Always best (dis-)connected: challenges to interconnect highly heterogeneous networks

Daniel Rodríguez-Fernández¹, Isaias Martinez-Yelmo^{1,2}, Ellen Munthe-Kaas¹, and Thomas Plagemann¹

¹ Department of Informatics, University of Oslo
Gaustadalléen 23 D, N-0373, Oslo, Norway

{dani, imyelmo, ellenmk, plageman}@ifi.uio.no

² Universidad Carlos III de Madrid
Av. Universidad 30, 28911, Leganés, Madrid, Spain
imyelmo@it.uc3m.es

Abstract. Wireless networks enable mobility, multihoming, and Delay Tolerant Networks. In such networking environments, the principles of the Internet, i.e., the end-to-end principle and the combination of location and identification in IP addresses, cannot be applied. In this paper, we propose a scalable application centric approach for mobility and multihoming that is able to interconnect highly heterogeneous networks where the networks may belong to different networking paradigms, e.g., IP based and Delay Tolerant Networks. Applications and users that aim to communicate, form communities. Community members might together have several network technologies available, and the community layer manages internetworking information for the members to seamlessly integrate these. Networking adaptation layers are used to provide a common interface for the different networks to the community layer. Addressing is based on names, cryptographic identifiers, and network locators and such that identifiers can also be created in infrastructure-less and disconnected situations.

1 Introduction

Mobile devices with multiple wireless networking capabilities, like IEEE 802.11, 3G and Bluetooth, have become mainstream devices in the recent years, and their popularity will increase also in the future. Applications running seemingly smoothly over these different networks give the impression that the Internet architecture is well suited for future wireless and mobile computing. Unfortunately, this impression is not correct. Basic mobility support required a patch to the Internet architecture, i.e., Mobile IP. Mobile IP enables the user to be reachable via the same IP address while roaming, through the help of a Home Agent. Since the Home Agent is a part of the network it conflicts with the fundamentals of the Internet architecture, i.e., the end-to-end principle. The basic problem that causes this conflict is the fact that IP addresses are used as locators and identifiers at the same time. The combination of identification and location causes also substantial problems for seamless integration of different networking technologies at the end host, i.e., for multihoming, because each network interface requires its own IP address. There are several ongoing efforts in the IETF to develop new patches to the

Internet architecture to handle multihoming, but so far no solution has been adopted and is mature enough.

Another challenge to the Internet architecture is introduced by Mobile Ad Hoc Networks (MANETs) and Delay Tolerant Networks (DTNs). DTNs are infrastructure-less networks, which are so sparse that end-to-end connections cannot be established. Therefore, DTNs apply the store-carry-forward principle, i.e., nodes carry messages for some time until they meet other nodes that can be used to forward the message one hop. Obviously, such an approach is contradicting to the end-to-end principle and IP cannot be used in DTNs. Thus, DTNs represent a new networking paradigm and are clearly separated from MANETs. MANETs are also infrastructure-less networks, but they obey the end-to-end principle and are based on IP. In MANETs, it is assumed that route breaks caused by mobility can be quickly fixed through discovery of new routes and re-routing. We believe that the strong separation between MANETs and DTNs is wrong, because both network types are infrastructure-less and are used either due to necessity, i.e., no network infrastructure exists, or due to the explicit choice of the user. Examples for the first case include sensor networks for wildlife and environmental monitoring, emergency and rescue operations in areas with destroyed infrastructure, and also military applications. Examples for the second case include Vehicular Networks which do not want to rely on cellular networks like 3G due to the unacceptable end-to-end delay between neighbouring vehicles, and social content distribution networks that either want to save money (roaming costs) or want to establish a network that nobody can control, like floating content [1]. The question whether MANETs or DTNs should be deployed in a certain setting depends on the density and mobility of the nodes. MANETs might become quickly partitioned and nodes in DTNs might be close enough to each other to form a MANET with end-to-end connection. Multihoming and a seamless integration of the different networking paradigms would enable applications that are always best (dis-)connected in highly heterogeneous networks. Therefore, full advantage and resource optimisation of available communication capabilities from underlying layers could be obtained for applications and services in any network (dis-)connected node.

The contributions of this work comprise the identification of the challenges introduced by mobility, multihoming and different network paradigms and a proposal to address these challenges. One of the main characteristics of our proposal is that it is application and user centric compared to related network centric solutions.

The rest of the paper is structured as follows. In Section 2, we detail the challenges that mobility and multihoming over highly heterogeneous networks present. The state of the art in mobility, multihoming and heterogeneous internetworking is discussed in Section 3. In Section 4, we present our proposal for solving the aforementioned problems. Section 5 concludes the paper, discussing the differences between the proposed approach and related work, and pointing out our future work.

2 Challenges

The proliferation of wireless technologies has broken many of the assumptions made in the design of the current Internet. On the one hand, the Internet was designed with a much more reliable wired network in mind, so some protocols as TCP may be inefficient

for wireless communication. On the other hand, wireless technologies have facilitated node mobility and multihoming through the access to different network technologies, such as e.g. IEEE 802.11, Bluetooth, and 3G. These two novelties have revealed a fundamental problem in the design of IP, i.e. the dual role of the IP address, which acts as node identifier and network locator. Due to this design issue, IP cannot provide native support for node mobility and multihoming. In Section 3, we present several approaches for identifier and locator split, which in turn introduce many new problems.

The application of wireless technologies has pushed to the limits the utilisation of networking technologies. In so-called challenged networks [2], it is necessary to deal with long delays and network disruptions, breaking some of the fundamental assumptions made during the Internet protocol suite design. In such networks, it is not possible to rely on the classical end-to-end principle widely used in the Internet. This has caused the apparition of a new networking paradigm, i.e. DTN. In this section we describe the challenges that arise from the conjunction of mobility and multihoming support when internetworking highly heterogeneous networks that might be based on different networking paradigms.

Wireless technologies have inherent problems which are not usually observed in wired networks. The transmission medium causes a higher probability of transmission errors due to interferences or collisions when several nodes try to transmit in a free transmitting slot. These effects are caused by the existing limitations in the wireless medium and can consequently not be removed. In addition, such undesirable effects might be increased by other phenomena such as hidden nodes. Since these effects were not taken into account in the design of the Internet protocol suite, the usage of TCP/IP on wireless technologies leads to a suboptimal utilization of the available resources. The current layered approach of the Internet does not offer enough flexibility to take full advantage of the available resources. This has resulted in a number of different cross-layering optimisation proposals [3].

Wireless technologies have enabled device mobility. At the same time, the diversity of wireless technologies has caused the proliferation of multihomed devices, i.e., devices which can use more than one wireless technology to communicate. Neither host mobility nor host multihoming are supported in the current IP. This is due to a fundamental architectural problem of the current Internet, where IP addresses have a dual role as both host identifiers and network locators. At the application and transport layers, IP addresses are used to identify a host, so all information flows are attached to them. At the network layer, IP addresses are network locators, indicating points of attachment to the network. In order to achieve global scalability, IP addresses are aggregated reflecting the hierarchical organisation of the network. Therefore, if host mobility causes a change in the point of attachment to the network, the host appears as a different node to the rest of the network. In this case, the host should reset the information flows in order to continue with the communication. A similar problem comes from multihoming, where a host appears as different nodes to the network. The hosts should decide upfront which IP addresses to use for communication. In order to provide native support in the network layer for mobility and multihoming, there are many proposals that rely on the separation of the roles that the current IP address possesses. In Section 3, we describe some of those approaches.

The Internet was designed to interconnect sub-networks, which may have different underlying networking technologies. During the design of the Internet, the end-to-end principle was proposed, where complexity is moved to the hosts, while the network is kept as simple as possible. The end-to-end principle works fine in presence of constant end-to-end connectivity, low probability of error, and small end-to-end delays. However, the end-to-end principle cannot be successfully applied in all kinds of networks, e.g. in challenged networks [2]. Such networks require a different networking paradigm, i.e. DTN. This is based on the store-carry-forward principle, where the intermediate nodes need to play a much more important role than in the Internet. It should be noted that from the DTN perspective, node mobility and multihoming could be considered more as an opportunity than as a challenge, since they allow new alternative paths to deliver information if they are properly handled.

The internetworking of heterogeneous networks is a complex problem that poses many issues. According to [4], networks can vary in type of service, routing and forwarding protocols, addressing scheme, packet size, support for QoS, flow and congestion control, security, accounting, etc. An internetworking solution needs to take into account all the differences. In this paper, we also consider the integration of different networking paradigms, i.e. IP networks and DTNs, which adds additional challenges. The internetworking should be aware of the networking paradigm of the underlying networks and perform according to them. For example, in order to provide interoperability between the Internet and a DTN, the internetworking system should be able to work in DTN mode, avoiding or replacing the usage of the end-to-end principle.

Current mobility and multihoming solutions relying on IP may not work properly over DTNs since IP assumptions are incompatible with the DTN paradigm. Any solution should nonetheless take into account all challenges that emerge when the identifier-locator split is done. The mapping between identifiers and locators must also work in DTN mode. One issue is that, at times, it may not be possible to update the mapping information due to network partitions. Thus, the mapping system should be able to work with partial knowledge. On the other hand, the mobility of nodes could be beneficial since node mobility may allow communication among otherwise isolated partitions (akin to message ferrying).

In the following section, we present the state of the art of solutions that consider some of the aforementioned challenges.

3 State of the art

In Section 2, we introduced the problems that the current IP has due to the dual role of the IP address. Unfortunately, IPv6 does not offer a native solution, having the same problems as IPv4. There has been an effort at the IETF to give solutions to issues like mobility, multihoming, security, etc. However, instead of providing a general solution, each of these solutions focuses on a specific challenge, leading to incompatible partial solutions. A network based solution for multihoming is LISP [5]. LISP does not require modifications on the hosts, only on network routers. The problem with LISP is that it presents scalability issues when deployed on the Internet. Other proposals are host based, requiring modifications only on hosts and a few nodes in the network. This

is the case of Mobile IP (MIP) [6, 7], SHIM6 [8], and HIP [9], which target mobility, multihoming and security, respectively, but with extensions that can be used to solve other issues. HIP provides a clear solution, not just to security but also to mobility and multihoming. HIP adds a new layer between network and transport, and introduces a new cryptography-based identity namespace. Identities are composed by a public-private key pair, where the public key is used as host identifier (HI). Upper layers use HIs, while IP continues using IP addresses which maintain their topological information. All these approaches are designed to solve problems of IP (IPv4 or IPv6), trying to patch the current Internet rather than giving new architectural designs. All assume that IP is used as network protocol. Thus, they are making the same assumptions as IP.

In parallel, the IETF DTN-RG has proposed the Bundle Protocol (BP) [10, 11], which aims to provide a general solution to DTN. The bundle protocol creates an overlay composed by BP agents that interconnect different DTN regions. A convergence layer is used to encapsulate the native network stack of each region. Security is an optional extension defined in [12]. The naming proposal is intentional naming [13]. It is very flexible, but it is unclear how to do routing and how to support multihoming and mobility. BP has been criticised that it lacks reliability and error detection support, relies on synchronised time among BP agents, lacks a clear naming scheme, and lacks support for application requirements [14].

There has also been a significant amount of effort on the definition of the Future Internet. Many of the proposals give a clean slate design, reviewing the fundamental principles of the Internet. A general survey in the area can be found in [15]. Due to space limitations we can mention only a few of them. The Ambient Networks [16] propose a flexible framework to create the so-called ambient networks, which are formed spontaneously by the nodes that have some access to networking technology. The project aims to fulfil the always best connected vision [17]. The proposal of SpoVNet is also interesting, by providing a generic framework to create overlays. The internetworking functionality is provided by the Underlay Abstraction, described in [18]. However, neither Ambient Networks nor SpoVNet consider the problems that may arise from the integration of DTN. Despite the large amount of work, there are not many proposals that consider the problem of internetworking of highly heterogeneous networks, where different networking paradigms need to be integrated.

PONA [19] proposes a separation layer that performs the ID-locator split at different levels, allowing host, user, and data IDs. It proposes a hierarchy of realms that reflects logical organisations, e.g. administrative or commercial organisations. A hierarchy of connectivity zones reflects the physical network connectivity. PONA does not propose a specific ID definition, allowing different types depending of the type of realm. It also allows users to specify requirements by providing policies. This vision also allows the possibility of supporting DTN relying on the realm structure.

MEDEHA/HENNA [20, 21] is a proposal for internetworking of highly heterogeneous networks. It can interconnect infrastructure networks and MANETs, and also can work in a DTN-like mode. MEDEHA was originally designed to work over IPv4. However, HENNA provides an ID-locator split which can be used to enhance MEDEHA. HENNA provides support for IPv4, IPv6 and DTN network locators. It works in a similar way to MIP, extending it to support DTN. The Location and Management Server

(LMS) plays a similar role to the Home Agent on MIP. HENNA has two main problems: (1) its lack of support for multihoming, allowing just one network locator for each node identifier, and (2) the construction of the identifiers. Node IDs are constructed by concatenating the LMS locator on the Internet and a local label managed by the LMS. This causes many security and scalability problems. HENNA makes many assumptions about the LMSs, which should be static and publicly available at all times. Furthermore, the node IDs depends on a specific locator (IP address) of the LMS. Thus, the ID-locator split problem still remains at the LMS level (e.g. making impossible to exploit LMS multihoming).

The Phoenix architecture [22] aims to fulfil the communication requirements after a disaster, on a so-called Day After Network (DAN). The idea of such networks is to integrate all available communication means after a disaster and use them to offer services to the rescue team, victims of the disaster, etc. The proposal relies on two protocols: the Phoenix Interconnectivity Protocol (PIP) which enables the creation on islands of connectivity through heterogeneous networks, and the Phoenix Transport Protocol (PTP) which enables delay tolerant communication between partitions. The proposed naming system is based on roles, such as firemen, police, etc. However, apart from the project description, there is not much information available about these two protocols.

4 Community Internetworking

In this section we present our proposal to solve the challenges presented in Section 2. In order to take full advantage of present and future networking heterogeneity, it is necessary to consider the following requirements. Firstly, it is vital not to make strong a priori assumptions about the underlying networks. Consequently, the internetworking system needs to be aware of the characteristics of these networks, by some kind of network description and corresponding ontology. The internetworking further needs to be aware of the changing status of the network and the nodes resources. It is also necessary to support node multihoming and mobility. An application should be allowed to influence the internetworking, describing the service that it expects from the network. Finally, security should be a main concern of the internetworking system.

Taking into account the aforementioned requirements, we propose community internetworking. The basic idea behind our proposal is to provide internetworking to a set of nodes that want to communicate with each other. This set of nodes is called *community* in our proposal. The nodes that compose a community may have access to several heterogeneous networks, so the community should handle this heterogeneity in order to perform the internetworking. Multihoming over highly heterogeneous networks can cause an explosion in routing information management, since nodes may be accessible through several networking interfaces. In order to avoid scalability issues, we propose that instead of managing all nodes in each network, just the community members are taken into account in each community. This allows performing a more sophisticated internetworking, because much less information needs to be managed. The communities are created to reflect a permanent or temporal relationship among the nodes that form them. The relationship can, e.g., reflect an organisation or a social network. It also can be created to support requirements of an application, as, e.g., a video streaming commu-

Fig. 1. (a) Community Internetworking example, and (b) Community Layer scheme

nity or a content sharing community. Another example can be a Personal Area Network, where the community integrates all the communicating devices currently used by a user (laptop, phone, PDA, etc.). We assume that the community receives a set of policies that reflects the requirements of the application and user. These policies determine the behaviour of the community. Node mobility and multihoming are supported through an identifier-locator split described in Section 4.1. Figure 1(a) illustrates how community internetworking reduces the size of the problem, as just the nodes that are members of the community are considered in the internetworking.

Our proposal relies on three main abstractions: *node*, *networking substrate*, and *community*. The *node* represents a device that has access to one or more networking substrata, and can be member of one or more communities.

The *community* abstraction represents a group of nodes that cooperate in order to fulfil some common objectives. Communities are formed to support the high level requirements of applications and users. These policies will be used to manage the community internals. Each community should perform the internetworking of its members over different networking substrata.

Finally, the *networking substrate* represents any kind of legacy or future network. Our proposal aims for minimising the assumptions about the internal characteristics of the networking substrata, which is just assumed to be able to provide some basic networking functionalities, such as transmission of packets and some kind of addressing. Thus, this abstraction can wrap many different types of networks, ranging from complex network overlays to simple communication services such as SMS, including direct communication over link layer protocols. The current Internet is one possible instance of a networking substrate. In order to be able to use the different networking substrata, we place on top of each of them a Networking Substrate Adaptation Layer (NSAL). The NSALs provide a common API to the communities by using the native primitives of each networking substrate. NSAL is further described in Section 4.4. However, the common NSAL API is not enough to perform a proper internetworking, it is also necessary to know the internal properties (addressing scheme, type of service, etc.) of each networking substrate. We obtain such knowledge from a self-description of the networking substrate, which is one of the key services provided by the NSAL.

4.1 Name-Identifier-Locator

Our approach for ID-locator split identifies three main entities: *name*, *identifier*, and *network locator*. Names constitute a high level abstraction, which can be resolved into identifiers, which are used to uniquely identify nodes and communities. Identifiers are mapped to the network locators that the nodes have on each of the networking substrata. The namespaces for names, identifiers and network locators are independent of each other.

Our proposal does not require names to be of a specific type. We believe that depending on the specific situation, different name systems might be used. The only requisite is that it must be possible to resolve names into identifiers. Two examples of name systems could be the fully qualified domain names provided by DNS, or the role based naming proposed in [22].

Network locator spaces are defined by each networking substrate. They may have meaning only inside the network substrate. For example, IPv4 and IPv6 would use IPv4 and IPv6 addresses respectively, while a DTN may use another type of network locators such as, e.g. GPS coordinates in geographic based DTN routing.

The common namespace for all the communities and nodes is the identifier namespace. We propose the usage of cryptographic based identifiers similar to what HIP proposes [9]. The identifiers are the public part of an *identity*, which is a public-private key pair. In our proposal, both nodes (hosts) and communities are associated with at least one identity, and therefore an identifier. Cryptographic identifiers have the following advantages: (1) They offer support for security and (2) can be auto-generated by each entity. In addition, (3) they are totally independent of names and network locators.

At the same time, they introduce some challenges since such identifiers do not provide any information about the location of the entities in the network substrata. Hence, we need a system that maps the identifiers to network locators in order to be able to start any communication with a node or community. An issue that arises from the usage of identifiers is privacy. In our proposal, nodes can have one or more identities. A node with two identities would appear to the rest of nodes as two different nodes.

4.2 Community Layer

The implementation of the community internetworking functionality is done in the community layer. As shown in Figure 1(b), this layer is introduced between layer 3 and layer 4. Nevertheless, the network substrate definitions allow the coexistence of very different networking substrata, ranging from link layer network technologies to complete protocol stacks. The community layer maintains on each node a registry that contains information about: (1) its identities, (2) its NSALs, and (3) the communities that the node is member of. The community layer is also responsible for coordinating the utilisation of shared resources on the node by different communities. The communities may, for example, have a priority that could be used to organise the usage of these resources.

As shown in Figure 1(b), a cross-layering information sharing component gives access to information from other components of the system. In order to achieve awareness, nodes, communities and networking substrata should be able to give a description of their characteristics. In addition to this description, information about the status of

the networking substrata and the node internal resources is needed. For this purpose, a monitoring framework component will be used. The monitoring framework is in charge of performing measurements about the status of nodes and networking substrata. This component is shared by all communities, leading to an efficient utilisation of the resources. For instance, if several communities require the same type of measurements, they need to be performed only once.

4.3 Community Management

Each node can create communities. We assume that, when a community is created, it is provided with a set of policies that reflects the preferences of an application or user. Community management and internal routing will be done according to those policies. The ownership of a community is given by the ownership of the private key of the community identity. The node that creates a community is therefore its owner. The owner may either delegate to other community members management responsibilities, or share the private key with other members of the community. This is decided by the community policies.

Another way to become member of a community is joining an existing community. Nodes can apply for joining a community, or they can be invited to join it. In both cases, community membership policies are checked before membership is granted.

The communities maintain the mapping between the identifiers of its members and their network locators within the different networking substrata. This mapping is performed by the mapping manager component. The implementation of this component depends of the community. For instance, in a small size community that should deal with disruptions, each node may maintain its own local copy. In a larger community where nodes are accessible through stable networks, the mapping manager may be implemented in a different form, e.g. by a distributed hash table, in order to achieve better scalability.

The inter-community routing is performed according to the community topology. This can be extracted from the mapping manager of each community. The selection of the next hop inside of the community is done by using the routing policies of the community. In this process, the NSAL and the network locators that will be used are selected. This last process might be too heavy for a per-packet processing. Thus, it may be possible to use a per-flow processing, where a temporal state about the open communication flows is used. This state may react to changes in the community topology and modify the next hop of the flow. Each hop at the community level may be done in a different networking substrate. This hop inside of a network substrate may potentially require more than one internal hop, which will be managed by the native routing and forwarding mechanisms of the networking substrate.

Communities might also consider the case when none of the available networking substrata can be used to send data to the destination. This could for example occur when nodes do not have a DTN substrate. In this case the community might decide, based on the community policies, to provide some basic DTN support, e.g. buffering data. The behaviour of this DTN support would be managed by the community policies.

4.4 Networking Substrate Adaptation Layer

The networking substrate provides a generic abstraction for any kind of legacy network. This allows the community layer to handle highly heterogeneous networks. There may be large differences between the technology, protocols, addressing, and even networking paradigm of the networking substrata. In order to reduce the complexity required at the community layer, we introduce the Networking Substrate Adaptation Layer, which wraps a particular networking substrate and provides a set of services through a common API to the community layer. One of the key elements is the networking substrate description. We assume a network definition language that can be used to communicate the networking substrate properties to the community layer. The last may be done by a network definition ontology.

The services provided by the NSAL API include services to (1) send data and (2) receive data, (3) resolve names to identifiers, and (4) resolve identifiers to network locators. These two functions can be combined to find communities by using a networking substrate. The first is used to, given a name, resolve the identifier of a node or a community, while the second is used to, given an identifier, resolve network locators to a node or a community in the networking substrate. And finally, (5) request a description of the properties of the path to the destination in the network substrate. The path description is used by the community layer to identify the best network substrate to forward data to. To implement the name-to-ID and ID-to-locator resolve services, the adaptation layer will use the native services offered by the network substrate, if available. For example, an IP NSAL could use DNS, while a DTN NSAL may use an epidemic dissemination. The send and receive services are mapped to the native send and receive primitives of the network substrate.

5 Conclusions

In this paper we present community internetworking, which is our proposal to support mobility and multihoming in highly heterogeneous internetworking. The basic idea behind our proposal is to provide internetworking just to the nodes that need to participate in a communication, following an application and user centric approach compared to other networking centric approaches. This decision reduces the size of the problem, allowing a better scalability. In our approach, we define three main abstractions: *node*, *community*, and *networking substrate*. A node may be member of several communities and have access to different networking substrata. Communities manage the internetworking among the nodes that are members of the community. The behaviour of a community is driven by a set of policies that reflects the preferences of an application and user. Finally, the networking substrate represents any kind of legacy or future network. It is integrated into the community internetworking by a Networking Substrate Adaptation Layer (NSAL) that offers a common API to the communities. In order to deal with the networking substrata heterogeneity we propose that each NSAL provides a description of the networking substrate that it wraps.

There have been several proposals to support mobility and multihoming, and to perform heterogeneous internetworking. However, just a few approaches consider the integration of Delay Tolerant Networking (DTN).

The DTN Bundle Protocol presents a set of problems as indicated in [14]. The current naming system [13] is quite imprecise about how mapping and routing should be done. Neither is there a clear way to support mobility and multihoming. Finally, the support for the application requirements is very limited.

The vision of the Future Internet provided by PONA [19] presents many similarities with community internetworking. For instance, the user requirements are specified through policies. Realms and zones have a similar function as communities and networking substrata. However, our concept of community is more general since it is not necessarily attached to an organisation. The same is the case of the networking substrata, which can be used to describe, not just the connectivity structure, but almost any kind of network, ranging from the Internet to a DTN, including MANETs. In our proposal we also present a clear identity system and a mechanism to support legacy or future networks by using NSAL and its network description functionality.

The MEDEHA/HENNA [20, 21] proposal for highly heterogeneous networks internetworking does not provide support for multihoming. Its ID-locator split scheme relies on the locators of a global Internet (LMS locator + a local label) and introduces the same ID-locator problems (mobility, multihoming, security) at the LMS level. Our proposal fully separates the identifier namespace from the specific network locator namespaces, in a similar fashion as HIP. MEDEHA/HENNA provides a limited support for the application requirements. There is no concept similar to the community, and therefore all the nodes need to be considered for the routing, which can lead to scalability issues.

The Phoenix architecture [22] targets DAN scenarios, while we propose a more general internetworking approach that covers a larger type of networks. Another difference is that Phoenix has a role based naming, whilst we propose the concept of community. Both role and community based approaches have as a goal to provide routing scalability. The community abstraction can be used to create role based communities, e.g. a police community. In our approach, application requirements are specifically supported through policies. Finally, [22] does not describe how node and network heterogeneity is managed, while our proposal relies on the node, community and networking substrate description, the last one provided by NSAL.

The ANA framework [23] provides very generic abstractions that can be used to implement different communication paradigms. It does not define how addressing, routing, etc. should be done, leaving this to the system designers.

Our ongoing and future work is to refine the design of the community internetworking presented in this paper and to implement a proof-of-concept prototype, for which we are considering to use the ANA framework. We plan to focus our first experimental studies on the integration of MANETs and DTNs.

6 Acknowledgements

This work has been funded by the VERDIKT Programme of the Norwegian Research Council through the DT-Stream project (project number 183312/S10).

References

1. Ott, J., Hyytiä, E., Lassila, P., Vaegs, T., Kangasharju, J.: Floating content: Information sharing in urban areas. In: Proceedings of the 2011 IEEE International Conference on Pervasive Computing and Communications (PerCom). PERCOM'11, IEEE Press (2011)
2. Fall, K.: A delay-tolerant network architecture for challenged internets. In: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications. SIGCOMM '03, ACM (2003) 27–34 ACM ID: 863960.
3. Srivastava, V., Motani, M.: Cross-layer design: a survey and the road ahead. Communications Magazine, IEEE **43** (2005) 112–119
4. Tanenbaum, A.S.: Computer Networks. 3rd edn. Prentice-Hall Inc. (1996)
5. Farinacci, D., Fuller, V., Meyer, D., Lewis, D.: Locator/ID separation protocol (LISP) (2011)
6. Johnson, D., Perkins, C., Arkko, J.: Mobility support in IPv6. RFC 3775 (2004)
7. Perkins, C.: IP mobility support for IPv4, revised. RFC 5944 (2010)
8. Nordmark, E., Bagnulo, M.: Shim6: Level 3 multihoming shim protocol for IPv6. RFC 5533 (2009)
9. Nikander, P., Gurtov, A., Henderson, T.: Host identity protocol (HIP): Connectivity, mobility, multi-homing, security, and privacy over IPv4 and IPv6 networks. Communications Surveys Tutorials, IEEE **12** (2010) 186–204
10. Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., Weiss, H.: Delay-tolerant networking architecture. RFC 4838 (2007)
11. Scott, K., Burleigh, S.: Bundle protocol specification. RFC 5050 (2007)
12. Symington, S., Farrell, S., Weiss, H., Lovell, P.: Bundle security protocol specification (2011)
13. Basu, P., Brown, D., Polit, S., Krishnan, R.: Intentional naming in DTN (2009)
14. Wood, L., Eddy, W., Holliday, P.: A bundle of problems. In: Aerospace conference, 2009 IEEE. (2009) 1–17
15. Paul, S., Pan, J., Jain, R.: Architectures for the future networks and the next generation internet: A survey. Computer Communications **34** (2011) 2 – 42
16. Niebert, N., Schieder, A., Abramowicz, H., Malmgren, G., Sachs, J., Horn, U., Prehofer, C., Karl, H.: Ambient networks: an architecture for communication networks beyond 3G. Wireless Communications, IEEE **11** (2004) 14–22
17. Gustafsson, E., Jonsson, A.: Always best connected. Wireless Communications, IEEE **10** (2003) 49–55
18. Bless, R., Hübsch, C., Mies, S., Waldhorst, O.: The underlay abstraction in the spontaneous virtual networks (SpoVNet) architecture. In: Next Generation Internet Networks, 2008. NGI 2008. (2008) 115–122
19. Paul, S., Jain, R., Pan, J., Bowman, M.: A vision of the next generation internet: A policy oriented perspective. In: British Computer Society (BCS) International Conference on Visions of Computer Science. (2008)
20. Rais, R.N.B., Mendonca, M., Turletti, T., Obraczka, K.: Towards truly heterogeneous internets: Bridging infrastructure-based and infrastructure-less networks. In: Communication Systems and Networks (COMSNETS), 2011 Third International Conference on. (2011)
21. Rais, R.N.B.: Communication Mechanisms for Message Delivery in Heterogeneous Networks Prone to Episodic Connectivity. PhD thesis, INRIA/University of Nice, Sophia Antipolis (2011)
22. Luo, H., Kravets, R., Abdelzaher, T.: The-Day-After networks: A First-Response Edge-Network architecture for disaster relief (2007)
23. Bouabene, G., Jelger, C., Tschudin, C., Schmid, S., Keller, A., May, M.: The autonomic network architecture (ANA). Selected Areas in Communications, IEEE Journal on **28** (2010) 4–14