A topic-based, publish-subscribe architecture for intermittently connected 802.15.4 networks

Abstract:
The small size and power consumption of IEEE 802.15.4 devices allows embedding them in GSM/UMTS U-SIM cards and/or SD cards. The availability of such a near field technology for data exchange within mobile phones is very useful to complement GSM/UMTS services, providing proximity services, such as chat and advertisements in a commercial center, configuration data, micro-payments, access control. However, appetite comes with eating and once that we have the availability of a free communication radio link, we can enlarge the assortment of offered services, supporting not only "direct" data exchanges between two users within the 802.15.4 connectivity range but also communication among intermittently connected users. This is a typical scenario of so-called Delay Tolerant Networks (DTN) and we argue that a communication paradigm well suited to this environment is publish-subscribe. On the other hand, publish-subscribe is also well suited to satisfy the requirements of a community of users such as the one of an university campus, accompanying other services such as voice and Internet access. The aim of this paper is to present a topic-based, publish-subscribe architecture for intermittently connected networks exploiting IEEE 802.15.4 devices, and taking into due account the severe constraints deriving from their physical characteristics. We describe the architectural model, the protocol design, the system analysis, and the implementation of our solution in a real test-bed, which we carried out in cooperation with Telecom Italia, that financed this work.

Keywords: Topic Based, Publish/Subscribe, Delay Tolerant Networks, Spray and Wait, IEEE 802.15.4

1. Introduction

Location-based services and proximity services are growing in numbers and importance, extending the range of possible application scenarios, see e.g. the diffusion of locality-based social networks such as "Facebook places" or "Foursquare". The importance of location and proximity information is not limited to specific context-aware, location-dependent applications used e.g. "to (a) adapt interfaces, (b) tailor the set of application-relevant data, (c) increase the precision of information retrieval, (d) discover services, (e) make the user interaction implicit, or (f) build smart environments"[1]. As a matter of fact, location and proximity information could be used also to ease or improve communication services among persons living or working in the same geographical area. Next-door neighbors, students following the same class, workers of the same company are likely to have common interests and need to share information. True they could use e-mail or any of the hundreds of Internet services available nowadays; however, in frequent cases there is the need of communicating with a neighbor exactly because she is a neighbor. For instance, think to a student attending a class, she would like to ask if someone has a copy of a given lesson note or if there is someone that can give her a lift to the University canteen which is far away from the classroom or if someone wants to join her to a social event. She could shout her question in the classroom, which is maybe not a good idea, or send an e-mail to all her colleagues, which would bother many, or use a web service, which should know the position of all her colleagues to work properly, and which is complex to implement and raises privacy concerns. It would be much easier if she could send a message to all students currently in that classroom or in its proximity. If this message is delivered by a near field technology, for free, so much the better. Along this line we can think to several application scenarios that would benefit of a tool allowing to reach users located in a specific geographical area, eventually for free. This we can do by using the short-range radio technology known as IEEE 802.15.4, supporting not only "direct" data exchanges between two users within the 802.15.4 connectivity range but also communication among intermittently connected users. In other words, we want to offer communication services to a community of users that are either directly connected to each other or that can reach the intended receiver by means of other users, or as they move and eventually get in direct contact with the target user. By the way, this approach allows implementing several location-based services without necessarily having to resort to external localization infrastructures such as GPS, which, in addition, can not be used indoor. We choose IEEE 802.15.4 as near field technology, for two reasons: i) its low energy consumption [4]: only 24.7 mA of current consumption during the transmission against 57 mA of Bluetooth and 219 mA of WiFi [2]; ii) the small size of IEEE 802.15.4 devices, needing an area of $7mm^2$ or less. These characteristics allowed Telecom Italia, the major Italian operator and a member of ZigBee Alliance, to integrate IEEE 802.15.4 devices within smart phones and mobile terminals. Telecom Italia envisages in [3] a scenario comprising mobile phones equipped with either the so-called ZSIMs, which are GSM/UMTS SIMs integrating a ZigBee node, or the so called ZSDs, which are micro or mini SD cards integrating a ZigBee node. In a previous work [6], funded by Telecom Italia as well, we added to ZSIMs and ZSDs the IPv6 protocol; this we did by exploiting the 6LoWPAN adaptation layer [7]. Such solution, comprising IEEE 802.15.4 plus 6LoWPAN plus IPv6 allows easy communications among mobile terminals within IEEE 802.15.4 connectivity range, and easily support applications designed to work over the TCP/IP suite. In this paper, we leverage on this previous work and we design a system, following Delay Tolerant Networks (DTN) principles, to support communications not only between users that are directly connected to each other (i.e. in IEEE 802.15.4. range) but also among intermittently connected users. In fact, given the mobility of nodes and the short radio coverage, most of the time a complete path between a source and a destination does not exist. For this reason we resort to DTN routing rather than to a MANET one. On top of the DTN we deploy a topic-based, publish-subscribe [9] communication paradigm, which we think is very well suited to this environment. The publish/subscribe interaction scheme provide a loosely coupled form of interaction. "Subscribers have the ability to express their interest in an event, or a pattern of events, and are subsequently notified of any event, generated by a publisher, which matches their registered interest. An event is asynchronously propagated to all subscribers that registered interest in that given event."[9] Publishers do not have knowledge of what, if any, subscribers there may be. Pub/sub schemes can differ in the ways of specifying the events of interest. The two most widely used schemes are topic-based and content-based. Participants of a topic-based publish/subscribe system can publish messages and subscribe to individual topics, which are identified by keywords. A message published on a given topic will be delivered to the subscribers of that topic. We argue that publish-subscribe is also well suited to satisfy the requirements of a community of users such as the one of an university campus; in this environment, a publish-subscribe service can complement other, more classical, communication services such as voice and Internet access. The result of this construction is our Campus++ system: a university campus service framework where students and university personnel, equipped with mobile terminals containing ZSIMs or ZSDs, can share and retrieve data exploiting the IEEE 802.15.4 technology by using a topic-based, publish-subscribe service model.

The paper is organized as follows. Section 2. describes the service and network framework, and presents some issues that drove the system design. Section 3. reports the main characteristics of the Campus++ system, and a system analysis, performed by means of simulations. In Section 4. we describe the implementation of our system in a real test-bed, carried out in cooperation with Telecom Italia. Finally, section 5. concludes the paper.

2. System framework and key issues

2.1 Service framework

The Campus++ system enables a community of users equipped with IEEE 802.15.4 devices (e.g., mini SD, ZSIMs) to exchange messages regarding specific "Topics". We call these messages data-samples: users that send data-samples regarding a topic are publishers of such topic; users interested in receiving data-samples of a topic are subscribers of that topic. Users exploit Campus++ services by running Campus++ software on their PDAs or mobile phones. The Campus++ PDA application allows visualizing the list of topics in the systems and the subscribed-to topics and publishing data-samples.

2.2 Networking framework

The delivery of data-samples toward subscribers is carried out by a Delay Tolerant Network (DTN) formed by 802.15.4 devices either belonging to users and infrastructure nodes. We choose "Spray and Wait"[8] as routing paradigm because of its simplicity, which suits the characteristics of IEEE 802.15.4 devices, and because of its native support for point to multi-point communication. Spray and Wait works as follows: when a new data-sample is published it is replicated on L different nodes of the network (including the source). Subscribers can retrieve the data-sample by coming in direct radio contact with one of these L nodes. Thus, the more the replicas in the system, the smaller the average time between the publication of a data-sample and its reception by subscribers (=mean delivery delay). We point out that "Spray and Wait" does not define a fully-fledged protocol but only a routing scheme. Thus, we had to design a publish-subscribe receiver-driven protocol that implements this routing scheme and whose details will be provided in section 3.2

2.3 Key issues

During the design of the Campus++ system architecture, some important issues arose. In the following we present these issues, which are typical of any topic-based publish-subscribe system supported by a DTN, with the exception of the last one.

- `Replication control`: several replicas of different data-samples have to share a limited distributed storage space. For this reason, the number of replicas must be evaluated as a trade-off between mean delivery delay and used storage space. This trade off must take into account not only the memory constraints but also the different popularity of different topics (i.e. the different number of subscribers per topic).

- `Data obsolescence`: Since memory is limited, oldest data-samples have to be removed from the system when new data-samples are published on the same topic. This means that we need a way to distinguish newer data-samples from older ones.

- `Distribution of control-data`: control-data are control information needed by Campus++ nodes for supporting the information distribution system. For instance, control-data include the list of topics supported by the system, the current number of subscribers per topic, the memory space available in the system, the current number of users, etc..

- `Ad-hoc mode`: to support the DTN functionality, we need a way to establish connections between intermittently connected nodes. Unfortunately the IEEE 802.15.4 Standard [4] does not provide any kind of ad-hoc mode. In particular, IEEE 802.15.4 defines the role of PAN coordinator as the principal controller of a personal area network (PAN) and states that there must be exactly one PAN coordinator in each PAN. However, such per-PAN unique role of the PAN coordinator, notwithstanding the fact that nodes are mobile, can be a problem.

3. The Campus++ System

3.1 Architecture

Figure 1a shows the Campus++ system architecture. The architecture is formed by 802.15.4 mobile user devices, one or more 802.15.4 infrastructure nodes called "way-servers", and an administrative server.

Way-servers are placed in strategic locations crossed by users when they enter the service area (e.g. campus entry points/doors/gates). The role of a way-server is: i) to be the source (publisher) of control-data; ii) to provide mobile terminals with a (loose) system clock reference; iii) to inform the administrative server about the current status of the systems parameters, such as the number of users, the number of subscribers per-topic, etc, by using a backbone communication system (e.g. ethernet, UMTS, etc.).

The administrative server provides the way-servers with the set of control-data (e.g., list of topics, overall system memory, number of subscriber per-topic, clock reference, etc.) that they have to distribute in the system. In this way, all way-servers transmit the same control-data.

To distribute control-data we use a special topic called "built-in topic" whose publishers are only the way-servers and whose subscribers are all user nodes. This technique is inherited from the DDS publish/subscribe architecture [11] and exploit the same publish-subscribe communication approach for diffusing both data and data control information.

3.2 Communication protocols

The distribution of the data-samples involves the MAC (802.15.4), the 6LoWPAN layer and DTN (Spray and Wait) functionality, which we describe in the following.

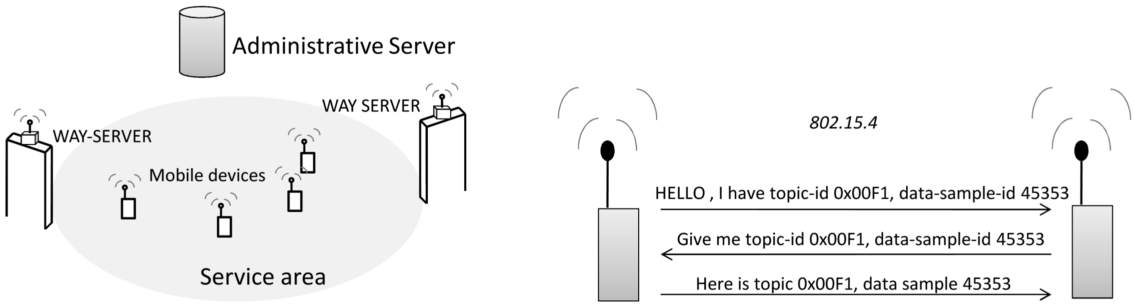

(a) System architecture (b) Receiver-driven approach

Figure 1

3.2.1 MAC and 6LowPAN

To support direct communications we use MAC 802.15.4 / 6LowPAN functionality. Nevertheless, to cope with the "ad-hoc" issue mentioned in section 2.3, we operate as follows. Given that security support is not considered in our scenario, and that we use the unique 64bit extended addressing mode (i.e. MAC addresses), we do not find any reasons to follow the rule of having a unique PAN coordinator. We force all nodes of the system to be PAN Coordinators. This could lead to cases in which there are more that one PAN Coordinator inside the same PAN, and this is a violation of the 802.15.4 standard. However, on the one hand we have verified that our hardware (Texas CC2430) allows this mode of operation; on the other hand we think that this can be accepted as an experimental way to use the 802.15.4 technology, and eventually start a process toward the modification of the standard to support an ad-hocmode in 802.15.4.

3.2.2 DTN

As regards DTN, we devised a receiver-driven Spray and Wait protocol. A node pulls data-samples when it either wants to retrieve subscribed-to data-samples or when it wants to replicate data-samples. To this end, nodes periodically emit Hello messages, which advertise the stored data-samples and if some of them need to be replicated [8].

The above idea is illustrated in figure 1b, which refers to the data retrieval operation. Figure 1b shows a node that advertises its stored data-samples to another node by means of a Hello message. The node on the right-hand side recognizes a new data-sample inside the ones listed in the Hello message and requests it. Data are requested by using the tuple {topic-id (0x00F1), data-sample-id (45353)}, where topic-id is an identifier of a specific topic and data-sample-id is an identifier of the data-sample. The association between topic-id and topic name is communicated to nodes by using the built-in topic.

To support data obsolescence, as discussed in 2.3, we (loosely) synchronize all devices with the clock reference provided by way-servers and use the the publishing time to set the data-sample-id (data-sample-id=publishing time). A node that is replicating a data-sample marks it as "removable" if the node sees a neighbor node with a newer data-sample of the same topic. Removable data-samples will be effectively removed from nodes' memory in case of new data-samples replication, so that the system memory is efficiently used. When the memory is full, only the newest data-samples will remain in the system.

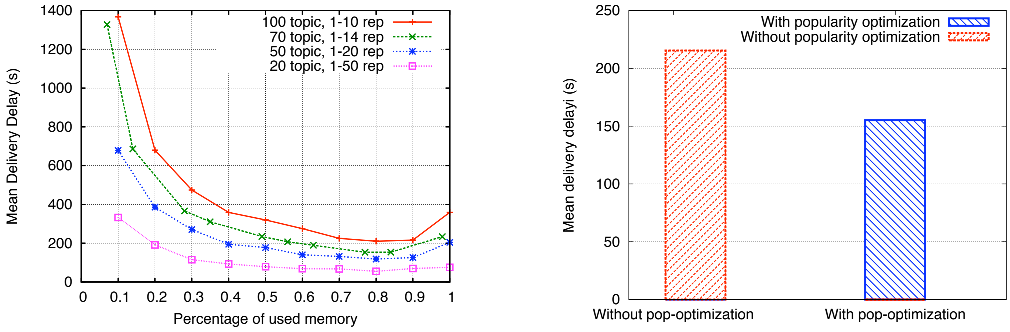

(a) Mean delivery delay as a function of the percentage of used memory

(b) Mean delivery delay with/without popularity optimization

Figure 2

3.3 Replication Control

The number of replicas for each data-sample must be chosen taking not account memory constraints, and minimizing the mean delivery delay. This is not an easy task. Memory constraints may show up either because a user allocates a limited amount of her device memory for DTN functionality or, as in our case, because the DTN functionality is implemented in embedded wireless chips, which have a small memory. For instance, the Texas CC2430 [10] system-on-chip provides 8kB of RAM memory that has to be shared among Operating System, 802.15.4 MAC, 6LoWPAN, and Campus++ functionality. All this implies that about 10 data-samples of 100 bytes can be stored on each node. Given these memory constraints, we determine how many replicas of each data-sample we should have in the DTN to minimize the *mean delivery delay*, defined as the mean interval between the instant in which a generic data-sample is published and the instant in which when it is retrieved.

The answer to this question could be very complex, given the different parameters that affect this optimization, such as the mobility model of nodes, the number of data sample that a node can store, and the topics popularity (i.e. the number of subscribers for each topic). For this reason we proceed by steps: i) first we find out which is the optimal number of replicas (workload) if all the topics have the same popularity; ii) then we consider the different popularity of topics. We use a simulative approach based on an event driven simulator written from scratches whose code is available in [12].

We simulate the behavior of 100 nodes that move according to a random waypoint (RWP) mobility model [13] on a 500x500 m^2 surface. Node speed and pause periods are constant and equal to 1 m/s and 5 s, respectively. The radio coverage range is set to 50m; the nodes' memory capacity is 10 data-samples; the simulation duration is set to 500000 seconds.

Figure 2a shows the mean delivery delay s a function of the percentage of used memory (or workload) for different numbers of topic in the system (100, 70, 50, 20). All the topics have the same popularity. Surprisingly enough, performance worsen as the workload goes above 80%. The reason of this behavior is due to the Spray Time. When we put too much replicas in the system, we decrease the probability for a spraying node to find other nodes with at least one free place in their memory to which pass half of its replica. For this reason, the spray time increases, and increases as well also the probability that subscribers retrieve data-samples when the spray operation is still in progress. In this way, subscribers see a smaller number of replicas and this leads to an higher mean delivery delay, despite the increment of the number of replicas in the system.

However, in the real-world topics have not the same popularity, typically there are very popular topics and un-popular topics. For this reason, it could be worthwhile to perform an optimization having the aim of advantaging popular topics with respect to unpopular one by giving to the former topics a bigger number of replicas with respect to un-popular topics. We can make subscribers of popular topics perceive a smaller delay, and subscribers of unpopular topic perceive a greater delay, so that the mean delivery delay is smaller that in the case in which all topics perceive the same delay. To evaluate the number of replicas following this idea, we start from the model in[8] and, using the same assumptions made in that work, we consider only the wait phase and an exponentially distributed inter-meeting time between nodes. Under these hypotheses the mean delivery delay grows as E_{mm}/L where L is the number of replicas and E_{mm} is the mean time required by two nodes to meet each other. Then we exploit a Lagrange optimization to derive the following "handy" formula, which gives the optimal number of replicas (L_i) for data-samples of the i-th topic, given the number of subscribers S_i of each of the T topics in the system and the available system memory C_{tot}.

$$L_i = C_{tot}\frac{\sqrt{S_i}}{\sqrt{S_1} + ...\sqrt{S_T}} \quad (1)$$

We experimented the effectiveness of this formula in the the same simulated scenario described above, for a workload of 80%. For this experiment, we considered 100 topics, and varied their popularity according to a Zipf distribution with parameter $\alpha = 0.8$. Figure 2b shows the performance improvements deriving from this optimization compared to the case of a "flat" distribution of replicas per topic.

Summing up, we found that to obtain a good trade-off between complexity of the optimization and performance, we must: (i) use up to 80% of the whole system memory; (ii) choose the number of replicas per data-sample according to the popularity of the related topic (see Equation 1).

4. Implementation of the system

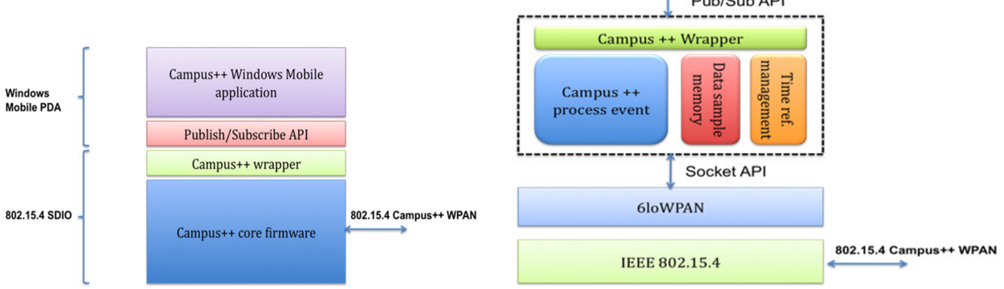

Figure 3: (a) Overview of Campus++ PDA protocol stack, (b) Campus++ CC2430 firmware architecture

We implemented the Campus++ system by using the TI CC2430 firmware [10], deployed in different devices: way-servers, user PDAs with SDIO/ZSD and mobile phones with ZSIM. Campus++ functionality are exploited by a Windows Mobile application, which interacts with the CC2430 by means of a publish/subscribe API (see Figure 3a).

Figure 3b shows the firmware architecture of a Campus++ device. Campus++ software exploits the Texas Instrument 802.15.4 MAC (TI-MAC) and our 6LoWPAN layer [6], and it is composed by three modules: i) Campus++ process event, which implements the state machine of the DTN; ii) Data Sample Memory, and iii) Time Reference Management, which is responsible for the time reference synchronization, necessary for handling data obsolescence. The total code-size of the firmware is about 60 KB.

5. Conclusions

In this paper we showed how IEEE 802.15.4, can be exploited by mobile terminals to deliver location-based services. We provide two main contributions: system design and performance optimization.

As regards system design, we showed that the coupling of DTN routing and topic-based publish-subscribe is a nice solution for the requirements stated in the Introduction. We also assessed the effectiveness of our proposed Campus++ by means of a real test-bed.

As regards performance optimization, we derived a trade-off procedure to dimension the number of replicas in the system. We showed that the system memory must not be completely filled up by replicas and we provided a relation between topics popularity and number of replicas of data-samples.

References

[1] Schmidt, A.; Aidoo, K.A.; Takaluoma, A.; Tuomela, U.; Van Laerhoven, K; Van de Velde W. *"Advanced Interaction in Context"*, 1st Int. Symp. on Handheld and Ubiquitous Computing (HUC99), Springer LNCS, Vol. 1707. pp. 89101.

[2] J.S. Lee, Y.W. Su, C.C. Shen, *A comparative Study of Wireless Protocols: Bluetooth, UWB, ZigBee and Wi-Fi*, Proc. IECON, 2007

[3] M. Turolla, E. Alessio, *ZSIM enabling innovative services to improve quality of life*, white paper on line available at www.zigbee.org/imwp/download.asp?ContentID=10403

[4] IEEE Standard 802.15.4-2006, *Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs)*, 2006.

[5] ZigBee Alliance, *ZigBee Specification*, October 2007.

[6] authors hidden for reviewing purpose, *Towards fully self-configuring and IP-enabled IEEE 802.15.4 LR-WPANs*, in the proceeding of Med Hoc Net 2010

[7] RFC 4944, *Transmission of IPv6 Packets over IEEE 802.15.4 Networks.*

[8] T. Spyropoulos , K. Psounis , C. S. Raghavendra, *Efficient routing in intermittently connected mobile networks: the multiple-copy case*, IEEE/ACM Transactions on Networking (TON), v.16 n.1, p.77-90, February 2008

[9] P. T. Eugster, P. A. Felber, R. Guerraoui, A. Kermarrec, *The Many Faces of Publish/Subscribe*, in ACM Computing Surveys (CSUR), Volume 35, Issue 2, 2003

[10] CC2430 web site, http://focus.ti.com/docs/prod/folders/print/cc2430.html

[11] Data Distribution Service (DDS) web site, http://portals.omg.org/dds

[12] http://stud.netgroup.uniroma2.it/ lorenzo/campus++/simulator.tar.bz2

[13] J. Broch, et al., *A performance comparison of multi-hop wireless ad hoc network routing protocols*, in Mobile Computing and Networking, pages 8597, 1998.