CoP3D: CONTEXT-AWARE OVERLAY TREE FOR CONTENT-BASED CONTROL SYSTEMS*

Mauro Caporuscio
Dipartimento di Elettronica de Informazione, Politecnico di Milano, Italy

Alfredo Navarra
Dipartimento di Matematica e Informatica, Università degli Studi di Perugia, Italy

Keywords: Content-based publish/Subscribe system, Sensor networks, Distributed computing, Overlay network.

Abstract: Publish/Subscribe systems are nowadays largely accepted for coordinating applications over wide-area networks, whereas they still suffer when applied to mobile environments. In fact, the typical tree overlay network implementation of the event service, which provides scalability in the context of wide-area networks, fails when a run-time reconfiguration is needed. This might happen when dealing with the dynamic and mobility characteristics inherent to the pervasive environments. In this proposal, we address such an issue by envisioning a context-aware approach in order to build and maintain the dispatching tree overlay with respect to the context sensed through the environment. The resulting distributed control system optimizes on the energy consumption aspect, which is the key cost measure for a mobile environment such as sensor networks.

1 INTRODUCTION

Continuous improvement of wireless network technologies and the miniaturization of electronic equipments are de facto enabling the exploitation of wireless sensor networks as distributed control systems for monitoring data within physical environments. In particular, a sensor network consists of a number of sensor nodes randomly disseminated within the environment. This requires for protocols and algorithms that let nodes to self-organize and cooperate with each other.

In this context, content-based publish/subscribe system is considered. Indeed, it is a well suited communication infrastructure for dealing with distributed applications over wireless networks. In fact, the loose coupling characteristic of publish/subscribe together with content-based routing protocol, where the message routing is determined by the interests of the receiver rather than by the explicit destination address, allow for dealing with the dynamism of such environments (Carzaniga and Wolf, 2001).

As depicted in Figure 1, a content-based publish/subscribe system is composed of two main entities: clients and servers. *Servers* are interconnected

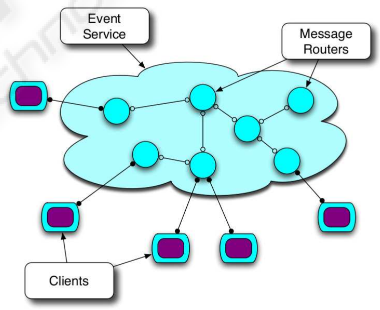

Figure 1: Tree overlay implementation of the event service.

in a distributed network, referred to as *event service*, and provide clients with *access points* offering an extended publish/subscribe interface. *Clients* are of two kinds: (*i*) *publishers* use the access points to *publish* events and (*ii*) *subscribers* use the access points to *subscribe* for events of interest by supplying a predicate, called *subscription*. A *subscription* is applied to the content of events and it allows *subscribers* to select the events they are interested in. The *event service* is responsible for selecting and delivering events of interest to subscribers via the access points. The idea of a publish/subscribe system is quite mature, and a fair number and variety of publish/subscribe systems have been proposed, including research prototypes (Carzaniga et al., 2001; Cugola et al., 2001;

*This research has been partially funded by the European Commission, Programme IDEAS-ERC, Project 227077-SMScom.

Caporuscio M. and Navarra A. (2009).
CoP3D: CONTEXT-AWARE OVERLAY TREE FOR CONTENT-BASED CONTROL SYSTEMS.
In *Proceedings of the 6th International Conference on Informatics in Control, Automation and Robotics - Intelligent Control Systems and Optimization*, pages 305-310
DOI: 10.5220/0002247703050310

Pietzuch and Bacon., 2002; Meier and Cahill, 2002), commercial products (TIBCO Inc., Palo Alto, CA, 1996), and attempts of standardization (Object Management Group (OMG), 2004; Sun Microsystems, Inc., Mountain View, California, 1999).

The typical implementation of the event service is a tree overlay network (as shown in Figure 1), which provides scalability in the context of wide-area networks (e.g., Siena (Carzaniga et al., 2001) and Elvin (Segall and Arnold., 1997)). In fact, having a unique path between any two nodes simplifies both the matching algorithm and the routing scheme, and avoids network flooding. Even though such an overlay topology has been successfully exploited in ad-hoc environments (see e.g., multicast tree (Royer and Perkins, 2000)) there are issues concerning its application to publish/subscribe systems in the context of mobile environments. In fact, the high degree of dynamics inherent to sensor networks requires to continuously reconfigure the event service and, consequently, to reconfigure data structures and routing tables. Such structures, distributed through the event service, are created by considering the interest expressed by the clients and evolve at runtime when subscriptions are added/removed/modified. This makes the reconfiguration process expensive (in terms of both message exchanges and computational operations) and hard to accomplish.

Many works have been published so far concerning publish/subscribe systems in mobile environments, such as (Huang and Garcia-Molina, 2004; Fiege et al., 2003). However, most of them focus on the clients mobility while the event service remains stable (Caporuscio et al., 2003). On the other hand, in (Huang and Garcia-Molina, 2003) and (Mottola et al., 2008), the authors concentrate on the event service topology reconfiguration in the field of ad-hoc networks.

This paper deals with publish/subscribe systems over sensor network by proposing a context-aware technique to build and maintain the dispatching tree overlay with respect to the context sensed through the environment. Our proposal is to exploit the Connectionless Probabilistic (CoP) protocol (McCann et al., 2005), originally designed for sensor networks, in the context of publish/subscribe systems. The CoP protocol assumes battery powered mobile devices that interact with each other in order to establish a virtual infrastructure for routing purposes. Namely, a virtual grid infrastructure is built by considering the physical position of the devices. Further, the protocol takes care of the energy consumption by benefiting from mobility and network dynamics.

We envision an extension of the CoP protocol in order to build a rooted tree on top of the virtual grid. This would realize a new protocol, called CoP3D, where the tree is used to accomplish publish/subscribe transmissions (i.e., the overlay network), whereas the underlying grid maintains its original aim. The main property of the proposed CoP3D protocol is that it builds the overlay network taking into account context information sensed through the environment. Namely, the devices representing the tree nodes will be selected according to their physical location, available energy, computational resources and expected mobility rate.

The paper is organized as follows. Section 2 describes the CoP protocol and the virtual grid construction. Section 3 presents CoP3D and describes how the tree overlay is built with respect to the sensed context. Section 4 points out interesting peculiarities of the considered context that must be carefully addressed. Section 5 discusses future work.

2 CoP: THE VIRTUAL GRID INFRASTRUCTURE

In this section, we outline the CoP protocol (McCann et al., 2005), which aims to manage communication for sensor networks. We have chosen this protocol since it efficiently performs in mobile ad-hoc environments and well fits our purposes.

We assume the devices can move over some given area A. From now on we refer to devices also by nodes. As it will be clear soon, our approach can be adapted to any kind of desired shape for A. It may be suitable for any building like offices, airports, companies and so forth where the network has to be deployed. For the sake of clarity, A is assumed to be simply a square area of side d. At the border of A there can be some fixed infrastructure, if required, that supports the survivability of the network. In the context of publish/subscribe systems, whenever a publisher wants to provide a service on the network, it has to inform the network about its publication. In the wired context this message was routed along a fixed tree responsible of the matching among publishers and subscribers. Since in our wireless and mobile context we cannot have such an infrastructure, the basic idea is to build a virtual one along which messages are routed.

Virtual infrastructure means that active nodes have to take care somehow of its feasibility and maintenance. The only thing that each node needs to know in order to participate in the network, by means of the CoP protocol, is its own physical position with respect to the area A. This can be achieved by means of either devices powered by GPS systems or some in-

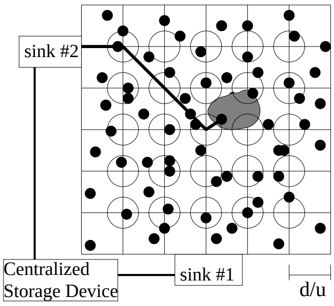

Figure 2: The area of interest covered by a virtual grid infrastructure of grid unit $\frac{d}{u}$. Nodes inside each circle of radius ds compete to represent the corresponding grid node.

stalled beacon external nodes that inform about relative positions (the mentioned infrastructure at the border of A), or other positioning techniques (see for instance, (Capkun et al., 2006; Caruso et al., 2005)). According to the expected population inside A, CoP builds a virtual infrastructure. Namely, a square grid G covering A is considered with grid unit d/u (see Figure 2), where u is determined according to the density of the nodes. From the "balls into bins" theory (see (Raab and Steger, 1998; McCann et al., 2005)), by choosing an appropriate distortion parameter ds, according to the density of the nodes in the area of interest, we can compute our desired value u in such a way that for each grid node there will be some device inside the circle of radius ds centered on it with high probability. We can choose for instance $u = \frac{d}{3ds}$. Note that u can change according to the established density of the network over the time. Its current value can be part of the beacon message sent by the external border nodes.[2]

Whenever a device in A starts its interaction with the network or just moves from its current position, it is aware about its location with respect to A and hence to G. This allows for determining whether it can be elected as representative of a grid node. For this step of the CoP virtual grid construction, we refer to standard leader election strategies suitably adopted in wireless environments like ad-hoc and sensors networks (see (Malpani et al., 2000)).

In doing so, at the expenses of some communication needed for the evaluation of a suitable value for u and the local leader elections, CoP builds a virtual grid G that can be used in order to correctly route desired communications. G covers all the given area A. One of the most important properties of such a construction is that it does not depend on the current nodes acting as leaders but only on the density of the nodes. The CoP protocol was originally applied in the field of sensor networks where sensed data needed to be route outside A to a fixed sink. Hence, the routing was performed in a multi-hop fashion over the grid. The main steps that must be performed by a sensor in order to participate to the CoP protocol can be summarized as follows.

1. It discovers its position according to the used technology, the area covered by the network and the size of the grid needed for the routing
2. If the received value for u is zero then it just waits
3. According to the previous information it evaluates its position with respect to the grid and decides whether it can represent any grid node or not. If not, it just assumes there is someone representing the closest grid node to its position and whenever needs to send a message it delegates such a node for the correct routing. If yes, it checks whether there is someone else playing or not that role
4. If it represents a grid node then it has to take care of the associated traffic that has to be route
5. One hop of the communication is made over the grid in such a way that the transmission can be received inside the whole circular area associated to the target grid nodes. This is due to the fact that a transmitting node does not know the exact position of its neighbors but it only assumes that someone is inside the circle of radius ds surrounding the target grid node
6. Whenever a device changes its position, i.e., either enters or exits a circular area, it has to change its role accordingly. If it was responsible for the routing, it has to take care to send the needed information to at least one of the current candidate to become its successor
7. If a device is running out of energy, it has to become a passive node even though still continues to receive its desired communications until it can.

Starting from this routing protocol and infrastructure, next section shows how this can be modified and suitably adapted in the case of mobile publish/subscriber systems.

3 CoP3D: THE TREE OVERLAY NETWORK

In this section, we describe our proposal concerning the management of a mobile publish/subscriber system. In particular, we exploit the CoP virtual grid construction in order to obtain a tree overlay implementing the event service. The main differences with re-

[2] $u = 0$ means there are not enough devices in the area to build the network.

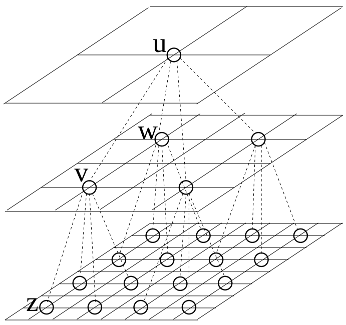

Figure 3: The tree construction based on the virtual grid infrastructure.

spect to a typical structured environment can be summarized as follows:

(a) *Energy Efficiency*: mobile devices have scarce power capacity and they can be required to stay alive for long periods without any support

(b) *Scalability*: it is desirable that any solution to manage such type of network would easily scale according to the number of devices and the corresponding area. Due to the mobility feature, manual deployment of devices is just not feasible

(c) *Fault-tolerance*: even though wired networks are also usually designed to cope with fault-tolerance issues, in a mobile environment faults are much more frequent. A device can easily disappear from the network due to either its movement or its low power level

(d) *Absence of the Infrastructure*: Due to their movements or to some handoff strategy, a given device is not always linked to the same set of neighbors at different time. This implies that assumptions about the topology of the network cannot be done unless they are straightforward from the area of interest (for instance networks along a street cannot differ too much from a line). The only fixed infrastructure that can be assumed can reside at the border of the area of interest

As shown in Figure 3, virtual grid nodes become the leaves of the tree overlay. Over the basic virtual grid infrastructure previously described, we build several levels of grids. What in the typical publish/subscribe system was the routing tree now is a 4-ary tree rooted at the center of the grid. Dividing the grid into four sub-grids we iteratively discover other 4 centers (one for each subgrid) to which the previous one is virtually connected. We iterate such a process $2\log\frac{1}{u}$ times in order to obtain a full coverage of the grid by means of a 4-ary tree of logarithmic height with respect to the number of grid nodes (see Figure 3). Starting from the bottom layer, once the basic virtual grid is built, all the nodes also know their location with respect to the overlay tree and hence, each one can play its corresponding role. By referring to Figure 3, for instance, node v knows that in the overlay tree it plays the role of parent for node z.

As it was for the virtual grid construction, many devices may cleverly compete in a leader election in order to become nodes of the tree overlay, hence responsible for routing publish/subscribe events and subscriptions. In particular, such an election is performed by considering context information. These include: (*i*) available energy and physical location to be compliant with the CoP protocol (see Section 2), (*ii*) mobility rate since the more a node is stable the more a node is qualified to play the router role, and (*iii*) computational resources since powerful nodes can better accomplish routing tasks.

In this way, the devices representing the virtual grid nodes of the 4-ary tree are responsible for the routing of the messages to realize the desired publish/subscribe system. Such nodes will be clearly more loaded than others and hence they will spend more energy. In order to distribute the energy consumption, changes to the grid structure can be done by means of shifting procedures or changing the grid unit u. It can also be implicitly obtained by means of a certain mobility ratio. Whenever a device changes its position, in fact, it has to evaluate if its role is also changing according to the distortion parameter ds. When a node moves out from the virtual circular area of radius ds defined around each grid node, it cannot be anymore representative of such a grid node hence it has to delegate someone else for doing its job. This means that mobility can play a central role in this routing process since it implies a more uniform distribution of the energy consumption. Once the tree overlay network is built, the event service can act as in the case of the typical structured environment. This allows us to apply such an approach to any publish/subscribe system that relies on a tree overlay network.

Concerning the properties outlined at the beginning of this section, we want to point out how our approach is suitable with respect to them. For (a) it is worth to note that the only resource of the devices is their energy. Moreover, both the virtual grid infrastructure and the tree overlay are heavily dependent on the devices survivability. On the other hand, this is the only resource that we can rely on in order to avoid the installation of typical structured environments. The tree overlay network construction is quite lightweight since it requires only some basic calculations. Furthermore the routing is well spread among all the devices and it is made in a multi-hop fashion. It takes into account both the maximum distance covered by means of a transmission (hence the maximum energy spent) and the time needed to deliver a message (that in the worst case requires only $4\log\frac{1}{u}$ steps). As last remark, our strategy makes use of the mobility rate as a mean to save energy since roles change accordingly.

Concerning (b) and (c), our approach is easily scalable. It is in its nature to adapt the network coverage based on the density of the devices with respect to the area that must be served. It also suitably copes with fault-tolerance issues. In fact, devices might continuously appear (disappear) in (from) the network without breaking virtual connections up to a certain threshold. Once reached such a threshold, communications among the devices become unfeasible.

The last property (d) is easily solved by means of the virtual grid infrastructure, hence avoiding the installation of both structured environments and more powerful devices.

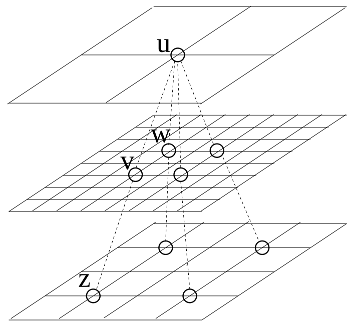

Figure 4: The tree construction based on the reversed virtual grid infrastructure. The aim is to better balance energy consumption among the devices.

4 FURTHER KEY-POINTS

We now focus the attention on three main properties that must be carefully addressed in our approach.

(*i*) As it was for structured event services, it is evident that starting from the root and descending until the leaves, nodes are differently loaded in terms of routing messages. As already outlined, one mean to avoid this unbalanced energy consumption might depend on the mobility rate or the shifting procedure. On the other hand it is worth noting that our construction can be easily modified in the case we prefer to not deploy the root at the center of the grid since we may have information about the distribution of the devices. The previous discussion is in fact based on the assumption of a uniform distribution of the devices inside A. Consequently, it is easy to understand that the area of interest can be subdivided into several areas according to the expected density, hence applying our methodology to construct different subtree overlays.

(*ii*) Another very important issue that must be addressed before practically apply our method concerns the maximum distance that devices' transmissions can cover. With our construction, in fact, not only the root is the most loaded node in terms of communications but also it is required to perform the most distant transmissions overall the network. In order to avoid such a situation and to better balance the energy consumption among the nodes of the network, we may think about a reversed structure. This implies that the more a node is far from the root (with respect to the tree overlay), the more must be its transmission distance. In Figure 4 it is shown how the tree overlay may become. In any case, when a device belongs to a layer in which its power cannot cope with the required transmission distance, it does not mean that it cannot perform the required communications. In the header of the messages can be added information in such a way that a long transmission, in terms of covered distance, can be emulated by multi-hop transmissions of shorter range. This recalls what was used in the field of ATM networks by the concepts of Virtual Channel and Virtual Path communications (see e.g., (Flammini and Navarra, 2009)). It is not in the aim of this paper to go into details of such discussion but we want only point out that the underlying virtual grid structure at each layer can be used for this purpose.

(*iii*) Finally, CoP was based on another assumption for which some communications can be wasted due to the absence of a sensor representative of some virtual grid node. According to the parameter ds, such an event can happen rarely but still there is the possibility. Since in our context the tree overlay must be somehow guaranteed until it is feasible, we need some strategy to cope with this possibility. Due to the underlying virtual grid structure, in fact, the transmissions directed to a missing node can be easily rerouted to some other grid path. For this issue we can refer to (Mostarda and Navarra, 2008) where CoP was modified in order to address security issues. Clearly, at the expenses of a higher energy consumption, such situations can be addressed.

5 FUTURE WORK

In this paper, we have envisioned a new approach to achieve content-based publish/subscribe control systems in the context of sensor networks. In particular, we have proposed a context-aware technique to build and maintain the tree overlay network for communication purposes. Specifically, our approach builds on the CoP protocol and takes into account context information sensed through the environment. In fact, the devices representing the tree nodes are chosen by considering their physical location, available energy, computational resources and expected mobility rate. Once the tree overlay network is built, the publish/subscribe system can act as in the case of the typical structured environment. This allows to apply such an approach to any publish/subscribe system that relies on a tree overlay network. It is of main interest to study both the theoretical feasibility and the effectiveness of our proposal. Moreover, our methodology might be also applied to other distributed systems rather than just the publish/subscribe ones (e.g., peer-to-peer) and to build other topology overlay networks.

REFERENCES

Capkun, S., Cagalj, M., and Srivastava, M. (2006). Secure localization with hidden and mobile base stations. In *Proc. of the 25th Annual Joint Conf. of the IEEE Computer and Communications Societies (INFOCOM)*. IEEE Computer Society.

Caporuscio, M., Carzaniga, A., and Wolf, A. L. (2003). Design and evaluation of a support service for mobile, wireless publish/subscribe applications. *IEEE Transactions on Software Engineering*, 29(12):1059–1071.

Caruso, A., Chessa, S., De, S., and Urpi, A. (2005). GPS free coordinate assignment and routing in wireless sensor networks. In *Proc. of the 24th Annual Joint Conf. of the IEEE Computer and Communications Societies (INFOCOM)*, pages 150–160. IEEE Computer Society.

Carzaniga, A., Rosenblum, D. S., and Wolf, A. L. (2001). Design and evaluation of a wide-area event notification service. *ACM Transactions on Computer Systems*, 19(3):332–383.

Carzaniga, A. and Wolf, A. L. (2001). Content-based networking: A new communication infrastructure. In *NSF Workshop on an Infrastructure for Mobile and Wireless Systems*, number 2538 in Lecture Notes in Computer Science, pages 59–68, Scottsdale, Arizona. Springer-Verlag.

Cugola, G., Nitto, E. D., and Fuggetta, A. (2001). The JEDI event-based infrastructure and its application to the development of the OPSS WFMS. *IEEE Transactions on Software Engineering*, 27(9):827–850.

Fiege, L., Grtner, F. C., Kasten, O., and Zeidler, A. (2003). Supporting mobility in content-based publish/subscribe middleware. In *Lecture Notes in Computer Science*, volume 2672, pages 103 – 122.

Flammini, M. and Navarra, A. (2009). Layouts for mobility management in wireless atm networks. *Discrete Applied Mathematics (DAM)*, 157(1):98–111.

Huang, Y. and Garcia-Molina, H. (2003). Publish/subscribe tree construction in wireless ad-hoc networks. In *Proc. of the 4th Int. Conf. on Mobile Data Management (MDM)*, pages 122–140.

Huang, Y. and Garcia-Molina, H. (2004). Publish/subscribe in a mobile environment. *Wireless Networks*, 10(6):643–652.

Malpani, N., Welch, J. L., and Vaidya, N. H. (2000). Leader election algorithms for mobile ad hoc networks. In *Proc. of ACM Joint Work. on Foundations of Mobile Computing (DIALM-POMC)*.

McCann, J. A., Navarra, A., and Papadopoulos, A. A. (2005). Connectionless Probabilistic (CoP) routing: an efficient protocol for Mobile Wireless Ad-Hoc Sensor Networks. In *Proc. of the 24th IEEE Int. Performance Computing and Communications Conf. (IPCCC)*, pages 73–77.

Meier, R. and Cahill, V. (2002). Steam: Event-based middleware for wireless ad hoc networks. In *Proc. of the 1st Int. Work. on Distributed Event-Based Systems (DEBS)*, Vienna, Austria.

Mostarda, L. and Navarra, A. (2008). Distributed idss for enhancing security in mobile wireless sensor networks. *Int. Journal of Distributed Sensor Networks (IJDSN)*, 4(2):83–109.

Mottola, L., Cugola, G., and Picco, G. P. (2008). A self-reparing tree topology enabling content-based routing in mobile ad hoc networks. *Transaction on Mobile Computing*, 7(8):946–960.

Object Management Group (OMG) (2004). Notification service version 1.1.

Pietzuch, P. R. and Bacon., J. (2002). Hermes: A distributed event-based middleware architecture. In *Proc. of the 1st Int. Work. on Distributed Event-Based Systems (DEBS)*, pages 611–618, Vienna, Austria.

Raab, M. and Steger, A. (1998). “Balls into bins” - a simple and tight analysis. In *Proc. of the 2nd Int. Work. on Randomization and Approximation Techniques in Computer Science (RANDOM)*, pages 159–170.

Royer, E. and Perkins, C. (2000). Multicast Ad hoc On-Demand Distance Vector (MAODV) Routing. IETF, Internet Draft: draft-ietf-manet-maodv-00.txt.

Segall, B. and Arnold., D. (1997). Elvin has left the building: A publish/subscribe notification service with quenching. In *Proc. of the Australian UNIX User Group*, pages 243–255, Brisbane, Australia.

Sun Microsystems, Inc., Mountain View, California (1999). Java message service.

TIBCO Inc., Palo Alto, CA (1996). TIBR+: a WAN router for global data distribution.