DTN content sharing among commuters

Armir Bujari and Claudio Enrico Palazzi*

Department of Pure and Applied Mathematics,
University of Padua,
Via Trieste 63 – 35123, Padua, Italy
Fax: +39-049-827-1426
E-mail: abujari@math.unipd.it
E-mail: cpalazzi@math.unipd.it
*Corresponding author

Marco Roccetti and Gustavo Marfia

Department of Computer Science,
University of Bologna,
Via Mura Anteo Zamboni 7 – 40127, Italy
Fax: +39-051-209-45-10
E-mail: roccetti@cs.unibo.it
E-mail: marfia@cs.unibo.it

Abstract: When discussing vehicular networking we have also to consider the multitude of commuters daily travelling to/from work through public transportation means such as buses, metros, and trains. Indeed, through their smartphones, commuters can generate traffic while moving, thus embodying a significant part of an urban communication scenario. In this context, we discuss a way to port the P2P paradigm into mobile disconnected networks composed by commuters' smartphones. Different from classical P2P file-sharing, users could exchange data in proximity with each other; they can also leverage on peer mobility and encounters in order to extend a requesting peer's reach-area to other local disconnected networks. This is achieved by implementing a DTN-like store-delegate-and-forward communication model over underlying social networks where a peer can delegate unaccomplished data downloads to frequently encountered peers. We show how involving frequently encountered peers/commuters improves the success ratio for the data search request while reducing useless communication overhead.

Keywords: mobile ad-hoc networks; MANETs; delay tolerant; DTN; opportunistic file sharing; mobility; working day movement model.

Reference to this paper should be made as follows: Bujari, A., Palazzi, C.E., Roccetti, M. and Marfia, G. (2012) 'DTN content sharing among commuters', *Int. J. Satellite Communications Policy and Management*, Vol. 1, Nos. 2/3, pp.138–151.

Biographical notes: Armir Bujari is a PhD student in Computer Science at the University of Padua, Italy, and is completing the second year of his programme under the supervision of Professor Claudio E. Palazzi. His research is focused on wireless networking and delay-tolerant communication for mobile networks. He previously completed a Master's degree in Computer Science, Summa Cum Laude, at the University of Padua.

Claudio Enrico Palazzi is an Associate Professor with the Department of Pure and Applied Mathematics, University of Padua, Italy. He received his MS in Computer Science from UCLA in 2005, his PhD in Computer Science from UniBO in 2006, and his PhD in Computer Science from UCLA in 2007. From 2007 to 2010, he is an Assistant Professor at the Department of Pure and Applied Mathematics of the University of Padua. His research is primarily focused on protocol design and analysis for wired/wireless networks, with emphasis on network-centric multimedia entertainment and vehicular networks.

Marco Roccetti received his Italian degree in Electronic Engineering from the University of Bologna, Italy in 1989. He has been a Professor with the Computer Science Department, University of Bologna since 2000. His research interests include digital audio and video, consumer multimedia and computer entertainment.

Gustavo Marfia received his degree in Telecommunications Engineering from the University of Pisa in 2003 and a PhD in Computer Science from UCLA, in 2009, respectively. He is currently a Postdoctoral Fellow at the Computer Science Department, University of Bologna. His research interests include vehicular, medical, multimedia networks and systems.

1 Introduction

Our cities' roads are complex ecosystems where cars have to coexist with public transportation means (i.e., bus and metro train), bicycles, and many pedestrians. Yet, most of current research in this scenario is focused just on car-to-car (C2C) or car-to-infrastructure (C2I) communication. Instead, to design new city 2.0 able to provide pervasive innovative services, it is crucial to broaden our vision, including and eventually integrating all entities and even social relationship among them (Marfia et al., 2011; McNamara et al., 2008; Rachuri et al., 2011). This work goes in this direction and aims at devising a new communication-based service specifically involving people daily travelling to work, i.e., commuters.

As most commuters travel at the same time of the day and along the same path, they often share the same transportation means, even if not knowing each other. Yet, this creates a potential scenario where their mobile devices could periodically establish communication. In this paper we show how this interesting property can be integrated with technology deriving from ad-hoc opportunistic networking, peer-to-peer (P2P) communication, and delay/disruptive tolerant network (DTN) to devise innovative services and communication paradigms for city users. It is also important to notice that our work represents a necessary step in creating a holistic vision of urban communication scenarios.

More in detail, thanks to the evolution of mobile device technology, smartphones have evolved from simple voice communication means into powerful devices able to handle complex multimedia documents. It is hence interesting to export a popular application such as P2P file-sharing into the new scenario of mobile networks. In this context, different from the classic wired P2P file-sharing approach, mobile users could exchange multimedia data when in proximity of each other through opportunistic, proximity-based P2P ad-hoc links (Ferdinando et al., 2003; Jung et al., 2007; Palazzi, 2004).

However, mobile network communication cannot always rely on the continuous end-to-end path between source and destination. Due to mobility, communication links are transient and short in time (Klemm et al., 2003); in sparse mobile ad-hoc networks (MANETs) this path might not even exist. Also, when considering our applicative scenario, low node density could imply lower data population in the overlay network, which could undermine the utility of a P2P file-sharing application.

In this context, it is interesting to consider mobile devices belonging to people realistically moving in an urban scenario. Some of these people will never meet or will meet just seldom: for instance, two customers entering in a shop almost simultaneously. Other users will repeatedly meet even without knowing each other: for instance, commuters utilising the same public transportation means (almost) every day. Indeed, as anticipated, commuters embody a particularly interesting case study as their daily home-work-home routing provides some sort of predictable link establishment among them, even in the considered heterogeneous mobile scenario.

We are inspired by DTNs for their ability in supporting communication with intermittent connectivity, long/variable delay, and high error rates, as in the wireless mobile scenario of our everyday lives (Cerf et al., 2007; Gaito et al., 2009). More in detail, we have created a smartphone application, namely M2MShare, which ports the DTN paradigm into the mobile world, addressing the node density issue by providing means for an asynchronous data exchange similar to that of DTNs. The idea of a DTN is modelled in an infrastructure-less environment where both the source of the request and the destination of the data are the same entities and where intermediary nodes (servants) can store-delegate-and-forward back the requested user data toward the source. In our application scenario the forward route consists of a single hop that is, from servant toward the request originator. Our focus is on exploring a new mechanism that allows peers to explore and download multimedia content available outside their reach area, provided in other local disconnected networks.

M2MShare tries to reach data in other local disconnected networks by leveraging on node mobility and existing, even if invisible, social relationships. Our application monitors periodic encounters among users regardless of the fact that they may or may not be aware of this social proximity. Content search is performed by both enquiring encountered nodes but also by assigning the file search request (delegations) to some of them. To avoid excessive transmission overhead, delegations are assigned only to nodes that will be probably met again in the future: heuristically, nodes that have a recent history of encounters.

In summary, our contribution regards a new example of opportunistic communications specifically designed for social DTNs composed by smartphones. Our solution also includes

1. a protocol that dynamically establishes forward routes (delegations) by exploiting users' real life periodic encounters
2. a smart criteria for download request delegation so as to speed-up file transfer through parallel operations, while reducing the transmission redundancy.

Commuters sharing part of their everyday paths between home and workplace embody a very representative example of users that frequently meet even if not knowing each other, sometimes even not noticing this frequent proximity. We have hence specifically included this kind of users in our experiments on M2MShare. Specifically, we have utilised a well-known tool for DTNs named THE_ONE and tested our M2MShare against a naïve model in a realistic scenario representing the town of Helsinki, its roads and its public transportation system. In this scenario, we employed the realistic working day movement model with nodes moving from one area to the other by feet, car, or bus, depending on the location and distance of departure and arrival.

The remainder of this paper is organised as follows. Section 2 summarises background information related to our work. In Section 3, we present M2MShare. Section 4 provides some insights on the servant election strategy, also we show some experiments which prove that the delegation technique serves its purpose. In Section 5, we introduce the file division strategy employed for delegations, demonstrating its efficiency against other strategies. Finally, concluding remarks are given in Section 6.

2 Background

In this section, we present some background information regarding the applicative domains concerning our work. We begin by stating some of the challenges encountered when deploying P2P on MANETs, further proceeding in our discussion with the DTN communication model.

2.1 P2P on MANETs

Mobile devices are providing functionalities previously restricted to desktop computers. Such devices can interconnect using broadband wireless network interfaces. Local storage and computing capacity are sufficient for storing and playing high-quality audio, movie files and complex multimedia presentations. Due to their low cost, they are growing in popularity and may at some time become the dominant mode by which users reach the internet, in a way we could consider them as an *opportunistic extension* of the today's internet. Consequently, they will have an important role in future P2P overlays. However, mobile networks have some important characteristics that differentiate them from the internet and that affect their interaction with the classical P2P wired overlays.

The static connection approach employed by such systems in the wired case is not suitable for MANETs; this due to the frequent mobility disconnections and to the maintenance overhead (Klemm et al., 2003; Wu, 2005). Also, in sparse mobile environments, there can be no path between the source and the destination, and when considering our applicative scenario low node density could mean lower data population in the overlay, thus undermining the functioning of a P2P file-sharing application (Ferdinando et al., 2003; Klemm et al., 2003). For these reasons, synchronous communications are not sufficient in such an environment and other communication approaches are needed.

Mobile devices have energy issues, thus the need to preserve energy is essential (Gurun et al., 2006; Wu, 2005). Power management involves a combination of techniques, including network adapters that can trigger power resume of the host while offloading certain network activity and network protocols that reduce power consumption. In current overlay designs, a mobile peer that goes into power-saving mode is treated as a node that has left the overlay. Another, not less important characteristic that differentiate the wireless from the wired world is the communication technology. The wireless medium, as opposed to the wired communication technology, is known to be error-prone and bandwidth limited. Limited bandwidth alone could mean lower data quantity transferred, adding to this the mobility factor and low peer density we might need other ways for reliable synchronous and asynchronous data transfer.

Concluding, mobile environments face multiple challenges and traditional solutions for P2P over fixed networks may need to be redesigned when applied in the wireless mobile world.

2.2 Delay/disruption tolerant network

The TCP/IP protocol suite has been a great success at interconnecting communication devices across the internet (Stevens, 2000). Yet, it is based on the assumptions of having connectivity on wired links, continuously connected end-to-end devices, and low-delay paths between the sources and destinations.

Nowadays, with the growing number of devices equipped with wireless interfaces, users increasingly find themselves in different types of networking environments than the aforementioned ones. These environments, spanning from globally connected networks such as cellular networks or the internet to the entirely disconnected networks of stand-alone mobile devices, impose different forms of connectivity. In particular, differently from the wired internet communication, MANETs composed by smartphones are characterised by intermittent connectivity, long/variable delays, asymmetric data rates and high error rates.

A DTN originally conceived for communication in outer space, is an overlay of networks, including the internet. DTNs support interoperability of networks by accommodating long delays between and within them, and by interfacing different network types. In providing these functions, DTNs accommodate the mobility and limited power of evolving wireless communication devices (smartphones). They employ a store-and-forward message switching: whole messages or pieces of such messages are forwarded from a storage place on one node to a storage place on another node, along a path that eventually reaches destination.

Another peculiarity of DTNs is that their routers need persistent storage for their queues as opposed to internet routers that use short-term storage provided by memory chips. This, for the following reasons:

- a communication link to the next hop may not be available for a long time
- one node in a communicating pair may send or receive data much faster or more reliably than the other node
- a message, once transmitted, may need to be retransmitted if an error occurred at an upstream (toward destination) node or link, or if an upstream node declined acceptance of a forwarded message.

The DTN paradigm is still an active area of research and a lot of other features would be worth mentioning but are outside of the scope of this work [for more insights refer to DTNRG (2002)]. However, as we discuss in this paper, it represents an interesting option for opportunistic communication among smartphones.

3 M2MShare: modus operandi

M2MShare uses the bluetooth to create a P2P overlay networks, which would allow the automatic exchange of multimedia content among smartphones (Bray and Sturman, 2001). Our software automatically initiates a search by broadcasting the user's query request toward other bluetooth enabled devices. Once an answer is received, data found to match the criteria will be automatically requested for download.

Given that energy consumption is a problem for handheld devices the notion of active sessions was introduced. An active session is a period of time in which the software is functioning and performing its duties; session periods are configurable through the graphical user interface. This way, the user can set the application to look for requested contents only in certain periods of the day; for instance, when commuting or during lunch time in the cafeteria, so as to be active only when it could be useful since the high number of other peers around.

As stated earlier, we leverage on mode mobility to reach multimedia content on other local disconnected networks. This is achieved by introducing an asynchronous communication model, store-delegate-and-forward where a peer delegates and unaccomplished/unsatisfied request to other peers in the overlay network. Delegating requests to all encountered peers is energy and bandwidth consuming. Also, it would be useless to assign requests to peers that will never be met again in the future. To avoid this M2MShare exploits social relations and delegates requests only to *frequently encountered peers*, peers whom are expected to be encountered again in the future (refer to Section 4.1).

We could say that we exploit social relations among users to determine possible request delegations. This is not a new approach as context and social relations are already studied in opportunistic data transmission (Boldrini et al., 2010). Yet, we utilise even unknown social relations by connecting users that have to pass by the same geographical location at the same time frequently enough. To assume that a candidate for a request delegation could be met again in the future (so that she/he will be able to deliver contents possibly found), we use the history of previous encounters.

4 Servant election

M2MShare implements an asynchronous communication model between peers, where a client peer can delegate an unsatisfied, unaccomplished task to a servant peer. While in DTNs there are pre-deployed entities that store-and-forward data along the destination path (routers), M2MShare achieves this functionality in an infrastructure-less environment, where forward routes are established dynamically. In this section, we state the underlying assumptions behind the DTN module design and show some experiments which demonstrate that the delegation technique serves its purpose.

4.1 Frequently encountered device

While research into routing in mobile environments is not new, researchers have for many years assumed node encounters to be random. In reality, mobile nodes are of course used by people, whose behaviours are better described by social models. This opens up new possibilities for routing, since the knowledge of behaviour patterns allows better routing decisions to be made (Keränen et al., 2009; Palazzi and Bujari, 2010; Rhee et al., 2007; Samal, 2003). In our work, we exploit this idea of social relations between users operating mobile wireless devices and provide a proof of concept implementation, adopting a DTN type solution for the mobile disconnected networks.

Delegating an unaccomplished task to all the peers in the established overlay network would be bandwidth and energy consuming; a smarter criterion is hence needed to choose one peer instead of others in order to decrease the number of transmissions when possible. It is sound to delegate tasks to peer devices operated by users expected to be encountered again in the future. In fact, it would be useless to assign a task to a node that will never be met in the future. As a heuristic to predict future probability of encountering a node we use the frequency of past encounters. This heuristic is useful for our purpose as it is able to detect, for instance, commuters sharing part of their everyday home-work-home path. Note that it is important that the considered commuters share only part of their path; this way they will both have the time to exchange tasks/files and have the chance to reach other town areas with different group of users (so as widen the search area).

In M2MShare, only frequently encountered devices can become servants. The concept of frequently encountered devices changes in time, adapting to the observed dynamics; this because the contact rate of a single device operating M2MShare might vary from day to day. Moreover, some devices frequently encounter many other devices, while others encounter a small number of them. In the first case we want to be more selective, choosing as servant candidates only those devices that are more frequently met. Therefore, we consider a device as 'frequently encountered' if met, for instance, several times during each day. In the second case, we cannot be such selective and, in order to have a certain number of devices being elected as 'frequently encountered', we have to lower the system expectations (thresholds) and consider nodes that are met, for instance, even just once every two or three days. The servant election algorithm takes these cases into consideration and tunes the parameters accordingly to the observed history.

4.2 Evaluation of delegation efficiency

We compare the efficiency of our system, which employs delegations, against the system where there are no delegations and file exchange is initiated only when a peer holding the requested data is found directly by the file requester. The metric we study is the average found time (Ft_{avg}) for a specific multimedia file. The found time (Ft) is the time interval between the first delegation made and the time an output return for that specific file is received. If no delegation is made and the first file request is satisfied directly between a file owner and the file requester, the Ft_{avg} is equal to zero.

To this purpose we implemented the two protocols in THE_ONE (Keränen et al., 2009), a well known DTN simulation environment. For sake of simplicity, delegations are one hop only, i.e., a servant peer cannot further delegate the task to other frequently encountered peers of its own. Also the software is always operational, there is only one active session configured that covers all day user activity.

We stated before that synthetic mobility models are not realistic and to this purpose we try to simulate a more realistic scenario provided by the working day movement model (Ekman et al., 2008) implemented for THE_ONE. In our scenario, we have a population of nodes (N) which emulate people operating M2MShare and are involved in their daily activities according to the cycle home-work-home. A node at home is inactive, thus the software is not operative. Nodes are uniformly distributed between the available districts in the Helsinki map available in THE_ONE and move along the streets from home to work and vice versa by feet, car, or bus, depending on the locations, distance and close availability of public transport lines. The simulation time is set to two days, during which there are two full cycles home-work-home; each cycle home-work-home amounts to 12 h. Therefore, the maximum simulation time during which a node can be active amounts to 24 h.

The file population (*Fp*) parameter denotes the percentage of nodes that posses the required file. The node requesting the file is randomly chosen between the population and we repeat the experiment 40 times in order to achieve more accurate results, independent from the initial user starting point.

In the first scenario (Figure 1), we consider *Fp* = 20%. The protocol not employing delegations (purple line in the chart) is not able to find any piece of the file during the simulation time when the considered nodes in the scenario are equal or less than 50. We have indicated this in the chart by assigning to Ft_{avg} a value of 24 h. This is due to the trivial strategy employed by the protocol and to the sparse environment. Even when able to find some node possessing the file (with $N \geq 100$), the time needed results bigger than when employing our solution for task delegation (blue line in the chart). Clearly, when increasing the *Fp*, even the number of nodes in the population that posses the data file increases; as a result, the time to retrieve the file decreases for both solutions.

A similar result is achieved also when considering a wider popularity for the required multimedia file (*Fp* = 50%, in Figure 2). However, in this case, the higher popularity of the requested file helps both solutions in finding the file possessor with a smaller Ft_{avg} than in the previous scenario.

Figure 1 Comparison between the trivial strategy (purple) vs. the strategy employing delegations (blue) and *Fp* = 20%

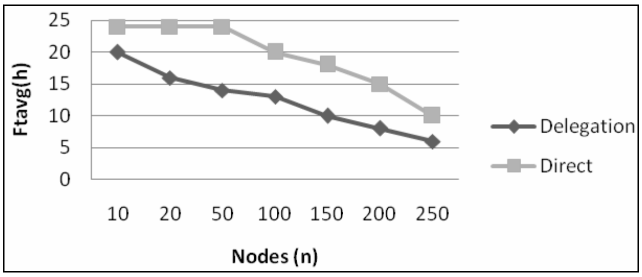

Figure 2 Comparison between the trivial strategy (purple) vs. the strategy employing delegations (blue) and *Fp* = 50%

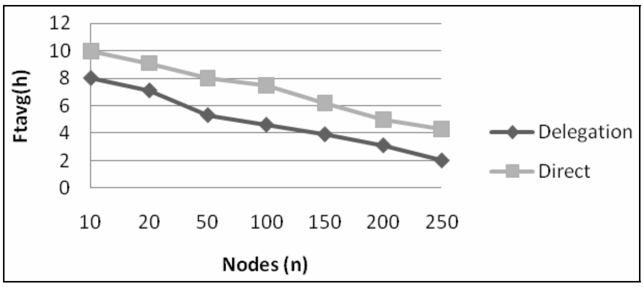

Finally, in Figure 3, the performances of the two compared solutions are very similar. This is due to the high file popularity among nodes (*Fp* = 80%): the chances of eventually finding a file possessor in a short time are clearly much higher.

Figure 3 Comparison between the trivial strategy (purple) vs. the strategy employing delegations (blue) and *Fp* = 80%

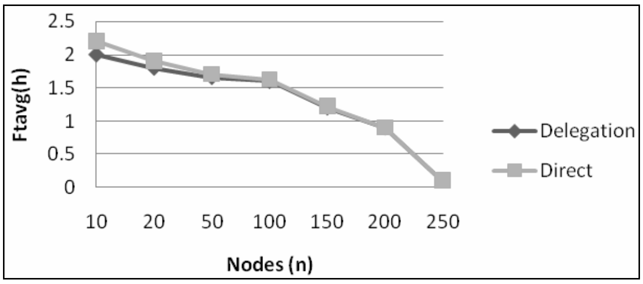

5 File division strategy

The majority of file transfer applications in the market follow a client-server paradigm where devices pair with each other for all the duration of data transfer. If a disconnection takes place, file transfer has to restart from the beginning and already downloaded data are of no use. Pairing for all the duration of file exchange is desirable if possible, taking in consideration the mobility of users in our scenario and that established connections are transient and short in time, the chances of this happening reduce drastically.

A common practise in P2P file-sharing is that of dividing the file into data chunks, which are the atomic transferable parts. This is a good starting point but taking into consideration the possibility of delegation of data search and the usage patterns mentioned before we require a more flexible file division strategy. M2MShare divides the file into data chunks. Although any node can retrieve the whole set of chunks, yet short time encounters between nodes may allow only for partial downloads of the requested file. Therefore, the file requester orders to each servant to start their possible downloads from a different chunk of the file. In this way, even if two delegated node will download just part of the file, there will still be the chance that the combination of the two parts at destination will reconstruct the whole file.

5.1 Different download starting points

For the sake of clarity, in the following, we utilise the term *file server* to denote both the device which is originally in posses of the file and any servant that has succeeded in finding a file possessor *and* is now carrying the file or part of it.

When a user chooses to initiate a file download, a task is created and scheduled for execution. Initially there is only one interval to be downloaded that is the entire file [*0, file length*] [Figure 4(a)]. Once a file server is in reach area, a chunk request is issued containing the missing data interval, in this case [*0, file length*]. If there is more than one file server in reach, a transfer might be initiated with each of them, starting from a different point in the file.

To illustrate how these intervals are computed, consider the potential scenarios in Figure 4, where *I* represents the total file length. In each case the whole file may be downloaded, yet the starting point for the download varies as follow:

- *Case Figure 4(a) and case Figure 4(b):* Two file servers are in reach and there are available resources to launch two parallel execution flows:
  - a chunks are requested from file server 1 starting from the beginning of the file
  - b chunks are requested from file server 2 starting from the middle of the file, i.e., point (1/2) * *I*.
- *Case Figure 4(c):* Four file servers are in reach area and there are available resources to launch four parallel execution flows:
  - a chunks are requested from file servers 1 and 2 as stated above
  - b chunks are requested from file server 3 starting from point (1/4) * *I* of the file
  - c chunks are requested from file server 4 starting from the point (3/4) * *I* of the file.

In essence, the starting point of the requested interval is calculated so as to halve the largest interval left undivided.

Clearly, when the end of the file is reached, the download continues from the beginning until the whole file is downloaded, if possible. However, in case all the parallel downloads are prematurely interrupted by disconnection, they will all have downloaded different parts of the file thus maximising the possibility to have cumulatively downloaded the whole file, instead of having redundantly downloaded the same limited part of the file.

Figure 4 M2M share file division strategy

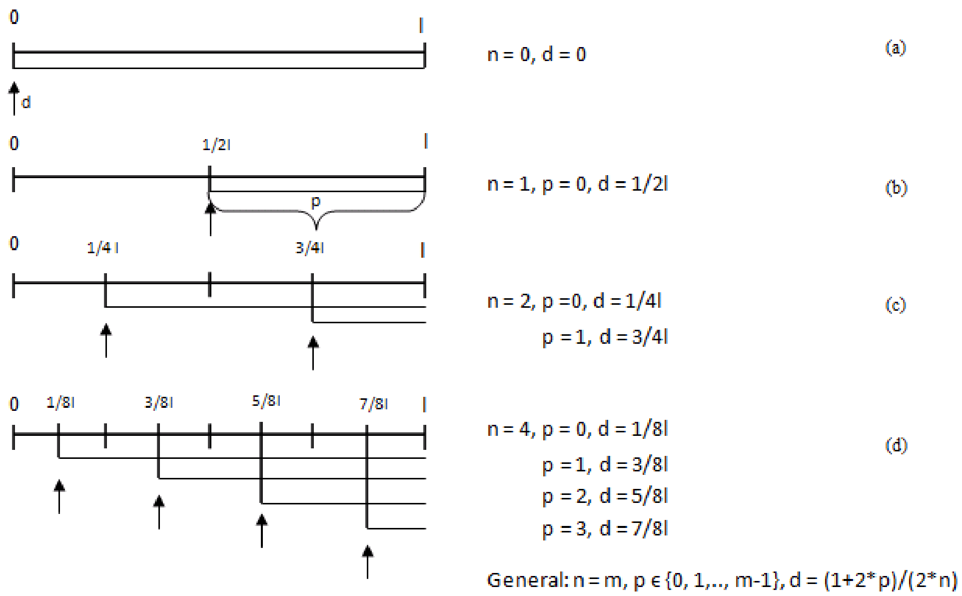

Figure 5 Data quantity downloaded employing different strategies for a video file of 25 MB

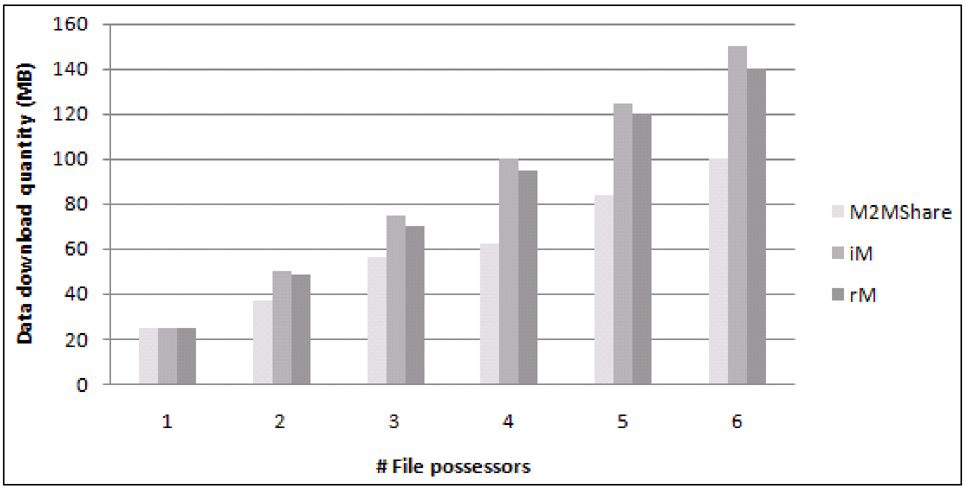

5.2 Evaluation of file division strategy

As mentioned earlier, DTNs often use message replication techniques along the destination path in order to increase the probability of data reaching the destination. In our case the file division strategy might add redundancy during data transfer as it can happen that overlapping data intervals are simultaneously downloaded by different servants. However, the fact that each servant is asked to download the file starting from different points allows reconstructing the whole file even if both downloaded just part of it.

To this aim, Figure 5 shows a simulative comparison of our file division strategy with two other division strategies:

- *iM:* a strategy which requests at each file server the entire file, always starting from its first byte
- *rM:* a strategy that randomly chooses the initial download point in the file request.

In the tested scenario, file possessors are in-reach area. If *M* is the number of file possessors (X-axis in Figure 5), the task handling the file download can initiate a maximum of *M* simultaneous transfers. Each experiment was repeated 40 times for each division strategy measuring the total amount of data transferred between any two nodes in the network (Y-axis in Figure 5). As we can see from the chart, our division strategy has the least redundancy (overlap among simultaneous file transfer) during data transfer, thus increasing the amount of useful data transfer and reducing the time required to retrieve the requested file (as previously demonstrated in Figures 1, 2 and 3) while reducing the total transmission overhead.

6 Conclusions and future works

End-user tends to shift toward wireless technologies and the device technology evolution has opened new application scenarios for the mobile device world. These challenging environments have fundamental characteristics that differ from the wired internet and that need to be addressed (e.g., mobility, low node density and energy preservation).

In this paper, we presented a delay tolerant solution for P2P multimedia file-sharing which uses delegations to reach data in other local disconnected networks. To this purpose, we were inspired by the asynchronous DTN communication model, and brought it into the wireless mobile world. With our solution, a peer can delegate unsatisfied file search tasks to other peers. However, in order to keep delegation overhead at a minimum while increasing the chances of eventually receiving the output back, only frequently encountered peers are chosen to this aim. We can hence state that this solution is particularly efficient when employed by commuters utilising public transportation as:

1 they share the part of their daily home-work-home path with the same people

2 they are generally toward different destinations thus reaching other disconnected networks in search for requested files.

We showed by experimental means employing a working day movement model that this technique serves its purpose with better performance with respect to traditional strategies.

Possible extensions of this work include the design of a file division strategy and local congestion control to increase the quantity of file downloaded per time unit (Marfia et al., 2007, 2008). Moreover, it would be interesting to test possible support of our solution to entertainment applications such as gaming, alert message propagation, and vehicular applications (Ferretti and Roccetti, 2005; Palazzi et al., 2007; Roccetti et al., 2007). Finally, we are currently planning a comprehensive set of simulations through THE_ONE that will compare realistic, even if different, users' mobility models (Gaito et al., 2011).

References

Boldrini, C., Conti, M., Delmastro, F. and Passarella, A. (2010) 'Context- and social-aware middleware for opportunistic networks', *Journal of Network and Computer Applications*, Vol. 33, No. 5, pp.525–541.

Bray, J. and Sturman, C.F. (2001) *Bluetooth 1.1: Connect without Cables*, 2nd ed., Prentice Hall, Saddle River, NJ.

Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K. and Weiss, H. (2007) *Delay-Tolerant Networking Architecture*, available at http://www.tools.ietf.org/html/rfc4838 (accessed on 18 September 2011).

DTNRG (2002) *Delay Tolerant Networking Research Group*, available at http://www.dtnrg.org/wiki (accessed on 21 September 2011).

Ekman, F., Keränen, A., Karvo, J. and Ott, J. (2008) 'Working day movement model', *MobiHoc 2008: Proc. of the ACM SIGMOBILE Mobile Ad Hoc Networking*, Hong Kong, China.

Ferdinando, A., McKee, P. and Amoroso, A. (2003) 'A policy based approach for automated topology management of peer-to-peer networks and a prototype implementation', *POLICY 2003: Proc. of the 4th IEEE International Workshop on Policies for Distributed Systems and Networks*, Como, Italy.

Ferretti, S. and Roccetti, M. (2005) 'Fast delivery of game events with an optimistic synchronization mechanism in massive multiplayer online games', *ACE 2005: Proc. of Advances in Computer Entertainment Technology*, Valencia, Spain.

Gaito, S., Maggiorini, D., Pagani, E. and Rossi, P.G. (2009) 'Distance vector routing for public transportation vehicular networks: performance evaluation on a real topology', *WD 2009: Proc. of the 2nd IFIP/IEEE Wireless Days*, Paris, France.

Gaito, S., Pagani, E. and Rossi, P.G. (2011) 'Strangers help friends to communicate in opportunistic networks', *Computer Networks*, Vol. 55, No. 2, pp.374–385.

Gurun, S., Nagpurkar, P. and Zhao, B.Y. (2006) 'Energy consumption and conservation in mobile peer-to-peer systems', *MobiShare 2006: Proc. of the 1st International Workshop on Decentralized Resource Sharing in Mobile Computing and Networking*, Los Angeles, CA, USA.

Jung, S., Lee, U., Chang, A., Cho, D.K. and Gerla, M. (2007) 'BlueTorrent: cooperative content sharing for bluetooth users', *Pervasive and Mobile Computing*, Vol. 3, No. 6, pp.609–634.

Keränen, A., Ott, J. and Kärkkäinen, T. (2009) 'The ONE simulator for DTN protocol evaluation', *SimuTools 2009: Proc. of ICST Simulation Tools and Techniques*, Rome, Italy.

Klemm, A., Lindemann, C. and Waldhorst, O. (2003) 'A special-purpose peer-to-peer file sharing system for mobile ad hoc networks', *VTC 2003: Proc. of the IEEE Vehicular Technology Conference*, Orlando, FL, USA.

Marfia, G., Lutterotti, P., Eidenbenz, S., Pau, G. and Gerla, M. (2008) 'Fair multi-media streaming in ad hoc networks through local congestion control', *MSWIM 2008: Proc. of the 11th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, Vancouver, Canada.

Marfia, G., Pau, G., Giordano, E., De Sena, E. and Gerla, M. (2007) 'VANET: on mobility scenarios and urban infrastructure: a case study', *MOVE 2007: Proc. of Mobile Networking for Vehicular Environments at INFOCOMM*, Anchorage, Alaska, USA.

Marfia, G., Roccetti, M., Palazzi, C.E. and Amoroso, A. (2011) 'Efficient vehicle-to-pedestrian exchange of medical data: an empirical model with preliminary results', *MobileHealth 2011: Proc. of the 2nd ACM International ICST Conference on Wireless Mobile Communication and Healthcare*, Paris, France.

McNamara, L., Mascolo, C. and Capra, L. (2008) 'Media sharing based on colocation prediction in urban transport', *Mobicom 2008: Proceedings of the ACM Conference on Mobile Computing and Networking*, San Francisco, CA, USA.

Palazzi, C.E. (2004) 'Buddy-finder: a proposal for a novel entertainment application for GSM', *NIME 2004: Proc. of the IEEE Global Telecommunications Conference and Exhibition*, Dallas, TX, USA.

Palazzi, C.E. and Bujari, A. (2010) 'A delay/disruption tolerant solution for mobile-to-mobile file sharing', *WD 2010: Proc. of the 3rd IFIP/IEEE Wireless Days*, Venice, Italy.

Palazzi, C.E., Roccetti, M., Ferretti, S., Pau, G. and Gerla, M. (2007) 'Online games on wheels: fast game event delivery in vehicular ad-hoc networks', *V2VCOM 2007: Proc. of IEEE Intelligent Vehicles Symposium on Vehicle to Vehicle Communications*, Istanbul, Turkey.

Rachuri, K.K., Mascolo, C., Musolesi, M. and Rentfrow, P.J. (2011) 'SociableSense: exploring the trade-offs of adaptive sampling and computation offloading for social sensing', *Mobicom 2011: Proceedings of ACM Conference on Mobile Computing and Networking*, Las Vegas, NV, USA.

Rhee, I., Shin, M., Hong, S., Lee, K. and Chong, S. (2007) 'Human mobility patterns and their impact on delay tolerant networks', *HotNets-VI 2007: Proc. of 6th ACM Workshop on Hot Topics in Networks*, Atlanta, GA, USA.

Roccetti, M., Gerla, M., Palazzi, C.E., Ferretti, S. and Pau, G. (2007) 'First responders crystal ball: how to scry the emergency from a remote vehicle', *NetCri 2007: Proc. of the 1st IEEE International Workshop on Research Challenges in Next Generation Networks for First Responders and Critical Infrastructures*, New Orleans, LA, USA.

Samal, S. (2003) 'Mobility pattern aware routing in mobile ad hoc networks', Unpublished Master thesis, Virginia Polytechnic Institute and State University, VI, USA.

Stevens, W.R. (2000) *TCP/IP Illustrated, Volume 1 – The Protocols*, Addison-Wesley, USA.

Wu, J. (2005) *Theoretical and Algorithmic Aspects of Sensor, Ad Hoc Wireless, and Peer-to-Peer Networks*, Auerbach Publications, USA.