

D3: Data-centric Data Dissemination in Wireless Sensor Networks

Maarten Ditzel

Netherlands Organisation for Applied Scientific Research
P.O. Box 96864, 2509 JG Den Haag, The Netherlands
Email: maarten.ditzel@tno.nl

Koen Langendoen

Delft University of Technology
Mekelweg 4, 2628 CD Delft, The Netherlands
Email: k.g.langendoen@ewi.tudelft.nl

Abstract—This paper presents a novel method to disseminate sensor data in a wireless sensor network, called D3 (Data-centric Data Dissemination). The method combines the advantages of data-centric routing like SPIN and directed diffusion and energy-efficient MAC protocols such as S-MAC and T-MAC. The protocol's strengths are its energy-efficiency and its simplicity. Messages are transmitted using broadcasting only, reaching as many nodes as possible with the least energy. Furthermore, D3 easily accommodates energy-dependent traffic balancing and data aggregation, crucial to prolong the lifetime of a sensor network.

I. INTRODUCTION

Wireless sensor networks [1], [2] have experienced increasing attention in academic, industrial and military environments over the past few years. These networks promise an easy-to-deploy, easy-to-use and moreover, low-cost means to remotely monitor environments. Furthermore, sensing accuracy can be improved significantly by processing and combining collected data within the network itself. Finally, the network can be made robust to the failure of individual nodes, which ensures that the lifetime and proper operation of the network is not limited to the lifetime of one node in particular.

Applications, either envisioned or already realized, are generally related to the remote monitoring of a, possibly inaccessible or hostile, environment. Examples are an aqueous surveillance system for a drinking water reservoir [3], a wildlife habitat observation system [4], and a network to monitor the behavior of glaciers [5]. Additionally, numerous military applications are envisioned such as a battlefield data collection network as described in [6].

In all the aforementioned applications, the network consists of tens to millions of tiny devices. Each device carries one or more sensors and has limited signal processing and communication capabilities. Usually, the devices are powered by batteries and can thus only operate for a limited time period. Key to implementing a network with such devices is that energy, computing power and communication bandwidth are scarce. Therefore, light-weight, scalable, energy-conserving communication protocols are essential to the successful operation of the network. Fast deployment of such a network and robustness against device failures require an ad-hoc network that is self-organizing. In general, radio communication (both transmitting and receiving) is generally the operation that consumes the most energy in a device.

Conventional ad-hoc *address-oriented* communication protocols, such as IEEE 802.11 [7], generally consume too much energy or poorly support multi-hop networks. In [8] and [9] the authors propose a new *data-centric* approach for the dissemination of data in sensor networks. At the same time, numerous

energy-efficient MAC protocols have been developed such as S-MAC [10] and T-MAC [11]. However, these protocols still operate in an address-oriented fashion. In accordance with the shift in paradigm to data-centric communication protocols, this paper proposes a novel method for communicating sensor data through a wireless sensor network. It is a simple and straightforward data-oriented method combining the ideas of data-centric routing with an energy-efficient MAC protocol.

II. DATA-CENTRIC ROUTING

To operate a wireless sensor network successfully, two closely related networking issues have to be addressed. The first is the media access protocol, i.e., how the nodes access and share the radio channel. The second is routing, i.e., how data is passed through the network. Dependent on the structure of the network different solutions may be found. In this paper we consider a homogeneous, randomly distributed wireless sensor network. All nodes are identical and offer the same processing and communication capabilities. Several routing schemes and MAC protocols can be found in literature (for an overview see [12]).

Due to the energy restrictions on both processing and communication, sensor data should be transmitted conservatively. Based on this restriction, the authors of [8] suggest the use of *data-centric* routing and introduced the family of the SPIN protocols. Here, sensor nodes broadcast an advertisement describing the available data and wait for a request of an interested neighbor, before sending the actual data. In [9] the authors propose the *directed diffusion* data dissemination paradigm. A node, called the *sink*, transmits an *interest*, describing the data it is interested in, to all nodes using flooding. Subsequently, if new data becomes available, the data is propagated using *gradients* that indicate where the interest originated from.

Although both directed diffusion and the SPIN protocols are data-centric routing protocols, they assume the use of an address-centric MAC protocol, as they rely on message passing to individual neighbors. As an alternative, this paper suggests to use data-centric message passing instead.

III. D3: DATA-CENTRIC DATA DISSEMINATION

The D3 protocol is based on two basic assumptions. First, transmitting information about data is cheaper in terms of energy than transmitting the actual data itself.¹ Second, if

¹The same assumption is made in [8].

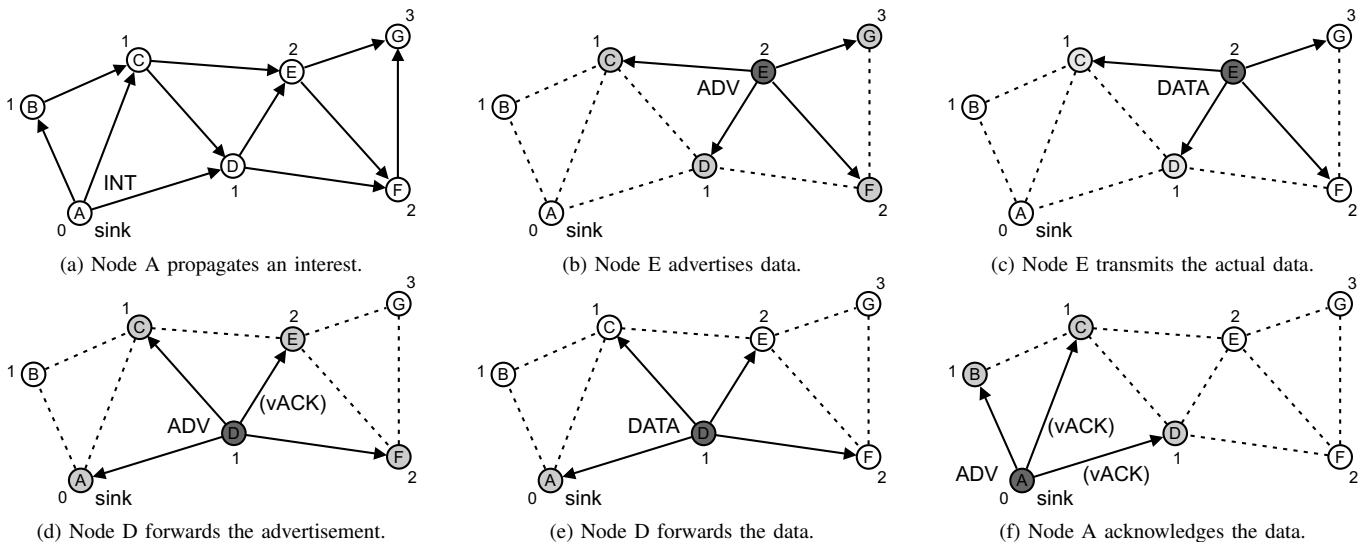


Fig. 1. The D3 protocol. After flooding an interest (INT), data is collected and forwarded using advertisements (ADV) to announce the transmission of the actual data (DATA). Nodes uninterested simply ignore broadcasted data messages.

a node has new data available, at least one of its neighbors is always interested.² Hence, the data must always be transmitted. Therefore, a node does not require a confirmation that a neighbor is interested and it can use broadcasting to communicate the data. The neighbors can decide individually whether to listen to the broadcast or not. Of course, these assumptions have to be verified.

Following the basic principles of the T-MAC protocol [11], the nodes adhere to an *active-sleep* regime. Periodically the nodes wake up from sleeping and enter the active state. In the active state nodes contend for the channel, applying collision avoidance by means of an initial random contention delay and carrier detection. During the sleep interval, messages are queued. The active-sleep regime circumvents the problem of *idle-listening*: the nodes are no longer required to keep their radio in receive mode all the time, since they now know when to expect a message from a neighbor.

The D3 protocol uses three types of messages to disseminate sensor data:

- INT – interest. Interests are propagated from a node to all other nodes in the network. It contains a unique identifier and a task descriptor that describes what kind of data the originating node (the *sink*) is interested in. Each sensor node stores the interest and the hop distance to the sink, also called the *depth* of the node with respect to the particular interest.
- ADV – data advertisement. When a node has data that suits a particular interest, it signals all its neighbors that it has relevant data ready using the interest identifier and its own depth. The data can either originate from its own sensors, or it has been previously received by other nodes and should be forwarded.
- DATA – data message. Data messages contain the actual sensor data.

The basic operation of the D3 protocol is depicted in Fig. 1. Node A starts with sending a data interest to all nodes using flooding (Fig. 1(a)). When node E has new data available (Fig. 1(b)) it advertises the data to its neighbors C, D, F, and

G, and claims a time slot for transmitting the data message. Subsequently, it sends the actual data (Fig. 1(c)) at the reserved time to all interested neighbors: C and D. Nodes F and G are not interested in the data and simply ignore the transmission by switching off their radios. This sequence of advertisement and data transmission (Fig. 1(d) and Fig. 1(e)) is repeated to forward the data until it reaches the sink where the interest originated from. Forwarded ADV messages are also used as virtual acknowledgements (vACK) (see section III-B). Finally, the sink acknowledges the reception of the data (Fig. 1(f)) using a further unused advertisement as acknowledgement.

A. Detailed Operation

The sensor network starts at rest. No data is collected by the nodes and no messages are communicated. To activate particular nodes and sensors an interest is injected into the network via a node, henceforth known as the *sink*. The interest is propagated to all the nodes in the network using flooding. Each node registers the interest's identifier and its corresponding depth. Once a node has received and registered an interest, it inspects the task descriptor and activates the corresponding sensors if necessary.

As soon as nodes have activated sensors they can start collecting data. If a node collects data relevant to a stored interest, it signals all its neighbors that it has data related to a certain interest available. It does so by broadcasting an advertisement containing the interest identifier, the node's depth with respect to the sink the interest originated from, the start time of the data message and the length of the data package. Furthermore, it attaches an identifier to the ADV message. The node also stores the identifiers in a history, to keep track of the recently sent ADV messages.

Upon receiving an advertisement, the node's neighbors check whether they are closer to the sink (lower depth) and, if so, prepare to receive the data message by switching on their radios at the specified start time. If not, they simply ignore the upcoming transmission.³ The same advertisement-transmission sequence is used to forward data messages. As

²Sensor data is only collected after receiving an interest (similar to directed-diffusion). Therefore, at least the neighbor the interest originated from must be interested.

³In this paper a node's depth is used as a decision criterion. However, other cost metrics (e.g., the physical distance to the sink) can also be used.

data is transmitted only to nodes with lower depths, the data will always propagate towards the sink node of the interest.

B. Virtual Acknowledgements

Apart from announcing that sensor data is available for transmission, the ADV message serves another purpose, namely as a virtual acknowledgement. In Fig. 1(d) node D broadcasts an advertisement in order to forward data originally coming from node E. When node E receives the broadcasted ADV message, it will ignore the upcoming transmission as usual, because the depth included in the advertisement is lower than its own depth. However, it also checks the advertisement's identifier to see where the data originated from. If the node can find this identifier in its own history of sent ADV messages, it implicitly knows the data package has been properly received. Hence, the ADV message also serves as an acknowledgement of the proper reception of a data message. If the node does not receive a virtual acknowledgement within a certain time after transmitting a data message, it can assume the data is lost and take appropriate actions.

Finally, the ADV message even has a third role. In case a node receives an ADV message from a node with the same depth (Fig. 1(d), node C), it will postpone the forwarding of the corresponding data messages for a certain time, in order to avoid duplication of data messages. However, if the node does not duly receive a virtual acknowledgement (Fig. 1(f)), it will reschedule the forwarding. In this way, energy waste due to unnecessary transmissions is kept to a minimum, while maintaining the flexibility to reach a sink along different paths.

IV. RESULTS

To verify the ideas and concepts introduced in this paper, simulations have been carried out using the OMNet++ simulation environment [13] (see Fig. 2). Moreover, the D3 protocol is implemented on a small scale experimental sensor network (see Fig. 3). Preliminary results show the proper operation of the protocol and, especially, its scalability. Apart from verifying the proper operation of the protocol, the first experiments focused on possible redundant paths from the source node to the sink, which may reduce the efficiency of the protocol.

Redundant paths may arise when nodes fail to receive virtual acknowledgements from neighboring nodes. Then a message will be multiplied and sent along different paths to the sink. A worst case scenario is depicted in Fig. 4, where the topology of the network is chosen such that every message is duplicated. For clarity, only the transmission of ADV messages is depicted.

Ideally, only one message should be received by the sink for every message sent by the source, i.e., a redundancy of 1. For multiple paths from the source to the sink, the redundancy R is calculated using

$$R = \frac{1}{d(n_1)} \sum_{i=1}^N s(n_i), \quad (1)$$

where n_i denotes the i^{th} node, with n_1 the source node, $s(n_i)$ the total number of data messages sent by node i , and $d(n_1)$ the depth of the source node (e.g., in the example of Fig. 4 $d(n_1) = 4$).

In the experiments the redundancy has been calculated for a large number of randomly generated networks with increasing

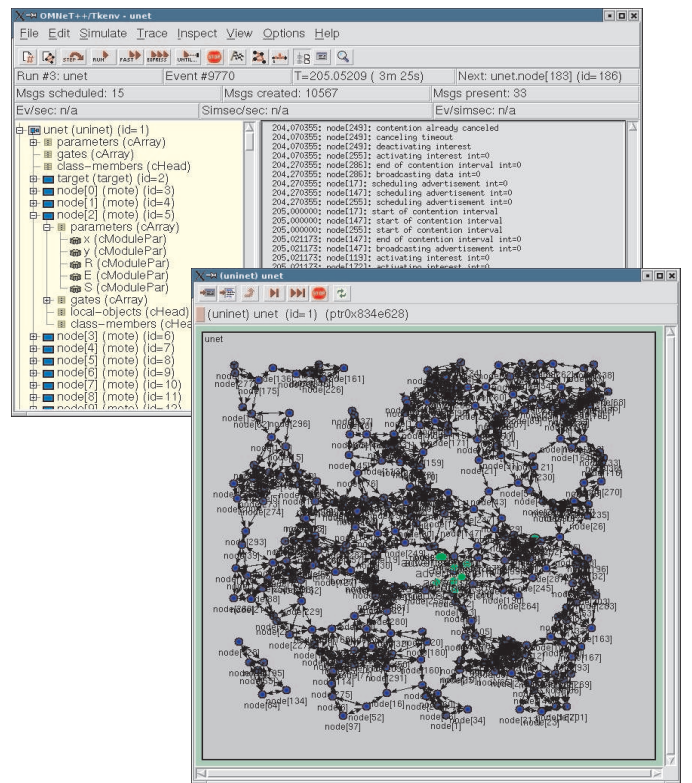


Fig. 2. A screenshot of a running simulation using the OMNet++ simulation environment.

connectivity. The connectivity of a network is the average number of neighbors per node. Each network consisted of over 250 nodes. In Fig. 5, this redundancy is plotted against the connectivity of the network. For each point in the plot, 100 simulations were run. In each simulation 100 messages were sent from the source node to the sink. The figure shows only a slight increase in the redundancy for networks with higher connectivity.

Nevertheless, more extensive simulations are needed to get sensible performance metrics in order to compare the performance of the D3 protocol against other protocols such as directed diffusion and the SPIN family of protocols in combination with, for instance, T-MAC.




Fig. 3. Picture of the TNode platform.

V. CONCLUSION

In this paper, the shift in paradigm towards *data-centric* routing is followed and incorporated into energy-efficient MAC protocols. This leads to a novel data dissemination protocol, called D3. The D3 protocol uses advertisement messages to announce that new data is available. Interested nodes listen to the subsequent transmission of the actual data, while uninterested nodes simply ignore the data message. All communication follows an active-sleep regime, alleviating the problem of idle-listening.

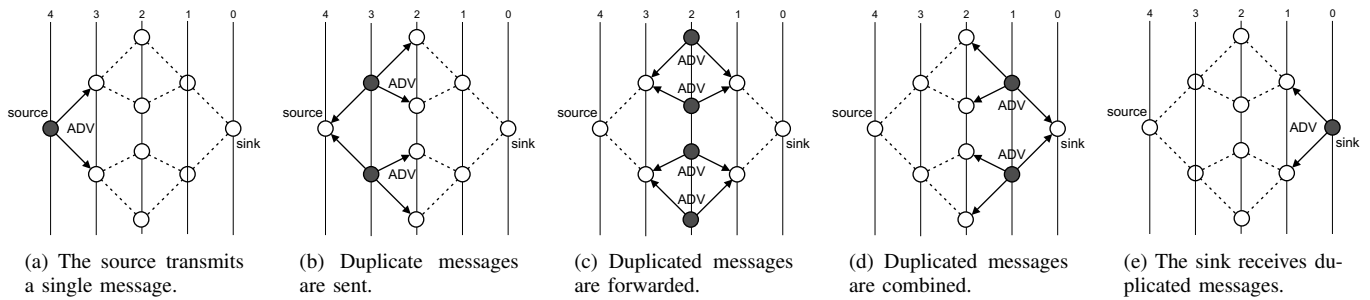


Fig. 4. As nodes along the vertical axes fail to receive each others advertisements, messages are duplicated as they propagate towards the sink node.

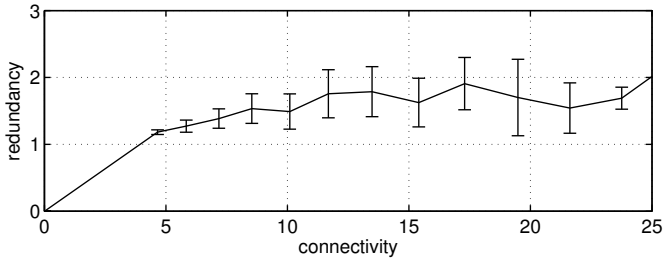


Fig. 5. The average number of redundant paths (connectivity) plotted against a network's connectivity.

The protocol combines the advantages of data-centric routing like SPIN and directed diffusion and energy-efficient MAC protocols such as S-MAC and T-MAC. Its strength is its energy-efficiency and its simplicity. Nodes only need to process local data and control messages are kept to an absolute minimum, both in number and in length. Furthermore, D3 can be easily extended to accommodate energy-dependent traffic balancing and data aggregation, vital for prolonging the lifetime of a battery-powered network.

ACKNOWLEDGEMENT

The authors would like to thank Wouter Vlothuizen and Tom Parker for their valuable comments and suggestions, which greatly contributed in improving the protocol.

REFERENCES

[1] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar, "Next century challenges: scalable coordination in sensor networks," in *Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking*. ACM Press, 1999, pp. 263–270.

[2] G. J. Pottie and W. J. Kaiser, "Wireless integrated network sensors," *Commun. ACM*, vol. 43, no. 5, pp. 51–58, 2000.

[3] X. Yang, K. G. Ong, W. R. Dreschel, K. Zeng, C. S. Mungle, and C. A. Grimes, "Design of a wireless sensor network for long-term, in-situ monitoring of an aqueous environment," *Sensors*, vol. 2, no. 11, pp. 436–472, November 2002.

[4] A. Mainwaring, D. Culler, J. Polastre, R. Szewczyk, and J. Anderson, "Wireless sensor networks for habitat monitoring," in *Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*. ACM Press, 2002, pp. 88–97.

[5] "Glacsweb," <http://www.glacsweb.org>.

[6] F. Ye, H. Luo, J. Cheng, S. Lu, and L. Zhang, "A two-tier data dissemination model for large-scale wireless sensor networks," in *Proceedings of the 8th annual international conference on Mobile computing and networking*. ACM Press, 2002, pp. 148–159.

[7] LAN MAN Standards Committee of the IEEE Computer Society, *IEEE Std 802.11-1999, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications*, IEEE, 1999.

[8] J. Kulik, W. Heinzelman, and H. Balakrishnan, "Negotiation-based protocols for disseminating information in wireless sensor networks," *Wirel. Netw.*, vol. 8, no. 2/3, pp. 169–185, 2002.

[9] C. Intanagonwiwat, R. Govindan, D. Estrin, J. Heidemann, and F. Silva, "Directed diffusion for wireless sensor networking," *IEEE/ACM Trans. Netw.*, vol. 11, no. 1, pp. 2–16, 2003.

[10] W. Ye, J. Heidemann, and D. Estrin, "An energy efficient MAC protocol for wireless sensor networks," in *In proceedings of the 21st Conference of the IEEE Computer and Communications Societies (INFOCOM)*, vol. 3, June 2002, pp. 1567–1576.

[11] T. van Dam and K. Langendoen, "An adaptive energy-efficient MAC protocol for wireless sensor networks," in *Proceedings of the first international conference on Embedded networked sensor systems*. ACM Press, 2003, pp. 171–180.

[12] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "A survey on sensor networks," *IEEE Communications Magazine*, vol. 40, no. 8, pp. 102–114, August 2002.

[13] "Omnet++: Objective modular network testbed in C++," <http://www.omnetpp.org>.