Wireless World Research Forum (WWRF)

Addressed objective: Identifying new research areas

Title of the research item: Making sensor networks intelligent

Authors

Andreas Köpke, Vlado Handziski[1], Holger Karl
koepke|handzisk|karl@ft.ee.tu-berlin.de
Telecommunications Network Group
Technische Universität Berlin
Einsteinufer 25, 10587 Berlin
http://www-tkn.ee.tu-berlin.de

Objectives of the required research

Our environment contains many physical processes that should be observed and/or controlled. Examples include climate control in intelligent buildings, factory floor automation and environmental control. These tasks require many individual sensors and actuators, which today are connected into fixed, wired networks. However, planning, installing and maintaining such a system is prohibitively complex and expensive.

In the future, this problem will become aggravated by the even wider spread of sensing and actuating facilities in our daily environment: not only buildings will have communication needs, but every day items like food packages, cars, etc. will all provide information and enable interaction. Wired communication infrastructures are evidently not up to this challenge.

Enabled by low-power electronics and efficient wireless communication, a new solution emerges in the form of wireless sensor networks. These sensor networks are based on ad-hoc networking capabilities. Since these networks are supposed to work for a long time without any maintenance, they can only consume a small amount of energy, possibly scavenged form the environment. The main focus of current research in this field is on technological issues like developing the necessary system parts and on basic networking protocols; very little heed has been paid so far to making these systems more dependable and useful.

We intend to highlight in this paper the currently ongoing research activities in the context of sensor networking and show why this research – while establishing the foundations upon which to build – is not sufficient to provide the additional functionalities and the usefulness that will be required from sensor networks in future ambient networking environments. We want to point out that the enrichment of sensor networks with semantic capabilities will be a major challenge in this context. After summarizing the ensuing research issues, we outline a possible research approach to solve this problem.

Related Work

Sensor networks have been an active research area for a number of years now. This section

[1] Corresponding author.

provides a very brief overview of the major research activities; the following section will then use this as a basis to highlight missing bits and pieces.

Hardware

The initial research has been mainly driven by hardware as the enabling technology: turning communication devices and microcontrollers both physically very small and very energy efficient. Examples for current sensor node prototypes include WINS [1], SmartDust [2], PicoRadio [3], Motes [4] and others.

MAC and Link Layer

On top of hardware, fundamental communication between two neighboring nodes is enabled by medium access control and link layer protocols. The situation for such protocols is similar to that for wireless ad-hoc networks (typical problems like hidden terminal have to be solved), but it is aggravated by the need to be highly energy efficient and to put nodes to sleep whenever possible. Examples for research in this direction include [5][6][7][8].

This is simple yet fundamental functionality for a sensor network. However, all of this research is targeted to one-to-one communication between two nodes. Very little research has been performed on "multicast-friendly" MAC and link layer protocols that can support the specific needs of communicating with nodes in the neighborhood of a specific node. We will later point out that such group-based communication is likely to become a major issue in wireless sensor networks; it is interesting to note in passing that MAC protocols for *wired* sensor networks are able to handle such requirements and use this capability in their programming philosophies (e.g., LON [9]).

Locationing

In many application scenarios for sensor networks, knowing the absolute or relative position of sensors is important, but explicitly providing this information is not acceptable. This information can be obtained from external sources like GPS, by purely RF-based solutions or by a combination of RF and ultrasound techniques. Examples include [10][11][12].

Clustering

Taming the complexity of large sensor networks can be tackled by introducing hierarchies. Some types of hierarchies exploit immediate vicinity of nodes by grouping them into clusters and exploit these clusters for a number of different tasks such as routing or data aggregation. Example for clustering protocols include [13][14]. Such vicinity-based groups are an important issue, but do not fulfill all the necessary requirements to support groups.

Routing and Non-Routing

The routing problem has probably received the lion share of attention: How to make sure that distant nodes can communicate with each other? Much of the ad-hoc routing research is to some degree applicable to sensor networks, but the crucial difference between sensor networks and ad-hoc networks is then not given due consideration: Sensor networks work in a data-centric fashion, ad-hoc networks are node-centric. Hence, some approaches forego routing in the

standard sense, e.g., directed diffusion [15][18]. While this is relevant, issues like observer or gateway mobility are not completely solved, as is the proper integration with data aggregation schemes.

Middleware

Perhaps the most important aspect to solve is the integration of sensor networks into existing Internet architectures. Some approaches like SINA [16] have been proposed, yet a lot of work remains to be done to make services rendered by sensor networks seamlessly accessible via existing Internet middleware architectures like .NET or CORBA.

Applications

The example applications that are currently discussed for sensor networks typically are "dumb" data acquisition applications or object tracking applications. Anything that requires a closed loop of sensing data, decision making and acting on these decisions is hardly possible with current approaches [17]. Such applications, however, can be found in many scenarios and have perhaps the most stringent requirements regarding both functionality and dependability.

Conclusions

Much research has been done to make moving of bits from A to B possible in an energy-efficient fashion. However, supporting truly distributed applications, e.g., closing the loop between sensors and actuators within a sensor network, is still a long way to go. For such a goal, distributed algorithms are necessary on a number of different levels.

In principle, a lot of the required algorithms and technology is available from theoretical computer science. However, it is not clear what can be expected when this theory is applied to sensor networks. Applying some of these algorithms to ad-hoc network, which are in many ways closely related to sensor networks, has already shown that many of the simple assumptions and abstractions used in theoretical computer science do not hold. Some of them will have to be revised to reflect the particular challenges of sensor networks (highly constrained energy and computation resources), leading to new approaches. For instance: One basic assumption in distributed algorithms is that failures are not correlated. This seems unrealistic in a sensor network where the movement of, e.g., sensors usually happens in a correlated fashion, making link breaks and reformations highly correlated events. Hence, directly applying such algorithms will certainly lead to inferior performance.

Research Issues

Providing such distributed functionality in a sensor network requires new innovations in a number of different fields. The long-term goal is to make sensor networks first-class participants in future ubiquitous ambient networks. Particular relevant appear to be new solutions in the following areas:

- Understand sensor networks' purposes
    - To begin with, the requirements of the application scenarios typical for sensor networks regarding dependability as well as basic abstractions have to be well

understood and formalized in an appropriate fashion. While this sounds fairly obvious, there is still only a limited understanding about how sensor network applications behave in the real world, what types of solutions they have to provide: Much is hypothesized, little is known.

  - As an important case, closing the loop between sensors and actuator will create requirements on dependability, timeliness and functionality that are very different from “dumb” data acquisition networks.

- Make sensor networks smart and functionally rich
  - How to built dependability, collaboration and semantics into the very sensor network itself? Providing such functionality will be important, but in order to achieve acceptable energy efficiency, it must be tightly integrated with lower layer protocols.
  - How to make these capabilities useable from an application programmer’s perspective? Simple building blocks for efficient distributed services are required; a set of necessary primitives that sensor networks have to provide to support meaningful distributed applications has to be found; a set of required basic distributed algorithms can also be derived to implement these service primitives. The goal is to develop a small set of building blocks that an application programmer can easily use.
    One possible approach could be to use group communication as a simple abstraction; actual distributed services like consensus could then easily built on top of it. The following section extends on this thought.
- Integration of sensor networks and conventional networks
  - Such “intelligent” sensor networks will enable enriched interaction patterns with their environments. How to enable semantically relevant interaction with sensor networks from without, bridging the gap over multiple, very diverse radio technologies and gatewaying between very different protocol design philosophies? On a higher layer, the question is how to translate the sensor network’s data-driven way of operation into the Internet’s node-address-driven way of operation.
  - How to integrate these different ways of addressing and handling data into middleware architectures? Making sensor network services accessible from/via existing middleware architectures like .NET or CORBA will be crucial to allow simple integration on an application level.
- Lower layer protocol network infrastructure for these services
  These higher-layer functionalities require some specific types of lower-layer communication protocol that are unlike those that are currently investigated. Specifically:
  - Bridging the gap between vicinity-based clusters and semantic groups is necessary to join energy-efficient communication with semantically relevant communication

- Supporting diverse types of mobility in data-centric routed networks: The mobility of sensor nodes themselves, observer / sink nodes, and data sources has to be treated in a seamless fashion
- Multicast-friendly MAC / Link Layer: As group communication is likely to become a basic paradigm for making enriched sensor network functionality accessible to an application programmer, it should be efficiently supported by the lower communication layers. In fact, this ties down to the very structure of the physical layer itself.

Possible Approach

We propose a top-down approach to make the open issues more concrete and to derive solutions for them later on.

As stated earlier, first the service primitives required by the applications have to be identified. The demanded dependability and service semantics influences the way these primitives have to be implemented. There is a trade-off between dependability and other metrics like delay or time needed to terminate such a basic algorithm on the one hand and energy efficiency on the other hand. This trade-off must be well understood to propose a suitable solution for a given set of requirements.

In order to make a distributed system dependable, the concept of a group and the communication in such a group is central. Furthermore, groups are a convenient tool for an application programmer. The established group communication algorithms have to be adapted to the particular nature of sensor networks. Identifying groups by nodes in their mutual immediate vicinity is to some extent natural in wireless sensor networks, since the wireless medium is a broadcast medium. But much more useful is the identification of group members based on their contribution to a common task: Identifying groups or a couple of sensors based on their task or sensed value introduces the concept of semantic addressing or intentional naming.

Supporting such semantics-based addressing and group communication will be a key element in future sensor networks, paving the way to a data-oriented sensor network architecture. It will open the door to new types of services such as collaboration within sensor networks. It will also be a major point of integration of sensor networks with outside, conventional networks by giving other applications the possibility to easily access data and services in their ambient environment without having to worry about the characteristics of the underlying network technology and topology. We are currently in the process of developing compact ways of expressing such semantic addressing in sensor networks and the group communication / collaboration functionality enabled by them. We believe that a major research effort is still required to easily and seamlessly integrate sensor networks into future ambient networks.

References

[1] G. J. Pottie and W. J. Kaiser. Wireless integrated network sensors. CACM, 43(5):51-58, may 2000.

[2] J. M. Kahn, Randy H. Katz, and K. S. J. Pister. Next century challenges: Mobile networking for smart dust. In Proc. ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 99), Seattle, WA, August 1999.

[3] J. M. Rabaey, M. J. Ammer, J. L. da Silva, D. Patel, and S. Roundy. Picoradio supports ad hoc ultra-low power wireless networking. IEEE Computer, 33(7):42-48, July 2000.

[4] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. E. Culler, and K. S. J. Pister. System architecture directions for networked sensors. In Proc. ASPLOS, pages 93-104, 2000.

[5] Alec Woo and David Culler. A transmission control scheme for media acces in sensor networks. In Proceedings of the Seventh Annual International Conference on Mobile Computing and Networking, Rome, Italy, july 2001.

[6] K. Sohrabi, J. Gao, V. Ailawadhi, and G.J. Pottie. Protocols for self-organization of a wireless sensor network. IEEE Personal Communications, 7(5):16-27, october 2000.

[7] Wei Ye, John Heidemann, and Deborah Estrin. An energy-efficient mac protocol for wireless sensor networks. Technical Report ISI-TR-543, USC/ISI, september 2001.

[8] Lizhi Charlie Zhong, Rahul Shah, Chunlong Guo, and Jan Rabaey. An ultra-low power and distributed access protocol for broadband wireless sensor networks. In IEEE Broadband Wireless Summit, Las Vegas, N.V., 2001.

[9] Dietmar Dietrich, Dietmar Loy, and Hans-Jörg Schweinzer, editors. LON-Technologie: Verteilte Systeme in der Anwendung. Hüthig Verlag, Heidelberg, 2. edition, 1999.

[10] C. Savarese, J. M. Rabaey, and J. Beutel. Locationing in distributed ad-hoc wireless sensor networks. In Proc. Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP 2001), Salt Lake City, Utah, May 2001.

[11] J. Li, J. Jannotti, D. S. J. De Couto, D. R. Karger, and R. Morris. A scalable location service for geographic ad hoc routing. In Proc. 6th ACM Intl. Conf. on Mobile Computing and Networking, pages 120-130, Boston, Massachusetts, August 2000. http://www.pdos.lcs.mit.edu/papers/grid:mobicom00/paper.pdf.

[12] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In Proc. 6th Intl. Conf. on Mobile Computing and Networking (Mobicom), 2000. http://nms.lcs.mit.edu/papers/cricket.html.

[13] Wendi Rabiner Heinzelman, Anantha Chandrakasan, and Hari Balakrishnan. Energy-efficient communication protocol for wireless microsensor networks. In Proceedings of the 33rd Hawaii International Conference on System Sciences, 2000.

[14] A. Cerpa and D. Estrin. ASCENT: Adaptive self-configuring sensor networks topologies. In Proc. Infocom, June 2002.

[15] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed diffusion: A scalable and robust communication paradigm for sensor networks. In Proc. Sixth Ann. Intl. Conf. on Mobile Computing and Networks, Boston, MA, August 2000. http://lecs.cs.ucla.edu/~estrin/papers/diffusion.ps.

[16] Chien-Chung Shen, Chavalit Srisathapornphat, and Chaiporn Jaikaeo. Sensor information networking architecture and applications. IEEE Personal Communications, 8(4):52-59, August 2001.

[17] A. Lim. Distributed services for information dissemination in self-organizing sensor networks. Special Issue on Distributed Sensor Networks for Real-Time Systems with Adaptive Reconfiguration, Journal of Franklin Institute, 388:707-727, 2001.

[18] W. Heinzelman, J. Kulik, and H. Balakrishnan. Adaptive protocols for information dissemination in wireless sensor networks. In *Proc. 5th Intl. Conf. on Mobile Computing and Networking (Mobicom)*, Seattle, WA, August 1999. An updated version of this paper is available at http://wind.lcs.mit.edu/papers/spin.ps.