

Reasoning about Multicast in Opportunistic Networks

Elena Pagani

Computer Science Dept., Univ. degli Studi di Milano, Italy
CNR-IIT, National Research Council, Italy
Email: pagani@ dico.unimi.it

Gian Paolo Rossi

Computer Science Department
Università degli Studi di Milano, Milano, Italy
Email: rossi@ dico.unimi.it

Abstract—Mobile computing services are becoming highly personalized and influenced by user location, social attitudes and interests. This trend is often synthesized with the term *behavior-awareness*. Behavioral services characterize most of the applications of Opportunistic Networks and advocate a single communication paradigm: *multicast*. Despite that, multicast is less studied than unicast and broadcast. We believe that there is not correspondence between problem relevance and research attention. This paper focuses on the multicast best effort service in opportunistic networks with the aim of better understanding the problem and defining the guidelines for the design of a novel multicast protocol.

I. INTRODUCTION

People move and meet while engaging in their social relationships. This provides the contact opportunities that Opportunistic Networks (ONs) leverage to obtain an unlicensed wireless communication infrastructure. In such a framework of *human contacts*, mobile computing services are becoming highly personalized and influenced by user location, social attitudes and interests. This trend is often synthesized with the term *behavior-awareness*. The understanding of human behavior is profitably adopted to enhance contact-based forwarding (e.g., [1]) and to seed users with broadly scoped contents that depend on a given *behavior-aware profile* (e.g., [2]). By rethinking the Opportunistic Network workload in terms of behavioral networking we envision specific services mediating between users' behavior and the actual resources matched in the neighborhood. For instance, a mobile user may be willing to share some geo-tagged contents with other users/friends with similar interests, or retrieve a content/resource among the set of neighbors that more likely possess it, or diffuse news to targeted destinations, and so on. Evidently, the nature of behavioral services is to aggregate people, and their handheld devices, in groups with homogeneous behavior. This class of applications represents the large part of ON applications and the hope to deploy them is still fueling the ON research. There is a single communication service advocated by this class of applications: *multicast*. Certainly, it is not *unicast*; in fact, no one is able to precisely address the node destination of the content or owner of the requested content. This because the user's query may have a wide range of vagueness and because in ONs we are moving

apart from classical server-centric approach of the wired Internet. It is not *broadcast* either; in fact, both users and the network are not willing to be flooded with unneeded contents or requests. Despite the common understanding of the above arguments, multicast is the less studied of the three mentioned communication paradigms. The huge of proposals on unicast and broadcast communications in ON is compensated by very few papers coping with the multicast issues. At the best of our knowledge, multicast is explored in detail only in [3], where new semantic models are defined taking into account timing aspects related with the long delays involved in ONs, and in [4], where a similar multicast semantics (*ProfileCast*) has been adopted. Other few provide multicast protocols under specific ON conditions. We believe that there is not correspondence between problem relevance and research attention.

This paper focuses on the multicast best effort service in ON with the aim of better understanding the problem and defining the guidelines for the design of a novel multicast protocol. We firstly define the novel multicast model and semantics, named *behavior-aware multicast*, or *BAM*, suitable for various application settings in ONs. Secondly, we characterize the notion of membership in a contact-based scenario and classify the problem in terms of changing levels of knowledge about group membership and contact events. Finally, we discuss how temporal constraints on message delivery can be enforced with different degrees of knowledge. The paper shows that, in *BAM*, the notion of membership is vaguely defined through similarity and the likelihood of capturing it depends on its stability and on the ability to extract some structure from human contacts. As a result, we are not plenty of opportunities to adopt multicast instead of broadcast and epidemic in practical scenarios. Nevertheless the space exists and should be followed with proposals to improve the feasibility of applications in ONs. According to our current understanding of the problem, however, there will not be a single general multicast solution to the *BAM* problem, but several different solutions for different settings.

II. MULTICAST ARCHITECTURE AND MODEL

We assume a two-layer architecture for the mobile nodes (fig.1). At the *network* layer, the multicast entity, *multicast-*

relay, receives messages from encounters, when a contact event occurs, stores them locally, takes forwarding decisions and may deliver a message to the upper application entity, (end-node), according to membership conditions. End-nodes agree on a super-set of metadata that can represent locations (e.g., [4]), resources, interests or social attitudes of users.¹

The behavior-aware multicast model we consider is different from classical IP-multicast adopted, e.g., in [6]. Let us consider fig.1 to show this point. Six people with their handheld devices (a through f) share a set of three (well known) metadata $\{m, n, p\}$. b, d and f share the common metadata m . Each end-node in $\{b, d, f\}$ has local awareness of its interest in m . This knowledge can be either explicitly declared or implicitly derived from human behavior. Moreover, it also knows that the same interest in m is possibly shared with a group of other end-nodes. However, neither the cardinality of the group nor the identity of the end-nodes are a-priori known and, unlike IP-multicast, the group is implicitly identified by the name of the metadata. End-nodes in $\{b, d, f\}$ may be willing to diffuse contents about m , or to query contents about it. This is performed by exploiting the service of the underlying multicast-relay entity. A multicast-relay is notified by upper entities about metadata of interest, i.e. their names, and is responsible for performing the behavior-aware multicast routing to enable delivery of the message to the intended recipients (b, d and f in our example). In practice, BAM is a sort of metadata-driven routing where each multicast-relay has no a-priori knowledge of other multicast-relays, if any, matching the specified metadata. As a consequence, a multicast-relay must go beyond its local awareness, achieve some richer knowledge of the membership aggregated around a given concept m , and learn how to get in contact with the members. These are all critical points in our highly mobile scenario and they have been in general skipped by most of the proposals in the literature. The required knowledge can be enriched only when a contact opportunity occurs, by exchanging relevant information among encounters (e.g. through beaconing). As we will discuss in Section III, the membership knowledge is a highly variable concept depending on the way information is diffused and on how contact events are utilized. In ONs, such a contact-based diffusion mechanism more or less performs the functionality of Join/Leave operations in IP multicast, although the relationship between a set of metadata and the named message is not regulated by a perfect matching but by similarity.

In the sketched metadata-driven multicast model, each multicast-relay is responsible for routing a message m to a group G on the base of the degree of membership knowledge it acquired. The problem it has to solve is as follows:

Definition 1 (Problem definition): Let G be a group of

¹As an alternative, methods for tag comparison can be adopted, such as the one proposed in [5].

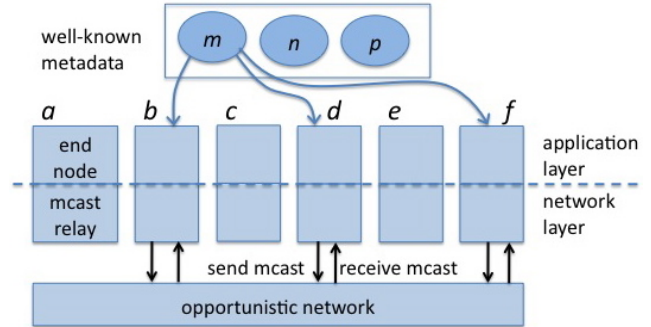


Figure 1. Architecture and example

end-nodes and m a message that matches the delivery rules of G (nodes b, d, f in fig.1). Let U be the set of all nodes in the system (nodes a through f). Then a multicast routing service must satisfy the following three properties:

- **correctness:** (i) m must be delivered to a subset $G' \subseteq G$ of nodes, but (ii) no node in $U - G$ must deliver m ;
- **multicast effectiveness:** the service should maximize coverage, i.e., to keep the cardinality of G' as close as possible to that of G ;
- **network efficiency:** the service should minimize network load, i.e., reduce the number of m copies that multicast-relays exchange with respect to an epidemic.

The likelihood for a multicast protocol to satisfy the problem above highly depends on the knowledge a multicast-relay has about membership and mobility patterns of mobile devices. Intuitively, without knowledge the only viable approach is to diffuse m through an epidemic protocol; despite the approach fully satisfies the multicast effectiveness, it fails in satisfying network efficiency and the correctness clause is satisfied by end nodes thanks to the separation between reception and delivery. By contrast, the full knowledge about both contacts and membership (i.e. each multicast-relay knows the identity of all group members anytime, and the set of contacts with their timing, to exploit to reach all recipients) enables the construction of the optimal multicast delivery tree. We indicate this type of knowledge with the term *complete knowledge*. In the next Sections, we explore the different conditions occurring between these two extremes. Time constraints on delivery are initially neglected due to the fact we are operating within a delay tolerant scenario. They will be discussed in Section IV.

III. ROLE OF KNOWLEDGE IN ON MULTICAST

With the aid of the ruler in fig.2, we identify different levels of membership knowledge that can be achieved in a contact-based communication scenario. Let us define K_M^∞ as the complete knowledge about membership, for any node and anytime. K_M^1 is the membership knowledge of one-hop neighbors, under the assumption that this information is obtained through beaconing and forgotten as soon as



Figure 2. Ruler for different degrees of knowledge about membership

the radio link breaks. K_M^1 can be used to provide multicast routing through direct contact. The lowest degree of knowledge a multicast-relay may achieve is K_M^0 or the local awareness that is never exchanged with encounters. Of course, this condition emphasizes privacy issues and leads to communication approaches different from multicast (see, for instance, [7]).

Similarly, K_C^∞ is the complete knowledge about contacts, for any pair of nodes and anytime, while K_C^1 corresponds to knowing of a contact only when it occurs (locally).² Both K_M^∞ and K_C^∞ can only be approximated in practical ONs. K_C^∞ can be achieved when relays follow mobility patterns that are regulated by a deterministic time schedule. This condition partially applies, for instance, to vehicles of a public transportation system (either in urban or rural area). K_M^∞ can be approximated under particular mobility assumptions (e.g., all vehicles eventually enter the bus depot, or all nodes move according to a uniformly distributed mobility model) and stability assumptions (e.g. the membership does not change until the required contacts occur). In both cases, each node is guaranteed to eventually encounter every other node and learn its membership.

Interestingly, most of the proposals in the literature (e.g., [8]-[13]) assume K_M^∞ conditions where the group membership is *static* and well known by the message source at the time of message generation. Similarly, as far as contacts are concerned, some approaches assume mobility patterns that approximate K_C^∞ such as, for instance, [14] that considers rural environments, or [15] that adopts message *ferries*. In some cases, mobility is assumed low enough to allow building tree-based or mesh infrastructures, such as in [16]-[21].

By contrast, under practical and more general ON conditions, only a partial degree K_M^P , with $1 < P \leq \infty$, of knowledge can be achieved. Multicast-relays can exploit contacts to learn about membership and to derive some utility function that captures human behavior to drive the forwarding decisions. The problem is complicated by the fact that, at a given time, each multicast-relay may have achieved only a partial view of the membership and that the set of members one knows is possibly different from the set of others. The higher the value of P in K_M^P , the more efficient and effective should be the multicast. Interestingly, at the network layer, contacts are exploited to propagate both the knowledge about membership and the multicast

messages. Should the membership be very stable we can assume that nodes have time to diffuse membership and exploit contacts to build internal utility functions before performing efficient multicast routing. By contrast, under dynamic membership conditions, we can envision a scenario in which nodes are contemporarily performing membership diffusion, construction of utility function and message routing. This will negatively affect the capability of satisfying the specified problem requirements. For instance, a set of sensors deployed to observe environmental parameters maintain location and functionalities over time. Similarly, students in a campus area maintain interests and mobility habits. In these two cases, while nodes should quite easily get acquainted of membership they still have the problem of how to reach the members. The understanding of the mobility patterns becomes crucial to define the proper utility function and then behavior-aware multicast. However, no useful and effective utility function exists if contacts do not reveal a behavioral, recurrent, structure. For instance, mobility scenarios where different nodes come back and forth frequently (e.g., such as people in a crowded urban area) are unsuitable to multicast; this is more likely the condition where an epidemic should rather be performed.

By browsing the literature, Mobility ProfileCast, [4], is probably the only approach fitting with the described multicast scenario and trying to deal with membership and contact knowledge. It assumes that relays know a common set of locations (that defines K_C), coincident with the common set of interests (that define K_M) shared by the nodes. It uses clustering algorithms and the assumption that node profiles do not change in order to approximate K_M^∞ , and uses the similarity among mobility profiles of both source and destinations to achieve $K_C^P > K_C^1$.

The intimate dependence of multicast on the underlying contact network has two implications. First, it makes difficult the design of a general purpose multicast protocol. Multicast in ONs, if any, will provide a dedicated solution to specific settings. Second, contacts are the surrogate of the multicast delivery tree in IP-multicast. But, unlike IP-multicast, contacts also reveal at the network layer, to multicast-relays, the identity of the group members.

In the considered scenario, we can easily show the following claims.

Claim 1: If no more knowledge than K_C^1 is available about contacts, then any K_M^P is useless for message forwarding, that is, relays can only behave as if they had K_M^1 .

Whatever is the knowledge that relays have about group membership, if relays have K_C^1 , then they cannot forecast anything about future encounters, that is, they are unable to discriminate between relays useful to reach the message destinations, and useless encounters. Hence, a relay can only forward the message blindly (as if it had K_M^1).

Claim 2: $K_M^P > K_M^1$ tends to K_M^1 , if the membership continuously changes.

² K_C^∞ and K_M^∞ corresponds to the *complete contact* and the *complete memberships* oracles respectively in [3].

Table I
IMPACT OF KNOWLEDGE ON FORWARDING POLICY

	K_M^1	K_M^P	K_M^∞
K_C^1	epidemic	epidemic	epidemic
K_C^P	broadcast	BAM	BAM
K_C^∞	optimal broadcast	optimal bcast / BAM	optimal multicast

If nodes continuously change their interests, then a relay r_1 may communicate its group membership to a one-hop neighbor r_2 , and then the local node changes it soon after having moved out of r_2 's range. As a consequence, the information held by r_2 becomes obsolete, that is, the unique correct knowledge is K_M^1 .³

A. Knowledge and Multicast Algorithms

Table I provides a classification of viable solutions to the multicast problem as a consequence of the available values of K_C and K_M . According to Claim 1, when K_C^1 is available only random message forwarding is possible with message filtering at the last hop. In order to maximize the multicast effectiveness, the best choice is an epidemic approach, which however congests the network. If K_C^∞ is available, then relays can build a temporal graph of all contacts and use it to compute optimal paths towards the message recipients. Should destinations be also known (K_M^∞ achieved, for instance, through oracles, as in [3]), then the optimal multicast tree can be built. By contrast, when K_M^1 is only available (for instance, because membership changes too quickly), then broadcast is the only viable choice because each relay potentially hosts a node interested in a multicast message. Under the remaining yet practical conditions, only a partial knowledge is available about contacts. If $K_M^P > K_M^1$ is available, then a behavior-aware multicast could be profitably designed to find the best trade-off between network efficiency and multicast effectiveness.

We have observed that most of the proposals in the literature adopt the scenarios with K_M^∞ and K_C^P , with a high value of P obtained from proper mobility settings. The condition with K_M^P and K_C^P is the most general and interesting, but, so far, (partially) covered only by [4]. This is actually a very open research area. Table I shows that there exists a narrow band of applicability of multicast solutions in ONs. However, when it is a viable solution it could actually provide significant service benefits making feasible the deployment of practical ON applications.

IV. ROLE OF TIME IN ON MULTICAST

Time constraints come to complicate the fragile scenario described above. In fact, in ONs the unpredictable latencies

³In fact, a node could change its membership between two successive beacons, but this inconsistency is healed at the successive beacon, and we assume that beacons are exchanged frequently.




Figure 3. Example for Claim 3

in message delivery generate several semantics issues. For instance, how to deal with a message received by a multicast-relay which does no longer host members of the group? How to deliver a query to a set of end-nodes that were members of a group at a given time in the near past? In this Section we will consider time constraints and evaluate the effects they produce on the previous considerations.

Time implications on multicast have been analyzed in depth in [3], where multicast messages are labeled with both temporal delivery and temporal membership constraints that determine the set of nodes that should deliver the message. The following models are proposed:

- **temporal membership (TM):** the intended recipients are the nodes that belong to a group G at any time during the specified membership interval, independently of any delivery interval;
- **temporal delivery (TD):** the intended recipients are those determined by TM that are also able to deliver the message during the specified delivery interval;
- **current-member delivery (CMD):** the intended recipients are those determined by TD that also belong to the group at the time of delivery.

Particular cases occur when the membership interval is (left and/or right) unbounded, and the same applies to the delivery interval. E.g., the multicast semantic considered in [4] could be casted to the CMD model, with CM flag `True`, delivery interval (t_{send}, ∞) and membership interval (t_{send}, ∞) .

In the above time constrained scenarios, the following effect may be derived:

Claim 3: If message diffusion takes place before or during the membership interval, then K_M^P with $P < \infty$ does not allow to bound diffusion, for any K_C . A message must be kept in buffers till the end of the membership interval.

The proof follows from fig.3: let r be a relay receiving a message m before the end of m 's membership interval t_{end}^M , and n its local node, which so far never belonged to G . If K_M^∞ is not owned, then r is not able to determine whether n will join G at some time $t_{end}^M - \epsilon$. Hence, r must buffer m till t_{end}^M ; otherwise, if n joins G after m has been discharged by the relay, m could never be delivered by n , thus weakening the multicast effectiveness (Def. 1). In the worst case, the argument applies to all the nodes in the network. Hence, m must be diffused to all relays, and they possibly have to buffer m till the end of the membership interval.

If by contrast the message diffusion is performed *after* the end of the membership interval, in case some $K_M^P > K_M^1$ and $K_C^P > K_C^1$ are available then BAM could be performed.

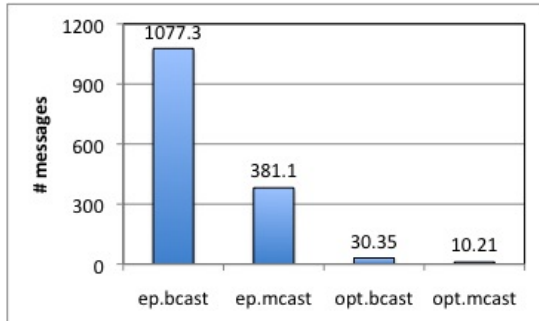


Figure 4. Network efficiency for one group

Besides [3] – whose models are also adopted by [6], [11] – only a few works consider temporal constraints. In [17], [19], late joins are managed by buffering messages till their lifetime expiration. In [17], investigation on stability requirements is put forward as an interesting area of research.

V. A FEW MEASUREMENTS

In Table I, we have proposed different problem solutions for different degrees of knowledge about membership and contacts. We observed that the practical conditions favoring the deployment of a *BAM* protocol are actually very few. The question is: is the effort of designing a multicast protocol worth in a real mobility scenario? We performed some measures with a data set of contacts obtained in a campus environment from 44 people equipped with wireless devices with 10 m. radio range [22].⁴ Groups have been chosen according to social relationships. Group size ranges from 8% to 20% of the number of nodes present in the considered days. With a custom simulator we performed the following measures: (i) network efficiency, in term of message copies; (ii) delivery latency; and (iii) multicast effectiveness. We used Dijkstra’s algorithm to compute the optimal paths under the K_C^∞ conditions and epidemic forwarding is adopted when K_C^1 . Fig.4 shows the number of message copies required to ensure the message delivery to each destination for a group G involving the 14.7% of the nodes present in the considered day. For epidemic multicast, we measured the number of messages generated upon delivery at each destination in G (i.e., K_M^∞ is owned); for epidemic broadcast (K_M^1 is owned) we wait to reach all the nodes in the system. Over the considered groups, epidemic multicast generates around 41% of the traffic generated by epidemic broadcast. Similarly, the optimal trees for broadcast and multicast provide an estimation of the gain that K_M^∞ may yield: the optimal multicast tree allows to exchange 30.4% of the messages exchanged with broadcast. Symmetrically, the comparison between epidemic and optimal broadcast (multicast) supplies a measure of the gain achievable moving

⁴The data set can be downloaded from the CRAWDAD archive (<http://www.crowdad.org/unimi/pmtr>).




Figure 5. Dynamics of membership learning

from K_C^1 to K_C^∞ . The optimal trees generate the 3.2% for broadcast (4.8% for multicast) of the traffic generated by the corresponding epidemics. A *BAM* protocol is supposed to approximate the optimal multicast performance under the hypothesis of achieving K_M^P close to K_M^∞ . This results in a drastic traffic reduction that corresponds to a significant saving of system batteries and resources.

All approaches obtain the same performance in terms of multicast effectiveness (epidemics just follow all paths, thus also those adopted by optimal broadcast and multicast) and latency. For all the algorithms and groups considered, group coverage is higher than 90% of the destinations, and around 95% in average. The mean delay to reach a destination – averaged over all destinations – ranges in between 35 minutes and 1 hour and 50 minutes (although, it may occasionally happen that reaching the last destination takes up to 7 hours). This leads to conclude that the specification of some delivery interval may hinder the multicast effectiveness.

We are unable to observe the performance of a *BAM* protocol for intermediate levels of knowledge (although the results provided in [4] are comforting). However, it is interesting to observe how long the membership learning procedure lasts in a real setting. In fig.5, the dynamics of group learning is shown for a group including the 20% of the nodes present in the considered day; dotted lines represent the learning process for each node present, while the solid line is their average. Relays learn membership both through contact with a destination, and via exchange of the owned information upon each encounter. Considering all groups, in average over a one-day time window the percentage of discovered destination ranges between 67% and 71%. The mean time to discover all destinations – when this is obtained – is in the order of 2-4 hours. This leads to conclude that K_M^∞ may require a non-negligible time to be achieved, also in restricted geographical areas and when all group members are present. On the other hand, if groups are stable, the learnt information can be re-used for a long time thereafter. Hence, K_M^∞ cannot be always assumed, but it rather depends on both

mobility, and dynamics of interest changes.

VI. CONCLUSIONS

In this paper, we investigate the impact of both group membership knowledge and contact knowledge on the feasibility of multicast content dissemination in opportunistic networks. The results show that: the membership may be unknown and unstable but can be discovered – with some delays – through communication. Though, the usefulness of learning the membership depends on membership dynamics. Moreover, the knowledge on membership is useless if no knowledge about contacts is available. The impact of time constraints on the feasibility of multicast is studied. Simulations show that knowledge may significantly help in performing a selective message forwarding. These results may help designers in understanding to what case a certain real setting can be reverted, in choosing trade-offs between forwarding efficiency and exchange of control information, and in characterizing viable solutions. Moreover, the results bring into evidence the open research issues, that is, (i) characterizing suitable contact models; (ii) investigating mechanisms to gain knowledge about nodes interests so as to evolve from the classical IP-multicast semantic, and (iii) designing behavior-aware multicast forwarding policies such that the greater the knowledge held, the better their performance. This is our future work.

ACKNOWLEDGMENT

This work was partially funded by the European Commission under the FP7 SCAMPI (258414) Project.

REFERENCES

- [1] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, *Impact of Human Mobility on Opportunistic Forwarding Algorithms*, IEEE Trans. Mobile Computing, 6(6), 2007, 606-620.
- [2] C. Boldrini, M. Conti, and A. Passarella, *Design and performance evaluation of ContentPlace, a social-aware data dissemination system for opportunistic networks*, J. Comput. Netw., 54(4), 2010, 589-604.
- [3] W. Zhao, M. Ammar and E. Zegura, *Multicasting in Delay Tolerant Networks: Semantic models and Routing Algorithms*, Proc. ACM SIGCOMM Workshop on Delay-tolerant networking (WDTN), 2005, 268-275.
- [4] W.-J. Hsu, D. Dutta, and A. Helmy, *Profile-Cast: Behavior-Aware Mobile Networking*, SIGMOBILE Mob. Comput. Commun. Rev., 12(1), 2008, 52-54.
- [5] G. Lo Justo, A.J. Mashhadi, and L. Capra, *Folksonomy-based Reasoning for Content Dissemination in Mobile Settings*, Proc. CHANTS 2010, 39-46.
- [6] M. Abdulla, and R. Simon, *A simulation analysis of multicasting in delay tolerant networks*, Proc. 38th Conf. on Winter simulation (WSC), 2006, 2234-2241.
- [7] G. Karlsson, V. Lenders, and M. May, *Delay-Tolerant Broadcasting*, IEEE Trans. on Broadcasting, 53(1), 2007, 369-381.
- [8] M. Chuah, and Y. Xi, *An Encounter-Based Multicast Scheme for Disruption Tolerant Networks*, Comput. Commun., 32(16), 2009, 1742-1756.
- [9] W. Gao, Q. Li, B. Zhao, and G. Cao, *Multicasting in delay tolerant networks: a social network perspective*, Proc. 10th ACM International Symposium on Mobile ad hoc networking and computing (MobiHoc), 2009, 299-308.
- [10] U. Lee, S. Y. Oh, K.-W. Lee, and M. Gerla, *RelayCast: Scalable Multicast Routing in Delay Tolerant Networks*, Proc. ICNP 2008.
- [11] Z. Narmawala, and S. Srivastava, *MIDTONE: Multicast in Delay Tolerant Networks*, Proc. CHINACOM 2009, 609-616.
- [12] Y. Wang, X. Li, and J. Wu, *Multicasting in Delay Tolerant Networks: Delegation Forwarding*, Proc. Globecom 2010.
- [13] J. Wu, and Y. Wang, *A Non-Replication Multicasting Scheme in Delay Tolerant Networks*, Proc. 7th IEEE Intl. Conf. on Mobile Adhoc and Sensor Systems (MASS), 2010, 89-98.
- [14] M. Musolesi, and C. Mascolo, *A Framework for Multiregion Delay Tolerant Networking*, Proc. WINSNR'08.
- [15] P. Yang, and M. Chuah, *Efficient Interdomain Multicast Delivery in Disruption Tolerant Networks*, Proc. 4th Intl. Conf. Mobile Ad-hoc and Sensor Networks, 2008, 81-88.
- [16] A. Afanasyev, K. Mayoral, Z. Zhu, and S. Y. Oh, *DTCAST: Delay and Disruption Tolerant Multicasting Protocol*, Proc. 11th Youth Technological Conf. on High Technologies and Intellectual Systems, 2009.
- [17] M. Demmer, and K. Fall, *The design and implementation of a session layer for delay-tolerant networks*, Comput. Commun. 32(16), 2009, 1724-1730.
- [18] J. Santiago, A. Casaca, and P. R. Pereira, *Multicast in Delay Tolerant Networks Using Probabilities and Mobility Information*, Ad Hoc & Sensor Wireless Networks, 7(1-2), 2009, 51-68.
- [19] S. Symington, R. C. Durst, and K. Scott, *Custodial Multicast in Delay Tolerant Networks – Challenges and Approaches*. MITRE Technical Paper, 2007.
- [20] P. Yang, and M. C. Chuah, *Context-Aware Multicast Routing Scheme for Disruption Tolerant Networks*, Int. J. Ad Hoc Ubiquitous Comput. 4(5), 2009, 269-281.
- [21] Q. Ye, L. Cheng, M. C. Chuah, and B. D. Davison, *Performance comparison of different multicast routing strategies in disruption tolerant networks*, Computer Communications, 32(16), 2009, 1731-1741.
- [22] S. Gaito, E. Pagani, and G.P. Rossi, *Opportunistic Forwarding in Workplaces*, Proc. 2nd ACM SIGCOMM Workshop on Online Social Networks, 2009, pp. 55-60.