data is committed and bound to any of the fixed nodes that lie in the smallest quadrant that fully encloses the data's coordinates or geographic extent. Therefore, in QTree, dataSource pipeline stages will process data acquired by any mobile node. In the deepest levels of the QTree aggregation tree, all the processed data pertains to the immediate neighborhood of the processing node. This enables QTree nodes to potentially detect highly localized phenomena more efficiently and sooner in the aggregation tree.

Query Dissemination and Processing QTree and RTree share a large number of features: overlapping queries are merged and split using the same re-publish strategy; a random tree containing all relevant infrastructure nodes is created by disseminating the query using a distributed algorithm; periodically re-publishing the query can be used to improve load-balancing; data dissemination uses dedicated trees spanning only nodes with matching queries.

The algorithms differ mainly on the query filter used and how queries are routed to build the aggregation tree. In QTree, the *query filter* corresponds

to the *minimum subdivision quadrant* it belongs to[1]. As for query (content-based) routing, it is performed by recursively subdividing space along quadrant boundaries. At each point of the routing tree, the current quadrant is divided into sub-quadrants, selecting for each of them, a child node, at random, from those enclosed and whose query filters intercept the query area. Along the way, the query path is updated and appended to the query as it travels along the dissemination tree, which will also serve as the data processing aggregation tree.

The role of homebase node in QTree is mainly that of a gateway to the fixed infrastructure. They do not necessarily store or process the data collected by their respective mobile nodes, but commit and route it to a suitable infrastructure node. By binding data to nodes close by, QTree allows for early detection of some data patterns (e.g., a certain number of samples in a given area has been reached), which can be published and notified to applications sooner, typically deep in the aggregation tree.

4 Evaluation

The evaluation of the proposed algorithms has been achieved through the simulation of an application for detecting congestion in city traffic. The application detects congestions in a given area, based on GPS data sampled by the simulated in-transit vehicles at periodic intervals. For this purpose, one virtual table is used: TraficHotspots, which supports querying for congestion detections, computed from average speeds, as explained in the previous sections.

In our simulation, we are using the Open Street Map (OSM) [17] vectorial representation of the road network of Lisbon city. This map data is used to map geographic coordinates to road segments, to determine the spatial extent of segments and their associated road type. The model used in the evaluation divides roads, as needed, into segments with a maximum of 1 km and uses separate segments for each driving direction. Each road is assigned an expected (uncongested) driving speed according to its type: highway, primary to tertiary and residential.

In the experiments, we simulated 5000 fixed nodes and 50000 mobile nodes. Fixed nodes are distributed randomly across the urban space with a minimum inter-node distance of 80 meters. Mobile nodes simulate in-transit vehicles, according to a traffic model, and report GPS readings every 5 seconds as they follow the assigned paths. They interact with the respective homebase counterpart directly, resulting in the delivery of raw GPS data with no latency. Communication among fixed nodes experiences latency and jitter. A common clock is used to timestamp readings; any effects of clock desynchronization are not considered.

Traffic is modeled by emulating a fleet of vehicles driving through random routes. The maximum speed for a given segment is the same value used for congestion detection and depends on the road type. An average speed, for each road segment at a given time, is determined by its current car density and used to generate the random speed individually for each vehicle, according to a normal

---

[1] QTree assumes space is sufficiently populated that it can be subdivided to achieve a uniform *minimum subdivision quadrant.*

Acquisition	278,969	96.3%
Aggregation	10,614	3.7%
Total	289,583	100%

Table 1: Workload using centralized processing (tuples/minute)

distribution. In the experiments performed, congestion occurs in segments with a density of at least 50 vehicles. Vehicle paths are determined by choosing a random start position and sequence of road intersections; a new path is assigned whenever a vehicle reaches its destination. Figure 2 shows a rendering of the traffic simulation.

In the following results, we have evaluated the effects of a single query covering 25% of the overall simulation area in an area with high density of mobile nodes. Note that having a query over a large area is what happens when multiple queries over nearby small areas (e.g., several points in the city downtown) are merged in our algorithms. The set of metrics captured was averaged over 5 runs, corresponding to different fixed node placements.

Workload Distribution We have started by evaluating how the effort required to evaluate a query is spread among the fixed nodes in the different distribution strategies. For comparison, we have also computed the effort that a centralized solution, with a single central node, would experience.

The workload is measured as the number of data tuples processed at each fixed node in both acquisition and aggregation steps. Specifically, the acquisition pertains to the number of GPS sensor readings received and processed by the acquiring node, while the aggregation refers to the number of inputs handled by the global aggregation stage and corresponds to the updates received for each segment aggregated by that node.

Table 1 presents the workload using a centralized approach, while Figure 3 shows the workload distribution using the different distribution strategies, which for RTree involved aggregation trees with fanout degree of 4 and 8, besides the default value of 6.

The results show that using a decentralized approach, the most loaded node only processes a small fraction of the tuples processed in a centralized approach - 0.8% for QTree and 11.5% for the best RTree configuration. Comparing the two distribution strategies, QTree balances load very effectively, while in RTree, irrespective of the configuration, there are some nodes with a much higher load. The reason for this is that the detection of congestion in QTree is done faster, in the lower levels of the tree, because tuples of the same area will be aggregated in some node sooner. In RTree, there is often the need to propagate the information up to the top of the tree to check if a congestion exists or not, as data for the same areas may be propagated in any path. This explains why the most loaded nodes, corresponding to the nodes near the root of the tree, have so much load.

Communication Load We have also measured the communication cost, as the number of messages exchanged during the execution of a query. The communication load includes one-hop DHT overhead (4.4 messages/min per node)

Fig. 3: Workload distribution - for the 50 most loaded nodes

Fig. 4: Total communication cost

[2], acquisition messages, data messages and binding events. Acquisition messages are those sent by the mobile nodes to a fixed node. Data messages carry the tuples exchanged between peers, relative to query data (incomplete aggregations that are forwarded up the tree). Binding events capture the additional overhead associated with uploading the data from mobiles to a fixed node in QTree, where the node the data is sent varies over time. In all cases, the number of messages sent by individual nodes is limited at 1 message per pipeline stage every 10 seconds by the use of the *timeWindow* operator.

Figure 4 presents the total communication cost in the centralized scenario and in both distribution strategies presented. As expected, decentralized approaches involve additional messages, as data with intermediate computations needs to be exchanged among nodes in the system. However, results show that these messages represent only a small fraction of the messages propagated from the mobile nodes to the fixed infrastructure in any case. The reason for this is that messages exchanged by the fixed node usually carry data that is computed from a large number of data items received from mobile nodes and data propagation can stop whenever a detection is performed at some tree node.

Comparing the two distribution strategies, RTree exchanges more data messages than QTree, as expected, because congestion computation tends to be computed in higher levels of the tree, requiring more messages to be propagated. However, QTree has an additional overhead related with binding events, as mobile nodes need to know to which fixed node to propagate a message when they move to a different area. This overhead is important, leading to an overall larger communication cost in QTree when compared with RTree. In the future, we intend to explore strategies to reduce the need for these binding events, e.g., by caching in mobile nodes the identity of fixed nodes responsible for a query in adjacent areas.

Detection Latency Detection latency measures the lag between the occurrence of a segment congestion and its detection by the system. Transient congestions (lasting less than 20 seconds) were not considered for the evaluation. As the traffic patterns produced by the traffic model are highly dynamic compared to real world conditions, with frequent short lived congestions, these experiments represent a challenging scenario to the system.

---

[2] Corresponding to a fixed cost of 7.5 join/leave DHT events per minute, cf. [5].

Fig. 5: Detection latency

Fig. 6: Success rate in the presence of node failures (at time = 100)

Figure 5 plots the detection latency for the accumulated level of successful detections. Results show clearly that QTree detects congestions faster than RTree. The reason for these results lie in the fact that in QTree, related data flow to the same nodes faster, making detections to occur in the lower levels of the tree. Up to 35% of successful detections, latency is similar for QTree and RTree. This is explained by the fact that, when a road is very congested, with much more cars than necessary for a detection, even in RTree enough data is aggregated in the lower levels of the tree.

Failures We implemented a simple failure recovery mechanism consisting in rebuilding the aggregation tree periodically or when failures exceed a certain threshold. Using this strategy, we observed that the loss of a small percentage of nodes has little impact on the query success rate, depending on the sensor acquisition rate used. Figure 6 summarizes those results, which involved sensor acquisition intervals of 5, 10 and 15 seconds and the sudden failure of 15% of the fixed nodes. As expected, detection success is adversely affected by node failures but afterwards it recovers to its original values. The lower the acquisition rate, the more pronounced the effect is. This behavior is explained by the fact that node failures cause the loss of some of the acquired data and, until a new tree is built, aggregation is impaired due to broken tree paths. If the acquisition rate is high enough, the effect can be negligible due to data redundancy. In the opposite case, the impact is more severe since the amount of lost data is proportionally higher. Comparing both distributed strategies, RTree suffers the most from node failures. In RTree node failures tend to have a global impact because data is spread among all nodes, so it requires enough data redundancy and tree integrity to operate properly. Due to its *a priori* reliance on data locality, QTree does not need to aggregate as much, so the effect of node failures is also much more localized and impacted less by problems in the aggregation tree.

5 Related Work

Participatory sensing has become popular recently [1, 4], but a large number of supporting systems have already been designed.

In BikeNet [6], users can share information about their bike rides. PEIR [16] computes estimates of environmental impact and exposure from location samples. CarTel [12, 7] focus on vehicle based sensing applications and served as

inspiration for our evaluation scenario. Although these systems are different, and some even rely on a delay-tolerant network substrate, they all are based on a centralized approach. Our work studies decentralized solutions that improves the scalability of the system, by spreading the load requirements by a large number of nodes, as shown in the results of our evaluation. Lacking the need for having a powerful server infrastructure also makes this approach more suitable for community-based sensing, with the needed resources being contributed by the community.

Some sensing systems present an architecture built entirely in the mobile nodes and ad-hoc coordination to support applications (e.g. [18] and [13]). A limitation of this approach is that mobile nodes have to spend computation, communication and energy resources in coordination efforts, regarding node and service discovery and context dissemination. This is specially relevant given that communication can be expensive and energy is scarce.

Other systems (e.g., [14, 10]) have an architecture more similar to ours, based on the combination of fixed and mobile nodes. For example, in SensorWeb [10], a set of fixed nodes act as gateways for sensor network and proxies for mobile devices. In this system, a coordinator node mediates and coordinates the access of applications to the different gateways and proxies. Unlike SensorWeb, in our system, nodes are also used to process information, thus allowing to more evenly distribute the load.

In wireless sensor networks, a large number of decentralized algorithms have been proposed [19]. However, these algorithms tend to focus on minimizing energy consumption and nodes can only communicate with nearby nodes. Thus, the trees that are formed to aggregate results have very different constraints when compared to our approach. Our algorithms are closer to distributed processing of data streams - e.g. [2, 3]. However, the application scenario makes our solution different, by having mobile nodes to feed geo-referenced data into the system periodically. In QTree, we explore this property for aggregating related data at lower levels of the dissemination tree, allowing to make detections faster.

6 Conclusions

In this paper, we presented and evaluated two algorithms for decentralized processing of Participatory Sensing data. Namely, RTree, which leverages random aggregation trees and personal data repositories, and QTree, which combines random trees with regular, *a priori*, subdivision of geographic space.

Their experimental evaluation using simulation and a case-study application, allow us to draw the following conclusions. Both decentralized solutions, in regards to data processing, show a small total message overhead compared to a centralized solution, and are clearly competitive. The most loaded node processes a small fraction of the data in both algorithms. QTree, however, performs better in this respect, spreading the processing load more evenly. QTree also shows that it can produce results earlier, exhibiting a smaller processing latency.

As for future work, we intend first to scrutinize how periodic update of the aggregation trees impacts on load-balancing, processing overhead and query suc-

cess and latency. Moreover, we want to evaluate inter-partition interactions and measure the impact on the overall system performance of inter-partition queries.

References

1. A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, R. A. Peterson, H. Lu, X. Zheng, M. Musolesi, K. Fodor, and G.-S. Ahn. The rise of people-centric sensing. *Internet Computing, IEEE*, 12(4):12–21, 2008.
2. M. Cherniack, H. Balakrishnan, M. Balazinska, D. Carney, U. Çetintemel, Y. Xing, and S. B. Zdonik. Scalable distributed stream processing. In *CIDR*, 2003.
3. T. Condie, N. Conway, P. Alvaro, J. M. Hellerstein, K. Elmeleegy, and R. Sears. Mapreduce online. In *NSDI2010: Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*, 2010.
4. D. Cuff, M. Hansen, and J. Kang. Urban sensing: out of the woods. *Commun. ACM*, 51(3):24–33, 2008.
5. S. Duarte, J. L. Martins, M. Mamede, and N. Preguiça. Catadupa: A load-balancing and fault-tolerant full-membership substrate for content-based routing. Technical Report UNL-DI-5-2010, Universidade Nova de Lisboa, August 2010.
6. S. B. Eisenman, E. Miluzzo, N. D. Lane, R. A. Peterson, G.-S. Ahn, and A. T. Campbell. The bikenet mobile sensing system for cyclist experience mapping. In *SenSys '07: Proceedings of the 5th international conference on Embedded networked sensor systems*, pages 87–101. ACM, 2007.
7. J. Eriksson, L. Girod, B. Hull, R. Newton, S. Madden, and H. Balakrishnan. The Pothole Patrol: Using a Mobile Sensor Network for Road Surface Monitoring. In *MobiSys '08: Proceeding of the 6th Int. Conf. on Mobile systems, applications, and services*, June 2008.
8. H. Ferreira. 4Sensing - Distributed Processing for Participatory Sensing Data. Master's thesis, DI - FCT - Universidade Nova de Lisboa, June 2010.
9. R. K. Ganti, N. Pham, Y.-E. Tsai, and T. F. Abdelzaher. Poolview: stream privacy for grassroots participatory sensing. In *SenSys '08: Proceedings of the 6th ACM conference on Embedded network sensor systems*, pages 281–294. ACM, 2008.
10. W. Grosky, A. Kansal, S. Nath, J. Liu, and F. Zhao. Senseweb: An infrastructure for shared sensing. *Multimedia, IEEE*, 14(4):8–13, Oct.-Dec. 2007.
11. A. Gupta, B. Liskov, and R. Rodrigues. One hop lookups for peer-to-peer overlays. In *Ninth Workshop on Hot Topics in Operating Systems*, pages 7–12, May 2003.
12. B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. K. Miu, E. Shih, H. Balakrishnan, and S. Madden. CarTel: A Distributed Mobile Sensor Computing System. In *4th ACM SenSys*, November 2006.
13. N. Kotilainen, M. Weber, M. Vapa, and J. Vuori. Mobile chedar: A peer-to-peer middleware for mobile devices. In *PERCOMW '05: Proc. of the 3rd IEEE Int. Conf. on Pervasive Computing and Communications Workshops*, pages 86–90. IEEE Computer Society, 2005.
14. Y. J. L. Marie Kim, Jun Wook Lee and J.-C. Ryou. Cosmos: A middleware for integrated data processing over heterogeneous sensor networks. *ETRI Journal*, 30(5), October 2008.
15. Mohan, Prashanth and Padmanabhan, Venkata and Ramjee, Ramachandran . Nericell: Rich Monitoring of Road and Traffic Conditions using Mobile Smartphones. In *Proceedings of ACM SenSys 2008*, November 2008.
16. M. Mun, S. Reddy, K. Shilton, N. Yau, J. Burke, D. Estrin, M. Hansen, E. Howard, R. West, and P. Boda. Peir, the personal environmental impact report, as a platform for participatory sensing systems research. In *MobiSys '09: Proc. of the 7th Int. Conf. on Mobile systems, applications, and services*, pages 55–68. ACM, 2009.
17. OpenStreeMap. http://www.openstreetmap.org, April 2010.
18. O. Riva and C. Borcea. The urbanet revolution: Sensor power to the people! *Pervasive Computing, IEEE*, 6(2):41–49, April-June 2007.
19. J. Yick, B. Mukherjee, and D. Ghosal. Wireless sensor network survey. *Comput. Netw.*, 52(12):2292–2330, 2008.