Narrowcasting: An Empirical Performance Evaluation Study

Franck Legendre
ETH Zurich
8092 Zurich, Switzerland

Vincent Lenders
Armasuisse
3000 Bern, Switzerland

Martin May
ETH Zurich
8092 Zurich, Switzerland

Gunnar Karlsson
KTH, Royal Inst. of Tech.
100 44 Stockholm, Sweden

ABSTRACT

This paper studies content dissemination (narrowcasting) in wireless ad hoc networks. We consider (i) the dissemination of general (non selective, broadcast) content (ii) with loose delay constraints (iii) aimed at an entire community confined in a restricted area (iv) served periodically by one single and stationary content emitter (i.e., an access point). These constraints target challenging environments where for practical and/or economical reasons content dissemination must be set up on the fly with the least effort. Examples are dissemination of information at a conference, in military operations, to team workers, and to inhabitants of remote villages. Based on real-world mobility traces, we analyze different dissemination strategies comparing cooperative versus non-cooperative behavior of nodes for storing and/or forwarding content. In particular, we evaluate how the performance is impacted by mobility. The main outcome is that node collaboration drastically increases the performances of content dissemination while the per-device overhead (or load) is very low and remains on average evenly distributed.

Categories and Subject Descriptors

C.2.1 [**Computer Systems Organization**]: Computer-Communication Networks—*Network Architecture and Design*

General Terms

Design, Performance, Measurement

Keywords

Ad Hoc, Multi-hop, Delay-tolerant, Testbed, Content dissemination, Narrowcasting

CHANTS'08, September 15, 2008, San Francisco, California, USA.

1. INTRODUCTION

Narrowcasting refers to broadcasting to a very narrow range of audience. This audience can be limited by requiring individuals to register such as with cable TV, electronic mailing lists, and electronic forums. It can also be restricted to audience in given areas such as the audience of local radio and television stations [1] or individuals in public transports or waiting queues watching prerecorded television programs. Hence, narrowcasting gathers individuals sharing similar interests or physically located in the same restricted area for whatever reason.

We believe there is a unfilled need for narrowcasting of local information (e.g., news, weather forecasts, forthcoming events, authoritative regulations) to communities gathered in a restricted area. This can be settled communities (e.g., work teams, military bases, remote villages, tactical operations) or transient communities gathering occasionally (e.g., conferences, exposition halls). With the advance of wireless technologies working in the ISM band and the increasing proportion of handheld devices, it is now feasible to offer a competitive alternative to terrestrial and satellite broadcasting/narrowcasting systems (e.g., TV, Internet, 3G). Current narrowcasting and broadcasting systems in general are highly regulated and costly insofar as they exist. For practical and/or economical reasons, such narrowcasting systems must be easy to set up and at low cost. For example, conference attendees or military troops are a perfect audience for local information (e.g., schedules, events, commands) as well as villagers of remote areas in developing countries where access to local news may not be assured by a traditional broadcasting/narrowcasting system.

In this context, we consider narrowcasting of content to a community (or broadcasting of content to a tacitly subscribing community) originating from one single and fixed (stationary) emitter such as an access point (AP) with possibly the cooperation of mobile nodes to participate in spreading. The main motivation for such a minimalistic setup is to target real life situations constrained by economical and/or practical factors. With such a system, we trade flexibility brought by new communication means (ad hoc, handheld devices) for some delay. The content we are considering is delay-tolerant in nature or must have at least loose constraints regarding the delay of delivery.

In this paper, we study to what extent and with what level of implication, nodes collaboration in forwarding and storing content can effectively increase content dissemination compared to traditional means such as broadcasting or flooding. We hence study dif-

[1]These are also called low-powered broadcasting systems, which transmit only to a small area such as a town, a district or a neighborhood and as such are not classified as narrowcast systems. The scope of narrowcasting systems is even smaller.

ferent strategies to disseminate general and delay-tolerant content i.e., (i) broadcasting from an AP where nodes request for new content when within transmission range, (ii) periodic multi-hop flooding initiated by the AP, and (iii) opportunistic broadcast dissemination involving store-and-forward epidemic spreading between nodes. As for content, we consider two types of content, (a) short messages (e.g., text messages such as press dispatches) and (b) large files (e.g., audio/video files such as news reports). We evaluate these strategies using mobility traces collected in an office environment at ETH Zurich during a whole week, which are replayed in a simulator. For each strategy and message type, we evaluate the performances of the dissemination i.e., the proportion of reached nodes (content reachability) and how fast content spreads. We also characterize the dissemination process per se i.e., through how many hops content is delivered and if particular nodes play a more important role in delivering content (i.e., carriers).

Much work has been done for content dissemination in the context of metropolitan areas with the concept of Infostation [12, 15, 16], vehicular [3, 11], and wildlife sensor networks [7]. Yet, we would like to point out the main differences of our work with related works. From a communication model and technical viewpoint, we consider narrowcasting. First, we consider only a stationary point from which content is originated to a large audience i.e., one AP whereas other approaches (e.g., Infostation, sensor networks) consider an infrastructure with more than one emitter of content. In the scenarios we target, one single emitter of information is considered for practical and/or economical reasons. Second, as regard to traffic, we consider periodic spreading of updated information relevant to an entire population or community. We assume localized interest of the information spread and its general scope. Last, we assume different levels of node cooperation for storing and/or forwarding content. From a methodological viewpoint, no previous work has performed a comparison between different content dissemination strategies in the particular setup we are targeting. Besides, our study is the first to replay real mobility measurement traces with such a fine-grained precision (i.e., 0.5s whereas other mobility measurement campaigns measured mobility at 7s [3] or 120s [5] intervals.).

The contributions of this paper are:

1. An empirical evaluation of different dissemination strategies.
2. A fair and repeatable comparison by replaying fine-grained mobility measurements from a real-world experiment.
3. We show in a realistic use case that node collaboration drastically increases the performances of content dissemination while fairly distributing the load on each node without enforcing particular mechanisms.

The remainder of this paper is organized as follows. Section 2 motivates our approach by describing several application scenarios. Section 3, describes the different dissemination strategies which are put into action in a community of team workers as well as the evaluation metrics. The evaluation methodology and results are presented in Section 4. Section 5 compares our approach with related work. Section 6 discusses deployment issues and Section 7 concludes our work.

2. APPLICATION SCENARIOS

We present three scenarios particularly adapted to narrowcasting as proposed in this paper:

Conference venue: At a conference (or any scientific gathering such as symposium or workshop), researchers gather from all over the world and form a transient community lasting a short period of time (e.g., from one to a couple of days). Activities include keynote speeches, paper presentations (in one or several parallel tracks), poster/demonstration presentations, and sometimes social activities. In this context, our setup can be used to disseminate ad hoc content from an emitter placed in a central location (e.g., lobby). Examples of content are the conference's schedule but also videos of keynotes speeches, videos of demonstrations, as well as papers and posters. In the case of a multi-track conference, one could get a presentation that occurred in a different track he/she couldn't attend. One could also watch videos of demonstrations after they have actually happened. Besides, proceedings materials would no longer be required to be printed on paper nor stored on CD for example. Eventually, our setup can be used to announce changes in the program and advertise at the right timing of a forthcoming social event (e.g., dinner, excursion). Of course, all this content could be available directly on the world wide web but in our case, anyone participating to the event, would be tacitly registered and get all the content as soon as it is available.

Military camp: Logistics and telecommunication means are the first facilities being setup when a military camp is installed. These camps or bases are often placed in areas where no fixed telecommunication access nor broadcasting systems are provided. Besides this, due to security issues, telecommunication solutions are proprietary military solutions (e.g., jamming-resistant). Narrowcasting as we propose in this paper, can be used in such context to disseminate news from the soldiers' homeland but also news related to the camp such as regular reports on its status of advancement when it is being setup or camp-life related news (e.g., menus, social activities). Note that emergency situations are akin to this scenario and narrowcasting can be very useful to inform endangered populations.

Remote village: In some remote areas of the developing countries, access to information is sometimes limited since broadcasting systems might not cover certain areas or even not exist (e.g., due to censorship). Consider a content emitter placed in a central location such as the main place of a village. This content emitter can be used to disseminate valuable and general information to villagers such as local news, goods market prices, weather forecasts, new regulations (e.g., social, agriculture), new facilities, etc. The information must be periodically updated and disseminated to the whole population which forms a settled community. Taking profit of node mobility, content can be spread from this kind of central location to everyone. We can expect the flow of content to spread when villagers gather in the village during markets, or rituals. For those not coming often to the village the flow of content will certainly follow the same route as mail or food.

3. NARROWCASTING STRATEGIES

We consider the case where new contents originate from a single fixed source such as an access point (AP) in ad hoc mode. We define three simple delay-tolerant dissemination schemes as follows:

- *Cell-based broadcast (Broadcast):* Content is emitted in broadcast mode from the fixed AP when a node is within transmission range and proactively requests for new content. It does not involve any cooperation between mobile nodes. This scheme is the exact sibling of traditional low-powered broadcasting systems (e.g., local radios/TVs) where content is broadcast to a small area except that the fixed AP only broadcasts content when requested. We benefit from the flexibil-

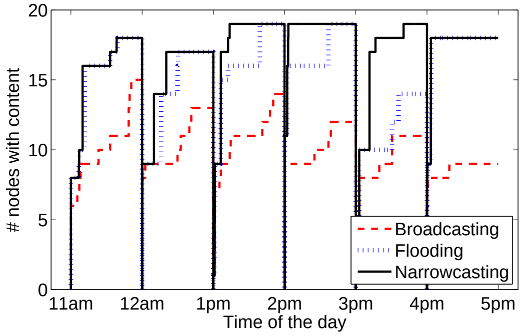

(a) day 1

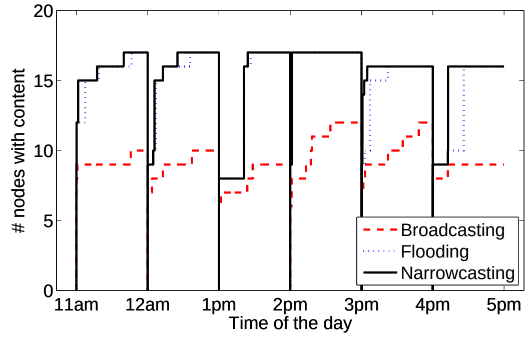

(b) day 2

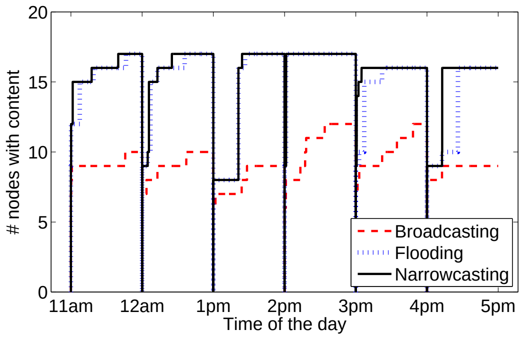

(c) day 3

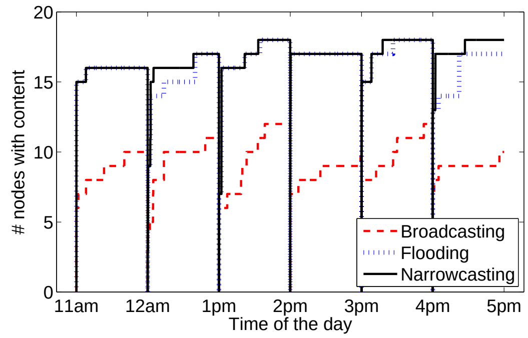

(d) day 4

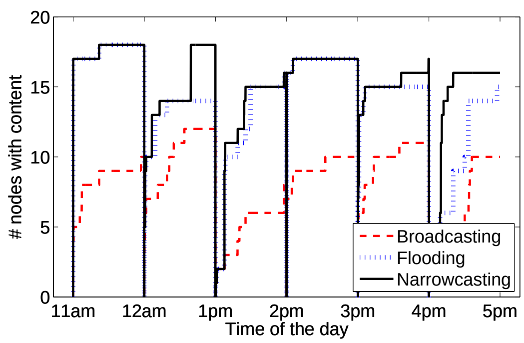

(e) day 5

Figure 1: Reachability vs. Delay: temporal evolution of the number of mobile nodes in possession of the latest short message (1500B) issued by a fixed content emitter (AP) on an hourly basis.

ity brought by the possibility to distribute any kind of digital content. The transmission range is, however, even more limited than analog low-powered systems (due to the ISM band transmission power regulations) i.e., typically 10-250m for IEEE 802.11 depending on the environment.

- *Multi-hop flooding (Flooding):* With this scheme, a message is flooded from the fixed AP over multiple hops composed only of mobile devices. The flooding scope is determined by the momentary placement of the nodes. If the network is partitioned, only the nodes that are in the same partition as the fixed AP can possibly receive a message. In order to distribute a message to newly arrived nodes, the flooding is repeated periodically. For our simulations, we have chosen a 1s period.

- *Delay-tolerant narrowcasting (Narrowcasting [2]):* This scheme relies on opportunistic communications between nodes that meet. Every time two or more nodes meet (they come in transmission range), they associate randomly with each other and solicit new content. Content is then broadcast so that other nodes that might not have the content can also benefit from the content exchange. When two associated nodes are done exchanging messages, they associate with another peer in range if there is any to make sure everyone has the content. Otherwise, they wait until a new node becomes reachable to associate and exchange potential new content. This approach can relay messages even when the network is partitioned with the help of mobility.

We have chosen these different dissemination strategies since they imply different node cooperation (in forwarding and storing) and complexity of implementation (in the network stack layers involved). On the one hand, the *broadcast* strategy does not require nodes' collaboration and only implies link-layer connectivity to the AP. On the other hand, the *flooding* and *narrowcasting* strategies rely on node collaboration. The former requires nodes to forward content to neighbors while the latter requires an additional buffering capability. From the implementation side, *flooding* requires to manage broadcast IDs while *narrowcasting* requires a more complex processing with a protocol managing content solicitations and retrievals.

The kind of content we consider to be spread with these strategies are general content relevant to an entire community such as news, weather forecasts, or schedules, which can take the form of short messages (e.g., press dispatch) or large files (e.g., recorded audio/videos such as news reports). In the sequel, we will evaluate the dissemination of two different workloads, a short message of 1500B fitting in the payload of one packet and a large file of 20MB. Note that we only consider flooding for short messages.

In order to compare these different strategies, we consider the following metrics:

- *Content Reachability (R):* This metric gives the number of nodes having the content after a defined time period. We are interested in the empirical mean as well as in the upper and lower bound of content reachability.

- *Delivery Delay (D):* The delivery delay gives the time from the publication of a content by the AP until it is delivered to to a node. We are interested in the upper limit of the delivery delay (D_{max}) i.e., the time required for the content to be distributed to all nodes as well as the mean delivery delay ($E[D]$) i.e., the average delivery delay seen by each node.

- *Hop count (H):* The hop count represents the number of nodes that carried or forwarded the content until delivery at a mobile node.

- *Carriers (C):* For a given content, this metric gives the number of times a node has delivered content to another node. It enables to evaluate if certain nodes have a higher importance in delivering content.

Note that the two last metrics are only relevant for the *flooding* and *narrowcasting* strategies.

4. NARROWCASTING IN AN OFFICE ENVIRONMENT

This section presents the methodology to evaluate and compare the different strategies in an office environment. We then present results in light of them being influenced by mobility.

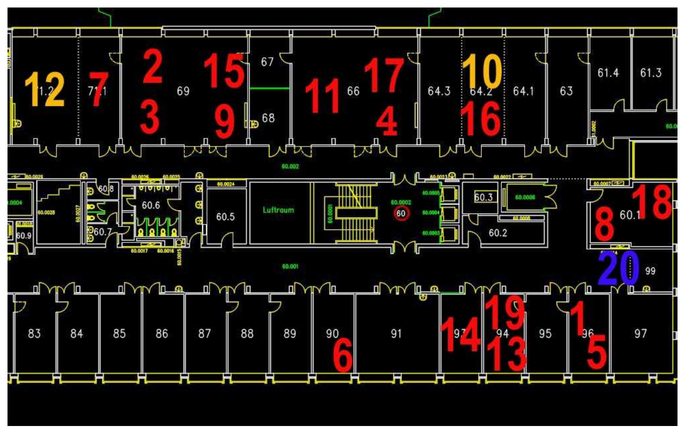

Figure 2: Map of the G-floor in the ETZ building at ETH Zurich. The numbers represent the users that are carrying the mobile devices and their respective desks (except for #20 which is the content emitter placement).

4.1 Evaluation methodology

We use real-world mobility traces that we collected at ETH Zurich. These traces are then replayed in a simulator we have developed. The setup used to collect these traces consists of 20 HP iPAQs probing their neighborhood every half a second using their integrated IEEE 802.11b chipset operated in ad hoc mode. [3] We asked nineteen test users to carry the devices during five consecutive working days (Monday-Friday from 11am to 5pm). The test users were researchers, staff members, and students of a networking research lab, all working on the same floor having a size of 100 meters x 30 meters. The map in Figure 2 shows the users' desks. The test users were instructed to carry the iPAQs with them throughout the day. A majority of the test users were researchers and spent most of the time at their desks. The users became mobile mainly due to lunch and coffee breaks, for going to the rest room,

[2] In the sequel of this paper, we use the term "narrowcasting" to refer to this strategy although all previous strategies are aimed at narrowcasting as defined in the Introduction.

[3] Note that a random time offset is added to the measurement period to avoid any undesired synchronization problem.

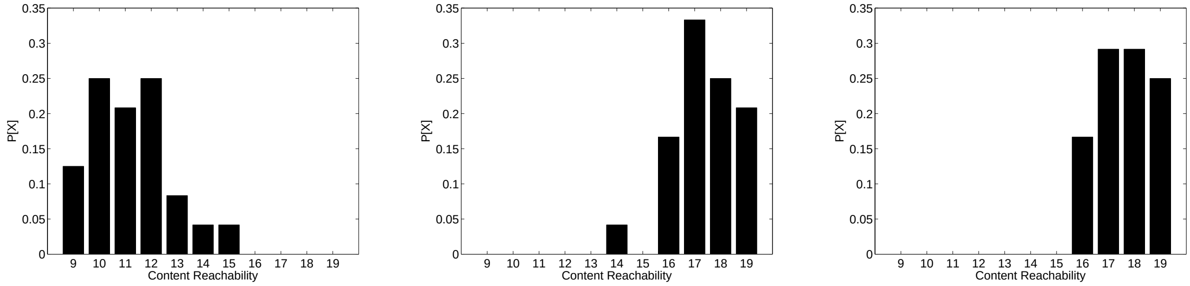

(a) Broadcast (b) Flooding (c) Narrowcasting

Figure 4: Reachability distribution after one hour for short messages ($19 = 100\%$).

Table 1: Reachability and delay statistics.

	Broadcast	**Flooding**	**Narrowcasting**
Reachability (short msg) $E[R \mid t = 3600]$	min=47%, max=79%, mean=59%	min=74%, max=100%, mean=91% (+32%)	min=84%, max=100%, mean=92% (+33%,+1%)
Reachability (large file)	min=37%, max=47%, mean=40%		min=42%, max=95%, mean=68% (+28%)
Delivery delay (short msg) $E[D_{max}]$, $E[D]$ (in seconds)	$E[D_{max}]$=2447.5 $E[D]$=483	$E[D_{max}]$=1571 (-36%), $E[D]$=304 (-37%)	$E[D_{max}]$=1264 (-48%,-20%), $E[D]$=227 (-53%,-25%)
Delivery delay (large file) (in seconds)	$E[D_{max}]$=2797, $E[D]$=1260		$E[D_{max}]$=1270 (-55%) $E[D]$=853 (-32%)

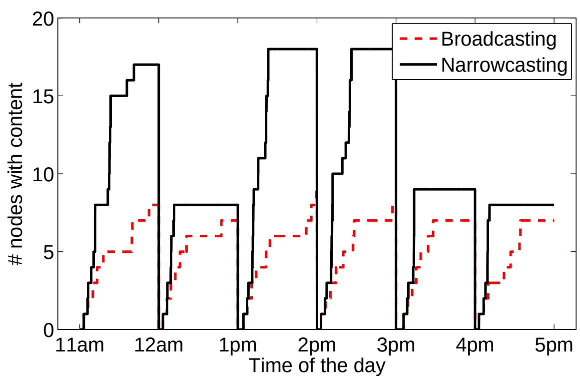

Figure 3: Reachability vs. Delay: dissemination of large files (20MB) updated every hour during day 1.

picking up printouts in the hallway, or meeting each other for discussions. Only few test users occasionally left the building or the campus for a limited period.

Having the collected traces for five consecutive days, we simulate the three different data dissemination schemes. One dedicated iPAQ that was installed next to the lab's coffee machine during the probing phase is considered as the content emitter (node 20 on the map). Our simulations evaluate how quickly and to what extent a new message that is updated on a hourly basis at this fixed iPAQ disseminates to the mobile devices carried by the test users. The radio model we use for our simulations is based on the reception probability we observe during the experiments. That is, we model the link delivery probability from node a to b as the fraction of received packets at b over the number of packets sent by a node a.

4.2 Reachability vs. Delivery Delay

We first investigate how the different dissemination schemes affect the delivery delays of the messages. We consider relatively short messages of 1500B that can be delivered in one single IP packet. Figure 1 shows for all one-hour period of each day the evolution of the number of device in possession of the latest message issued by the access point. Considering day 1 in this Figure, the resulting distributions reveal quite much about the network topology and the node mobility. In the *broadcast* scheme, every node has to be in range of the access point in order to retrieve the message. The sharp increase after the moment a message is published reflects the number of nodes in transmission range of the access point at this time. Later retrievals are related to mobility, i.e., when a node moves in range of the access point. Interestingly, this strategy manages to make more progress around noon as later in the day. This reflects the increased mobility of the people during lunchtime. The best distribution is observed in the *narrowcasting* scheme. This scheme outperforms simple *flooding* due to its ability to relay messages when the network is partitioned. With *flooding* only nodes that are in the same partition as the access point can possibly receive a message. In contrast, *narrowcasting* manages to relay messages between partitions, mainly when one node moves from one partition to another. The best example for this is between 3pm and 4pm. After around ten minutes, one node must have moved to a partition with seven nodes that is isolated from the access point since *flooding* was not able to make progress at that time. When both *flooding* and *narrowcasting* manage to make progress at the same time like a couple minutes after 4pm, a node must have moved in such a way that an isolated partition merged with the partition of the access point.

As explained previously, our methodology allows replaying the same scenario with different workloads. For example, we can answer the question on how the distribution would have been if the messages created at the access point were 20MB instead of 1500B. The resulting distribution times are plotted in Figure 3 for day 1 (except for the *flooding* that we omit since *flooding* fails to deliver an entire message that exceeds one IP packet's payload in the presence of links with intermediate loss rates). As we can observe with the *broadcast* scheme, the distribution time is now limited by the large message size. The message retrieval with no relaying does no longer significantly depend on the time of the day. Since the mobility and the resulting link qualities are identical as before, we can hence conclude that short transfer opportunities resulting from a node traversing the wireless range of the access point are not sufficiently long any more to transfer a message of 20MB. *Narrowcasting* manages to distribute well the message created at 11pm, 1pm, and 2pm, but comparatively fails for the messages of 12pm, 3pm, and 4pm. Again, this is because short-lived contacts are not sufficient to relay messages between partitions as it was possible with the 1500B message.

We now look in more details at the performances of delivery averaged over the five days. Figure 4 gives the content reachability distribution for all strategies. Clearly, we observe the higher dissemination for *flooding* and *narrowcasting*. We can also remark that *narrowcasting* performs slightly better than *flooding*.

Table 1 presents numerical values of reachability and delay statistics. Statistics are computed for each of the strategies and the different workloads and averaged over each one-hour period for all days. We give the proportion of nodes with content (R) one hour after it has been updated (minimum, maximum, and mean reachability). We also compute the mean delay to reach the maximum content reachability ($E[D_{max}]$) [4] , and the mean delay for all deliveries ($E[D]$). We observe that the content reachability reaches on average 92% of the nodes with the *narrowcasting* strategy, i.e., 33% better than the *broadcast* strategy but only 1% better than *flooding*. Actually, the main difference between *flooding* and *narrowcasting* is the delivery delay. Considering the all network of diffusion, the delay until the maximum content dissemination is reached is 20% less with *narrowcasting* than with *flooding* and 48% less than *broadcast* for the short message case. Now from an individual perspective, the mean delay for the content to reach any node is less than 4 minutes with *narrowcasting* of the text message and 14 minutes for the file. This is 53% and 32% faster than simple *broadcast* for the message and the file respectively.

4.3 Characterization of content propagation

We now look in more detail how the content is disseminated for the *flooding* and *narrowcasting* strategies, that is, for how many hops the content is stored and forwarded before delivery. In Figure 5, we observe a monotonous decrease of the number of hops for the *flooding* strategy while with *narrowcasting* we observe that nearly 70% of the content is delivered directly from the AP or from two hops (i.e., from a node which retrieved the content from the AP). This clearly shows the benefit of node mobility with *narrowcasting*. This also explains why the content is disseminated more rapidly with this scheme. Actually, the *flooding* strategy requires certain specific nodes to be in a certain location to merge different partition in order to increase the content reachability. With *narrowcasting*, content dissemination follows a spraying process. It exploits the mobility of nodes retrieving the content from the AP.

We also study if some nodes play a more important role in delivering content. Figure 6 shows for each individual nodes the number of messages that it has delivered averaged over all days for each one hour period. We observe that node 20 which corresponds to the AP has a higher number of deliveries, between 5 and 6 for both schemes, compared to other nodes. This value can be explained by looking at the map in Figure 2 where we can clearly observe 4 nodes in the vicinity of the AP (i.e., nodes 1, 5, 8, and 18) when they are at their desks. Flooding performs similarly with this respect. More surprising is that the load is relatively distributed between the mobile nodes. We cannot distinguish clearly nodes having a higher role in distributing content (i.e., carriers). Each node delivers on average the content to at most two other nodes and the mean delivery is below 1 (without accounting for the AP). This means that nodes have only to carry a message and deliver it to at most two other nodes. This last results clearly indicates that no specific content dissemination policy must be enforced to reach a fair cooperation. Now coming back to previous results, we have observed that cooperation in the *narrowcasting* strategy not only allows a better content reachability but also a shorter time to delivery.

5. RELATED WORK

Despite its potential growth in the coming years, the networking community has not considered narrowcasting except for commercial aims. Harrison et al. [6] propose to enhance the diffusion of advertisements in public places using wireless networks.

Research on delay tolerant networking is very related to this work. We do not aim for unicast communications in our work but for narrowcast communications i.e., broadcast communications to communities. Multicast routing in DTN has been addressed in [17]. While the goal of our system is also to deliver data to a group of people, our approach is decoupled from any multicast semantics (such as group memberships, et cetera). The infostation concept [12, 15, 16] is akin to our proposal. [16] by Yuen et. al. considers a similar setting as ours with one infostation. It assumes noncooperation of nodes and stipulates nodes will follow a social contract where exchange of data occurs if only each entity has interesting data to exchange to the other. Papadopouli et al. [12] also studies noncooperation in the Infostation context along with power conservation and wireless coverage. We differ in that we don't consider the same (non)cooperation concepts. And besides, both works consider mobility models that are tractable for analytical performance analysis. In our case, we use real-world mobility traces. Cooperation is also studied in a similar context to ours. In [18] by Helgason et al., an analytic study of delay-tolerant narrowcasting is presented. It compares the broadcasting scheme to narrowcasting; also when narrowcast forwarding is limited to a single copy or to a limited time period.

Vehicular ad hoc networks are also quite compelling and highly studied as a way to communicate over opportunistic contacts between vehicles [3, 11].

In the Zebranet [7] project, an opportunistic network was used to collect zebra mobility information in the desert of Kenya where there is no available communication infrastructure by relaying data between devices attached to the zebras.

The Delay Tolerant Network Research Group (DTNRG) [1] has proposed an architecture [4] to support messaging that may be used by delay tolerant applications. The architecture consists mainly of the addition of an overlay, called the bundle layer, above a network transport layer. Messages of any size are transferred in bundles in an atomic fashion that ensures node-to-node reliability. Since, several works have extended this concept. TACO-DTN [13] is a content-based dissemination system for delay tolerant networks implemented as a publish/subscribe system. It is mainly designed to

[4]Note that the maximum content reachability varies between hours and days.

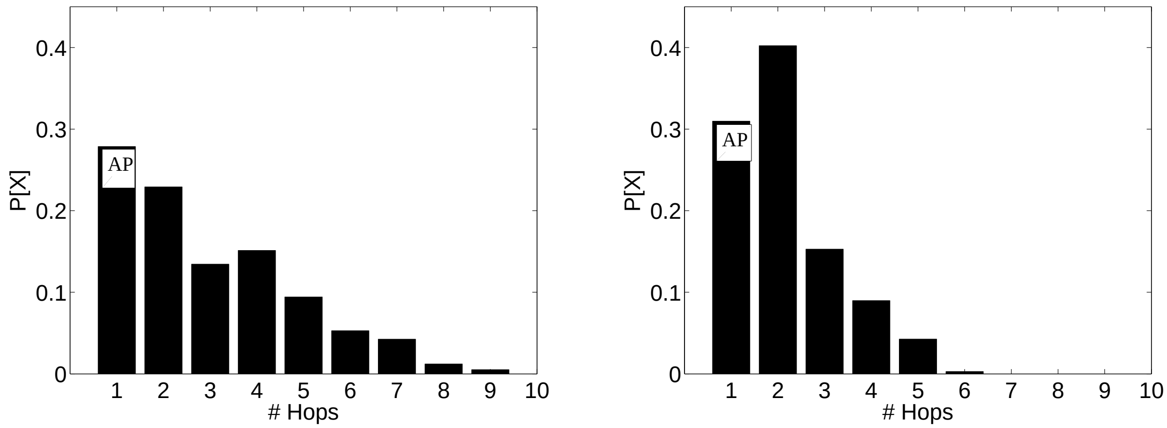

(a) Flooding (b) Narrowcasting

Figure 5: Hop count distribution for short messages deliveries.

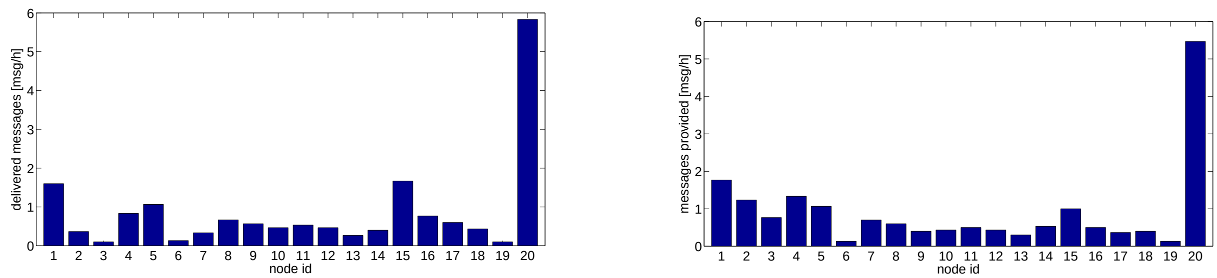

(a) Flooding (b) Narrowcasting

Figure 6: Messages delivered by each node during a one-hour period averaged over all days.

be time-aware which adds a level of complexity we do not need to consider in our case. [2, 8] by Lenders et. al. introduces the paradigm of delay-tolerant broadcasting, which is a receiver-driven content dissemination solution for unregulated content. Although having a different purpose, the narrowcasting strategy can use the same technology. Narrowcasting is actually a particular case of delay-tolerant broadcasting with only one source generating content, where dissemination is *forced* (and no more receiver-driven) and where the caching strategy is set to cache all content.

Contact patterns of human mobility have been analyzed in the Haggle project [5]. This project aims at developing an application-independent networking architecture for delay-tolerant networks. De Lara et al. [14], Leguay et al. [9], and Lindgren et al. [10] have studied different content distribution schemes using Bluetooth and Wifi connectivity traces. These works provide useful insights into the impact of human mobility on the opportunistic contacts in real environments. However, these works only focus on unicast communications whereas we target the narrowcasting of content. Besides and most important of all, the mobility traces used in these works are either coarse-grained in time (i.e., interval of 2 to 5 minutes) for Bluetooth measurements or coarse-grained in space with WiFi measurements where nodes are associated to an AP. In our case, we use Wifi measurements in ad hoc mode performed at beacon intervals of half a second (0.5s).

6. DISCUSSION

The content emitter placement is an issue, which might greatly influence the performance of the different studied schemes. In traditional broadcasting systems, the radio emitter is usually mounted in a central and high point to deserve a large area. With our strategies, we heavily rely on node mobility to spread content. As a consequence, the content emitter must also be centrally located but from a mobility viewpoint. We do not study this issue hereinsince for the same reasons mentioned earlier, the location of the AP might be imposed by external constraints (location of electrical power plug and/or wired Internet access plug) or left to ones own judgment based on the intuitive assumption that the more centrally located, the better dissemination will perform. In our case study, we have chosen the coffee corner at our floor as placement since it is one of the most visited locations where people meet.

The proposed flooding and narrowcasting mechanisms are rather simple. Yet, we wanted our study to be simple in order to understand the fundamental behavior of narrowcasting. Many optimizations could be brought to these schemes. One optimization could be to allow content transfers to be resumed after disruptions due to node mobility for example. This would have been particularly profitable for large files dissemination since with our schemes, exchanged had to be started again from the beginning if interrupted before the end. Possible solutions could be to use network coding strategies or use the principle of random dissemination of chuncks. This latter technique is used for delay-tolerant broadcasting [2, 8] and can be easily adapted for narrowcasting. Going further, chuncks could be disseminated in a smart way instead of random dissemination to optimize delivery. Such kind of optimizations are the subject of further investigations we intend to carry.

7. CONCLUSION

We have studied content dissemination/narrowcasting in mobile ad hoc environments where general and delay-tolerant content aimed at an entire community (e.g., local news, special events) is regularly emitted from a single emitter (e.g., daily, hourly). We provide a comprehensive comparison of different cooperative and non-cooperative dissemination strategies. Overall, we show empirically using real-world mobility traces the benefits of nodes cooperation (i.e., relaying and caching) combined with node mobility to efficiently spread content. The proposed setup and communication traffic is particularly adapted to communities sharing a common interest and requiring a common source of information. This can be settled communities (e.g., work teams, military bases, remote villages) or transient communities gathering occasionally (e.g., conferences, exposition halls, tactical operations). Future works will consider the case of several emitters to span larger communities such as campuses, airports or metropolitan areas.

8. REFERENCES

[1] Delay Tolerant Network Research Group (DTNRG). http://www.dtnrg.org.

[2] The PodNet Project. http://http://podnet.ee.ethz.ch/.

[3] John Burgess, Brian Gallagher, David Jensen, and Brian Levine. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. In *Proceeding of IEEE INFOCOM*, Barcelona, Spain, April 2006.

[4] S. Burleigh, A Hooke, L. Torgerson, K. Fall, V. Cerf, B. Durst, and K. Scott. Delay-tolerant Networking: An Approach to Interplanetary Internet. *IEEE Communications Magazine*, 41(6):128–136, 2003.

[5] Augustin Chaintreau, Pan Hui, Jon Crowcroft, Christophe Diot, Richard Gass, and James Scott. Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms. In *Proceedings of IEEE INFOCOM*, Barcelona, Spain, April 2006.

[6] J.V. Harrison and A. Andrusiewicz. Using wireless networks to enhance narrowcasting in public spaces. In *Proceedings of the IEEE GLOBECOM*, 2004.

[7] Philo Juang, Hidekazu Oki, Yong Wang, Margaret Martonosi, Li-Shiuan Peh, and Daniel Rubenstein. Energy-Efficient Computing for Wildlife Tracking: Design Tradeoffs and Early Experiences with ZebraNet. In *Proceedings of the Tenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-X)*, San Jose, CA, USA, October 2002.

[8] Gunnar Karlsson, Vincent Lenders, and Martin May. Delay-Tolerant Broadcasting. *IEEE Transactions on Broadcasting*, 51(1), March 2007.

[9] Jeremie Leguay, Anders Lindgren, Timur Friedman, James W. Scott, and Jon Crowcroft. Opportunistic Content Distribution in an Urban Setting. In *Proceedings of the ACM SIGCOMM workshop on Challenged Networks (CHANTS)*, Pisa, Italy, August 2006.

[10] Anders Lindgren, Christophe Diot, and James W. Scott. Impact of Communication Infrastructure on Forwarding in Pocket Switched Networks. In *Proceedings of the ACM SIGCOMM workshop on Challenged Networks (CHANTS)*, Pisa, Italy, August 2006.

[11] J. Ott and D. Kutscher. A Disconnection-Tolerant Transport for Drive-thru Internet Environments. In *Proceedings of IEEE INFOCOM*, Miami, USA, March 2005.

[12] M. Papadopouli and H. Schulzrinne. Seven degrees of separation in mobile ad hoc networks. In *Proceedings of the IEEE GLOBECOM*, 2000.

[13] Giuseppe Sollazzo, Mirco Musolesi, and Cecilia Mascolo. TACO-DTN: A Time-Aware COntent-based dissemination system for Delay Tolerant Networks. In *Proceedings of the First International Workshop on Mobile Opportunistic Networking (Mobiopp)*, Puerto Rico, June 2007.

[14] Jing Su, Ashvin Goel, and Eyal de Lara. An empirical evaluation of the student-net delay tolerant network. In *Proceedings of MOBIQUITOUS*, San Jose, California, July 2006.

[15] W. H. Yuen, R. D. Yates, and S. C. Mau. Exploiting data diversity and multiuser diversity in mobile infostation networks. In *Proceedings of the IEEE INFOCOM*, 2003.

[16] W. H. Yuen, R. D. Yates, and S. C. Mau. Noncooperative content distribution in mobile infostation networks. In *Proceedings of the IEEE WCNC*, 2003.

[17] W. Zhao, M. Ammar, and E. Zegura. Multicasting in delay tolerant networks: Semantic models and routing algorithms. In *Proceedings of the Sigcomm Workshop on Delay Tolerant Networking*, August 2005.

[18] Ólafur Ragnar Helgason and Gunnar Karlsson. On the effect of cooperation in wireless content distribution. In *Proceedings of the IEEE/IFIP WONS 2008*, Garmisch-Partenkirchen, Germany, January 2008.