

A Network Architecture for Distributed Event-Based Systems in an Ubiquitous Sensing Scenario

Cristina Muñoz, Pierre Leone
 Department of Computer Science
 University of Geneva
 Carouge, Switzerland
 Email: {Cristina.Munoz, Pierre.Leone}@unige.ch

Abstract—Ubiquitous sensing devices frequently disseminate their data between them. The use of a distributed event-based system that decouples publishers of subscribers arises as an ideal candidate to implement the dissemination process. In this paper, we present a network architecture which merges the network and overlay layers of typical structured event-based systems. Directional Random Walks (DRWs) are used for the construction of this merged layer. Our first results show that DRWs are suitable to balance the load using a few nodes in the network to construct the dissemination path. As future work, we propose to study the properties of this new layer and to work on the design of Bloom filters to manage broker nodes.

Keywords—Distributed Event-Based systems; Directional Random Walks; Bloom filters; Wireless Sensor Networks.

I. INTRODUCTION

Ubiquitous or pervasive computing [1] uses many sources and destinations to gather and process data related to physical processes with the aim of making possible human-computer interaction. In the process of dissemination, some devices generate the data, while others are waiting for the sensing data. In this context, the use of a distributed event-based system [2] arises as an ideal candidate to implement the model of communication on the reception or transmission of events.

The main characteristic of an event-based system is that publishers and subscribers are decoupled. This means that they do not have any information about each other. The element in charge of matching notifications with subscriptions is called the event notification service. In distributed networks, the event notification service is implemented using a network of brokers nodes (see Figure 1). It is considered that a broker is any node in the network that has information about any single or set of subscriptions. The complexity of designing this type of systems usually lies on the way to elect the nodes which will act as brokers because of the decentralized nature of a distributed network.

In our research, we assume that a node can be a publisher, a subscriber, a broker or a combination of these three possibilities. We also assume that all the nodes in the network are able to participate in it without the requirement to adopt the specific role of publisher or subscriber. Nodes that are actively participating in the network but do not take any specific role will be considered as part of the overlay layer. Those nodes of the overlay layer that are able to redirect messages will be considered as brokers.



Figure 1: Distributed notification service using a network of brokers.

Event-based systems are classified as topic-based or content-based [2]. Topic-based systems take into account the subject of messages in order to match publications with subscriptions. Content-based systems use filters to specify the value of subscriptions' attributes to redirect notifications. A filter is a boolean function built taking into account the set of subscriptions. In our proposal, we plan to deal with a content-based system that uses Bloom filters at broker nodes in order to save memory resources and speed up routing decisions.

Sensor networks frequently use tiny devices with limited battery capabilities that make unsuitable the use of a Global Positioning System (GPS) to disseminate information according to the coordinates of nodes. In addition to this, the use of virtual coordinates to substitute real coordinates requires the use of sinks or landmarks to structure the network. For these reasons, the use of coordinates in an unstructured sensing scenario is not recommended. We assume that we work in an unstructured scenario in which no routing protocol provides communication between the nodes of the network.

The constraints of the network's infrastructure lead us to the design of a network architecture for distributed event-based systems that must use as less resources as possible (i.e., battery, memory, etc.). In this paper, we present a solution that avoids implying all the nodes of the network in the dissemination process by using a distributed notification service defined by Directional Random Walks (DRWs).

The rest of this paper is organized as follows: Section II analyzes the state of the art. Section III points out the approach to solve the problem specified in this section. Section IV presents the research efforts already done for the approach specified in Section III. Finally, Section V summarizes our proposal and suggests further work.

II. STATE OF THE ART

A. Distributed and Structured Event-based Systems

Distributed and structured event-based systems use three layers on the top of a bottom layer (see Figure 2), which provides data link functionalities, to facilitate topology control:

- 1) The network layer is in charge of providing data forwarding between the different nodes involved in the network. A network protocol, such as the Multicast Ad-hoc On-demand Distance Vector (MAODV) [3] is needed to provide point-to-point communication.
- 2) The medium layer is called the overlay layer. It is a virtual layer that builds the event notification service by providing a network of brokers that redirect notifications to the corresponding subscribers.
- 3) Finally, on the top layer the event-based protocol is implemented.

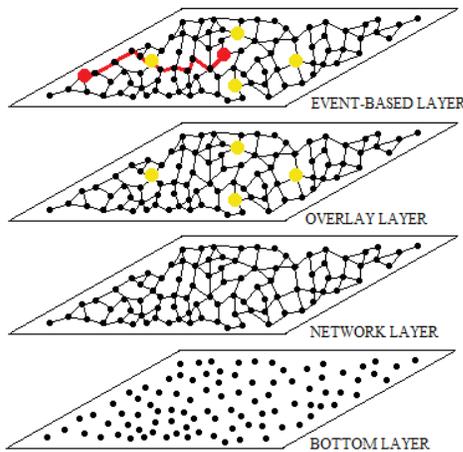


Figure 2: Decomposition in layers of the typical design of a distributed and structured event-based system.

One strategy to construct the overlay layer is to use a tree. In TinyMQ [4], which is designed specifically for wireless sensor networks, a multi-tree overlay layer is maintained.

Another strategy is to clusterize the network and use cluster heads to manage messages as in Mires [5], which is a middleware for sensor networks. The Gradient Landmark-Based Distributed Routing (GLIDER) [6] organizes the network using some defined landmarks to compute the Delaunay graph for network partition. Then, the Landmark-Based Information Brokerage scheme (LBIB) [7] uses an overlay layer based in GLIDER to match publishers with subscribers.

A typical solution is to build the overlay layer using Distributed Hash Tables (DHTs). In these systems, a key is mapped to a particular node with storage location properties. In some DHT architectures, rendez-vous nodes depend on the node ID as in Pastry [8]. In others, as the Content Addressable Network (CAN) [9], a region of the space is used to map a key. Some efforts have been made to apply this solution to sensor networks [10]. When coordinates are available, sensor networks use Geographic Hash Tables (GHT) instead of a typical DHT. Currently, technology companies as Ericsson Research, are making an effort to develop applications that use GHTs in wireless sensor networks [11].

B. Distributed and Unstructured Event-based Systems

The main characteristics of distributed and unstructured event-based systems is that they do not maintain an overlay layer. This fact makes easier to deal with network's changes. The distributed notification service may be built using flooding, gossiping or random walks.

Most of the algorithms proposed deal with the unstructured nature of wireless communications using flooding to

build a tree. A typical solution is to use the On-Demand Multicast Routing Protocol (ODMRP) [12], which is based in the forwarding group concept. Groups are constructed and maintained periodically when a multicast source has data to send. This task is done by broadcasting the entire network with membership information. An extension for ODMRP has been proposed [13] to adapt a content-based system by adding subscriptions to Bloom filters. Trees also may be configured to self-repair themselves in base to brokers' dynamicity [14]. These solutions are reliable but increase the traffic of the network because they use flooding at some point.

Flooding may also be used to continuously exchange subscription information clusterizing the network [15]. Then, notifications are sent to the appropriate cluster, improving the efficiency of the network. Other mechanisms can be used as the combination of a DHT and random walks [16]. Cluster heads manage the DHTs while random walks help to connect the different cluster heads of the network. The cluster concept in the network of brokers can be improved in a dynamic scenario by enriching the topology management with predictions based on location [17].

1) *Probabilistic approaches:* Probabilistic approaches are suitable to deal with dynamicity but they do not offer reliability. Some solutions propose that all the nodes in the network implement a broker that forwards messages to neighbors depending on the estimation of potential subscribers [18]. Other solutions [19], propose to flood subscriptions in a small area and then use random walks to reach that areas. In Quasar [20], subscriptions of a certain area are able to attract or reject notifications, that are propagated with a random walk, using an attenuated Bloom filter [21]. A probabilistic solution that uses a random walk specifically designed to go deep into the network is CoQUOS (Continuous Queries on Unstructured OverlayS) [22]. Continuous queries are launched to the network using random walks. Peers compute the overlap between their neighbor's lists and use this information to forward the random walk to avoid remaining in a cluster. Then, some peers register the query with a probability that depends on the number of hops.

III. NETWORK ARCHITECTURE

Due to the unstructured nature of our network, we propose the development of a dissemination algorithm that merges the network and the overlay layers of a typical distributed and structured event-based system (see Figure 2). This means that no network protocol is needed. The main advantage of not using a network protocol is that there is no necessity of maintaining a network topology. This implies that most nodes of the network, which do not actively participate in the process of dissemination, do not have to keep any information about topology. The main consequence is that nodes not involved in the system are able to save energy and computing resources.

Our design (see Figure 3) uses two layers on the top of a bottom layer that provides data link functionalities:

- 1) The overlay layer is in charge of providing the distributed network of brokers and, at the same time, provides point-to-point communication between publishers and subscribers. The main objective of this strategy is to avoid the use of global information of the network which is costly to get and maintain.
- 2) As in Figure 2, the event-based protocol is implemented at the top layer.

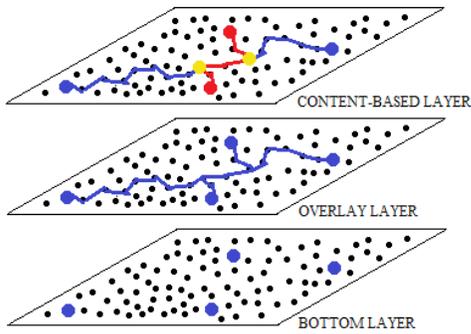


Figure 3: Decomposition of the architecture of our design in layers.

As Section I mentions, we assume that a node can be a publisher, a subscriber, a broker or a combination of these three possibilities. Moreover, our design takes advantage of some nodes in the network which want to collaborate. Nodes that participate in the system are considered part of the overlay layer (blue path of Figure 3). The overlay layer is formed by the intersection of different publishers and subscribers (blue nodes). Publishers and subscribers implement a DRW until intersecting other DRW. Broker nodes (yellow nodes) are the meeting point between two DRWs.

A DRW is a probabilistic technique able to go forward into the network following a loop-free path. The principle assumed in this strategy is that two lines in a plane cross (see Figure 4). It is unclear how to construct a straight path of relaying nodes in ubiquitous unstructured networks without requiring global information and without making use of geo-coordinates. In this research, two different methods have been proposed in order to build DRWs [23][24].

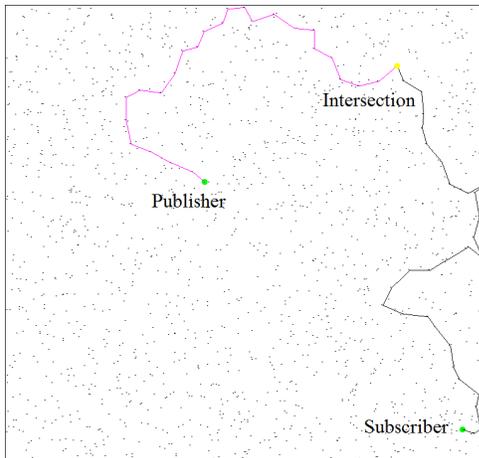


Figure 4: Directional Random Walks intersecting using a Java simulator.

The matching of publishers and subscribers will be done using a special architecture of Bloom filters [21] implemented at broker nodes. Bloom filters are probabilistic data structures that efficiently manage membership of a certain number of elements. The content related to membership is hashed using different hashing algorithms. Then, the positions of the Bloom structure corresponding to the hashes are set to one. The maximum number of elements to be inserted to the filter is fixed in order to maintain a certain probability of false positives. When searching for elements of a certain membership, the correspondent positions of the data structure are checked. The main advantage of Bloom filters is that they do not require

much memory space and processing resources; so its use is very convenient in a sensing scenario in which devices have limited capabilities.

It is remarkable to mention that in our event-based system no advertisement table is required because filters just manage information about subscriptions.

IV. RESEARCH EFFORTS

In this section, we present the efforts already made in order to build the overlay layer proposed.

Based on [23], a first method to build DRWs is proposed. It is based in the addition of different nodes to the DRW by pre-computing different weights at each node that take into account the two hops path. A weight is defined as follows:

$$n_{xz}^y = |N(x) \cap N(z)| \quad (1)$$

where y is the last node added to the DRW; x is the penultimate node added to the DRW; z can be any node of the set $N(y)$ and $N(a)$ is the set of neighbor nodes of node a . Furthermore, in this method, a penalty is added to the weight when a node is added to the DRW.

Some properties about our heuristics were found using extensive simulations. The first property claims that DRWs decrease the time to intersection compared to pure random walks. The second property states that cooperation also decreases the time to intersection. Cooperation refers to synchronicity between publishers and subscribers. Finally, it is shown that DRWs are good at balancing the load of the network.

Based on [24], a second method to build DRWs is proposed. The main difference with the first design presented for DRWs is that nodes of the first and second neighborhoods of nodes added to the DRW are marked. In addition to this, the cost is not pre-computed, but it is computed when selecting a node as follows:

$$c(v) = \alpha|N(v) \cap N(DRW)| + \beta|N(v) \cap N^2(DRW)| \quad (2)$$

where α and β are parameters used as weights; v can be any node of the set related to the neighborhood of the last node added to the DRW; DRW is the set of nodes that are part of the DRW; $N(a)$ is the set of neighbor nodes of node a and $N^2(DRW)$ is the set of neighbor nodes of $N(DRW)$.

In the first part of this research, the properties associated with a DRW were assessed. Implementations of DRWs of one branch or two branches were studied. The main results show that the use of one branch is as efficient as the use of two branches. Moreover, it is shown that the use of second neighborhoods to forward the DRW does not improve the Euclidean distance traversed in the network. It is also shown that shorter paths are obtained when using higher densities of nodes in the network. In the second part of this research, an information brokerage system was evaluated using a double ruling method. As in the first paper, it is shown that the algorithm is good at balancing the load using a few nodes of the network. In fact, we can state that the method proposed is as good as a traditional Rumor Routing algorithm [25] with an infinite memory.

V. CONCLUSION AND FUTURE WORK

In this paper, a novel network architecture for distributed event-based systems that use sensing devices has been proposed. Our first results, validated through extensive simulations, show that DRWs are suitable for the construction of an

overlay layer that provides point-to-point communication and a distributed notification service. This is due, mainly to the good properties of DRWs to balance the load using a few nodes of the network for the establishment of paths.

Currently, we are working on the first simulations of the overlay layer to study the impact of having different number of publishers and subscribers. Figure 5 shows a simulation of the overlay layer in which yellow squares represent the distributed network of brokers while publishers and subscribers are represented using green circles.

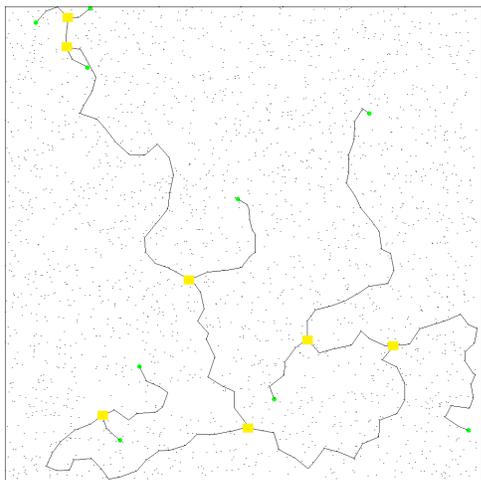


Figure 5: Overlay layer using Directional Random Walks in a Java simulator.

At the same time, we are working on the design of Bloom filters for broker nodes that will be specifically fitted for sensing constrained devices.

Finally, the deployment of a sensing testbed that will use wireless sensor nodes, will evaluate the suitability of the proposed solution under real conditions.

ACKNOWLEDGMENT

This work has been developed as part of the POPWiN project (Parallel Object Remote Programming for Heterogeneous Wireless Networks over IPv6) that is financially supported by the Hasler Foundation in its SmartWorld - Information and Communication Technology for a Better World 2020 program.

REFERENCES

- [1] D. J. Cook and S. K. Das, "Review: Pervasive computing at scale: Transforming the state of the art," *Pervasive Mob. Comput.*, vol. 8, no. 1, Feb. 2012, pp. 22–35.
- [2] G. Mühl, L. Fiege, and P. Pietzuch, *Distributed Event-Based Systems*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006.
- [3] S. Roy, V. Addada, S. Setia, and S. Jajodia, "Securing maodv: attacks and countermeasures," in *Sensor and Ad Hoc Communications and Networks, 2005. IEEE SECON 2005. 2005 Second Annual IEEE Communications Society Conference on*, Sept 2005, pp. 521–532.
- [4] K. Shi, Z. Deng, and X. Qin, "Tinymq: A content-based publish/subscribe middleware for wireless sensor networks," in *SENSORCOMM 2011, The Fifth International Conference on Sensor Technologies and Applications, 2011*, pp. 12–17.
- [5] E. Souto et al., "Mires: a publish/subscribe middleware for sensor networks," *Personal and Ubiquitous Computing*, vol. 10, no. 1, 2006, pp. 37–44.
- [6] Q. Fang, J. Gao, L. J. Guibas, V. Silva, and L. Zhang, "Glider: Gradient landmark-based distributed routing for sensor networks," in *Proc. of the 24th Conference of the IEEE Communication Society (INFOCOM, 2005*, pp. 339–350.
- [7] Q. Fang, J. Gao, and L. J. Guibas, "Landmark-based information storage and retrieval in sensor networks," in *In The 25th Conference of the IEEE Communication Society (INFOCOM06, 2006*, pp. 1–12.
- [8] A. I. T. Rowstron and P. Druschel, "Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems," in *Proceedings of the IFIP/ACM International Conference on Distributed Systems Platforms Heidelberg*, ser. *Middleware '01*. London, UK, UK: Springer-Verlag, 2001, pp. 329–350.
- [9] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker, "A scalable content-addressable network," in *Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications*, ser. *SIGCOMM '01*. New York, NY, USA: ACM, 2001, pp. 161–172.
- [10] G. Fersi, W. Louati, and M. Ben Jemaa, "Distributed hash table-based routing and data management in wireless sensor networks: a survey," *Wirel. Netw.*, vol. 19, no. 2, Feb. 2013, pp. 219–236.
- [11] H. Mahkonen, T. Jokikyyny, J. Jiménez, and S. Kukliński, "M3: Machine-to-machine management framework," in *Proc. 3rd Intl. Conference on Sensor Networks, SENSORNETS, 2014*, pp. 139–144.
- [12] S. J. Lee, W. Su, and M. Gerla, "On-demand multicast routing protocol in multihop wireless mobile networks," *Mob. Netw. Appl.*, vol. 7, no. 6, Dec. 2002, pp. 441–453.
- [13] E. Yoneki and J. Bacon, "An adaptive approach to content-based subscription in mobile ad hoc networks," in *Proceedings of the Second IEEE Annual Conference on Pervasive Computing and Communications Workshops*, ser. *PERCOMW '04*. Washington, DC, USA: IEEE Computer Society, 2004, pp. 92–97.
- [14] L. Mottola, G. Cugola, and G. P. Picco, "A self-repairing tree topology enabling content-based routing in mobile ad hoc networks," *IEEE Transactions on Mobile Computing*, vol. 7, no. 8, Aug. 2008, pp. 946–960.
- [15] S. Voulgaris, E. Rivire, A.-M. Kermarrec, and M. V. Steen, "Sub-2-sub: Self-organizing content-based publish subscribe for dynamic large scale collaborative networks," in *In IPTS06: the fifth International Workshop on Peer-to-Peer Systems, 2006*.
- [16] R. Tian et al., "Hybrid overlay structure based on random walks," in *Proceedings of the 4th international conference on Peer-to-Peer Systems*, ser. *IPTPS'05*. Berlin, Heidelberg: Springer-Verlag, 2005, pp. 152–162.
- [17] F. Abdennadher and M. Ben Jemaa, "Accurate prediction of mobility into publish/subscribe," in *Proceedings of the 11th ACM International Symposium on Mobility Management and Wireless Access*, ser. *MobiWac '13*. New York, NY, USA: ACM, 2013, pp. 101–106.
- [18] J. Hailiot and F. Guidec, "Content-based communication in disconnected mobile ad hoc networks," in *Proceedings of the 8th international conference on New technologies in distributed systems*, ser. *NOTERE '08*. New York, NY, USA: ACM, 2008, pp. 21:1–21:12.
- [19] P. Costa and G. Picco, "Semi-probabilistic content-based publish-subscribe," in *Distributed Computing Systems, 2005. ICDCS 2005. Proceedings. 25th IEEE International Conference on*, 2005, pp. 575–585.
- [20] B. Wong and S. Guha, "Quasar: a probabilistic publish-subscribe system for social networks," in *Proceedings of the 7th international conference on Peer-to-peer systems*, ser. *IPTPS'08*. Berkeley, CA, USA: USENIX Association, 2008, pp. 2–2.
- [21] S. Tarkoma, C. Rothenberg, and E. Lagerspetz, "Theory and practice of bloom filters for distributed systems," *Communications Surveys Tutorials*, IEEE, vol. 14, no. 1, First 2012, pp. 131–155.
- [22] L. Ramaswamy and J. Chen, "The coquos approach to continuous queries in unstructured overlays," *IEEE Trans. on Knowl. and Data Eng.*, vol. 23, no. 3, Mar. 2011, pp. 463–478.
- [23] P. Leone and C. Muñoz, "Content based routing with directional random walk for failure tolerance and detection in cooperative large scale wireless networks," in *Proc. 2nd Intl. Workshop on Architecting Safety in Collaborative Mobile Systems, SAFECOMP, 2013*, pp. 313–324.
- [24] C. Muñoz and P. Leone, "Design of an unstructured and free geo-coordinates information brokerage system for sensor networks using directional random walks," in *Proc. 3rd Intl. Conference on Sensor Networks, SENSORNETS, 2014*, pp. 205–212.
- [25] D. Braginsky and D. Estrin, "Rumor routing algorithm for sensor networks," in *Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*, ser. *WSNA '02*. New York, NY, USA: ACM, 2002, pp. 22–31.