Semana de Engenharia 2010
Guimarães, 11 a 15 de Outubro

EFFICIENT AGGREGATE COMPUTATIONS IN LARGE-SCALE DENSE WSN

Nuno Pereira

Abstract— **This work addresses scenarios where even a small area may contain several tens of nodes, in this context, we addresses the problem obtaining aggregate quantities (e.g., the minimum, maximum or median of the values proposed by all nodes) with a time-complexity that is independent of the number of nodes, or grows very slowly as the number of nodes increases. This is achieved by co-designing the distributed algorithms for obtaining aggregate quantities and the underlying communication system. In this paper, we summarize the results of the research developed around this problem. These results are significant because often networks of nodes that take sensor readings are designed to be large scale, dense networks and it is exactly for such scenarios that the proposed distributed algorithms for obtaining aggregate quantities excel. The implementation and test of these distributed algorithms in a hardware platform developed has shown that aggregate quantities in large-scale, dense wireless sensor systems can be obtained efficientlly.**

Keywords—**Medium Access Control (MAC), Data Processing, Wireless Sensor Networks, Cyber-Physical Systems, Data aggregation.**

Motivation and Objectives

Microprocessors are everywhere. Nowadays, we can find computing capabilities in everyday physical objects as diverse as mobile phones, digital personal assistants, gaming platforms, household appliances or cars, just to name a few examples.

Computing-enabled physical objects often have to deal with physical processes and tightly integrate computing with the physical world via sensors and actuators. The integration of physical processes and computing is not a new problem. Embedded systems, which have been in place long ago, often combine physical processes with computing. However, with the massive deployment of networked embedded computing devices, we are observing the next step in the evolution of embedded computing. The term Cyber-Physical Systems (CPS) has been used to describe these pervasive computing systems, where emphasis is put on the physical, real-time and embedded aspects [1].

Such large-scale, sensor-rich networked systems will generate an enormous amount of sensor data [2], and handling such amounts of data introduces significant challenges. One approach to deal with the amount of data generated in these systems is to perform in-network data aggregation. Instead of collecting data from all nodes to a central point, in-network data aggregation applies a data-reduction function to the data traveling through the network such that the total number of messages transmitted is reduced [3].

Despite the previous research developed in the field of data aggregation, its performance is limited by the fact that, nodes in the same radio broadcast range cannot transmit in parallel, hence the time-complexity still depends on the number of sensor nodes. If we envision scenarios where even a small area may contain several tens of sensor nodes, the advantages of typical data aggregation solutions found in the literature are significantly impaired.

It is in this context that the research work described addresses the problem of performing scalable and efficient aggregate quantities (e.g., the minimum, maximum or median of the values proposed by all nodes) in dense networks. Here “efficient” means that the desired computation should be performed while consuming very little resources (such as energy, communication links, memory and processor) and “scalable” means that the consumption of resources should increases slowly or not at all as the number of sensor readings to be processed and/or the number of embedded computer nodes increases.

To illustrate this concept, consider a large-scale dense networked sensor system, whose nodes have a common sensing goal to measure temperature. Now consider the problem of computing a simple aggregate quantity: the minimum (MIN) sensed temperature among the nodes at some given moment. Computing MIN seems trivial, but for dense and large-scale systems, it poses an important problem: communicating sensor data individually makes the time-complexity of computing MIN a function of the number of nodes. This is true for any data aggregation mechanism employed.

This research work aims at being able to validate and explore the following hypothesis:

Is it possible to efficiently obtain aggregate quantities with a time-complexity that is independent of the number of sensor nodes?

In other words, and taking the example of MIN, we aim at computing MIN with a time-complexity that is equivalent to the time of transmitting a single message, even if tens or thousands of nodes share the same radio broadcast range.

Obtaining scalable and efficient aggregate quantities in large-scale dense networked sensor systems requires tight integration between the data aggregation techniques and communication mechanisms. This is a key observation underlying this research work, where

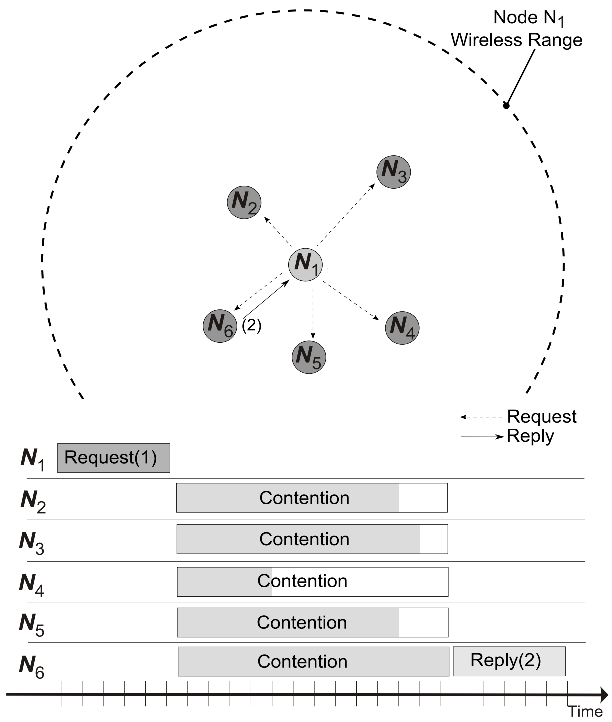

Fig. 1. Data Aggregation by Exploiting a Prioritized MAC.

the approach to obtain scalable and efficient aggregate quantities in large-scale dense networked sensor systems is co-designing (i) distributed algorithms to obtain aggregate quantities and (ii) the underlying communication services.

The main objective of this thesis is to demonstrate that is possible to obtain aggregate quantities efficiently by co-designing distributed algorithms for data aggregation with the underlying communication services. The approach to achieve this includes developing a prioritized MAC protocol and design distributed algorithms that exploit this MAC protocol. The next subsection presents a summary of the contributions made by this research work.

Research Approach

This research work explores mechanisms for obtaining aggregate quantities that are efficient, even in very dense networks. The efficiency of traditional data aggregation mechanisms results from applying data reduction functions to data coming from different sources, and from exploiting the opportunities for parallel transmissions. In the extreme case where all nodes are in the same broadcast domain, nodes cannot transmit in parallel and there are no opportunities for traditional data aggregation techniques to apply a data reduction function.

The novel approach explored in this thesis is based on the adaptation to wireless media of a family of medium access control (MAC) protocols. This family of protocols is known as dominance or binary countdown protocols [4] and is already present in wired networking solutions: the Controller Area Network (CAN) technology [5]. We then design distributed algorithms that exploit the MAC protocol to efficiently obtain aggregate quantities.

Dominance/binary countdown protocols can be exploited to efficiently obtain a range of aggregate quantities. Let us briefly exemplify, to give further intuition, the case of MIN, which can be obtained with a time-complexity that is equivalent to the time of transmitting a single message. This is illustrated in Figure 1, where all nodes are in the same broadcast domain. Suppose that the temperature values are coded as n-bit integers. Starting with the most significant bit first, let each node send the temperature reading bit by bit. Consider also that, for each transmitted bit, nodes read the resulting value in the channel (something straightforward in a wired medium) and the channel implements a logical AND of the transmitted bits. Furthermore, if a node reads '0' and is transmitting a '1', it stops transmitting. Then, at the end of the transmission of the n bits, the "observed" value in the channel will correspond to the MIN. It is as if all temperature readings were transmitted in parallel at the same time, and the resulting value of this non-destructive collision is a useful aggregate quantity.

It is based on this concept that the novel distributed algorithms developed in this work are designed upon. To accomplish this proposal, first it is necessary to design a MAC protocol that implements dominance/binary countdown in wireless environments, and then develop the algorithms to exploit that MAC protocol. However, designing and implementing such MAC protocol for wireless media is not trivial. For this reason, a substantial part of this work is focused on the development of such MAC protocol.

First, the problem is tackled assuming that all nodes belong to a *single broadcast domain* (SBD). Nodes are in a SBD when (i) a wireless broadcast made by one node reaches all other nodes in the same broadcast domain and (ii) if a node transmits a packet, then it can be correctly received by another node in the same broadcast domain only if the transmission of the packet does not overlap in time with another packet transmission. Note that, in this research work, there is no assumption about regular propagation patterns of the radio signals.

Achieving dominance in the wireless domain is challenging. To begin with, it is not possible to directly translate the behavior of wired protocols, as these require that nodes are able to transmit and receive at the same time. This is not possible in common radio transceivers, because the transmitted energy is much higher than the received energy. For this reason, dominance in wireless systems was achieved using a simple principle: when the transmitted bit is dominant, a pulse of a carrier wave is transmitted and there is no need to sense the medium. Conversely, when the bit to transmit is recessive, nothing has to be effectively sent, instead only the medium state has to be sensed.

Although the concept and approach sounds simple, a number of difficulties must be solved when proposing the design of a correct dominance protocol for wireless networks. These include achieving proper synchronization between the nodes, defining the parameters of the protocol such that clock inaccuracies, time-of-flight and other real-world effects are dealt with and how to perform reliable carrier detection. These aspects are addressed in this work and an implementation of dominance protocols in wireless media – WiDom (short for *wireless dominance*) – is presented.

This work also deals with the case of networks with *multiple broadcast domains* (MBD). Considering MBD is important because it will be difficult to make the SBD assumption hold in a large number of networks deployed in the real-world. Nodes are in a MBD network if it holds that a wireless broadcast made by one node cannot reach all nodes in the network. Such networks (MBD networks) suffer from the well known hidden node problem. This is a challenge that needs to be solved when considering the extention of WiDom to MBD.

While a significant effort of this research work is put into designing novel distributed algorithms to obtain aggregate quantities in a SBD, local aggregation between nodes in geographic proximity can be used as an intermediate step to compute aggregated quantities among all nodes in a multihop network; hence the solution to the problem of computing aggregated quantities in a SBD forms a relevant building block for large-scale data aggregation in multihop networks. This challenge is also tackled in this work.

A final note on the MAC protocol that implements dominance/binary countdown in wireless media (WiDom). One important property of WiDom is that it allows enforcing static priorities. Therefore, it enables, for the first time, static priority scheduling over wireless media. This is also a relevant characteristic in emerging embedded systems because these systems deal with the physical world, therefore one important requirement to be met is that their data services are able to meet timing constraints [1]. The research approach also takes this property into account, and a response-time analysis for the proposed MAC protocol is also developed.

Contributions Developed

In this work we reasoned on the design and implementation of a prioritized MAC protocol (WiDom) which enforces strict priorities over wireless channels and on the design of algorithms that efficiently exploit this MAC protocol to obtain, in a SBD, the minimum (MIN), maximum (MAX) and interpolation of sensor values with a time-complexity that is independent of the number of sensor nodes (it depends only on the sensor value range). These techniques also enable to efficiently obtain estimates of the number of nodes (COUNT) and the MEDIAN. For MBD, the time-complexity of the proposed distributed algorithms developed also depends on the network diameter.

Overview

Adaptation of Dominance Protocols to Wireless Media. This research work introduced an adaptation of a dominance protocols for wireless media, which existed previously only for wired media. The implementation of a dominance protocol for wireless media was named WiDom and was initially proposed under the assumption of a SBD [6], [7]. WiDom can be exploited to efficiently obtain aggregated quantities, and it is also useful to provide pre-runtime guarantees for sporadic messages streams. A schedulability analysis for WiDom was developed accordingly.

Extension of WiDom to Support Multiple Broadcast Domains. To cope with larger geographical areas, networks with *multiple broadcast domains* (MBD) need to be considered. An extension of WiDom for wireless networks with MBD was also proposed. The proposed solution is the first prioritized and collision-free MAC protocol designed to successfully deal with hidden nodes without relying on out-of-band signaling [8].

Improving the Reliability of WiDom in SBD. The techniques employed to solve the hidden node problem in [8] can also be adapted to improve the reliability of the protocol in a SBD. The proposed solution has a result that is similar to a cooperative relaying scheme, where several nodes can participate in the transmission of the priority bits [9].

Scalable and Efficient Aggregate Quantities.

By exploiting dominance protocols it is straightforward to demonstrate that, in a SBD, the minimum value (MIN) can be obtained with a time-complexity that is $O(npriobits)$, where *npriobits* is the number of bits used to represent the sensor data (the same technique can be applied to obtain the maximum value (MAX);). techniques to efficiently compute more complex aggregate quantities such as the number of nodes (COUNT), MEDIAN and interpolation by exploiting dominance protocols were also implemented in the wireless domain [10]. Finally, the techniques employed to obtain aggregate quantities in a SBD were also extended for multihop networks [11]. The algorithms proposed for MBD have a time-complexity that only depends on the network diameter and on the value range of the sensor readings.

Discussion

There are two important sets of contributions in this research: (i) the development and implementation of WiDom and the support for static priority scheduling over wireless links; and (ii) the algorithms for efficient data aggregation based on WiDom. Let us now review and briefly discuss each of them.

WiDom Development and Implementation.

WiDom was proposed, a novel wireless MAC protocol inspired in dominance/binary countdown protocols which existed previously only for wired media [4]. This achievement is non-trivial. Firstly, implementations of dominance protocols for a wired medium are based on a wired–AND behavior of the bus, where the dominant signal overwrites the recessive signal. Secondly, these implementations require that nodes are able to monitor the medium while transmitting. Clearly this does not easily extend to the case of wireless channels. Moreover, due to non-idealities of transceivers and the nature of the wireless medium, it was not obvious how a dominance protocol could be achieved.

WiDom supports a large number of priorities. Although this number of priorities introduces overhead, the application developer has the freedom to choose the number of priority levels required, and thus possibly reduce the overhead introduced. Nevertheless, such a large number of priorities can be supported by other prioritized protocols (see e.g., [12], [13]) but at the cost of an overhead several orders of magnitude higher.

The initial design of WiDom was created under the assumption of a SBD. An extension of WiDom for wireless networks with MBD was also developed. In such scenario, the hidden node problem must be dealt with. The proposed solution is the first prioritized and collision-free MAC protocol designed to successfully deal with hidden nodes without relying on out-of-band signaling.

The idea of retransmitting priority bits used to solve the hidden node problem can also be adapted to improve the reliability of WiDom in a SBD. In the presence of several nodes in the same broadcast domain, various nodes can cooperate in the transmission of the priority bits. This simple modification of the protocol can result in a substantial gain in the reliability, as the number of nodes increases.

WiDom was implemented and evaluated experimentally using Commercial-Off-The-Shelf (COTS) technology. The implementation of WiDom using COTS technology suffered however from a significant overhead. Therefore, a platform with better characteristics to implement dominance protocols was also studied and developed. This platform is the proof of concept that highly scalable aggregate computations in wireless networks are competitive in practice.

The experimental evaluation of WiDom shows that the probability that a message is transmitted collision-free, correctly prioritized and received (neither lost nor corrupted) by all other nodes is high and this reliability justifies the study of schedulability analysis techniques for sporadic messages in wireless networks; WiDom is an enabling technology allowing schedulability analysis (for example to exercise in practice the analysis proposed by [14]) in wireless multihop networks with multiple broadcast domains. For the case of SBD, a response-time analysis for WiDom was developed and tested as well.

Efficient Data Aggregation.

The research on efficient data aggregation in this work is motivated by scenarios where even a small broadcast domain may contain several tens of sensor nodes. In these scenarios, the advantages of data aggregation solutions found in previous research are lost, since it is neither possible to apply data reduction functions to data coming from different sources nor is it possible to exploit the opportunities for parallel transmissions. In this thesis was demonstrated that it is possible to exploit a dominance-based MAC protocol to efficiently compute aggregate quantities with a time-complexity that is equivalent to the time of transmitting a single message, even if hundreds of nodes are in the same broadcast domain.

Concretely, the present work demonstrated that, in a single broadcast domain (SBD), the minimum value (MIN) can be obtained with a time-complexity that is $O(npriobits)$, where *npriobits* is the number of bits used to represent the sensor data. In this case, the message complexity (and thus, the time-complexity) is independent of the number of sensor nodes. The same technique can be applied to obtain the maximum value (MAX).

Based on these techniques (of obtaining MIN or MAX), more elaborate aggregated quantities can be obtained. In this work, useful examples such as the number of nodes (COUNT) and MEDIAN were also addressed.

Often, it is required to know how physical quantities (such as temperature) vary over an area. Clearly, the physical location of each node must be known then. For such systems, an algorithm that produces an interpolation of the sensor data as a function of space coordinates was proposed. The resulting interpolation is a compact representation of sensor data at a moment and is obtained efficiently.

The algorithms (to obtain aggregate quantities) were initially designed with the assumption of a SBD network. Nevertheless, in practice, most networks are not SBD networks. Therefore, solutions for MBD networks were also studied and proposed. An algorithm for computing the MIN (or MAX) of sensor readings in a multihop network was proposed. That algorithm has the particularly interesting property of having a time-complexity that does not depend on the number of sensor nodes; only on the network diameter and the range of the value domain of sensor readings matter. Other more sophisticated algorithms were demonstrated also feasible for MBD networks.

These results are significant because often networks of nodes that take sensor readings are designed to be large scale, dense networks and it is exactly for such scenarios that the proposed algorithms (designed in close articulation with the MAC protocols) excel. The implementation of these algorithms in the hardware platform developed shows that such highly scalable aggregate computations in wireless networks are indeed competitive in practice.

Discussion on Important Assumptions

There are two scenarios in this research work that are important to fully grasp the relevance of the contributions presented.

Dense Wireless Networks

The algorithms designed in this research are developed for large-scale, dense networks. In particular, it is in the presence of many nodes in the same broadcast domain that the advantages of an algorithm whose complexity does not depend on the number of nodes becomes evident. Note that, it was show that we can see advantages with as few as two nodes [10].

Some may object accepting such scenario. Indeed, there are a few results which might advise against deploying dense networks. For instance, previous results on the capacity of ad-hoc wireless networks [15] show that, under certain assumptions, the capacity per node approaches zero as the number of nodes increases. However, several facts show that this is not the case in the context of WSN [16], [17], and other solutions can be devised (see, for example, the results reported in [18]).

One reason to search for other solutions is that data in sensor network is correlated, and this can be explored in several ways (e.g., [16], [18]). In this work, the spatial correlation between sensor readings was exploited to perform a weighted-average interpolation and select only a subset of nodes.

A second reason differentiating WSN is that the communication pattern is often from several source nodes to a sink. This enables several techniques such as in-network data aggregation or other techniques such as antenna sharing [17].

Another question is *why deploy a dense sensor network if we know that we will be gathering a lot of redundant data*? Deploying a dense sensor network might be convenient for several reasons. First, when considering the case where the individual cost of each node is negligible, then deploying redundant nodes might not be a primary factor in the cost of the system. Redundant nodes are useful for fault-tolerance (under certain fault assumptions) and noise immunity. Deploying redundant nodes also allows for a very fine spatial resolution in the sampling of the phenomena being observed. More importantly, deploying a dense network allows a better resolution of how the physical world is perceived; for example, when we are interested in high-resolution sampling in a certain region of interest, but we do not know in advance where that region is. This is expected to be essential in forthcoming innovative applications in cyber-physical systems.

Timeliness Guarantees in Wireless and Reliability

An important contribution of this research work is also a MAC protocol that supports static priority scheduling in wireless networks. This contribution encompasses its full impact, when assuming that *it is relevant to analyze the problem of providing timing requirements within a hard real-time context*.

It is often indicated that wireless links are unreliable, thus it is not meaningful to approach the problem from a guaranteed timeliness perspective.

It is true that designing a protocol with an upper bound on queuing time is not sufficient to guarantee that hard real-time deadlines are satisfied in practice. However, it is a necessary step towards that goal. It is important to note that part of this problem is a technological one. The reliability of wireless networks has evolved noticeably over the last few years; it is safe to assume this evolution will continue. The experimental evaluation performed in this research work suggests that deadline misses due to noise are rare, so this provides evidence that it is still useful to consider hard real-time bounds on queueing delays. Obviously, any kind of guarantees are subject to some assumptions.

For the case of WiDom in SBD networks, there is a range of techniques originally developed for the CAN bus that can be applied, and this is a very interesting research path to explore. These techniques include the schedulability analysis as presented in [6], [7], which are easily extendable to consider acknowledgements and retransmissions, and also other techniques such as stochastic approaches to model faults [19], for example. Moreover, the scheme to improve the reliability of WiDom in a SBD results in a substantial gain in the reliability, as the number of nodes increases. This fact alone can enable hard real-time deadlines to be satisfied with a very high probability in practice.

Even if not employing a guaranteed framework, being able to enforce strict priorities is useful in general. One recent example can be found in [13], where the authors built a streaming audio application employing a prioritized MAC protocol previously proposed in the literature [20]. It was found that, although the overhead of this MAC protocol is high, it still offers better throughput than normal CSMA/CA protocols because such prioritized MAC protocols eliminate the overhead of back-off after collisions. The scalability of the work reported in [13] is partially limited to the small number of priorities supported by the MAC protocol adopted and it is exactly in this point that WiDom stands out. Compared to that MAC protocol, WiDom offers a much higher number of priorities for a given similar overhead.

Future Work

There are several opportunities that can be identified for further research:

- development of other mechanisms to obtain other aggregate quantities or perform other distributed computations;
- enclosing in a query processing system;
- integration with other communication protocols;
- power management schemes/duty cycling;
- improvement of current implementations;
- further development of radio hardware;
- maximize the number of parallel transmissions of WiDom-MDB.

The following paragraphs briefly discuss each one of these possible future research topics.

The fact that MIN can be obtained efficiently by exploiting WiDom can serve as a building block for other computations like COUNT, MEDIAN or Interpolation. It is possible to foresee that other computations may eventually be devised out of similar ideas and this is a research topic to be explored.

The computations enabled by WiDom could be encapsulated in a query processing system, similar TinyDB [21]. The query processing would receive the query specifications and map these into primitives that exploit WiDom adequately. Eventually, queries that cannot be more efficiently carried out by exploiting WiDom could be mapped into other primitives that do not work by directly exploiting WiDom.

One interesting possibility is to integrate WiDom in other communication protocols. For example, a TDMA protocol can benefit from WiDom by allowing several nodes to share the same timeslot. Inside that timeslot, contention is resolved using WiDom. This combination would allow reducing the TDMA cycle. Another similar example is to integrate WiDom in the IEEE 802.15.4 beacon-enabled mode, where one timeslot could be reserved for WiDom. In this case, during that timeslot, all nodes are allowed to access the medium using WiDom. Both examples can improve the schedulability of the system.

WiDom is energy efficient as it can avoid packet collisions and enable very efficient computations, but it lacks energy efficiency in the sense that it requires that nodes constantly monitor the medium. However, this does not need to be the case. WiDom is compatible with power management schemes proposed in WSN, such as coordinating activity/sleep schedules between the nodes (e.g. [22]). This is an important topic for further research.

At this point, the implementations of WiDom (and distributed algorithms) available are considered only as proof-of-concept prototypes. This means that the implementations were not developed for general use and they require some effort of understanding the details of the implementation in order to be used. This is not a topic for research, but it is important to consider this fact in order to asess the amount of effort needed to experiment with WiDom and/or further develop it.

The development of an efficient implementation of WiDom is still an important subject for future work. While a platform specifically designed for implementing WiDom efficiently in provides an indication that the overhead of the protocol can be low such that the algorithms based on exploiting WiDom can be very competitive [10], there is still a lot of work to be done in this aspect. The platform presented is a prototype developed from components commercially available at the time. The development and research of better radios for the execution of WiDom is a topic that would benefit the algorithms for obtaining aggregate quantities presented and would also enable low overhead static priority scheduling in wireless systems.

Finally, the proposed protocol for WiDom in a network with MBD [8], does not maximize the number of parallel transmissions. This is a problem that can be difficult to solve, but, in practice, strategies as planning the priorities in the nodes such that multihop competing is avoided can provide interesting solutions.

Conclusion

This paper presented the research efforts developed around the following hypothesis: *Is it possible to compute aggregate quantities with a time-complexity that is independent of the number of sensor nodes?*

To achieve this goal, WiDom was designed in close articulation with the data aggregation mechanisms. By exploiting the properties of a prioritized MAC protocol, the formulated hypothesis can be supported in SBD networks. Effectively, the time-complexity of the algorithms for obtaining aggregate quantities in SBD networks only depends on the sensor value range. In the case of a MBD network, the time-complexity of the algorithms developed also depends on the network diameter.

One important part of this research was dedicated to make this approach effective in wireless networks. This included the design and implementation of a prioritized MAC protocol in wireless media. Several implementations [6], [7], [8], [10] have shown that such approach is viable. While evidence was presented that the approach is feasible and the algorithms based on exploiting WiDom can be very competitive, there is a big gap to fulfill until such mechanisms are useful in more general settings. Filling this gap involves a wide range of aspects, from development of a framework to ease the use of such techniques by application developers, maturation of radio hardware or integration with other communication protocols.

This research work has demonstrated innovative mechanisms for obtaining certain aggregate quantities in dense wireless sensor networks. The work included the development of a prioritized MAC protocol that enables these mechanisms and can also efficiently schedule sporadic messages. Despite the dificulties yet to overcome, these constitute an attractive set of solutions for emerging Cyber Physical Systems.

Acknowledgment

This research work was partially developed at the Real-Time Computing Systems Research Centre (CISTER), from the School of Engineering of the Polytechnic Institute of Porto (ISEP).

References

[1] John A. Stankovic, Insup Lee, Aloysius Mok, and Raj Rajkumar. Opportunities and obligations for physical computing systems. *Computer*, 38(11):23–31, 2005.

[2] D. Estrin, D. Culler, K. Pister, and G. Sukhatme. Connecting the physical world with pervasive networks. *IEEE Pervasive Computing*, pages 59–69, January-March 2002.

[3] Bhaskar Krishnamachari, Deborah Estrin, and Stephen Wicker. Modelling data-centric routing in wireless sensor networks. In *21st Annual Joint Conference of the IEEE Computer and Communications Societies INFOCOM'02*, New York, USA, 2002.

[4] Aloysius K. Mok and Steve Ward. Distributed broadcast channel access. *Computer Networks*, 3:327–335, 1979.

[5] CAN specification, ver. 2.0, 1991. Bosch GmbH, Stuttgart, Germany.

[6] N. Pereira, B. Andersson, and E. Tovar. Implementation of a dominance protocol for wireless medium access. In *12th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA'06)*, pages 162–169, Sydney, Australia, 2006.

[7] N. Pereira, B. Andersson, and E. Tovar. WiDom: A dominance protocol for wireless medium access. *IEEE Transactions on Industrial Informatics (TII)*, 3(2):120–130, May 2007.

[8] Nuno Pereira, Bjrn Andersson, Eduardo Tovar, and Anthony Rowe. Static-priority scheduling over wireless networks with multiple broadcast domains. In *28th IEEE Real-Time Systems Symposium (RTSS'07)*, pages 447–456, Tucson, Arizona, USA, 2007.

[9] N. Pereira, B. Andersson, E. Tovar, and P. Carvalho. Improving the reliability of widom in a single broadcast domain. In *IEEE Symposium on Industrial Embedded Systems (SIES'09), Work-in-Progress Session*, pages 144–147, Lausanne, Switzerland, 2009.

[10] N. Pereira, R. Gomes, B. Andersson, and E. Tovar. Efficient aggregate computations in large-scale dense WSN. In *15th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS'09)*, pages 317–326, San Francisco, California, USA, 2009.

[11] N. Pereira, B. Andersson, E. Tovar, and P. Carvalho. Efficient computation of min and max sensor values in multihop networks - by exploiting a prioritized MAC protocol. *Springer Lecture Notes on Electrical Engineering (LNEE)*, 38, Intelligent Technical Systems:233–246, 2009.

[12] Joo L. Sobrinho and A.S. Krishnakumar. Quality-of-service in ad hoc carrier sense multiple access networks. *IEEE Journal on Selected Areas of Comm.*, 17(8):1353–1368, 1999.

[13] B. D. Bui, R. Pellizzoni, M. Caccamo, C. F. Cheah, and A. Tzakis. Soft real-time chains for multi-hop wireless adhoc networks. In *13th IEEE Real Time and Embedded Technology and Applications Symposium (RTAS'07)*, pages 69–80, 2007.

[14] T. F. Abdelzaher, S. Prabh, and R. Kiran. On real-time capacity limits of multihop wireless sensor networks. In *IEEE Real-Time Systems Symposium (RTSS'04)*, pages 359–370, Lisbon, Portugal, 2004.

[15] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.

[16] Daniel Marco, Enrique J. Duarte-melo, Mingyan Liu, and David L. Neuhoff. On the many-to-one transport capacity of a dense wireless sensor network and the compressibility of its data. In *2nd International Workshop on Information Processing in Sensor Networks (IPSN'03), Lecture Notes in Computer Science*, pages 1–16, Palo Alto, CA, USA, 2003.

[17] H.E. Gamal. On the transport capacity of the many-to-one dense wireless network. In *IEEE 58th Vehicular Technology Conference (VTC'03-Fall)*, pages 2881–2885, Vol.5, Oct. 2003.

[18] Himanshu Gupta, Vishnu Navda, Samir Das, and Vishal Chowdhary. Efficient gathering of correlated data in sensor networks. In *6th ACM Intl. symposium on Mobile ad hoc networking and computing (MobiHoc'05)*, pages 402–413, 2005.

[19] Ian Broster, Alan Burns, and Guillermo Rodríguez-Navas. Probabilistic analysis of can with faults. In *23rd IEEE International Real-Time Systems Symposium (RTSS'02)*, pages 269–278, Austin, TX ,USA, 2002.

[20] Joo L. Sobrinho and A. S. Krishnakumar. Real-time traffic over the IEEE 802.11 medium access control layer. *Bell Labs Technical Journal*, 1(2):172–187, 1996.

[21] Samuel R. Madden, Michael J. Franklin, Joseph M. Hellerstein, and Wei Hong. TinyDB: an acquisitional query processing system for sensor networks. *ACM Trans. Database Syst.*, 30(1):122–173, 2005.

[22] Wei Ye, John Heidemann, and Deborah Estrin. Medium access control with coordinated adaptive sleeping for wireless sensor networks. *IEEE/ACM Transactions on Networking (TON)*, 12(3):493–506, 2004.

Nuno Pereira is a researcher at CISTER/ISEP and a Lecturer at the Computer Engineering Department of ISEP, the School of Engineering of the Polytechnic of Porto. He received a licentiate degree in Computer Engineering from the School of Engineering of the Polytechnic Institute of Porto and an MSc Degree from the University of Minho in 2002 and 2005 respectively, and has also recently presented his doctoral thesis. Nuno worked in a number of research projects related to several aspects of real-time systems and has authored and co-authored more than 20 research papers in international refereed conferences and journals.