

NASDI – Naming and Service Discovery for DTNs in Internet Backbones

Sebastian Schildt, Wolf-Bastian Pöttner, Oliver Ohneiser, and Lars Wolf

Institute of Operating Systems and Computer Networks,
Technische Universität Braunschweig, Germany
{schildt,poettner,ohneiser,wolf}@ibr.cs.tu-bs.de

Abstract. Delay Tolerant Networking (DTN) approaches based on the Bundle Protocol are commonly used within mobile IP based networks. Instead of being isolated applications, the Internet is often used to provide additional services or to route through other parts of the DTN network. A major drawback is that current DTN routing and discovery protocols are not generally applicable in the Internet as there is no common protocol to resolve DTN node names to convergence layer addresses outside a local network.

We present NASDI, an approach based on Distributed Hash Tables which can support naming, routing, notifications and service discovery in a heterogeneous DTN linked by the Internet. We present the architecture and initial evaluations of a NASDI prototype system we built for the IBR-DTN software.

1 Introduction

Delay Tolerant Networking (DTN) approaches replace the end-to-end semantics of common protocols such as IP with a hop-by-hop store and forward architecture [1]. Originally devised for interplanetary networks where nodes might see each other only occasionally, this approach has also been widely applied for ad-hoc networks with high mobility such as vehicular networks [2]. The Bundle Protocol [3] is a standardized widely used DTN protocol. It supports optional end-to-end acknowledgements on top of the hop-by-hop approach. In fact, the Bundle Protocol can be seen as a superset of IP (and TCP, as it includes elements from both the networking and the transport layer): In a continuously connected network it works much like the former, while in addition it is able to deal with disruptions of the network. This leads to the use of the Bundle Protocol in networks with mobile IP devices such as smartphones, which are regularly connected to the Internet [4] [5].

However, the Bundle Protocol ecosystem has a major shortcoming when it comes to operating in large-scale fully interconnected networks such as the Internet: There is no standard mechanism to find a node or the next hop for a specific DTN nodename. Compared to the standard IP architecture there is no standardized naming system such as DNS and there is no usable routing protocol to find the next hop to a destination if that hop is located in a so far unknown

network across the Internet. The various proposed routing protocols for DTNs normally assume an ad-hoc scenario relying on different forms of flooding and network discovery, both of which are not applicable in the Internet.

In fact the “DTNBone” [6], which is a collection of DTN nodes operated in the Internet by different institutions for DTN testing purposes, consists mainly of a webpage. This page contains (often inaccurate) information about which node can be reached at which IP address using which transport protocols. A DTN administrator who wants to connect with the DTNBone takes this information in order to configure a static route within his DTN server.

Therefore, we propose NASDI, a mechanism that allows for naming, service discovery and routing between DTN nodes operated in a large backbone network such as the Internet. NASDI is able to integrate peripheral networks and nodes which are only intermittently connected to the backbone and it allows nodes which are unaware of the approach to take advantage of the benefits. As DTNs contain intermittently connected nodes that may enter or leave the network at any time, it is our goal to be notified about such events even if the node in question is not located in our direct network neighborhood. To facilitate this NASDI offers an asynchronous notification mechanism. NASDI is based on Distributed Hash Tables (DHT) and is specifically adapted to the needs of DTNs.

In Section 2 we present some related work. Section 3 describes our system architecture followed by Section 4 describing the proposed algorithms. Section 5 introduces our implementation. Finally, in section 6 we wrap up our experiences with NASDI and line out future work.

2 Related Work

For LAN or small-scale ad-hoc DTNs there exists a standard mechanism to detect neighboring nodes and services: IP Neighbor Discovery (IPND) [7]. IPND works by regularly broadcasting beacons and thus it does not scale to the Internet. In [8] Waldhorst sketches Arriba, a general architecture for routing in overlay networks spanning heterogeneous technologies based on generic node ids. That work focusses on routing, but it does not specify how to create unique node ids and assign them to device names or underlay addresses. Closely related to the problem of node discovery is the problem of routing: Most general routing protocols proposed for DTNs are designed for ad-hoc type scenarios and as such are often variants of flooding like Epidemic routing [9] or PROPHET [10]. Other approaches exploit domain specific knowledge [11] and are thus not generally applicable.

Earlier DTN specifications included the concept of “regions”. DTN regions were a hierarchical naming concept for DTN nodes based on their network affiliation [12]. Current DTN specifications have abandoned this concept in favour of a more flexible flat URI based namespace. The idea is, that different networks can be identified by different URI schemes, but generally the usage of

URIs is meant to be much more open. The specification suggests things such as “expressions of interest” based URIs [1].

To allow hosts to find out the network layer address for a host name in the Internet the Domain Name System [13] is used. In addition to a number of shortcomings of traditional DNS in the Internet [14], it is also not optimal in a DTN. DNS systems are partitioned assuming a hierarchy of hostnames, which does not exist in the flat URI based namespace of DTNs. Furthermore DNS is not self-organizing but instead it involves significant organizational overhead. Also, DNS assumes that the network of servers is static and rather stable - a property, which can not necessarily be found in DTN. To overcome some of these problems, DDNS [15] has been proposed. It is based on a DHT as data storage and embeds a hierarchical namespace in the flat key space of a hash table. DDNS tries to mimic the behavior of DNS and is destined to be a DNS replacement. Although it is also based on a DHT, it has some significant differences to our approach. It just transfers DNS semantics to a distributed DHT store and thus lacks an asynchronous notification mechanism and support for groups.

As suggested in [4], DHTs might be a feasible way to tackle the naming problem in DTNs. DHTs are a robust way to store data in a distributed fashion. A DHT is a key-value store which is distributing the load evenly across participant nodes while still providing good lookup performance. Generally DHTs are resilient against node failure, have excellent scalability and support a flat name space. DHT implementations differ in the topology in which they organize participant nodes, and in the metric used to determine which key belongs to which node. One of the earliest well-known DHT is Chord introduced in the seminal paper by Stoica et al. [16]. Chord uses a ring topology. In addition to successor information, each node has a routing table which enables $O(\log(n))$ communication complexity while searching for a key. Among other well known DHT variants are Pastry [17], which tries to exploit local neighborhood information and CAN [18] which organizes data in a d-dimensional grid. A more recent DHT is Kademlia [19] which is widely used in the EMule and BitTorrent P2P networks. Thus, from all DHTs Kademlia has best proven the DHT’s alleged scalability and performance in real world scenarios [20]. For a more thorough DHT introduction see [21].

As lined out above, one of NASDI’s goals is to provide an asynchronous event notification mechanism. Several papers about publish-subscribe solutions based on DHTs have been published, such as the SCRIBE [22] system. It is based on Pastry [17] and allows topic-based subscriptions. The XEvent [23] publish-subscribe system is based on Bamboo [24] and allows either topic-based or event-content-based subscription using XPath expressions as filter. Both approaches target networks with a large number of subscribers and events and are therefore focused on strategies to efficiently and reliably deliver these events. Our use case has only a limited number of events that does not need such sophisticated structure. In addition, neither SCRIBE nor XEvent support the notion of asynchronous, cached events that will be delivered whenever a node rejoins the network.

3 System Architecture

We propose NASDI, a distributed system that can provide naming, routing and service discovery for Internet connected DTNs. NASDI’s goal is to help connecting to other DTN nodes from which only the name is known. As we will show, NASDI does not necessarily find the destination itself, but possibly only a suitable next hop that can be used to route to the destination. Due to the various, sometimes very application specific, routing protocols available for DTNs, unlike [8], NASDI does not try to replace former routing protocols in DTNs. It rather augments them by bridging the gap between separate DTN networks connected through the Internet that cannot discover each other by the conventional link local discovery mechanisms. For example, to reach a so far unknown node in another network, NASDI provides the connectivity information of a suitable node. This node can then be contacted and sent the bundle. Additionally routing packets for a mechanism such as PRoPHET might be exchanged with the newly discovered node to learn about the network behind it.




Fig. 1. Scenario Overview

An overview of a NASDI System is depicted in Figure 1. Apart from providing naming services to well connected IP capable DTN nodes, the NASDI architecture also allows to integrate peripheral networks that are not directly reachable or which use a non-IP transport layer for the Bundle protocol. These networks can be transparently proxied by gateway routers as explained in section 3.2. Additionally NASDI can also improve the connectivity of legacy nodes not supporting NASDI.

3.1 Assumptions

The following standard DTN terminology is required to understand the information managed by NASDI: Each DTN Node has one or more EIDs (Endpoint Identifier). EIDs identify a node or a service and have the form of URIs [1]. To connect to an EID within a DTN it is necessary to know which convergence layer, and which convergence layer address can be used to connect with the

node. Common convergence layers are TCP [25] and UDP [26] or the Licklider Transport Protocol [27].

Generally, any DHT structure can be used. For the notification we do however assume, that the location of a node in the DHT topology is deterministic, and that this position can be determined by any DHT member. For example a node's location within the DHT's topology can be determined by that nodes name. More specifically, for the notification to work we assume that each DHT member is responsible for a range in keyspace containing its own node id.

As mentioned before the EIDs used in Bundle Protocol [3] have the form of an URI. Commonly the *dtn://* schema is used. An example EID would be *dtn://node1/echo* which identifies the application "echo" on the node "node1". For the remainder of the text we assume that we use the EID's scheme and authority parts [28] as keys for the DHT, i.e. *dtn://node1* represents a node id and is used as key for the DHT. Other schemes might use a different mechanism to derive the DHT keys depending on the scheme's semantics.

3.2 DTN Node Roles

When NASDI is deployed each node can assume different roles in the NASDI system. The different roles are depicted in figure 1. First, it is to be expected that not all nodes will support NASDI:

- **Nasdi aware nodes:** Nodes, which implement NASDI (filled circles and diamonds in figure 1). These nodes can be DHT members, i.e. storing information for the DHT or query the DHT.
- **Nasdi unaware nodes:** Nodes, which do not know about the NASDI mechanism (empty circles in figure 1). Lots of unaware nodes are to be expected before this approach is widely adopted within the DTN community. NASDI unaware nodes can still be announced or proxied within the DHT by a DHT member. They cannot, however, query the DHT for information. They can benefit from NASDI by routing through a DHT member.

NASDI aware nodes can choose to become a **DHT member** (diamonds in figure 1). A DHT member is responsible for storing data which is assigned to it by the DHT implementation. Additionally it is responsible for regularly refreshing the DHT information of nodes it proxies or announces in the DHT. Therefore a node with high mobility or insufficient network connection might choose not to join the DHT. For nodes which have contact information in the DHT, two kinds of information can be stored in the DHT:

- **Announced nodes:** Announced nodes are nodes, whose convergence layer information is stored within the DHT. The convergence layer information stored in the DHT points to the node itself, i.e. contains its current IP address. A node can announce itself in the DHT if it is a DHT member, or it can ask a DHT member to announce it.
- **Proxied node:** Nodes in networks that are not accessible via IP from the Internet. They need a DTN router in order to participate in the global

network. Reasons for unreachable nodes might be firewall or NAT router. Proxied nodes are nodes which have the convergence layer information of another DTN node stored in the DHT. The proxy can also be used as a gateway between different underlying network technologies such as IP and ZigBee. Another rationale behind proxying nodes is, that a node might only be intermittently available which leads to frequent DHT updates and inaccurate information. Instead, storing the convergence layer information of a node that is more likely to be online, allows other DTN members to route bundles in the correct direction while the proxy node is in a good position to relay the information to the target as soon as it is available.


Please note, that a node may be announced directly and at the same is being proxied by others.

4 DHT Information Management

This section details the information stored in the DHT and the steps needed to maintain and query the DHT. We assume that the DHT provides a method `DHT_ROUTE(MSG, KEY, VALUE)` which delivers a message of type *msg* with content *value* to the node(s) responsible for the partition of key space containing *key*.

4.1 Information Stored in the DHT

The storage at DHT member nodes is assumed to be a set val_{stored} of $(key, value)$ tuples. For a given key a number of values can be concatenated, which is needed for group management and is also a way to deal with duplicate names: As the Bundle Protocol forces no structure on the EID naming space, it is valid and to be expected that for example multiple `dtm://test` nodes will join the network. Replacing is not an option because we do not want malicious or similarly named nodes to expunge valid entries from the DHT. Security critical applications which need to certify the identity of the other node, can use the Bundle Protocol



(a) *val_{stored}* single entry

(b) *notify_{pend}* entry

Fig. 2. Data stored in the DHT

Security extension [29]. Therefore, spoofing other nodes in the DHT does not pose an additional risk.

Figure 2a shows an example entry for a stored node. The *key* is the key used as address in the DHT and it is derived from a node's id by SHA-1 hashing it. The *type_entry field* denotes whether this is a group or single node entry (the Bundle Protocol allows the same form of id to be used for either a group or a node). The *type_information_list* contains the IP address and port numbers of the TCP and/or UDP convergence layer for a single node entry, or a list of hashed node ids for a group entry.

Different timers are used to determine when to expunge an entry and to assess the freshness of the data:

- time_to_live (*tll*): This timer determines how long this entry is considered to be valid. The initial value is determined by the node publishing the key into the DHT. The node that stores this entry decrements it. This entry is measure how long the contact information in this entry is assumed to be valid
- time_since_last_seen (*tls*): Even if the *tll* is very high (e.g. for announcing a stationary node), the entry should be refreshed periodically. The *tls* counts how many seconds have past since the entry was last updated.
- time_refresh_passive (*trp*): The *trp* value is constant and indicates how often the publishing nodes intends to refresh the entry. If $tls > trp$ this means that an entry was not refreshed within the expected time. This can indicate that the publishing node has connectivity problems.

4.2 DHT RPCs

On top of the DHT the following message types have to be implemented by DHT members. We assume that a *key* (which is used for DHT routing) and a *value* is associated with each message.

- GET: Standard DHT operation. Returns all entries for *key*. The *value* parameter is ignored for this call.
- STORE: Standard DHT store. Stores *value* associated with *key*. Existing entries for *key* are extended.
- JOIN_GROUP: Allows to augment information stored for a *key* describing a group. Creates a new val_{stored} group entry if the group does not exist so far.
- LEAVE_GROUP: Deletes the node in *value* from the group entry designated by *key*. Does not touch other entries for *key*.
- NOTIFY_REQUEST: Indicates that the node id contained in *value* wants to be notified when a modification is done to key *key*. A user should be able to specify whether this should be a “one-shot” notify, i.e. whether the notification request should be cleared after the notification is fired the first time or whether this should be a permanent notification request. See also section 4.3.
- NOTIFICATION: This message contains a notification for the key *key*. If the current node's id is *key*, the notification is forwarded to the application layer. Otherwise the (*key*, *value*) tuple is stored.

The processing of these messages is shown in algorithm 1.

Algorithm 1. Process messages

```

1: procedure PROCESSMESSAGE(msg, key, value)
2:   if msg = GET then
3:     return  $\{(k, v) \mid (k, v) \in val_{stored} \wedge k = key\}$ 
4:   else if msg = STORE then
5:     entry  $\leftarrow \{(key, v) \mid (key, v) \in val_{stored}\}$ 
6:     if entry ==  $\emptyset$  then
7:       entry  $\leftarrow \{(key, value)\}$ 
8:     else
9:       MERGE_SINGLE(entry, value)
10:    end if
11:    valstored  $\leftarrow val_{stored} \setminus \{(k, v) \mid k = key\}$ 
12:    valstored  $\leftarrow val_{stored} \cup entry$ 
13:  else if msg = JOIN_GROUP then
14:    entry  $\leftarrow \{(key, v) \mid (key, v) \in val_{stored}\}$ 
15:    if entry ==  $\emptyset$  then
16:      entry  $\leftarrow \{(key, value)\}$ 
17:    else
18:      MERGE_GROUP(entry, value)
19:    end if
20:    valstored  $\leftarrow val_{stored} \setminus \{(k, v) \mid k = key\}$ 
21:    valstored  $\leftarrow val_{stored} \cup entry$ 
22:  else if msg = LEAVE_GROUP then
23:    entry  $\leftarrow \{(key, v) \mid (key, v) \in val_{stored}\}$ 
24:    if entry  $\neq \emptyset$  then
25:      entry  $\leftarrow$  REMOVE_FROM_GROUP(entry, value)
26:      valstored  $\leftarrow val_{stored} \setminus \{(k, v) \mid k = key\}$ 
27:      valstored  $\leftarrow val_{stored} \cup entry$ 
28:    end if
29:  else if msg = NOTIFY_REQUEST then
30:    notifypend  $\leftarrow notify_{pend} \cup \{(key, value)\}$ 
31:  else if msg = NOTIFICATION then
32:    if key = myid then
33:      NOTIFY_APP(value)
34:    else ▷ Indirect notification
35:      valstored  $\leftarrow val_{stored} \cup \{(key, value)\}$ 
36:    end if
37:  end if
38:  if msg  $\neq$  NOTIFY
39:    and msg  $\neq$  NOTIFICATION then
40:      CHECK_NOTIFY(key) ▷ See algorithm 2
41:    end if
42: end procedure

```

4.3 Asynchronous Notification

In a mobility enabled DTN network nodes might not be reachable at all times. This is a standard case in DTN networks and participating nodes keep bundles for an unreachable destination until a suitable next hop becomes available or the bundle expires. However, it is beneficial if the node storing the bundle is notified as soon as the destination becomes available again. This can be implemented using the DHT. To support notifications a DHT member node maintains a second set *notify_{pend}*. A *notify_{pend}* entry is depicted in figure 2b containing the following items:

- key: This is the DHT key of the node, we want to receive notifications about
- number_of_notifications: How often this notification should fire. Typical values are 1 or ∞ . For 1 the event fires once, and afterwards the *notify_{pend}* entry will be deleted, for ∞ the event will fire every time its triggering conditions are met.
- event: Defines, which kind of event triggers this notification. Possible triggers are the reappearance of a node, the change of any value in the *val_{stored}* entry for *key*, or the change of a specific value.
- interest: The key of the node that wants to receive a notification when this events fires.

To demonstrate the different steps of establishing a notification request and the further processing, we will look at an example from figure 1:

Assume that in Figure 1 the node *dtm://mobile1.y.dtm* is proxied by *dtm://gw.y.dtm*. When *mobile1.y.dtm* becomes unavailable this will be detected by *gw.y.dtm* and the corresponding entry will expire in the DHT. However, foreign nodes might still be sending bundles for *mobile1.y.dtm* to *gw.y.dtm* because they used a cached older entry with a higher *tll*, or the bundles have been send before the DHT entry expired. Thus *gw.y.dtm* stores a NOTIFICATION_REQUEST for *mobile1.y.dtm* into the DHT, issuing the DHT command

```
DHT_ROUTE(NOTIFY_REQUEST, dtm://mobile1.y.dtm, dtm://gw.y.dtm)
```

This request will be routed to the same node that is responsible for the key *dtm://mobile1.y.dtm* in the DHT keyspace. Whenever a DHT member receives a call that creates or modifies tuples in its *val_{stored}* it will check whether there are any notification requests pending for the modified key (see algorithm 2). In our example *mobile1.y.dtm* joins another network and gets itself proxied by *dtm://gw.z.dtm*. This means *dtm://gw.z.dtm* will issue a STORE to the DHT:

```
DHT_ROUTE(STORE, dtm://mobile1.y.dtm, conv_layer(dtm://gw.z.dtm))
```

The node responsible for the key *dtm://mobile1.y.dtm* will check its pending notifications for this key and finds *gw.y.dtm*. Instead of directly trying to contact *gw.y.dtm* it will use the DHT:

```
DHT_ROUTE(NOTIFICATION, dtm://gw.y.dtm, conv_layer(dtm://gw.z.dtm))
```

If *gw.y.dtm* is online, this has the same effect as contacting *gw.y.dtm* directly, because we assumed that we use a DHT where each node is responsible for a range in keyspace containing its own node id (see sect. 3.1). If *gw.y.dtm* is currently not available, the notification is stored on another node currently responsible for the key *dtm://gw.y.dtm*. Once *gw.y.dtm* becomes available again and

Algorithm 2. Check and transmit pending notifications

Check whether there are pending notifications for EID

```

1: procedure CHECK_NOTIFY( $EID$ )
2:    $targets \leftarrow \{target \mid (EID, target) \in notify_{pend}\}$ 
3:    $data \leftarrow \{(k, v) \mid (k, v) \in val_{stored} \wedge k = EID\}$ 
4:   for all  $target$  in  $targets$  do
5:     DHT_ROUTE(NOTIFICATION,  $target$ , ( $EID$ ,  $data$ ))
6:   end for
7: end procedure

```

rejoins the DHT, depending on the DHT, the mechanism of the underlying DHT will hand over the data for its chunk of the keyspace, including the notification. This ensures, that receivers are notified as early as possible.

5 Implementation

While NASDI is not a routing protocol in a strict sense, it could be implemented as such using the routing module interface of DTN2 [31]. For IBR-DTN [30] a new discovery module is the best choice for integration. IBR-DTN allows different submodules to plug into its event-based core. DTN2 offers an XML based interface for external routing implementations. In our implementation we opted to use the Maidsafe library¹ which provides a Kademlia implementation including NAT traversal capabilities. The NASDI implementation is an external program using Maidsafe which communicates via TCP IP with a new IBR DTN discovery module that acts as a wrapper for the external NASDI implementation. This setup allows for great flexibility while developing NASDI and should make it relatively easy to connect NASDI to DTN2's external interfaces later.

5.1 DHT Functionality Tests

While the performance of large scale Kademlia deployments has already been examined, e.g. in [32] we performed some small scale tests, to verify that the NASDI implementation is performing as expected. We used virtual machines running instances of NASDI and IBR-DTN. The first NASDI instance is always started standalone, while the following instance get one of the running instances as bootstrap partner.

Lookup Test. Figure 3 shows the time to lookup a value in the DHT. The diagram includes the min max and median values as well as the $Q_{0.25}$ and $Q_{0.75}$ quantiles. We modified the number of NASDI nodes and the amount of tuples stored in the DHT. When the number of stored elements is increased from 1 to 1000 the average response time goes up from ~ 70 ms to ~ 90 ms, which shows

¹ <http://code.google.com/p/maidsafe-dht>




Fig. 3. Lookup times

the additional processing overhead in the NASDI instances. The variance for each measurement is due to the fact, that the DHT structure is different between runs, so that the responding Maidsafe instance might be nearer or further away. The rightmost plot shows the situation, when a node can answer the query from its local storage without the need to contact other DHT members.

Notification Delay Test. For this test we used 9 NASDI instances. The IBR-DTN node *node1* was started sending a bundle to *node2*, which was currently not available. This leads NASDI to store a notification request. Subsequently we started IBR DTN node *node2*. The NASDI instance for *node2* announces its contact information in the DHT, which leads to a notification being dispatched to *node1*, which in turn connects to *node2*, delivering the stored bundle. We measured the time between starting of *node2* and the reception of the bundle. This took around 3 seconds. A breakdown of the used time can be seen in figure 4. As can be seen in this case the notification itself is nearly instant, while the biggest amount of time is spent in the IBR-DTN daemon getting the cached bundle from storage and preparing it for transmission.




Fig. 4. Notification latencies

6 Conclusions and Future Work

We presented NASDI, an approach that allows DTN nodes connected to the Internet to efficiently store and retrieve convergence layer information of other DTN nodes in a distributed manner. NASDI supports building groups and integrates an asynchronous event notification mechanism. While NASDI's DHT member nodes should be located in the Internet and should be chosen in such a way that a long uptime can be expected, peripheral, possibly non-IP, intermittently connected networks can benefit from NASDI whenever they have a connection to the backbone. Those peripheral networks can be transparently bridged through the Internet. NASDI can be easily implemented as a routing module for widely used DTN implementations. Deploying NASDI is simple, because it can coexist with NASDI unaware nodes and still provide its benefits.

During the implementation and evaluation of NASDI we identified some areas for further improvement: We think that the services offered by NASDI can be very beneficial when they are used on a large scale in DTN implementations. To reach this goal, NASDI functionality should be shipped and enabled by default in Bundle Protocol implementations. We choose Maidsafe as basis for our DHT, which proved to be a very versatile library. However, its huge size and various dependencies may be a negative point when trying to integrate it with IBR-DTN or DTN2. This is especially a problem for IBR-DTN which is optimized for embedded devices. Therefore, we are currently looking into more lightweight DHT implementations.

We have only been able to perform tests with 9 nodes. While this shows the systems works as advertised, it makes it hard to predict how the system would perform with thousands of nodes. This also reveals a problem when the system is deployed initially: If there are only a few nodes operating in the NASDI DHT the overall resilience and reliability of the system might be less than in our controlled lab experiments. To solve this, we are currently looking into the possibility of leveraging the DHT subsystems of popular filesharing applications such as BitTorrent or the eMule network for NASDI. This means having less control over the DHT implementation, which could mean that some NASDI functionalities can not be implemented as efficiently or fully featured as outlined in this paper. On the other hand at any given time thousands of nodes will be online and available in the DHT, which should make the overall system very robust.

Acknowledgments. This work has been supported by the NTH School for IT Ecosystems.

References

1. Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., Weiss, H.: RFC4838 - Delay-Tolerant Networking Architecture. RFC (2007), <http://tools.ietf.org/pdf/rfc4838.pdf>

2. Lahde, S., Doering, M., Pöttner, W.-B., Lammert, G., Wolf, L.: A practical analysis of communication characteristics for mobile and distributed pollution measurements on the road. *Wireless Communications and Mobile Computing* 7(10), 1209–1218 (2007)
3. Scott, K., Burleigh, S.: RFC5050 - Bundle Protocol Specification. RFC (2007), <http://tools.ietf.org/pdf/rfc5050.pdf>
4. Ott, J.: Application protocol design considerations for a mobile internet. In: *Proceedings of First ACM/IEEE International Workshop on Mobility in the Evolving Internet Architecture, MobiArch* (2006)
5. Caini, C., Cornice, P., Firrincieli, R., Livini, M., Lacamera, D.: DTN meets smartphones: Future prospects and tests. In: *5th IEEE International Symposium on Wireless Pervasive Computing, ISWPC* (2010)
6. Delay Tolerant Networking Research Group, DTN-Bone, <http://www.dtnrg.org/wiki/DtnBone>
7. Ellard, D., Brown, D.: DTN IP Neighbor Discovery (IPND). Internet-Draft (2010), <http://tools.ietf.org/pdf/draft-irtf-dtnrg-ipnd-01.pdf>
8. Waldhorst, O.P.: On Overlay-Based Addressing and Routing in Heterogeneous Future Networks. In: *2010 Proceedings of 19th International Conference on Computer Communications and Networks (ICCCN)*, pp. 1–8 (2010)
9. Vahdat, A., Becker, D.: Epidemic Routing for Partially-Connected Ad Hoc Networks. Duke University, Tech. Rep. CS-200006 (May 2000)
10. Anders, L., Avri, D., Olov, S.: Probabilistic Routing in Intermittently Connected Networks. *SIGMOBILE Mobile Computing and Communication Review* 7(3), 19–20 (2004)
11. Doering, M., Pögel, T., Wolf, L.C.: DTN Routing in Urban Public Transport Systems. In: *ACM MobiCom 2010 Workshop on Challenged Networks (CHANTS)*, Chicago, USA (September 2010)
12. Fall, K.: A delay-tolerant network architecture for challenged internets. In: *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)* (August 2003)
13. Mockapetris, P., Dunlap, K.J.: Development of the domain name system. *SIGCOMM Computer Communication Review* 18(4), 123–133 (1988)
14. Balakrishnan, H., Lakshminarayanan, K., Ratnasamy, S., Shenker, S., Stoica, I., Walfish, M.: A Layered Naming Architecture for the Internet. In: *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, SIGCOMM* (August 2004)
15. Cox, R., Muthitacharoen, A., Morris, R.T.: Serving DNS Using a Peer-to-Peer Lookup Service. In: Druschel, P., Kaashoek, M.F., Rowstron, A. (eds.) *IPTPS 2002*. LNCS, vol. 2429, pp. 155–165. Springer, Heidelberg (2002)
16. Stoica, I., Morris, R., Liben-Nowell, D., Karger, D., Kaashoek, M., Dabek, F., Balakrishnan, H.: Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications. *IEEE/ACM Transactions on Networking (TON)* 11(1), 17–32 (2003)
17. Rowstron, A., Druschel, P.: Pastry: Scalable, Decentralized Object Location, and Routing for Large-Scale Peer-to-Peer Systems. In: Guerraoui, R. (ed.) *Middleware 2001*. LNCS, vol. 2218, pp. 329–350. Springer, Heidelberg (2001)
18. Ratnasamy, S., Francis, P., Handley, M., Karp, R., Shenker, S.: A scalable content-addressable network. In: *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)* (August 2001)

19. Maymounkov, P., Mazières, D.: Kademia: A Peer-to-Peer Information System Based on the XOR Metric. In: Druschel, P., Kaashoek, M.F., Rowstron, A. (eds.) IPTPS 2002. LNCS, vol. 2429, pp. 53–65. Springer, Heidelberg (2002)
20. Wang, C., Yang, N., Chen, H.: Improving Lookup Performance Based on Kademia. In: Proceedings of Second International Conference on Networks Security Wireless Communications and Trusted Computing (NSWCTC), vol. 1 (2010)
21. Steinmetz, R., Wehrle, K. (eds.): Peer-to-peer systems and applications. Springer-Verlag New York, Inc., Secaucus (2005)
22. Castro, M., Druschel, P., Kermarrec, A.-M., Rowstron, A.: SCRIBE: A large-scale and decentralized application-level multicast infrastructure. *IEEE Journal on Selected Areas in Communications* 20(8), 1489–1499 (2002)
23. Wang, R., Rao, W., Zhang, C.: XEvent: An Event Notification System over Distributed Hash Table (DHT) Networks. *IEEE Intelligent Informatics Bulletin* 6(2), 19–25 (2006)
24. Rhea, S., Geels, D., Roscoe, T., Kubiawicz, J.: Handling churn in a DHT. In: Proceedings of the USENIX Annual Technical Conference (USENIX) (June 2004)
25. Demmer, M., Berkeley, U., Ott, J.: Delay Tolerant Networking TCP Convergence Layer Protocol. IETF Draft (2008), <http://tools.ietf.org/pdf/draft-irtf-dtnrg-tcp-clayer-02.pdf>
26. Kruse, H., Ostermann, S.: UDP Convergence Layers for the DTN Bundle and LTP Protocols. IETF Draft (2008), <http://tools.ietf.org/pdf/draft-irtf-dtnrg-udp-clayer-00.pdf>
27. Ramadas, M., Burleigh, S., Farrell, S.: Licklider Transmission Protocol - Specification. Experimental RFC (2008), <http://tools.ietf.org/pdf/rfc5326.pdf>
28. Berners-Lee, T., Fielding, R., Masinter, L.: RFC3986 - Uniform Resource Identifier (URI): Generic Syntax, RFC (2005), <http://tools.ietf.org/pdf/rfc3986.pdf>
29. Symington, S., Farrell, S., Weiss, H., Lovell, P.: Bundle Security Protocol Specification, IETF Draft (2010), <http://tools.ietf.org/pdf/draft-irtf-dtnrg-bundle-security-17.pdf>
30. Schildt, S., Morgenroth, J., Pöttner, W.-B., Wolf, L.: IBR-DTN: A lightweight, modular and highly portable Bundle Protocol implementation. *Electronic Communications of the EASST* 37, 1–11 (2011)
31. DTN2 implementation, <http://sourceforge.net/projects/dtn/>
32. Ou, Z., Harjula, E., Kassinen, O., Ylianttila, M.: Performance evaluation of a Kademia-based communication-oriented P2P system under churn. *Computer Networks* 54(5), 689–705 (2010), <http://dx.doi.org/10.1016/j.comnet.2009.09.022>