A Framework for Multi-region Delay Tolerant Networking

Mirco Musolesi
Department of Computer Science
Dartmouth College
6211 Sudikoff Laboratory
Hanover 03755, New Hampshire,USA
musolesi@cs.dartmouth.edu

Cecilia Mascolo
Computer Laboratory
University of Cambridge
15 JJ Thomson Avenue
Cambridge CB3 0FD, United Kingdom
cecilia.mascolo@cl.cam.ac.uk

ABSTRACT

Almost all the existing work on routing in delay tolerant networks has focussed on the problem of delivery of messages inside a single region, characterized by the same network infrastructure and namespace. However, many deployment scenarios, especially in developing regions, will probably involve routing among different regions composed of several heterogeneous types of network domains such as WiMAX or satellite networks and ad hoc networks composed of short-range radio enabled devices, like mobile phones with Bluetooth interface.

In this paper, we introduce a proposal for inter-region routing based on both probabilistic and deterministic forwarding mechanisms, embedded in an architectural framework able to support it. We also compare our solution to existing approaches in delay tolerant networking, discussing the main requirements and possible solutions, and outlining the open research problems.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Distributed networks, Network communications, Wireless communication; C.2.2 [**Network Protocols**]: Routing protocols

General Terms

Algorithms, Design

Keywords

Delay tolerant networking, multi-region protocols, data replication, push/pull architecture, persistent caching, communication paradigms

1. INTRODUCTION

Networked systems in developing regions are usually characterized by intermittent and unreliable connectivity, especially in rural areas [2]. For this reason, the idea of using store-and-forward protocols as a basis for communication systems in developing countries has been widely investigated in the recent years [7, 10, 22].

However, almost all the existing approaches for routing in this class of networks, generally called Delay Tolerant Networks (DTNs) [3, 9], have focussed on the problem of message delivery inside a single region[1], characterized by the same network infrastructure and a single namespace. Many approaches have been proposed to solve the problem of intra-region routing [28]. Yet, many realistic deployment scenarios for DTNs may involve routing among different regions which are likely to be composed of heterogeneous networks.

Many application scenarios can be envisioned for multi-region delay tolerant networked systems. As a first scenario, let us consider the classic example of users wishing to send email from a mobile device or a laptop (not connected to any Internet access point) to a user in the Internet. Messages should be routed firstly to Internet gateways and then forwarded in the usual way through the Internet. In other words, the routing process consists in finding an inter-region gateway that connects the two regions. If no cellular coverage is available, one solution is the exploitation of opportunistic forwarding schemes in the ad hoc network among mobile devices based on short-range radio communication interfaces such as Bluetooth to reach an Internet gateway. An alternative solution is the design of a delivery system based on scheduled messages ferries. The routing protocol in delay tolerant mobile networks can be deterministic (e.g., a person in the village collects all the messages and then moves close to a wireless gateway and downloads all the messages) or probabilistic (e.g., the delivery can be based on an opportunistic forwarding scheme using the probability that a person gets in contact with the gateways, given that she travels around that area). A unified routing framework should allow for the integration of these different routing schemes.

Multi-region routing schemes can also be considered of key importance for designing systems to support healthcare initiatives. Many systems based on mobile phones for remote healthcare interventions have been presented in the recent years [8, 15]. This kind of systems can be extended to support and integrate local communication inside communities of people. Alerts or advices can be sent firstly to

[1]The current DTN architecture specification [3] has replaced the concept of region with the notion of URI scheme (only related to the naming aspects): in this paper, we use the generic term of region to indicate a network domain, characterized by a homogeneous network infrastructure, underlying routing mechanisms, and namespace.

WiNS-DR'08, September 19, 2008, San Francisco, California, USA.

e-kiosks in villages by means of the cellular infrastructure or by means of carriers such as vehicles providing a transitive link from the source (such as non-governmental organizations or healthcare authorities) to the e-kiosks that act as local distribution centers. Messages can be downloaded to the mobile devices of people passing nearby, using short-radio distance communication technologies. Then, opportunistic routing mechanisms can be used inside communities to spread these messages to the population. Possible solutions may include epidemic protocols [27] or socially-aware routing schemes [4, 12], i.e., forwarding protocols based on the social network structure of the individuals carrying the devices. These protocols can be very effective in places where social ties among members of communities are traditionally very strong. Moreover, in this case, the system should also support persistent caching and broadcasting of the messages for a certain interval of time on the e-kiosks in order to be able to spread the messages to the devices of people in their proximity (for instance, for the entire day).

Multi-region delay tolerant networked systems can also be used to support communication in case of emergencies or natural disasters. For example, central gateways can distribute message alerts to phones in coastal area in presence of tsunami (e.g., as an extension of the current Tsunami Warning System [20]). The alert can be broadcasted to all phones using the cellular network. These messages can be successively further spread by using epidemic protocols among the mobile phones, in a peer-to-peer manner. It would also make sense for the authorities to be able to use infostations connected to the Internet and with additional network interfaces such as Bluetooth and WiFi, given the increasing number of phones equipped with these technologies. Mobile phones nearby these infostations can then receive updates sent through the Internet and the Bluetooth link. Other network infrastructures may be employed at the same time, like WiMAX.

Finally, another application scenario can be a sensor network for environmental monitoring in rural areas for rain-fed agriculture [21]. Let us suppose that each sensor is equipped with a Zigbee interface. The sensor data need to be transmitted to farmers and researchers located in distant villages and cities. Sensors are both spread in the environment and attached to people (that may use vehicles) or animals. A certain number of specialized devices, also carried by humans or animals, can be used to retrieve data from the sensors. These mobile sinks will eventually transfer their data to infostations in the villages, for example by means of satellite communication. The infostations act as gateways between the two types of network. Alternatively, transport vehicles with scheduled routes can be used to transmit the data to the infostations. The infostations can be connected to a local wireless network and send this information to all the farmers in the village and, by using a possibly intermittently available connection, to researchers in nearby cities. An alternative practical scenario is the design of urban sensing systems, for example composed of a wireless network of sensors mounted on a public transportation system and the network of data sinks connected to the Internet [31].

In this paper, we propose the Network-aware Opportunistic Multi-region Asynchronous Delivery (NOMAD) framework to address the common design issues of these systems. The goal of this work is to discuss the requirements of a generic architecture that can encompass different types of communication paradigms and delivery mechanisms. In general, the design of a routing architecture between multiple regions can be decomposed into the problem of delivering a message to one or multiple hosts inside a certain region (*intra-region routing*) or to hosts inside other regions (*inter-region routing*). In the latter case, the routing problem consists in reaching the gateways (e.g., the mobile sinks in the example of the sensor network in the rural area) that can transfer the message to another region. Inside a region, a connected path may exist between the source and the destination: in this case the delivery process may rely on existing synchronous routing mechanisms, if available. However, given the inherent characteristics of these networks, if a connected path does not exist, the delivery process has to rely on *store-and-forward* mechanisms.

Another key aspect is the choice of the best message custodian for the message delivery if a connected path between the sender and the receiver does not exist. The choice of the best custodian can be based on the calculation of delivery probabilities. The framework that we present is protocol-agnostic; in other words, we assume that these delivery probabilities are computed using different routing mechanisms that may be present in the regions. As observed in the application scenarios discussed above, there is also a need for persistent caching and broadcasting.

The remainder of this paper is structured as follows. The general requirements and design issues of an architecture to support multi-region routing are presented in Section 2. In Section 3, we compare our approach with the existing work in this area. In Section 4 we outline the open research questions. Section 5 concludes the paper summarizing the contribution of this work.

2. MULTI-REGION ROUTING FOR DTNS

In this section, we discuss the problem of multi-region routing in DTNs. We introduce a model of the problem and decompose it into *intra-region routing* (i.e., routing inside a region) and *inter-region routing* (i.e., routing across regions). We then describe how we use delivery probabilities to implement the forwarding process inside a region and we discuss the naming system in NOMAD.

2.1 Modeling of the Routing Problem

In the context of DTNs, we reformulate the routing problem for both inter and intra-region communication as the selection of the best host(s) for the temporary storage of messages. We refers to these hosts as message *custodians*. Gateways can be fixed or mobile. Examples of gateways include mobile sinks, like DataMules [14], or wireless access points with storage connected to the Internet. The problem of enabling inter-region communication can be seen as the problem of reaching these gateways and then forward the messages beyond these points.

We assume that regional forwarding mechanisms are provided by existing underlying protocols like Internet protocols in the Internet domain and protocols designed for specific types of networks, like AODV [23] in MANETs or DTN protocols for intermittently connected networks. This model also includes the simple case where underlying protocols are not present, i.e., only one hop communication is possible.

We now illustrate all the possible cases of inter-region communication:

- **The sender and the recipient(s) are in the same region** In case the sender and recipient are (temporarily) disconnected, the best custodian is the host with the highest delivery probability (see algorithms used for example in PRoPHET [17] or CAR [19] for the calculation of the delivery probabilities).

- **The sender and the recipient(s) are in different regions** In this case, the routing can be seen as a two-step process:

  1. selection of the region(s) through which the recipient(s) can be reached;
  2. direct delivery of the message, if a connected path exists between the sender and the gateway for that region, or selection of the best host(s) which would allow to reach the gateway(s). Through the gateway, the region of the message recipient can eventually be reached. Clearly, the recipient can be in a neighboring region or in a non-adjacent region.

The problem of selecting the region (i.e., the gateway) can be solved through a rule-based selection mechanism. Mechanisms for the automatic exchange of routing information are not necessary, since the deployment scenarios are composed of a limited number of regions[2].

2.2 Deterministic vs Probabilistic Delivery

A recipient i of a message m can be a single host or a group of hosts. In other words, the name i refers to a *class* of hosts which can be a singleton set, in the case of unicast, or a set with a higher cardinality, in the case of multicast.

We associate a delivery probability $P_i(h)$ to every host h, i.e., every potential custodian. $P_i(h)$ provides a measure of the probability of h of being co-located with hosts belonging to a class i. Every host maintains information about the sets of potential custodians for a certain number of classes. This model can be applied to the delivery of a message to a class of final recipients in the region or a class of gateways used to reach another region. An example of a class of gateways is a set of mobile sinks used to collect data from a sensor network.

The calculation of $P_i(h)$ is based on routing mechanisms that may be different in different regions. We refer to the delivery of the message as deterministic, when it is not based on a predicted or potential connectivity, but it is rather "certain" (given the characteristics of the network infrastructure or, more in general, the *a priori* known or planned network connectivity) For example, the delivery probability $P_i(h)$ of a host h connected to a host i by means of a satellite link available during scheduled transmission slots will be equal to 1, since we know in advance that the delivery of the message will be possible in the future.

More formally, we identify two types of transitive delivery mechanisms: *probabilistic* delivery based on predicted or potential connectivity (with $P_i(h) \leq 1$) and *deterministic* delivery based on an expected and planned connectivity between the custodian and the recipient(s) (with $P_i(h) = 1$). With this model, we can treat and unify deterministic and probabilistic delivery mechanisms.

Let us consider a possible inter-region communication deployment scenario for data collection of medical or environmental information in a rural area. The communication can be supported by the following infrastructure: smart phones carried by people produce data messages that are collected by mobile sinks (e.g., buses moving in a rural area). These are downloaded onto Internet gateways that allow for delivering them to back-end servers on the Internet. This scenario is composed of three regions: i) the PDA network and the mobile sinks (that act as gateways), ii) the mobile sinks and the Internet gateways that communicate using a wireless connection and iii) the Internet itself.

For example, the data collection process by means of the mobile sinks in the PDA network can be based on the probabilistic CAR protocol [19]. The mobile sinks just upload the messages when they are in the transmission range of the Internet gateways (i.e., no multi-hop routing, just one hop communication). The communication model in the CAR region is 1-to-N with at least one semantics. Instead in the mobile sinks/Internet gateways region, the communication model is typically N to M. In the Internet region, we can assume that we have unicast communication between the gateway that receives the message from a mobile sink and the recipient of the message. In this region, we have a direct delivery of the message[3], without using transitive mechanisms.

The CAR protocol (used in the PDA region) is exploited to calculate the utility of each host to act as a carrier for messages to the mobile sinks (i.e., the delivery probability $P_i(h)$ described in Section 2). This allows to determine which host is the best custodian for reaching the mobile sinks. Instead, in the mobile sinks/Internet gateway region there are no intermediate hosts. In this case, the intermediate recipient of the messages is set to the name of the class of the Internet gateways. Once a message is with the mobile sinks, it will eventually get delivered to the Internet when the sinks travel to locations in reach of a gateway.

In the following section, we will discuss the possible intra-region communication models and how their composition helps in modeling more general inter-region communication in NOMAD.

2.3 Communication Models

In general, *intra-region message delivery mechanisms* can be classified according to the following paradigms:

- **1-to-1 Communication Model** This is the case of the classic unicast communication. The receiver may be a generic host inside the region but also a gateway. The class i is a singleton set.

- **1-to-N Communication Model** In this case we distinguish two cases based on the different communication semantics: 1 to (*at least 1 of N*) model when the message has to be successfully received at least by one of the recipients (like in the case of data sinks in sensor networks) and 1 to (*all of N*) model when the

[2] If this was instead necessary, the exchange of routing information could be based on the same communication framework described in this paper, with messages containing routing tables addressed to the other gateways.

[3] We assume that the message can be dispatched by means of a TCP connection or by using email protocols with a high reliability, also using retransmission mechanisms, if necessary. We note that also email services actually rely on store-and-forward mechanisms.

message has to be received by all the recipients (like in the case of multicast). A possible scenario is the delivery of emergency alerts generated in the Internet to be spread through the population beyond an Internet gateway with various means (e.g., Bluetooth).

- **N-to-1 Communication Model** This corresponds to the case where messages are collected by a host that acts as a sink/recipient of data generated by N sources. The typical example is a sensor network with one sink. The sink is the element of the singleton set of the class i.

- **N-to-M Communication Model** This can be seen as an extension and generalization of the previous models with multiple recipients. Also in this case, we can have either an at least one or all of M semantics.

By composing these intra-region communication models, different architectures for *inter-region routing* can be derived. For example, environmental data can be collected using one sink (N-to-1 communication) that acts as a gateway in a rural area. Then the data can be transmitted by a satellite link (1-to-1 communication) and then forwarded to group of M farmers. The overall inter-region communication model is N-to-M.

2.4 Naming and Region Transition

In NOMAD, each message contains the name of the recipient(s) (i.e., the name of the class of the final recipient(s)), the name of the region which the recipient(s) belongs to, and the name of the intermediate destination (gateway), if the recipient(s) is not in the current region. We assume that the names of the regions are unique, whereas the names of the hosts must be unique only inside the region, since the unique name of the region acts as namespace.

When a message is generated, if the recipient of the message is in the same region of the sender, the name of the intermediate region is left blank. Otherwise, this is set to the name of the gateway that allows for reaching the next region in the path to the destination. The name of the gateways can be set manually by the developers or retrieved by using discovery protocols i.e., the gateways can broadcast their names in the region periodically together with the names of the regions that can be reached through them.

When a gateway receives a message, if the neighboring region contains the recipient(s), it sets the intermediate recipient to blank. Instead, for instance, in the case of routing from region A to region C through an intermediate region B, the gateway through which the message enters region B inserts, as intermediate recipient, the name of the class of the gateway(s) that allows for reaching region C.

Finally, the communication between gateways can happen either through intermediate hosts that carry the messages (using the delivery probability to the gateway(s)) or, in case of intermittently connected gateways, this can happen directly without exploiting intermediate custodians.

2.5 Persistent Caching and Broadcasting

The system should be able to provide mechanisms for delivering the same information for a given period of time, such as an alert to all the devices of people walking in the proximity of a WiFi access point or a Bluetooth enabled access point. In order to support this functionality a caching mechanism should be provided by the gateways. Clearly, this kind of services can be implemented at application level, but we believe that the support of this design abstraction at network level can enormously simplify the development and, especially, deployment of these systems. Given the low cost and level of miniaturization of storage devices, the limitations in terms of buffer space does not represent a problem at least for fixed hosts. The framework should also allow for the specification of temporal constraints associated to messages. For example, it should be possible to indicate temporal validity and timeouts like in [25].

2.6 Push vs Pull Delivery

The forwarding algorithms discussed in the previous sections can be considered typical *push* delivery mechanisms. However, *pull* delivery mechanisms are very useful in certain types of deployment scenarios involving DTNs (especially those based on infostations with local storage).

A typical example is a digital kiosk with multiple interfaces connected to the Internet and equipped with various network interfaces, like Bluetooth and 802.11. Clients passing by can connect to pull information using different interfaces. Messages can be sent in *push* mode to a particular network, stored *temporarily* in the gateways and retrieved (*pulled*) by the hosts that request them.

A possible extension in this sense includes infostations inside the regions to which messages are sent from the gateways (push-mode)[4]. The infostations can act as temporary repositories of the messages which are then pulled by mobile hosts. The communication process between gateways and infostations can be based on the same deterministic and probabilistic mechanisms that are used to dispatch the messages inside the same region.

2.7 Replication Mechanisms

To increase the fault tolerance of the delivery process, messages can be replicated. Two kinds of replication mechanisms are possible:

- **Intra-region Replication** The set of potential custodians can be sorted according to their delivery probabilities and replicas of messages can be sent for example to hosts that are in the higher positions in the list or above a given threshold.

- **Inter-region Replication** This replication mechanism allows for multi-path routing through different regions. Copies of the messages are sent to gateways connecting the current region to a set of different regions. An example of application of this strategy can be the broadcasting of alerts and information bulletins in emergency situations by means of multiple network technologies. For instance, if governmental authorities need to reach the population after a terrorist attack, alerts can be replicated and sent using the cellular network, email and infostations equipped with Bluetooth and WiFi to try to reach all the people through the network infrastructures that are still in place.

[4]In other words, we envisage the possibility of having logical regions inside a physically homogeneous region, i.e., the framework allows for multiple logical regions from the naming (and routing) point of view in a single physical network infrastructure.

3. RELATED WORK

The NOMAD framework is inspired by existing solutions, in particular the DTN reference implementation [1]. A number of approaches dealing with DTNs [9, 13] and trying to overcome the limitations of synchronous forwarding have been presented, both in infrastructured and infrastructureless networks. In the area of mobile ad hoc networking, for instance, epidemic routing protocols [27] form the basis for much of the following work in the area. As discussed above, alternative solutions are based on probabilistic delivery based on previous contact patterns, such as in PRoPHET [17] or CAR [19] or based on scheduled ferries [29]. A survey of routing protocol for delay tolerant networks can be found in [28].

With respect to the application of DTNs for developing regions, in [5] Demmer and Fall present the DTLSR routing protocol, an adaptation of classic link state routing techniques for deployments in these geographical areas. The authors discuss multi-region routing as future work: their proposal is to use a protocol based on the exchange of summaries of the set of the reachable endpoints identifiers in order to extend their single-region routing solution to the case of multiple regions according to a typical link state paradigm, whereas in our approach we do not assume that every host of the network keeps information about how to reach endpoints in other regions.

In general, most of the DTN research has concentrated on intra-region and not on inter-region forwarding. An exception is the architecture proposed by Khaled and Almeroth in [11] addressing the problem of routing among different regions; more specifically, the authors consider a scenario with hosts belonging to several disconnected networks and communicating by means of scheduled messengers. This scenario can be modeled as part of NOMAD by associating at least one messenger to each region with a single intermediate host (i.e., the messenger itself) with a delivery probability to that region. However, as we have described, NOMAD is more general and it can be used in scenarios characterized by the presence of multiple regions. More recently, Kutscher et alii in [16] presented an architecture to connect different regions using unidirectional communication links such as satellite links.

4. OPEN ISSUES

In this paper, we have presented the requirements of a general framework for inter-region communication. This can be considered as an initial result of our investigation of the problem of communication in heterogeneous DTNs. In fact, our research agenda includes many issues:

Selection among Multiple Regions The selection among multiple regions in order to choose the best route through can be rule-based (i.e., static) or can rely on the evaluation of the current state of the system (i.e., dynamic) in terms of current resource availability, connectivity and security. An intelligent choice of the region can also be based on economic aspects. In fact, users may want to use the cheapest transmission media (for example, a buffer in a vehicle instead of a satellite link). Moreover, the selection of the regions (and, possibly, of particular custodians) may be based on the evaluation of the security requirements of the sender of the message. This information can be included in the bundles and evaluated when a routing decision has to be made. For example, hosts may not be selected even if they have the highest delivery probability for economic or security reasons according to the sender requirements.

Device Profiles Another dimension of the problem of message routing in DTNs is the description of the available resources. The description of the capabilities of the gateways is fundamental in terms of both current and future connectivity, available bandwidth and buffer space. We believe that a *device profile* should be associated to each host, and, in particular, to the gateways. In the case of mobile devices, the battery level should also be considered to avoid, for example, to select mobile sinks with low residual energy. Gateways with limited resources (buffer space or battery) or temporarily disconnected from other regions can advertise their current situation by means of these device profiles that can be broadcasted periodically to all the hosts in the network. The description of the available resources is also essential for the calculation of the delivery probabilities. The device description can also contain security and trust information (for example, the membership to a particular social group). We are currently investigating these aspects, also considering a possible declarative routing approach [18].

Metadata for Content Description This design aspect is, in a sense, complementary to the definition of device profiles described above. In fact, the bundle abstraction used in delay tolerant networks allows for very expressive description of data that can be used for various goals. Metadata can be used to specify the type of the data, their confidentiality or the priority in order to make routing decisions. For example, a security level can be associated to the data and highly confidential data might be routed only using specific types of networks (for instance, by avoiding untrusted carriers). In case of scarce resources, such as in presence of constrained memory space or bandwidth, higher priority data might be selected to be transmitted first.

Possible Integration with the DTN Architecture Our work is inspired by the DTN architecture and we believe that nodes based on the DTN reference implementation can be integrated easily in the framework. Nodes implementing the DTN standard specification [3] can be used, as gateways, since they implement the abstraction of convergence layers and, for this reason, they can be exploited to bridge different communication regions. Moreover, the DTN architecture specification distinguishes between routing and forwarding. According to the DTN specification, routing refers to the execution of an algorithm for computing routes so to achieve a certain objective defined by a function. Forwarding refers to the act of moving a message from one DTN node to another. Nodes based on the DTN architecture could be integrated by assuming that the forwarding decision is made by the gateways. Routing is provided inside the region using the delivery mechanisms described in Section 2. Our model allows for the integration of nodes for intra-region routing that can be based on different software architectures: the intermediate custodians do not need to be DTN architecture compliant, but they should only be able to communicate with DTN nodes (i.e., they must implement the bundle standard [24]).

Support for Multicast The semantics of multicast communication in delay tolerant networks presents many open research questions both from theoretical (e.g., naming and group membership issues) and implementation points of view. For example, the definition of the multicast bundle format, possibly using a bundle-in-bundle solution, is still an open issue. Interesting considerations are presented in [30]. In some cases, content based routing approaches represent an effective and interesting solution to this problem, which we are investigating [4]. In a multi-region perspective, the definition of the multicast semantics is even more complex, since it may involve the delivery to hosts or to groups in different regions with different time constraints [25].

Rethinking the API Finally, it is apparent that we have to rethink the programming interface for this class of networks in order to support the different communication paradigms presented in Section 2.3. In [6] Fall and Demmer propose an API based on the publish-subscribe abstraction that looks promising for specifying system calls for transmitting messages to multiple recipients, capturing the multicast semantics discussed above. However, we argue that a more expressive API is needed in order to specify other aspects such as persistent caching and push/pull delivery. A new programming layer should be introduced; a solution might be a clean-slate approach like that at the basis of Haggle [26].

5. CONCLUDING REMARKS

We have analyzed the design requirements of several application scenarios of multi-regions delay tolerant networked systems for developing regions. We have identified the key design aspects of this class of systems and we have shown how to integrate different routing schemes based on both probabilistic and deterministic forwarding mechanisms considering various communication paradigms in a single framework. Finally, we have outlined the open research problems for the mobile networking community in this area which also constitute our future work path.

Acknowledgments We acknowledge the support of EPSRC through Project CREAM. Mirco Musolesi is currently supported from a research program in the Institute for Security Technology Studies at Dartmouth College, under award 60NANB6D6130 from the U.S. Department of Commerce. The statements, findings, conclusions, and recommendations are those of the authors and do not necessarily reflect the views of the National Institute of Standards and Technology (NIST) or the U.S. Department of Commerce.

6. REFERENCES

[1] DTN2 Reference Implementation. http://www.dtnrg.org.

[2] E. Brewer, M. Demmer, M. Ho, R. J. Honicky, J. Pal, M. Plauche, and S. Surana. The challenges of technology research for developing regions. *IEEE Pervasive Computing*, 5(2):15–23, 2006.

[3] S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H. Weiss. RFC 4838: Delay-Tolerant Networking Architecture, April 2007.

[4] P. Costa, C. Mascolo, M. Musolesi, and G. P. Picco. Socially-aware Routing for Publish-Subscribe in Delay-tolerant Mobile Ad Hoc Networks. *IEEE Journal on Selected Areas in Communications (JSAC)*, 26(5), June 2008.

[5] M. Demmer and K. Fall. DTLSR: delay tolerant routing for developing regions. In *Proceedings of NSDR'07*, pages 1–6, Kyoto, Japan, August 2007.

[6] M. Demmer, K. Fall, T. Koponen, and S. Shenker. Towards a Modern Communications API. In *Proceedings of HotNets-VI*, November 2007.

[7] B. DeRenzi, Y. Anokwa, T. Parikh, and G. Borriello. Reliable data collection in highly disconnected environments using mobile phones. In *Proceedings of NSDR'07*, Kyoto, Japan, August 2007.

[8] B. DeRenzi, N. Lesh, T. Parikh, C. Sims, W. Maokla, M. Chemba, Y. Hamisi, D. Shellenberg, M. Mitchell, and G. Borriello. E-imci: improving pediatric health care in low-income countries. In *Proceedings of CHI'08*, pages 753–762, Florence, Italy, April 2008.

[9] K. Fall. A delay-tolerant network architecture for challenged internets. In *Proceedings of SIGCOMM'03*, Karlsruhe, Germany, August 2003.

[10] S. Guo, M. Falaki, E. Oliver, S. U. Rahman, A. Seth, M. Zaharia, and S. Keshav. Very Low-Cost Internet Access Using KioskNet. *ACM Computer Communication Review*, October 2007.

[11] K. Harras and K. Almeroth. Inter-regional messenger scheduling in delay tolerant mobile networks. In *Proceedings of WoWMoM'06*, Niagara Falls, NY, June 2006. IEEE Computer Society Press.

[12] P. Hui, J. Crowcroft, and E. Yoneki. BUBBLE Rap: Social-based Forwarding in Delay Tolerant Networks. In *Proceedings of MobiHoc'08*, Hong Kong S.A.R., China, May 2008.

[13] S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. In *Proceedings of ACM SIGCOMM'04*, Portland, OR, USA, August 2004.

[14] D. Jea, A. A. Somasundara, and M. B. Srivastava. Multiple Controlled Mobile Elements (Data Mules) for Data Collection in Sensor Networks. In *Proceedings of DCOSS'05*, pages 244–257, April 2005.

[15] W. A. Kaplan. Can the ubiquitous power of mobile phones be used to improve health outcomes in developing countries? *Globalization and Health*, 2(9), May 2006.

[16] D. Kutscher, J. Greifenberg, and K. Loos. Scalable DTN distribution over uni-directional links. In *Proceedings of NSDR'07*, pages 1–6, Kyoto, Japan, August 2007.

[17] A. Lindgren, A. Doria, and O. Schelen. Probabilistic routing in intermittently connected networks. *Mobile Computing and Communications Review*, 7(3), July 2003.

[18] B. T. Loo, J. M. Hellerstein, I. Stoica, and R. Ramakrishnan. Declarative routing: Extensible routing with declarative queries. In *Proceedings of ACM SIGCOMM'05*, Philidelphia, PA, USA, August 2005.

[19] M. Musolesi and C. Mascolo. CAR: Context-aware Adaptive Routing for Delay Tolerant Mobile Networks. *IEEE Transactions on Mobile Computing*. To Appear.

[20] NOAA. The Tsunami Warning System. http://www.prh.noaa.gov/ptwc/.

[21] J. Panchard, P. S. Rao, M. Sheshshayee, P. P. Papadimitratos, and J.-P. Hubaux. Wireless Sensor Networking for Rain-fed Farming Decision Support. In *Proceedings of NSDR'08*, Seattle, WA, USA, August 2008.

[22] A. Pentland, R. Fletcher, and A. Hasson. DakNet: Rethinking Connectivity in Developing Regions. *IEEE Computer*, pages 4–9, January 2004.

[23] C. E. Perkins and E. M. Royer. Ad hoc On-demand Distance Vector routing. In *Proceedings of WMCSA '99*, February 1999.

[24] K. Scott and S. Burleigh. RFC 5050: Bundle Protocol Specification, November 2007.

[25] G. Sollazzo, M. Musolesi, and C. Mascolo. TACO-DTN: A Time-aware Content-based Dissemination System for Delay-Tolerant Networks. In *Proceedings of MobiOpp 2007*, Puerto Rico, USA, June 2007. ACM Press.

[26] J. Su, J. Scott, P. Hui, J. Crowcroft, E. de Lara, C. Diot, A. Goel, M. Lim, and E. Upton. Haggle: Seamless networking for mobile applications. In *Proceedings of Ubicomp'07*, pages 391–408, 2007.

[27] A. Vahdat and D. Becker. Epidemic routing for partially connected ad hoc networks. Technical Report CS-2000-06, Department of Computer Science, Duke University, 2000.

[28] Z. Zhang. Routing in intermittently connected mobile ad hoc networks and delay tolerant networks: overview and challenges. *Communications Surveys & Tutorials, IEEE*, 8(1):24–37, 2006.

[29] W. Zhao, M. Ammar, and E. Zegura. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *Proceedings of MobiHoc'04*, Tokio, Japan, May 2004.

[30] W. Zhao, M. Ammar, and E. Zegura. Multicast routing in delay tolerant networks: Semantic models and routing algorithms. In *Proceedings of the SIGCOMM WDTN Workshop*, Philidelphia, PA, USA, August 2005.

[31] K. D. Zoysa, C. Keppitiyagama, G. P. Seneviratne, and W. W. A. T. Shihan. A public transport system based sensor network for road surface condition monitoring. In *Proceedings of NSDR'07*, pages 1–6, Kyoto, Japan, August 2007.