A Case for Peer-to-Peer Network Overlays in Sensor Networks

Muneeb Ali and Koen Langendoen
Delft University of Technology
The Netherlands
{m.ali, k.g.langendoen}@tudelft.nl

Abstract

One of the primary roadblocks in pushing sensornet technology into the everyday lives of people is to provide mechanisms for *co-existence* and *seamless integration* of sensornets with other (more established) networks like the Internet, ad-hoc networks, wireless infrastructure-based networks, and personal area networks. In this paper, we argue for P2P overlays in sensornets as we believe that they hold the potential to provide powerful common abstractions for such seamless integration, eliminate the need for specific infrastructure/proxy support, and could help moving towards a general-purpose architecture for a future world-wide sensor web. We debunk some myths about why P2P overlays are not feasible in sensornets, propose a Chord-based P2P protocol called *Tiered Chord* (TChord), and show how TChord could seamlessly integrate sensornets with IP networks.

1 Introduction

Culler et al. [14] argue that one of the main factors limiting research progress in sensornets is the lack of an overall sensornet architecture. The cross-layer designs in sensornets have lead to monolithic, vertically integrated solutions, which might work independently, but are not really useful for other research groups. Developing a sensornet architecture would be a growing and organic process. In this paper we focus on the interaction of sensornets with IP networks from an architectural point of view.

When sensornets go on-line at a large scale, the number of sensors and actuators connected to the Internet would outnumber the traditional Internet hosts. These tiny networked sensors connected to the Internet can not be treated as peripheral devices, as proxies or sinks connecting them will eventually become a bottleneck on performance and scalability. Sensornets will need to become a part of the Internet core itself and there is a need to eliminate sinks and proxies from the design of sensornets [13]. We propose using Peer-to-Peer (P2P) overlays over (traditional) sensornets to eliminate the need for proxy support and to enable flexible access to sensed data.

P2P design, for Internet-like environments, has been a very active research area and there are many P2P Internet protocols and systems available like CAN [31], Pastry [33], and Chord [34]. In spite of the potential overlap of the sensornet and P2P communities, to date, they remain isolated from each other and researchers in both communities have been investigating their fields separately without considering possible cross-cutting issues.

In this paper we outline the benefits of using P2P design in sensornets (Section 3.1), present popular arguments against P2P approaches in sensornets, show why these arguments are not valid (Section 3.2), present early hypothesis about design tradeoffs (Section 3.3), and propose a Distributed Hash Table (DHT) based P2P overlay for sensornets called *Tiered Chord* (TChord). We discuss TChord in the light of moving towards a sensornet architecture (Section 3.5), and present how our design can seamlessly integrate sensornets with the Internet and other networks (Section 3.6).

Gerla et al. argue for the applicability and transfer of wired P2P models and techniques to MANETs [19] - we extend their discussion to sensornet scenarios. Readers not familiar with P2P systems and DHTs are encouraged to refer to Lua and Crowcroft et al.'s recent survey on P2P [25].

2 Motivating Application

This paper looks beyond the "dumb" data collection applications of sensornets. In the last six years sensornet research primarily focused on data collection (e.g., environmental monitoring, and infrastructure monitoring). In typical deployments sensornets were treated as peripheral networks with the focus on getting data out of sensornets into the mainstream domain of the Internet or local area networks, for making it useful for end-users. Today, networking research in sensornets is entering a new era marked by a wide variety of applications (e.g., urban sensing, body sensor networks, disaster management, and healthcare) where end-users not only consume data but are actively engaged in *generating* data as well.

Sensornets will not only be deployed in remote or isolated areas (with a sink collecting data), but will enter the everyday lives of people in offices, residences, and public areas. This people-centric [6] focus on future sensornet applications changes several underlying assumptions of typical sensornets for example urban (instead of isolated) deployment, no pre-defined sink nodes, rechargeable nodes, more powerful (master) sensor nodes (Intel Imote2, and Sun SPOT), a variety of sensor types (mobile phones, video cameras, etc.), possible high data-rate applications, mobility (humans, cars, and buses), and multiple users of sensed data. In this paper, we address some challenges introduced by these new assumptions and focus primarily on providing mechanisms for *co-existence* and *seamless interaction* of sensornets with other (more established) networks like the Internet, ad-hoc networks, wireless infrastructure-based networks, and personal area networks (PANs).

3 P2P Overlays in Sensornets

In this section we motivate the use of P2P overlays in sensornets, and introduce our DHT-based P2P protocol for sensornets called *Tiered Chord*.

3.1 Benefits

P2P overlays in sensornets can be useful in several ways:

3.1.1 Efficient Data Lookup

The main purpose of sensornets is to collect data about a phenomena of interest; making sense out of this sensornet data, coming from potentially thousands or millions of sensornet nodes, is like trying to drink from a firehose. Looking up or reporting data of interest is one of the main operations in sensornets and optimizing this operation is a fundamental research problem. DHT-based P2P protocols provide near-optimum data lookup times for queries made on networks of distributed nodes [31, 33, 34] and mapping such protocols to sensornets could help solve the efficient data lookup problem in sensornets.

3.1.2 Guaranties on Lookup Times

To date, the primary design goal for sensornets has been energy efficiency. However, as new applications of sensornets emerge, other optimization criteria (or Quality-of-Service parameters) such as latency and compliance with real-time constraints (e.g., monitoring and control in industrial environments), or reliable data delivery (e.g., medical applications) gain importance [1]. So far, little attention has been paid to them in the context of sensornets. With DHT-based P2P overlays, bounded times for data lookups can be achieved in sensornets and guarantees can be provided to the applications running atop. Such bounds on data lookup could be used to provide QoS guarantees on other parameters like latency and reliable data delivery. One could argue that limits on number of overlay hops does not necessarily translate to latency bounds, however in sensornets (unlike the Internet) a P2P overlay can be structured according to the actual routing topology [2]. Furthermore, providing guarantees above the routing layer when the underlying sensornet layers (routing, link) are unreliable by nature, is possible, but will not be straight forward.

3.1.3 Location Independence

In sensornets Geographic Hash Tables (GHTs) [32] have been considered as an alternative to DHTs. With GHTs, as data within the network is stored according to the type of data, the queries for data could be directly sent to the node storing the named data. GHTs are inherently dependant on localization information which may or may not be available in sensornets. Also, GHTs require foreknowledge of the geographic boundaries of the physical application area and keys could potentially hash outside the geographic limits. We believe that foreknowledge of geographic boundaries is not a very realistic assumption for most future sensornet applications (where sensornets are formed on-the-fly in urban environments instead of isolated deployments in pre-defined areas). DHTs, unlike GHTs, are location independent.

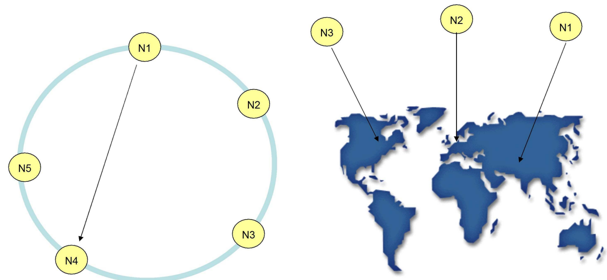

Figure 1. Virtual DHT ring and physical location

3.1.4 Overlay Applications & Services

DHT based data lookup operation is the only *core function* of P2P overlays. The data lookup operation of P2P protocols, like Chord [34] or Pastry [33], has been used to build a wide variety of successful applications and services including file systems [26], event notification [10], fair sharing of resources [27], and content distribution [9]. Some, if not all, of these applications and services (resource sharing, event notification, etc.) could be of great use in sensornets and the DHT data lookup function could serve as a natural *abstraction point* in the sensornet protocol stack; the applications running on top of the P2P overlay do not need to know how the data lookup function and the layers below it are implemented.

3.1.5 Proxies Considered Harmful!

In Section 1 we argued that sensornets will need to become a part of the Internet core itself and to achieve this we need to eliminate proxies/sinks from the design of sensornets. Furthermore, in urban sensing applications it is likely that a mobile end-user is well within communication range of some sensors of interest, but is out of communication range of the sink or proxy. To enable flexible access to sensed data, we will have to eliminate the need for pre-defined sinks and proxy support for sensornets; P2P overlays in sensornets could help achieve this goal as there is no concept of any "central authority" in P2P system design.

3.1.6 Limited Broadcast

Sensornets that do not require supporting infrastructure (sinks or proxies) are attractive because they would work even if the single supporting infrastructure fails or when it is not feasible (due to environment conditions or cost factors) to deploy permanent supporting infrastructure. To replace supporting infrastructure (like proxies) network designers generally use network-wide broadcasts to implement services - which is not efficient. Castro et al. argue that DHT solutions evenly distribute the load of building applications and services amongst participating nodes and could help achieve the goal of building services, with no central point of failures or need for supporting infrastructure, without global broadcasts [11].

3.2 Debunking Some Myths

In this subsection we present and debunk the main arguments against P2P designs in sensornets.

3.2.1 *Logical Topology $\neq$ Physical Topology*

A generic mapping of DHT-based P2P protocols to sensornets is considered difficult; as DHT protocols typically interconnect nodes independently of their proximity in the physical network topology (Figure 1), which is not suitable for energy-constrained sensornets as neighbors in the DHT logical identifier space may actually be far apart and each logical hop within a DHT based overlay may cost energy of many packet transmissions [32].

However, connecting overlay nodes independently of physical proximity is *not* a fundamental requirement for constructing P2P overlays. There have been works on assigning the Internet P2P neighbors according to proximity [8]. Ali et al. [2] show how to construct Chord-like virtual rings, in sensornets, with each neighbor being the physically closest node.

3.2.2 *Route Maintenance Overhead*

It is believed that in sensornets, particularly in large-scale sensornets, maintaining routing information among all pairs of nodes becomes expensive [32]. We give a counter example to this belief; VRR [5] demonstrates an implementation of a DHT-inspired routing protocol directly on top of the link layer in sensornets. VRR enables both (traditional) point-to-point routing and DHT routing to nodes responsible for respective DHT keys. VRR nodes are organized into a virtual ring (not mapped to physical location) and every node maintains a small number of routing paths to its neighbors in the ring. Experiments on a Mica2dot testbed indicate that VRR outperforms other routing protocols (e.g., BVR [18]) while nodes are able to route packets between any pair of nodes in the network [5].

3.2.3 *Sensor Nodes are Not Named*

DHT-based protocols use the IP-address of each node in the Internet, for obtaining unique node identifiers, whereas the sensornet literature presents a view that individual sensors are generally not named but data/application attributes are used to identify nodes [22]. However, we observe that in deployments (both testbed and real) generally a large (e.g. 64 bit) unique address is almost always used [24, 35]. The literature suggests that the globally unique address is only used for administrative and debugging tasks (configuring the network, monitoring individual sensors, and downloading binary code to specific nodes) but we observe that, in order to reduce system complexity, sensornet developers use unique addresses for "normal" operations (data collection, event notification, etc.) of the network as well.

A popular argument against using unique addresses in sensornets is energy waste by including (large) addresses with each packet. However, it is possible to reduce the address size of network-wide unique sensornet addresses [28] or to construct network-wide unique addresses dynamically from small, locally-unique addresses [3].

3.2.4 *DHTs are Computationally Intensive*

Contrary to popular belief, computing DHTs is not a very computationally intensive process and is well within the processing capabilities of sensornet platforms, for example VRR [5] was implemented on Mica2dot - which is one of the early generation sensornet platforms. The next generation sensornet platforms (e.g., Intel Imote2 and Sun SPOT) are moving towards 32-bit processors, memory in the order of 512KB, and storage in the order of 1024KB [4] - which is more than enough to support distribued hashing operations needed for P2P overlays as demonstrated by Caesar et al. [5]. Gupta et al. [21] demonstrate that even Elliptic Curve Cryptography (which requires more complicated hashing operations than P2P overlays) is feasible on standard Mote hardware. Furthermore, programming abstractions like Protothreads [16] can reduce the code size required to implement P2P overlays on sensor nodes with limited hardware resources.

3.3 Some Early Hypothesis

There are many open questions in designing P2P overlays for sensornets (see Section 4). We hypothesize that, instead of unstructured P2P systems like Gnutella and variants [12], mapping the design of structured DHT-based P2P systems, in general and Chord [34] in particular, to sensornets may prove more suitable. Unlike many P2P systems, Chord provides strong guarantees that are important for QoS parameters, and its lookup function runs in predictable time and always results in a success or definite failure. Some other systems, e.g. Oceanstore [23], provide stronger guarantees, but Chord is substantially less complicated and limited processing capabilities of sensornet nodes favor simpler protocols. Furthermore, urban sensing applications could have a significant number of mobile sensors and/or users and Chord is better in handling concurrent node joins and failures.

We believe that successfully recovering from node failures is an important performance metric (in urban environments with high network dynamics) and Chord helps in achieving that goal. Furthermore, the scalability of Chord is another factor that gives it an edge over systems like Gnutella and variants [12], that make widespread use of broadcasts.

One of the main arguments against structured approaches, like Chord, is that they perform worse than unstructured ones under Churn (high network dynamics, node joins and failures, mobility of nodes). However, recent research indicates that coping with churn is not a fundamental problem for structured overlays [7]. Structured overlays could achieve results better than, or equal to, unstructured ones even under a lot of Churn - while providing efficient and bounded data lookup times, unlike unstructured approaches [7].

3.4 TChord Protocol

Tiered Chord (TChord), is a simplified mapping of Chord [34] onto sensornets and the main purpose of this proposal is to serve as a *starting point* in designing P2P overlays for sensornets; with TChord we can begin experimenting with P2P overlays in sensornets, evaluate the effects of various design tradeoffs, and evolve TChord accordingly. Like Chord, at the heart of TChord is *one* main operation; the lookup operation. Given a set of sensor nodes, we hash, using SHA-1 [17], the unique address of each sensor node (e.g., 64-bit MAC address of Telos [30]) to obtain *node identifiers*. Meta-data keys, generated from the data stored on the nodes, are hashed to obtain *key identifiers*. Figure 2 shows

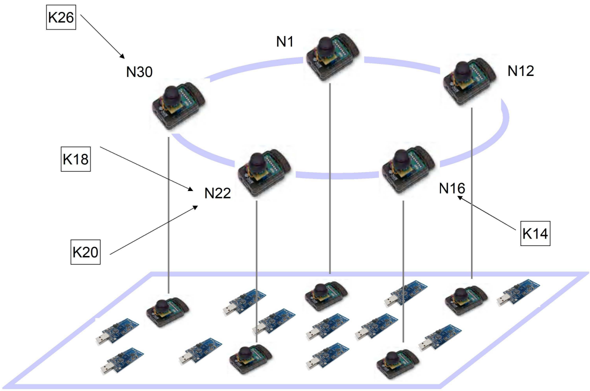

Figure 2. A TChord Master Ring in Sensornets

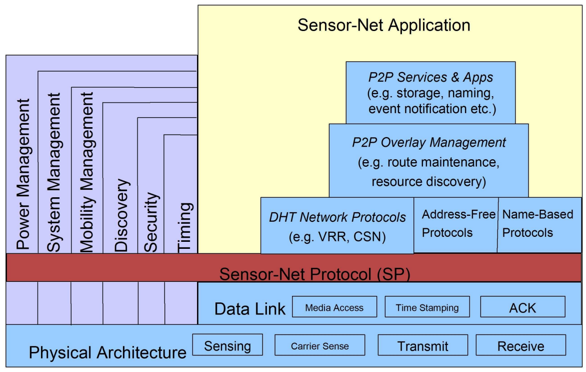

Figure 3. P2P Overlay in SP Architecture

a TChord ring arrangement of four master nodes with four keys mapped onto them. As meta-data keys are basically information about data, they are much smaller than the actual data itself and replicating meta-data keys amongst neighbors of a sensor node will not require a lot of storage.

The key identifiers are assigned to nodes in a manner, that *key identifier* $\leq$ *node identifier*. The sensor nodes are connected in a ring arrangement (Figure 2) and all messages are routed clock-wise. For details on how to form virtual DHT rings of sensor nodes, with each node being the physically closest neighbor, see [2]. In every ring it is necessary to have at least one high-powered master node. The master node maintains information (in its local finger table) about all its slave nodes and $O(log\, N)$ other master nodes. All queries are resolved in a distributed manner with a bound of $O(log\, N)$ messages. When a master node receives a query it first checks its own keys to resolve the query, if the lookup is not successful (note this means that the data element is not at the master node or any of its slaves) the master node then checks its local finger table. The finger table contains information about $O(log\, N)$ other master nodes and if the key can be located according to the information stored in the finger table, the query is directly forwarded to the master node storing the data. If the lookup on the local finger table also fails then the master node routes the query to the master node closest to the target according to the finger table.

When a master node joins a ring then the respective successor pointers need to be updated. Similarly, if a node voluntarily leaves the ring or fails then the node failure is detected and successor pointers are updated. Fault tolerance is handled by replicating the data of master nodes on neighboring masters.

Slave nodes do not store information about their neighbors. If a slave node directly receives a query, it checks its own data and if the lookup fails it simply forwards the query to its master node. For simplicity, in the TChord proposal we opt for not connecting the slave nodes in a ring arrangement and DHT lookups are not implemented in slave nodes (unless future experiment results prove otherwise). The master nodes of our proposal could be thought as “virtual sinks” with a DHT overlay between these virtual sinks.

In recent work, CSN [2] and VRR [5] take a comprehensive first step at designing DHT protocols for sensornets and TChord is similar to, and inspired by, CSN. However, unlike CSN our design is more generic (to support a variety of applications and services on top instead of just serving incoming data queries), and we assume the presence of more powerful master nodes (a tiered approach [20]). DHT lookups are performed only between the master nodes and not for every node. We do not opt for hierarchical clustering (which adds needless complexity to the design), and we consider parameters other than just energy-efficiency (e.g., latency). In VRR [5] the logical overlay is not constructed according to the physical location and this could lead to redundant packet transmissions and energy wastage (current experiments do not report on power consumption of VRR [5]).

3.5 Towards a Sensornet Architecture

P2P overlays in sensornets can help in the evolution of a general purpose sensornet architecture.

3.5.1 Co-existence with SP

The sensornet protocol (SP) by Polastre et al. [29] allows different MAC and link-layer technologies to co-exist by providing a standardized “narrow waist” interface to MAC, and provides an important step towards building a larger sensornet architecture. Unlike IP in the Internet, SP is not at the network layer but instead sits between the network and data-link layer (because data-processing potentially occurs at each hop, not just at end points).

Figure 3 shows how P2P overlays can be implemented on top of SP. The P2P overlay (shown as *P2P Overlay Management* in Figure 3) could be build on top of any generic network protocol. An underlying DHT routing protocol (e.g., VRR [5] or CSN [2]) is not necessary but recommended as it simplifies the job of overlay management and Caeser et al. show that it might be more efficient to build DHT-based routing directly on top of the link layer instead of implementing it as an overlay on top of a routing protocol [5]. *P2P Services & Applications* (e.g. event notification, resource allocation, and file systems) can then be built on top of the P2P overlay and sensornet applications could either use these services or communicate with the P2P overlay themselves.

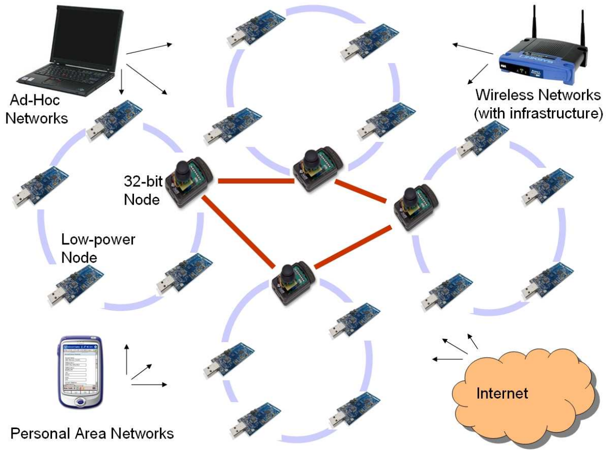

Figure 4. Interaction of sensornets, ad-hoc networks, and the Internet

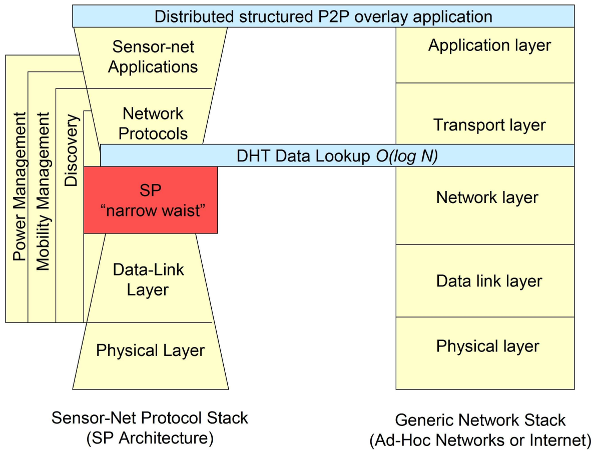

Figure 5. Integration of Sensornets with IP Networks

3.5.2 *Tiered Sensornets*

Early research in sensornets focused primarily on energy-efficiency and lead to application-specific designs, which perform data fusion as close to the source of data as possible. Gnawali et al. argue that performing such in-network processing on every sensor node increases the complexity of the sensornet; resulting in a system that is hard to program, debug, and maintain [20]. Instead of implementing complicated data fusion and application logic on every node, they argue for a *tiered architecture* where only a select group of more powerful nodes (masters) implement complicated in-network processing techniques [20]. As discussed in Section 3.4, our TChord P2P overlay design complies with this tiered approach to network design.

3.6 Putting the Pieces Together

Figure 4 shows our vision of future sensornets where P2P overlays (like TChord) seamlessly integrate sensornets with other networks. Dunkels developed a micro-TCP/IP stack (uIP) for 8-bit architectures [15] (a full implementation of TCP/IP) and tiny sensor nodes could run uIP to directly communicate with IP networks. However, uIP only provides means for communication between sensornets and other networks; we need a *common abstraction* that makes the underlying implementations and technologies transparent to distributed applications and services running over a variety of devices and networks like sensornets and the Internet. We propose that the DHT lookup operation could serve as this *single* common abstraction and a variety of applications and services could be build across sensornets, ad-hoc networks, and the Internet using this single common abstraction.

For example, in Figure 4 a PDA user (bottom left) within communication range of a sensor of interest could directly query sensors for data. With TChord, the sensor node will first look for the key in its own data elements (and in most cases reply with the data) or if the lookup fails then it will forward the query to its master node. Queries coming from the Internet (bottom right) or infrastructure-based wireless networks (top right) would be handled by the P2P overlay of master nodes (as described in Section 3.4). In this example, there are no proxies, no single points of failure, and as long as the device (be it a sensor node, a laptop with 802.11 connection, a remote Internet user, a PDA, or a small embedded device communicating over ZigBee) supports the DHT lookup operation it could participate in the network regardless of the underlying implementations and link-layer technologies.

In Figure 5 we present a general purpose framework for integrating sensornets with various IP networks. The left hand side of Figure 5 is a simplified SP stack [29], which was presented in more detail in Figure 3. The right hand side of the figure is the generic network stack used in various IP networks.

In IP networks the DHT data lookup operation is normally implemented as an overlay on top of the network layer [34] whereas in sensornets the DHT data lookup operation could be implemented directly on top of the link-layer [5] or on top of the network layer [2]. Regardless of how and at which layer the lookup operation is implemented, various distributed applications and services could be built above the DHT data lookup operation in a transparent manner. For example, geographic maps of metropolitan areas could be updated with traffic information (from wireless networks or GPS satellites), weather information (from the Internet or sensornets), parking spots (from live camera feeds or sensor information) and the distributed geographic map application does not care from which network the data is coming from; it simply uses the DHT lookup operation to locate the data of interest.

4 Open Questions

There are several open questions in the area of P2P for sensornets; P2P overlays (even if mapped to the actual routing topology) would have some overhead - do the benefits outweigh the costs? Which changing environmental parameters (e.g., query rate, replication rate, and QoS requirements), and network parameters (e.g., mobility, number of nodes, and network dynamics) to consider when designing lookup protocols for sensornets? Is structure required in sensornets?, and if yes, how much of it? Do general DHT approaches provide the right abstraction for a large class of sensornet applications, or do we need application-specific solutions? Can some applications, methods, and tools that are currently implemented for P2P systems be directly used on and/or mapped to sensornets? What new methods and tools are needed? If routing state is maintained in the network how does this scale under Churn? How do we evaluate the different approaches under realistic environmental parameters and workloads? The state-of-art in sensornet simulation (specially for simulating P2P characteristics in sensornets) is far from perfect; how do we achieve better validation of simulators and the underlying models?

5 Conclusion & Future Work

We showed that, contrary to popular belief, P2P overlays are viable for sensornets and their benefits (freedom from sinks/proxies, efficient data lookup, bounds on query times, location independence, limited broadcast) make them an attractive solution, especially when the goal is to integrate sensornets with other IP networks. We are planning an implementation and experimental evaluation of the proposed TChord protocol, on our TNOdes testbed, in order to answer some of the open questions presented in this paper.

Acknowledgements

The work was supported by Project "Relate" No. 013790 funded by the European Commission (FP6 IST Programme). The authors thank Thiemo Voigt for his helpful comments.

6 References

[1] M. Ali, U. Saif, A. Dunkels, T. Voigt, K. Römer, K. Langendoen, J. Polastre, and Z. A. Uzmi. Medium access control issues in sensor networks. *ACM SIGCOMM CCR*, 36(2):33–36, Apr. 2006.

[2] M. Ali and Z. A. Uzmi. CSN: A network protocol for serving dynamic queries in large-scale wireless sensor networks. In *2nd CNSR'04*, pages 165–174, Fredericton, N.B, Canada, May 2004.

[3] M. Ali and Z. A. Uzmi. An energy-efficient node address naming scheme for wireless sensor networks. In *IEEE INCC'04*, pages 25–30, June 2004.

[4] J. Beutel. Metrics for sensor network platforms. In *ACM RealWSN'06*, Uppsala, Sweden, June 2006.

[5] M. Caesar, M. Castro, E. B. Nightingale, G. O'Shea, and A. Rowstron. Virtual Ring Routing: Network routing inspired by DHTs. In *ACM SIGCOMM'06*, pages 351–362, 2006.

[6] A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, and R. Peterson. People-centric urban sensing. In *2nd ACM/IEEE WICON'06*, Boston, MA, Aug. 2006.

[7] M. Castro, M. Costa, and A. Rowstron. Debunking some myths about structured and unstructured overlays. In *2nd ACM/USENIX NSDI'05*, Boston, MA, May 2005.

[8] M. Castro, P. Druschel, Y. C. Hu, and A. Rowstron. Proximity neighbor selection in tree-based structured peer-to-peer overlays. Technical Report MSR-TR-2003-52, 2003.

[9] M. Castro, P. Druschel, A.-M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh. Splitstream: High-bandwidth multicast in a cooperative environment. In *19th ACM SOSP'03*, Bolton Landing, NY, Oct. 2003.

[10] M. Castro, P. Druschel, A.-M. Kermarrec, and A. Rowstron. Scribe: A large-scale and decentralized application-level multicast infrastructure. *IEEE JSAC*, 20(8), Oct. 2002.

[11] M. Castro, G. O'Shea, and A. Rowstron. Zero servers with zero broadcasts. In *ACM MobiShare'06*, pages 36–41, 2006.

[12] Y. Chawathe, S. Ratnasamy, L. Breslau, N. Lanham, and S. Shenker. Making gnutella-like P2P systems scalable. In *ACM SIGCOMM'03*, pages 407–418, Karlsruhe, Germany, Aug. 2003.

[13] D. Clark, C. Partridge, R. T. Braden, B. Davie, S. Floyd, V. Jacobson, D. Katabi, G. Minshall, K. K. Ramakrishnan, T. Roscoe, I. Stoica, J. Wroclawski, and L. Zhang. Making the world (of communications) a different place. *ACM SIGCOMM CCR*, 35(3):91–96, 2005.

[14] D. Culler, P. Dutta, C. T. Ee, R. Fonseca, J. Hui, P. Levis, J. Polastre, S. Shenker, I. Stoica, G. Tolle, and J. Zhao. Towards a sensor network architecture: Lowering the waistline. In *USENIX HotOS-X*, Santa Fe, NM, June 2005.

[15] A. Dunkels. Full TCP/IP for 8 bit architectures. In *USENIX/ACM MobiSys'03*, San Francisco, CA, May 2003.

[16] A. Dunkels, O. Schmidt, T. Voigt, and M. Ali. Protothreads: Simplifying event-driven programming of memory-constrained embedded systems. In *4th ACM SenSys'06*, Nov. 2006.

[17] FIPS 180-1. Secure hash standard. Technical report, U.S. Department of Commerce/NIST, National Technical Information Service, Springfield, VA, Apr. 1995.

[18] R. Fonseca, S. Ratnasamy, J. Zhao, C. T. Ee, D. Culler, S. Shenker, and I. Stoica. Beacon vector routing: Scalable point-to-point routing in wireless sensornets. In *NSDI'05*.

[19] M. Gerla, C. Lindemann, and A. Rowstron. P2P MANET's - new research issues. In *Dagstuhl Seminar Proceedings*, number 05152, Dagstuhl, Germany, Apr. 2005.

[20] O. Gnawali, K.-Y. Jang, J. Paek, M. Vieira, R. Govindan, B. Greenstein, A. Joki, D. Estrin, and E. Kohler. The tenet architecture for tiered sensor networks. In *4th ACM SenSys'06*, pages 153–166, Boulder, CO, Nov. 2006.

[21] V. Gupta, M. Millard, S. Fung, Y. Zhu, N. Gura, H. Eberle, and S. C. Shantz. Sizzle: A standards-based end-to-end security architecture for the embedded Internet. In *3rd IEEE PERCOM'05*, pages 247–256, 2005.

[22] J. Heidemann, F. Silva, C. Intanagonwiwat, R. Govindan, D. Estrin, and D. Ganesan. Building efficient wireless sensor networks with low-level naming. In *ACM SOSP'01*.

[23] J. Kubiatowicz, D. Bindel, Y. Chen, S. Czerwinski, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, C. Wells, and B. Zhao. Oceanstore: an architecture for global-scale persistent storage. In *ACM ASPLOS-IX*, pages 190–201, Boston, MA, Nov. 2000.

[24] K. Langendoen, A. Baggio, and O. Visser. Murphy loves potatoes: Experiences from a pilot sensor network deployment in precision agriculture. In *14th Int. Workshop on Parallel and Distributed Real-Time Systems*.

[25] E. K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim. A survey and comparison of peer-to-peer overlay network schemes. *IEEE Communications Surveys*, 7(2):72–93, 2005.

[26] A. Muthitacharoen, R. Morris, T. Gil, and B. Chen. Ivy: A read/write peer-to-peer file system. In *5th USENIX OSDI'02*, Boston, MA, Dec. 2002.

[27] T. Ngan, D. S. Wallach, and P. Druschel. Enforcing fair sharing of peer-to-peer resources. In *2nd IPTPS'03*, Feb. 2003.

[28] E. Ould-Ahmed-Vall, D. M. Blough, B. S. Heck, and G. F. Riley. Distributed unique global ID assignment for sensor networks. In *IEEE MASS'05*.

[29] J. Polastre, J. Hui, P. Levis, J. Zhao, D. Culler, S. Shenker, and I. Stoica. A unifying link abstraction for wireless sensor networks. In *3rd ACM SenSys'05*, pages 76–89, 2005.

[30] J. Polastre, R. Szewczyk, and D. Culler. Telos: Enabling ultra-low power wireless research. In *IEEE SPOTS/IPSN'05*, Los Angeles, CA, Apr. 2005.

[31] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Schenker. A scalable content-addressable network. In *ACM SIGCOMM'01*, pages 161–172, San Diego, CA, 2001.

[32] S. Ratnasamy, B. Karp, S. Shenker, D. Estrin, R. Govindan, L. Yin, and F. Yu. Data-centric storage in sensornets with GHT, a geographic hash table. *MONET*, 8(4):427–442, 2003.

[33] A. Rowstron and P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *IFIP/ACM Middleware'01*, pages 329–350, Nov. 2001.

[34] I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup protocol for internet applications. *IEEE/ACM Trans. Netw.*, 11(1):17–32, 2003.

[35] R. Szewczyk, J. Polastre, A. Mainwaring, and D. Culler. Lessons from a sensor network expedition. In *EWSN 2004*, Berlin, Germany, Jan. 2004.