

TACO-DTN: A Time-Aware Content-based dissemination system for Delay Tolerant Networks

Giuseppe Sollazzo
Dept. of Computer Science
University College London
Gower Street
London WC1E 6BT, UK
g.sollazzo@cs.ucl.ac.uk

Mirco Musolesi
Dept. of Computer Science
University College London
Gower Street
London WC1E 6BT, UK
m.musolesi@cs.ucl.ac.uk

Cecilia Mascolo
Dept. of Computer Science
University College London
Gower Street
London WC1E 6BT, UK
c.mascolo@cs.ucl.ac.uk

ABSTRACT

Content dissemination applications are becoming more and more popular on fixed infrastructure: in this paper we introduce TACO-DTN, a content dissemination system which, by virtue of being time-aware in terms of subscriptions and events, is appropriate for delay tolerant networks, where a number of nodes act as infostations, enjoying some form of connectivity to the backbone, and other nodes are mobile devices, reachable sometimes only through intermittent connectivity of carriers. Examples of applications benefiting from such a system could be travel information dissemination systems in large cities (exploiting infostations at bus stops) or on highways, advertisements dissemination at specific times, and information dissemination to remote villages. The approach is based on a novel concept of *temporal utility* of subscriptions and events. The temporal utility is used to govern the routing of the events to the right infostation (i.e., the one reached by the interested subscribers at the right time), avoiding unnecessary information transfer on slow links and the buffer management, in case buffer limitations are an issue. We give a description of our protocol and discuss its validation through simulation.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Distributed networks, Network communications, Store and forward network, Wireless communication; C.2.2 [Network Protocols]: Applications, Protocol architecture

General Terms

Algorithms, Theory, Design

Keywords

Delay tolerant networking, publish/subscribe systems, content-based routing, event dissemination, infostations, buffer management.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiOpp'07, June 11, 2007, San Juan, Puerto Rico, USA.

Copyright 2007 ACM 978-1-59593-688-2/07/0006 ...\$5.00.

1. INTRODUCTION

The ability to push relevant information (or advertisement) to people on the move on a highway or while they are reaching a point of interest, e.g., a tourist attraction, a bus stop or a mall, is one of the many recently targeted goals for location based services. Infostations [5] positioned in the interested areas act as forwarding points from where content is delivered to interested subscribers reaching that location. However, behind its appeal, this scenario offers many challenges including:

- *management of possibly limited resources* of the infostations, often intermittently connected to a backbone through slow links (e.g., at bus stops in a city, on a stretch of highway, at a kiosk in a remote village) and which might have to store content for quite an extensive period of time;
- *routing of the right content* from the backbone to selected infostations, depending on the analysis of which subscribers will roam close to the infostations in the near future (e.g., events related to travel information for trains going out a city should be downloaded onto an infostation in a train station around 5pm);
- *relaying of the right content* from infostation to more remote areas *through the use of selected carriers*, which will take the selected content to subscribers located in that area at the right time (e.g., an infostation mounted on a bus which reaches a remote village with a periodic schedule can be loaded with the right content depending on the time the bus is due to the remote village).

In this paper we argue that these challenges, which are inherently linked to the delay tolerant nature of the network, can be tackled by the use of temporal information associated with both the validity of the content and of the subscriptions (i.e., the expressions of interest). We have developed a *Time-Aware Content-based dissemination system for Delay Tolerant Networks (TACO-DTN)*, according to the following idea: a node can subscribe to publications of certain content for a certain period of time; these subscriptions can be *periodic*, i.e., have a validity which is dependent on the time of the day/week/year (e.g., travel news towards downtown at 8am). On the other hand, every publication of content has a temporal specification of its validity. We make use of this temporal information to enforce some kind of *temporal event matching*: publications are matched

against subscriptions according to their temporal correspondence rather than on content only. We introduce the concept of *temporal utility* that is used for intelligent buffer management and for content routing. Depending on the interest of subscribers and their periodicity in the area of an infostation, only publications which will be relevant in the next future, as predicted by the subscription profile of the infostation, are routed onto this infostation. This mechanism is also used for buffer management on infostations with buffer constraints, such as maybe mobile carriers.

Recent work on Delay Tolerant Networks (DTN) [4, 1] has focused on unicast routing [8, 10] and on opportunistic communication [14]. To our knowledge the only other work concentrating on multicast routing and temporal issues for delay tolerant networking is [19], where an approach to multicast is introduced trying to account for temporal group membership. Previous work in this area has been based on the exploitation of epidemic-style techniques [17, 11].

Some approaches which tackle location based dissemination with temporal constraints in specific scenarios like sensor networks from a pure algorithm perspective have been recently presented [7]. However, these approaches mainly focus on the problem of tuning the replication of messages in order to reach a certain group of nodes, given some spatio-temporal constraints, rather than on supporting decoupled communication with temporal semantics using a publish/-subscribe architecture. In [6], Huang et al. investigate the problem of delivering messages to a large set of nodes in a manner that satisfies a potentially dynamic set of spatio-temporal constraints without any infrastructure.

The solution that we are presenting in this paper is the first attempt of designing a content-based routing mechanism that allows for an intelligent buffer management and forwarding of content *based on the evaluation of the temporal validity and constraints of subscriptions and publications*.

The paper is organised as follows: Section 2 introduces a number of scenarios to motivate our work further. Section 3 describes our approach in detail and Section 4 our evaluation of the performance through simulation. In Section 5 we compare our solution with the state of the art, while Section 6 concludes the paper illustrating possible future work.

2. APPLICATION SCENARIOS

While developing our approach we have kept in mind a number of scenarios, such as:

Advertising: Let us consider the case of a person regularly travelling downtown to reach her workplace during week days. She may register an interest in receiving information to her mobile device about both travel news on certain routes around her specific travelling times and of restaurant promotions for lunchtime. Her subscriptions, then, will be *periodic* ones, around 8am and 6pm for travelling on certain trains and around lunch time for lunch promotions (only Mondays to Fridays).¹ Infostations around town, with quite limited/slow connectivity to the backbone network², will register her presence in particular locations in periods of the day and of the week: this will allow the *temporal profiling*

¹Some of these subscriptions may happen implicitly by the system registering a mobility pattern and linking that to some sort of content interest.

²An example may be infostations that communicate periodically with a central server using a GSM network.

of the infostation in terms of which kind of subscribers are in the area at the different times of the day or of the week: this information is useful for both the selective download of the content from the backbone (given the slow/limited connectivity) and the buffer management, if this is limited.

In general, we envisage a scenario where infostations are in charge of delivering the information to the subscribers passing by using short-range wireless technologies such as Bluetooth or 802.11. The business case is an interesting one: service providers could sell specific time slots in specific places to different advertising/travel companies according to the content interests of the users seen in the past in those places at those times. The billing of the service may be based on the number of the publications actually delivered to the users.

Remote Areas: Let us consider a remote village connected to a bigger town through a bus route. Information can be carried through the busses (with the help of mobile infostations installed on the busses) between the town/Internet and the village. A number of papers have focussed on this model such as DakNet [12] and the system developed by the University of Waterloo for delay tolerant communication support in India [13]. Information could for instance be related to stock values or market prices for goods. As the memory on the infostation is reasonably limited and the network link to it is quite slow, it makes sense to have techniques to selectively store only the information which is going to be relevant in the village both in terms of content and on timing.

3. TACO-DTN

In this section we illustrate our approach. We first define the general topology of the systems we consider, then we illustrate the algorithmic details of TACO-DTN. Our purpose is to enhance the typical pattern matching mechanism used in content-based routing with a notion of time-awareness that improves delivery in scenarios characterised by temporal constraints and delay tolerant traffic.

3.1 System Topology

We assume a system composed of fixed infostations, mobile (subscriber) nodes and, in some cases, mobile infostations which act as carriers (e.g., infostations on busses).

- *Fixed infostations* are connected (in some cases intermittently or with slow links) to a backbone and possibly to the Internet. Content is sent to infostations and stored temporarily in their (potentially limited) buffer.
- *Mobile subscriber nodes* can be PDAs, mobile phones, laptops, embedded devices for instance in cars, and the like. The nodes are intermittently connected to the infostations, as described in the scenarios in Section 1³.
- *Mobile infostations* are infostations which move around (e.g., busses) and can potentially carry information from fixed infostations to subscribers in a different (potentially remote) location. They will often have limited buffers and intermittent, or slow, connectivity.

³We assume no multi hop routing between the mobile nodes: routing is only considered for forwarding among infostations or the Internet and the infostations.




Figure 1: Examples of subscription/event temporal functions

3.2 Content Subscriptions

The entities of the system are content *events* and content *subscriptions*. Mobile nodes express interest in content by emitting subscriptions. We represent a generic subscription s as $\langle \text{SID}, \text{topic}, m_s(t), \text{TTL} \rangle$ where **SID** is the subscription identification number, **topic** is the type of content, $m_s(t)$ is the *subscription temporal function* describing the time validity of subscription and *TTL* the expiration time of the subscription. A *subscription temporal function* is expressed as a discrete function, allowing to define periodic time intervals as well as validity in the past and in the future within a given temporal range. The concept of temporal validity defined by $m_s(t)$ can also be seen as a temporal group membership. The value of $m_s(t)$ at instant t is set to 1 if the user is subscribed to the topic, to 0 otherwise. The subscription may either be in the past or in the future or both. Users may want to receive content published in the recent past and/or content published in the coming hours. The expiration time *TTL* may exceed the temporal validity of the subscription, since users may want to receive events published recently, even if, at the current time, they are no more subscribed: the *TTL* only binds the delay with which the event can be received (and it is an important parameter in networks with intermittent connectivity).

To give an example of subscription, let us consider a subscription which could be emitted by a person interested in sport updates on Saturdays from 1pm to 6pm, around the time matches are played. The subscription may have a *TTL* set to 11pm, since the user may be interested in receiving information also after the sport events, but not after Saturday nights. Temporal validity has semantics contents, whereas expiration time is used to define the constraints in terms of delivery delay of past events.

Subscriptions can be formulated by users manually or, in some cases, can be generated automatically by extracting information (for example, by using machine learning algorithms) from their electronic agendas, calendars (like Google Calendar), in-vehicle systems, blogs and so on. The automatic generation of the subscription is outside the scope of this paper.

3.3 Content Events

Similarly to a subscription, a content event e is a tuple: $\langle \text{PID}, \text{topic}, \text{data}, v_e(t), \text{TTL} \rangle$, where **PID** is the event identification number, **topic** is the content type, **data** is the actual content, $v_e(t)$ is the *event temporal function* and *TTL* expresses the expiration time for the event. Event temporal functions have a form that is similar to the subscription tem-

poral functions: the value of $v_e(t)$ at the instant t is set to 1 if the publication is valid in that particular instant, 0 otherwise. Periodic events can be represented with this model, e.g., advertisement of restaurants at the same time of the day for a week. The *TTL* is used as an expiration time, for example to allow for an asynchronous (delayed) delivery of the publication to subscribers that may have been disconnected when the event was published.

3.4 Event Matching

When a mobile node (subscriber) reaches a fixed or mobile infostation, an exchange happens which allows the infostation to register the interest of the subscriber. The mobile node sends all its content subscriptions, which will be stored in the infostation. The infostation sends all the events that *match* the subscriptions. An event matches a subscription if they are related to the same topic and if there is overlap between the length of time the subscription temporal function and the event temporal function validity (i.e., when they are 1). An example is reported in Figure 2.




Figure 2: Example of event matching.

3.5 Temporal Profiles

Each subscription sent by a subscriber is stored and is used by an infostation i to update its *temporal profiles* \mathbf{P}_i . Given a set of topics, a profile $p_{k_i}(t)$ is maintained for each topic k . A profile stores the information about the popularity of k as a function of time, both in the past and in the future. In other words, it expresses quantitatively the importance of keeping and receiving events of type k . If a subscription related to k is received for the first time, the

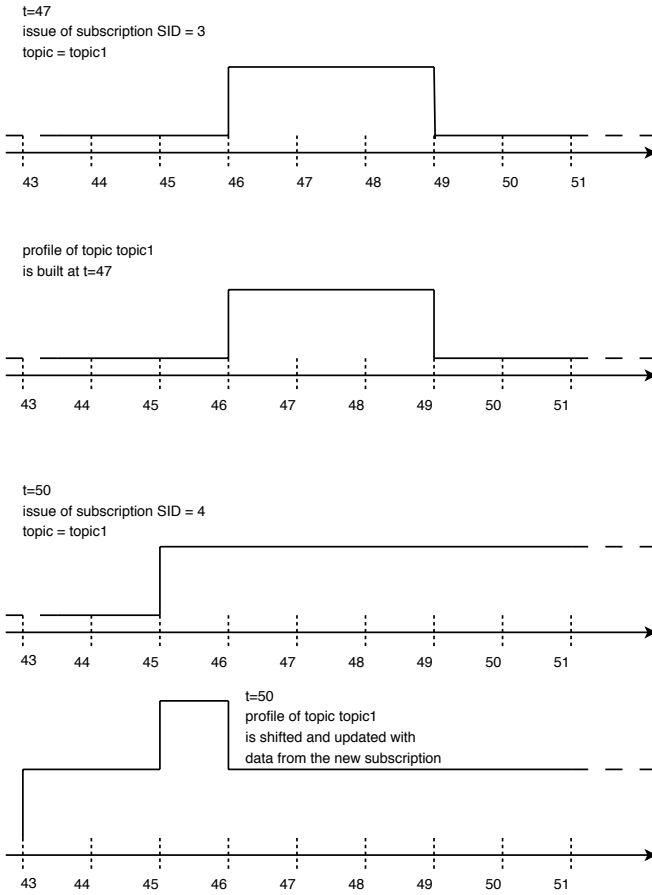


Figure 3: Infostation temporal profile construction.

infostation creates a temporal profile equal to the received subscription temporal function for that topic. If the profile for that topic already exists, the subscription temporal profile is just *added* to the current profile function. When a subscription expires, the temporal function is *subtracted* from the corresponding profile. An example is shown in Figure 3.

The temporal profile maintains a record of the subscribers for a certain period of time in the past and in the future. More formally, the temporal function of the profile $p_{k,i}$ for topic k at time t represents the number of nodes that are interested in k at time t in the neighbourhood of infostation i , considering all the subscriptions received. A profile is deleted when all its values are set to 0. More complex methods of analysis of the expected number of subscribers could be implemented, for example by using forecasting techniques and by assigning different weights to different class of users. We will soon show how temporal profiles are used to efficiently route and store events to/in the right infostations.

A profile of a fixed infostation is also updated when the infostation is reached by a mobile infostation (i.e., a carrier). The carrier will have collected subscriptions from other mobile nodes and created a set of temporal profiles for some of the topics. When reaching the fixed infostation, the carrier would update the temporal profile of the infostation with its profiles. When performing this update we keep into account the carrier mobility pattern (i.e., with which periodicity the

subscribers have been seen) in the update of the temporal profile, so to account for the delay in delivering the information to these subscribers.

3.6 Event Routing

Events are usually generated by servers other than infostations. TACO-DTN uses the temporal profiles of the infostations to decide which events to route to which infostation through the backbone or to which mobile infostation from a fixed one. This is particularly useful if slow and intermittent links are present, for instance, in case of fixed infostation at remote bus stops or mobile infostations on busses [13].

A utility function $u_{e_k,i}$ is associated to each event e_k related to topic k , for infostation i . The utility function of an event e_k , given its event temporal function $v_e(t)$ and the profile $p_{k,i}(t)$, is calculated as follow:

$$u_{e_k,i} = \sum_{t=t_{evalMin}}^{t_{evalMax}} v_e(t)p_{k,i}(t) \quad (1)$$

In other words, $u_{e_k,i}$ is defined as the sum of the products of each value of the event temporal profile $v_e(t)$ and the profile $p_{k,i}(t)$ over an *evaluation window* defined by the range $[t_{evalMin}, t_{evalMax}]$. The evaluation window is used to delimit the interval of evaluation of the utility. The choice of the upper and lower bounds of the range is dependent on the type of information that has to be delivered. The utility function u_{e_k} can be used to make informed decisions on the selection of the infostations on which to publish events. The publisher may retrieve the profiles of the reachable fixed infostations for each topic of the events that it is publishing and then calculate the corresponding utility function of the events. Possible selection criteria can be applied: for example, events may be forwarded only to all the infostations for which the utility of those events is higher than a certain threshold. The same approach can be applied when a mobile infostation is requesting events from a fixed one. If bandwidth is an issue, this technique allows some savings as only relevant events are forwarded.

This approach allows TACO-DTN to implicitly account for patterns of mobility of subscribers, by constructing a temporal profile of the subscribers seen by the infostation. An alternative use of the utility function is to send all the events to all the infostations (this is possible if bandwidth is not an issue): these then will discard or insert the events in the buffer depending on the calculated utility. The next section illustrates how this same utility function can be used for buffer management.

3.7 Buffer Management

The same utility function described in the previous section can be used to make informed decisions on which events to keep in the buffer, if space is limited. A simple example is shown in Figure 4. The events which are more likely to be requested in the near future will be kept.

Furthermore, different importance can be attributed to events published at specific times (e.g., all the events published in the future). To achieve this, a weighted function is used to assign a different importance to different units of time:




Figure 4: Calculation of the event utility with evaluation window [44,48]. If the buffer is full, publication 2 may be replaced by publication 1, since it is characterised by a lower temporal utility. In fact, in this particular case, the temporal utility of event 1 is 2, whereas the utility of event 2 is 0.

$$u_{e_{k,i}} = \sum_{t=t_{evalMin}}^{t_{evalMax}} w_{k,i}(t)v_e(t)p_{k,i}(t) \quad (2)$$

Assuming that more importance should be given to the events in the future, for example, we may assign weights to the future time slots equal to twice the weights related to the past instants of time. Using these weights, for instance, the system manager can decide to store all the events related to the future, discarding all the other ones.

The utilities of the events in the buffer are re-calculated periodically. When an event has to be added to the buffer and this is full, the event is inserted in the buffer only if its utility is less than the lowest utility of the events stored in the buffer. If this is the case, the event of the lowest utility is replaced and the event is inserted in the buffer. Otherwise, the event is simply discarded. Priority fields can also be used to define events that must be inserted in the buffers in any case (for example emergency updates about terrorist attacks and so on).

4. EVALUATION AND DISCUSSION

We have evaluated the performance and the scalability of TACO-DTN in large scale scenarios using the OMNeT++ simulator ([18]).

4.1 Description of Simulations

Our simulation scenario is composed of 100 nodes, 10 of which are infostations. Mobile nodes do not have buffering capabilities. The first set of graphs refer to scenarios without carriers, which are instead evaluated in the last part of this section. Infostation are initially placed on the $1000m \times 1000m$ area randomly, as mobile nodes are. The nodes move in the space in two possible ways: according to the *Random WayPoint* standard model, or to a modified version of the former in which every mobile node moves along a Hamiltonian path between all the infostations: this allows us to evaluate the impact of mobility patterns on TACO-DTN. The nodes move at variable speed, $1 - 6m/s$, among the infostation.

Functions and profiles have a granularity of information of 60s (i.e., each function “slot” describes one minute of interest). Variable parameters are first of all buffer size, mobility model and temporal function shape. The unit for the buffer size is an event. The duration of each simulation run is 20000s for simulation concerning buffer management and 6000s for the others. The difference is due to the different amount of time needed to deliver at least 80% of the messages for scenarios built to show different simulation goals. We made 10 run of each simulation scenario; a 10% confidence interval is shown in the diagrams. These parameters describe a very specific scenario, but they provide insight in the factors that influence the performance of the protocol and its behaviour compared to random mechanisms.

4.2 Simulation Results

4.2.1 Evaluation of Buffer Management

The first set of simulations shows the behaviour of TACO-DTN in terms of buffer management. We studied the behaviour of our algorithm using a temporal function with the following temporal pattern: all the subscriptions are valid only in the past, whilst 50% of publications are valid only in the past and 50% only in the future. In each of these scenarios other parameters are changed, especially buffer size.

In the simulations of this section, random selection is used to decide on which infostation to publish. Firstly, we consider a scenario where we generate subscriptions with interest in 1 topic with a temporal function valid only for a limited period of time: we identify a time in the simulation, i.e., $t = 3000s$, and we shape the temporal function so that the subscriber will be interested in the topic for $t \in [0, 3000]s$. Publications are issued on one topic between $t = 1000s$ and $t = 2000s$, with temporal function shaped in one of the following ways:

1. validity for $t \in [0, 3000]s$, or
2. validity for $t \in [3000, 20000]s$

that is, if we identify the moment *now* with $t = 3000s$ (i.e., the moment of publication), subscriptions are only in the past whilst publication can be either in the past or in the future.




Figure 5: Partial match, with RWP.

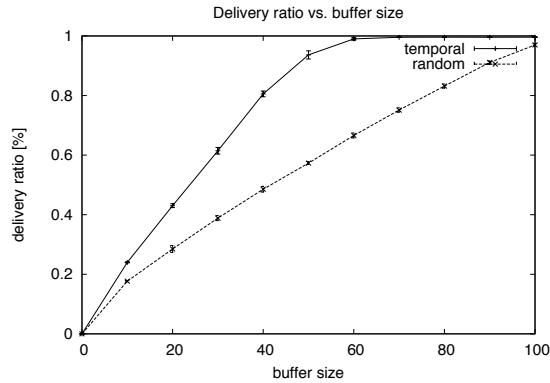


Figure 6: Partial match, with modified RWP.

In Figure 5, we show performances when standard RWP is used. The *temporal* line refers to the buffer management policy following temporal utility as shown in Section 3, whilst the *random* one refers to a random deletion policy for the buffer. We can clearly see that the performance of both random and temporal selection improve as the buffer size increases. The improvement of the temporal selection over the random one, though, increases as the buffer size increases. When buffer size is around 100, however, probabilistically, there will be sufficient space for all the publications and, obviously, both kinds of selection start to exhibit the same performances.

In Figure 6, we evaluate the same scenario with our modified version of RWP using Hamiltonian paths between infostations for mobile nodes. As we may expect, this scenario exhibits better performance, with a delivery approaching 95% for a buffer of 50 with temporal profiling.

4.2.2 Evaluation of Event Routing

In this section we describe some performance results related to infostation selection for publications. The goal of infostation selection is to issue a publication on the infos-

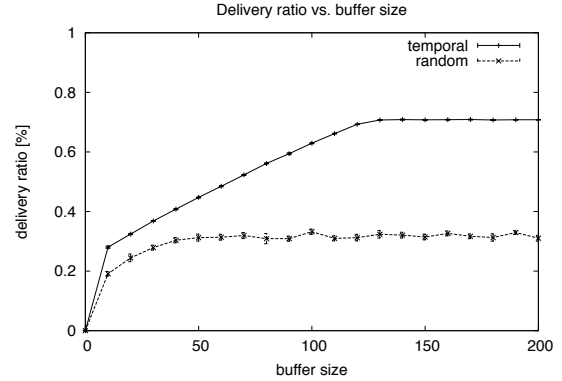


Figure 7: Event routing to infostations. 500 messages sent.

tation or the group of infostations that is/are likely to be visited by mobile hosts interested in the topic of the publication. As described in Section 3, infostations on which to publish an event are selected based on their temporal profile (built observing the subscriptions of the subscribers getting in reach over time).

To evaluate the effectiveness of the approach we divide the infostations into groups and mobile hosts into sets, and let each set travel only between infostations of one group and to subscribe to a single topic. Groups and sets are disjoint. Groups are used to simulate mobility models based on the notion of groups of interests. In other words, users with similar interests move according to the same pattern. In this scenario, we experimented with 4 groups of the same size. All the hosts are subscribed to the same topic, but a different temporal function is associated to each group. We do not consider different topic, to isolate the impact of the selection mechanism purely based on the temporal aspects of the subscription functions.

We studied this scenario with only RWP modified, obviously, to get the notion of group of mobility (that would not make any sense with standard RWP). What we expect from these simulations is better performance with the temporal algorithm, due to its ability to select infostation with higher temporal utility. Since not all nodes reach all infostations, the infostations will build temporal profiles of the nodes they see (actually, of the subscriptions they see). These profiles will allow routing of the right events to these infostations, increasing delivery ratio. On the other hand, the random protocol will probabilistically route (with a uniform distribution) events to infostations, not taking advantage of the subscription patterns.

Figure 7 shows that a random selection of infostations is not able to exploit information about groups inferred from the utility of topic. These simulations exhibit an expected behaviour: publications are placed on “good” infostations with higher probability by the temporal selection algorithm.

4.2.3 Impact of Temporal Information

We also investigated the impact of temporal information on the global performance of our system. We now present a




Figure 8: Impact of evaluation window length on delivery ratio.

scenario to assess how the length of the evaluation window, used to compare topic profiles and the temporal function of an event in order to calculate the utility value of that event, may affect performance.

The evaluation window length may affect delivery ratio as the longer it is, the more information is taken into account when calculating the temporal utility value of a publication. This value is used when deciding whether to store a publication in the local buffer, and which publication is to be replaced when the new one has a higher utility value. Subscriptions and events are relative to 1 topic, and each mobile node is subscribed to that topic. That is, the same scenario settings as the previous evaluation. As far as temporal functions are concerned, we use functions valid from a certain, random time slot on both for events and publications. This choice caters for several situations in which there is only a partial match between subscriptions and events. As shown in Figure 8, the longer the evaluation window is, the higher the delivery ratio. We did not reach 100% delivery because of the duration of the simulation. A 100% delivery is reached after about 10000 seconds.

4.2.4 Impact of Presence of Mobile Infostations

The last set of results refer to a scenario where mobile infostations (or carriers) are present. Mobile infostations are nodes with limited buffer capabilities that are able to carry publications from parts of the network to others. In a network with carriers, temporal information could be useful to decide which publication to carry from one side to another; for example, there should be cases of messages issued when the subscriber is present on the infostation, but which temporal function is defined far in the future.

The simulated scenario describes a network with only one infostation and three groups of mobile nodes moving in confined areas, without reaching the infostations. The nodes are deployed in a simulation area of $1000m \times 1000m$. There are 60 nodes in the three groups and each group occupies a $200m \times 200m$ area near one of the edges and are out of reach with respect to each other and to the infostation. There is no communication between the groups and nodes in each

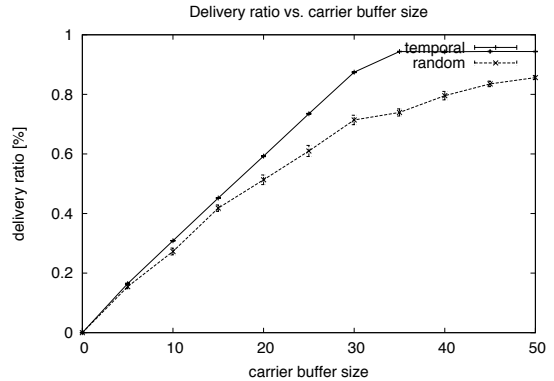


Figure 9: A simple scenario with mobile infostations.

group are confined to very constrained areas. All nodes are subscribed to only 1 topic, but each group has a different temporal characterisation, i.e., the temporal functions of subscriptions present 3 different patterns, one for each group. Publications are routed to the infostations, and each one has its temporal function consistent with only one of the 3 groups. Communications between the infostation and the mobile nodes are provided by three carriers: each carrier travels between one mobile group and the infostation. The goal of this simulation is to show that the carrier selects messages for the group according to temporal utility. Each carrier would only carry messages relative to one temporal function pattern and deliver them to the node in the relative group. As Figure 9 shows, carriers make correct choices, and delivery ratio increases with the carrier buffer size. Moreover, carriers using the temporal selection mechanism perform better than those with random selection. With a buffer of around 35 slots, TACO-DTN is able to already reach 95% delivery ratio, while the random protocol is around 70%.

5. RELATED WORK

Research projects about routing in Delay Tolerant Networks (DTN) [4] have focused on unicast [8, 10] and on opportunistic infrastructure-less communication [14, 1]. Other solutions exploit epidemic-style techniques [17]. In [11] an approach based on epidemic models considering network topologies and temporal constraints, in terms of maximum acceptable delay, is presented. However, these works concentrate on broadcast communication to all the nodes or a subset of the nodes. The problem of broadcasting in delay tolerant networks has also been studied in [9], where an analysis with different mobility models is presented. To our knowledge the only other work concentrating on multicast routing and temporal issues for delay tolerant networking is [19], where an approach to multicast is introduced trying to account for temporal group membership.

In the context of publish/subscribe systems [3], some solutions designed for location based dissemination with temporal constraints in specific scenarios like sensor networks from a pure algorithm perspective have been recently pre-

sented [7]. The main focus of these approaches is the problem of tuning the replication of messages in order to reach a certain group of nodes, given some spatio-temporal constraints rather than to support communication according to a decoupled publish-subscribe communication paradigm. In [6], the authors investigate the problem of delivering messages to a large set of nodes satisfying a potentially dynamic set of spatio-temporal constraints by exploiting geographic and topology information.

With respect to the activities of the IRTF Delay Tolerant Networking Research Group [2], a proposal for the extension of the Bundle Protocol for multicast communication has been discussed [15]. In this context, multicasting is defined as the ability of a source bundle node to transmit a bundle to a destination multicast endpoint without necessarily having to originate a separate bundle for each bundle node that is registered in that endpoint. However, this initial effort is focussed on the definition of the extension of the Bundle Protocol itself and not on the design of the anycast/multicast routing protocols implementing this semantics [16].

In this work, we have presented a solution for content based event dissemination exploiting a hybrid architecture composed of fixed infostations and mobile devices. With respect to the existing work, we have introduced a novel concept of temporal interest expressed using profiles with an associated temporal utility for intelligent event forwarding and buffer management.

6. CONCLUSIONS

This paper has described a time-aware approach to delay tolerant content based dissemination. Temporal profiles are associated to each subscription and allow the construction of temporal profiles of infostations. Events also have temporal validity. Temporal profiles are used in two main tasks: buffer management, in order to decide which events to store when buffer space is limited, and event routing, to select the right infostation or carrier on which to publish content. Our preliminary evaluation has described the advantages of the temporal profiling over random mechanisms.

This work is a first step towards the implementation of a working system for delay-tolerant networks using J2ME technologies for the mobile devices. Our next steps will be the execution of further simulation and the definition of more realistic scenarios and the implementation of a prototype version using access points acting as infostations connected to a central server.

7. REFERENCES

- [1] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Impact of human mobility on the design of opportunistic forwarding algorithms. In *Proceedings of INFOCOM'06*, Barcelona, Spain, April 2006.
- [2] Delay Tolerant Networking Research Group Homepage, 2007. <http://www.dtnrg.org>.
- [3] P. T. Eugster, P. A. Felber, R. Guerraoui, and A.-M. Kermarrec. The many faces of publish/subscribe. *ACM Computing Surveys*, 35(2):114–131, 2003.
- [4] K. Fall. A Delay-Tolerant Network Architecture for Challenged Internets. In *Proceedings of SIGCOMM'03*, pages 27–34, New York, NY, 2003. ACM Press.
- [5] R. H. Frenkiel, B. R. Badrinath, J. Borres, and R. D. Yates. The Infostations challenge: balancing cost and ubiquity in delivering wireless data. *IEEE Wireless Communications*, 7(2):66–71, 2000.
- [6] Q. Huang, C. Lu, and G.-C. Roman. Mobicast: Just-in-Time Multicast for Sensor Networks under Spatiotemporal Constraints. In *Proceedings of IPSN'03*, pages 442–457, 2003.
- [7] Q. Huang, C. Lu, and G.-C. Roman. Spatiotemporal Multicast in Sensor Networks. In *Proceedings of SenSys'03*, pages 205–217. ACM Press, 2003.
- [8] S. Jain, K. Fall, and R. Patra. Routing in a Delay Tolerant Network. In *Proceedings of SIGCOMM'04*, pages 145–158. ACM Press, 2004.
- [9] G. Karlsson, V. Lenders, and M. May. Delay-tolerant broadcasting. In *Proceedings of the 2006 SIGCOMM workshop on Challenged networks (CHANTS'06)*, pages 197–204, New York, NY, USA, 2006. ACM Press.
- [10] M. Musolesi, S. Hailes, and C. Mascolo. Adaptive Routing for Intermittently Connected Mobile Ad Hoc Networks. In *Proceedings of the 6th International Symposium on a World of Wireless, Mobile, and Multimedia Networks (WoWMoM'05)*. Taormina, Italy. IEEE press, June 2005.
- [11] M. Musolesi and C. Mascolo. Controlled Epidemic-style Dissemination Middleware for Mobile Ad Hoc Networks. In *Proceedings of MOBIQUITOUS'06*, San Jose, California, July 2006.
- [12] A. Pentland, R. Fletcher, and A. Hasson. DakNet: Rethinking Connectivity in Developing Regions. *IEEE Computer*, pages 4–9, January 2004.
- [13] A. Seth, D. Kroeker, M. Zaharia, S. Guo, and S. Keshav. Low-cost communication for rural internet kiosks using mechanical backhaul. In *Proceedings of the 12th Annual International Conference on Mobile Computing and Networking (MobiCom'06)*, pages 334–345, New York, NY, USA, 2006. ACM Press.
- [14] J. Su, A. Goel, and E. de Lara. An Empirical Evaluation of the Student-Net Delay Tolerant Network. In *Proceedings of MOBIQUITOUS'06*, San Jose, California, July 2006.
- [15] S. Symington, R. Durst, and K. Scott. Delay-tolerant networking custodial multicast extensions draft-symington-dtnrg-bundle-multicast-custodial-02. DTN Research Group Internet-Draft. Expires: October 22, 2007.
- [16] S. Symington, R. Durst, and K. Scott. Custodial multicast in delay tolerant networks. In *Proceedings of IEEE Consumer Communications and Networking Conference (CCNC'07)*, January 2007.
- [17] A. Vahdat and D. Becker. Epidemic Routing for Partially Connected Ad Hoc Networks. Technical Report CS-2000-06, Department of Computer Science, Duke University, 2000.
- [18] A. Varga. The OMNeT++ discrete event simulation system. In *Proceedings of the European Simulation Multiconference (ESM'2001)*, Prague, June 2001.
- [19] W. Zhao, M. Ammar, and E. Zegura. Multicasting in delay tolerant networks: semantic models and routing algorithms. In *Proceedings of the 2005 ACM SIGCOMM workshop on Delay-tolerant networking (WDTN'05)*, pages 268–275, New York, NY, USA, 2005. ACM Press.