

Efficient multicasting for delay tolerant networks using graph indexing

Misael Mongiovì, Ambuj K. Singh,
Xifeng Yan and Bo Zong
Department of Computer Science
University of California
Santa Barbara, CA 93106, US
Email: {misael,ambuj,xyan,bzong}@cs.ucsb.edu

Konstantinos Psounis
Department of Electrical Engineering and
Computer Science
University of Southern California
Los Angeles, CA 90089 US
Email: kpsounis@usc.edu

Abstract—In Delay Tolerant Networks (DTN), end-to-end connectivity between nodes does not always occur, due to limited radio coverage, nodes mobility and other factors. Long-range communication may assist in guaranteeing delivery. Since long-range communication has a considerable cost, minimizing it is an important task. For multicast routing, this problem is NP-hard, and naive approaches are infeasible on large instances.

In this paper we define the problem of minimizing the long-range communication cost for multicast in DTNs. Our formulation handles the realistic scenario in which a data source is continuously updated and nodes need to receive recent versions of data. We analyze the problem in the case of scheduled trajectories and known traffic demand, and propose a solution based on a novel graph indexing system. We also give an adaptive extension that handles limited knowledge in node mobility. Our method significantly reduces the search space and finds an optimal solution in reasonable time. An extensive experimental analysis on large real and synthetic datasets shows that the proposed method responds in less than 10 seconds on datasets with millions of encounters and thousands of messages, with an improvement of up to 100 times compared to a naive approach.

I. INTRODUCTION

Delay Tolerant Networks (DTN) are communication networks that lacks continuous connectivity due to node mobility, failures or other factors. They experience frequent partitioning, and end-to-end paths between two nodes may even never exist. Routing in DTNs uses a *store-carry-forward* approach [11], where intermediate nodes delay the transmission of messages until new links are available and the message is “eventually” delivered with some delay. When the lack of connectivity is due to node mobility, the movement of nodes can be exploited to carry the messages.

In recent years, routing protocols for multicast in DTNs have received considerable attention [28], [13], [8]. Multicast protocols optimize the transmission cost by sharing routing paths among multiple destinations. Recent advances allow us to achieve a good tradeoff between minimization of the transmission cost and maximization of the delivery rate. However, due to the nature of DTN networks, proper delivery cannot always be guaranteed.

To guarantee the connectivity, nodes can be equipped with a *long-range communication* device [27] to be used when end-to-end communication cannot be otherwise achieved. Since

long-range communication is expensive, its utilization should be limited as much as possible. Providing an extra long-range communication to guarantee delivery introduces the new challenging problem of minimizing its cost.

In this paper we formulate the problem of optimizing the long-range communication cost in a network of moving nodes and call it the *demand cover problem*. Our model considers a set of *moving nodes* (e.g. people or vehicles) that are equipped with a device providing two kinds of communication. A *short-range* communication (e.g. radio), considered non-costly, can occur between two nodes when they are close to each other. A remote (long-range) communication (e.g. cellular or satellite), which involves a cost, can occur at any time independently of the node positions. In our model, a set of continuously-updated *data objects* need to be shared among nodes. Each data object needs to be received by a subset of destinations. For each destination, its *deadline* (time instant at which the object is needed) is specified. To avoid receiving non-recent copies of objects, a *latency* is also specified. We aim to find the set of remote transmissions that minimizes the communication cost subject to the aforementioned delay constraints. To our knowledge, we are the first to analyze the problem of optimizing the long-range communication cost for multicast in a DTN.

The described problem has several practical applications. For instance, consider a network of city buses, in which the transportation agency wants to provide passengers with personalized news that depends on their position, traveling plan etc. Each bus can obtain the updates by the cellular network (costly). However, it is more convenient to share the information among various buses via radio communication (non-costly). Another example considers soldiers or military vehicles that move following a specific strategy. They need to access certain information related to their location (e.g. satellite images). In this case, the only options available are satellite communication (highly costly) and node-to-node communication. The proposed approach also has applications in data ferrying [27], [21]. A set of moving nodes (ferries) is charged for gathering data. Depending on time constraints on data delivery, the ferries may decide whether to use short-range or long-range communication. Note that in these examples, the

node trajectories and the traffic demand are known in advance.

Solving the demand cover problem introduces new challenges due to the temporal constraints. A data object may need to be transmitted remotely more than once, due to either lack of connectivity or the need of satisfying a latency constraint. For instance, consider the four nodes in Fig. 1(a) that move following certain trajectories. Initially, nodes 2 and 3 are in contact. At time t_1 , nodes 2 and 4 enter in each other's radio range and a new contact begins. At time t_2 , the contact between nodes 2 and 3 ends since they move away from each other. Three of these nodes (shown with triangles) need to receive the same data object. Each of them has a given deadline (t_a , t_b and t_c respectively) and a latency (δ_a , δ_b and δ_c respectively). A remote transmission to node 3 at time t_1 covers the data needs r_a and r_b . Although r_c can be reached by transmitting the object to node 4 at any time after t_1 , the latency δ_c cannot be satisfied. Therefore, an updated copy of the object needs to be transmitted after $t_c - \delta_c$.

In this paper, we prove that the demand cover problem is NP-hard and present a baseline graph-based approach for it. In order to make this problem feasible on large datasets, we formulate it as a query processing problem and develop a novel graph-indexing-based solution. Due to the index, we are able to handle thousands of destinations on a network with millions of encounters in less than 10 seconds, with an improvement of up to 100 times compared to a naive approach.




Fig. 1. An example of DTN among moving nodes. (a) The four solid lines represent four trajectories. To simplicity we use the x-axis for indicating the time. The big dashed circles represents the radio range of nodes. Nodes that are involved in a transition (contact beginning or contact end) are filled. Three data needs (r_1 , r_2 and r_3) are represented with filled triangles. Their deadlines are t_a , t_2 and t_3 respectively while their latencies are δ_a , δ_2 and δ_3 respectively. (b) The corresponding space-time graph. Snapshots of the connectivity graph at three different times are depicted within big ovals. Temporal links joining contiguous snapshots are represented with dotted lines.

Our contribution can be summarized as follows: (i) We define a novel problem (named *demand cover*) that formalizes the problem of optimizing the long-range communication cost in a network of moving nodes. (ii) We prove that demand cover is NP-hard. (iii) We develop a compact graph-based representation of a *demand cover* instance and give a baseline algorithm for it. (iv) We propose a novel indexing system for quickly solving demand cover on graphs optimally. Our system further compresses the compact graph and uses an efficient filtering approach to retrieve a small portion of vertices that are relevant for achieving an optimal solution. (v) We evaluate the proposed approach on two real and one synthetic datasets

and show that an exact solution can be found in reasonable time in datasets with millions of encounters.

The paper is structured as follows: Sect. II describes related work. Sect. III introduces some basic concepts in DTNs. Sect. IV defines the problem, provides a graph representation and gives a simple solution. Sect. V describes our indexing system for demand cover. Sect. VI presents an extensive experimental analysis on real and synthetic datasets. Finally, Sect. VII concludes the paper with some future directions.

II. RELATED WORK

Previous works on DTNs focus on three types of contacts including scheduled, predicted and opportunistic contacts. Scheduled contacts result from applications of known trajectories, such as deep-space communication and data service in developing regions [11], [2], [21], [15]. Predicted contacts are considered when there exist mobility patterns in applications [26], [14], [16], [3], [9], [23], [22]. Opportunistic contacts deal with completely uncertain circumstances where mobile nodes meet each other by random chance. Our work falls into the category of scheduled contacts [21], [11].

Graph representations are widely applied in studying routing strategies. In [17], [5], [21], evolving graphs are employed to model topological mutations in DTNs. In our work, we use compression and indexing techniques to efficiently explore evolving graphs with the purpose of minimizing the communication cost.

Multicast for DTN has recently drawn considerable attention. In [28], semantic models are proposed to unambiguously describe multicast in the context of DTNs. The throughput of multicast is discussed in [13] and mobility-assisted routing is used to improve the throughput bound of wireless multicast. In [8], multicast problems in DTNs are considered in a social network setting where centrality and community in DTNs are employed to help determine the appropriate selection of relays, with the objective of minimizing the delay of multicast message transmissions. In this paper, we study a novel optimization problem which is similar to the traditional multicast problems. However, instead of minimizing the delay of message transmissions, we are interested in minimizing the communication cost subject to some time constraints. To this end, the question of whether a node is reachable from another node is more important than the question of how a message flows in the network. In our problem setting, the DTN is a medium to propagate information, and our goal is to maximize the role it plays in information sharing.

Graph indexing systems have been widely studied by the database community. The most common approaches aim to efficiently solve problems as graph matching [19], [25] or reachability test [6], [12]. The closest to our work are systems for reachability tests, which aim to efficiently check if two vertices are reachable from each other (a path that connects them exists) in a directed graph. Some systems use chains [10] (generalization of paths) decomposition or path-tree [12] decomposition. The underlying idea is that if a vertex u of a chain (or a tree-of-paths) is reachable from another vertex v , all

the vertices downstream in that chain are reachable from v . We use a similar idea, but our system is designed to fast identify the regions of the graph that can reach a given destination instead of verifying the reachability between pairs. Moreover, we give a method for quickly identifying a small subset of representative vertices that allows us to solve the demand cover problem optimally and with reasonable efficiency on large datasets.

III. PRELIMINARIES

In this section we describe some basic concepts concerning DTNs that introduce our approach. We consider a network of moving nodes whose *trajectories* are known in advance or can be predicted. When two nodes enter each other's radio coverage area, a link between them is formed and a *contact* (or encounter) begins. A contact between two nodes terminates when they lose radio connectivity as they move away from each other. Contact beginnings and contact ends are also called *transitions*. The status of the network at a certain time instant can be described by a *connectivity graph*, whose vertices represent moving nodes and a link is placed between two nodes if their distance is within a given threshold d , called *radio range*. The network dynamics can be described by a series of snapshots of the connectivity graph over the time horizon [5], [11]. All the snapshot graphs are aggregated in a unique composite graph (named *space-time graph*) where vertices corresponding to the same moving node in two consecutive connectivity graphs are joined by a *temporal link*. In contrast to *spatial links*, temporal links are directed. A message can travel across a so called *space-time path*. If some spatial links toward the destination are available, the message is forwarded, otherwise the message is *carried* by the moving node (a temporal link is traversed) and forwarded when another suitable node is encountered. In the following, we refer to *route* for indicating the space-time path that a message traverses.

Fig. 1(b) shows the space-time graph corresponding to Fig. 1(a). Three snapshot graphs are represented, each of them describing the connectivity of the network at the time intervals $[t_0, t_1)$, $[t_1, t_2)$ and $[t_2, t_{MAX})$ respectively. Contiguous snapshots are joined by temporal links (in dotted line). Each snapshot is associated with its *lifetime*, i.e. the extent in time that it refers to. Message routes travel across paths of the space-time graph that can include both spatial and temporal links.

IV. PROBLEM DESCRIPTION

In this section we define formally the demand cover problem and give some baseline approach for it. We consider a set of n nodes (numbered $1, 2, \dots, n$) that move following certain *trajectories* (T_1, T_2, \dots, T_n respectively). A trajectory associates each time instant t in the range $[t_0, t_{MAX})$ with the position in the space (typically a plane) that the corresponding node occupies at time t . At a specific time, two nodes can communicate with each other through a so called *contact*

transmission (short-range, typically radio) if they are *in contact*, i.e. their Euclidean distance is within a fixed threshold d . The contact transmission (between nodes) does not involve any cost. Each node can also communicate at any time with a *data source* (Internet or a central server) through a costly *remote transmission* (cellular or satellite), with a certain cost. Our approach can be extended to a decentralized scenario where a central server is not available and data are distributed among nodes.

We consider the problem of delivering *data objects* to multiple *destinations*. In contrast to other multicast approaches in which static messages are sent to multiple destinations, we consider the problem of sharing data objects that are continuously updated over time. The dynamic character of data objects introduces new constraints: each destination needs to receive the object before a given *deadline* and with a delay that is limited by a given *latency*. We define the *data demand* I of a data object as a set of *data needs*, i.e. triples of the form (i, t, δ) , where i , t and δ represent the destination, the deadline and the latency respectively. We call the instant $t - \delta$ *release time*. It represents the earliest time instant in which an object can leave the data source.

The data flow is modeled by two kind of transmissions: a *remote transmission* is denoted by a pair (i, t) , where i represents the node that receives the object and t is the time instant in which the transmission occurs; a *contact transmission*, is represented by a triple (i_1, i_2, t) where i_1 , i_2 , t represent the node that transmits the object, the node that receives the object and the time instant at which the transmission occurs respectively. For simplicity, all the transmissions are considered instantaneous. Although it may not be true in all cases, the movement between nodes is usually very slow compared to the speed of transmission. Therefore, in most cases all the necessary objects can be transmitted before the contact terminates. We say that a remote transmission (i_s, t_s) covers a data need (i_d, t_d, δ) if there exists a sequence of contact transmissions $(i_0, i_1, t_1), (i_1, i_2, t_2), \dots, (i_{k-1}, i_k, t_k)$ with $i_0 = i_s$ and $i_k = i_d$, such as $t_s \leq t_1 \leq t_2 \leq \dots \leq t_k \leq t_d$ and $t_d - t_s \leq \delta$. The set of data needs covered by a remote transmission is also called *coverage* of the remote transmission. The demand cover problem is defined as follows:

Problem definition: Given a set of trajectories and a data object with demand I , detect the minimum set of remote transmissions that covers all the data needs in I .

The formulated problem can be shown to be NP-hard (even in the 2D plane) by reduction from the well known Set-cover problem. Given a family of sets $\mathcal{S} = \{S_1, S_2, \dots, S_m\}$ of elements taken from a set C , Set-cover calls for finding the minimum sub-family of S that covers all the elements of C . A proof is given in our extended technical report [18].

A. ILP formulation

The demand cover problem can be formulated in ILP (Integer Linear Programming) and solved by a standard solver. Here we give an ILP formulation and show that solving it on large datasets is infeasible.

We consider a set of n moving nodes numbered 1 through n and a special node that represents the central server, numbered 0. We write $i \rightarrow_t j$ if node i can communicate with node j at time t (i.e. they are within distance d or $i = 0$). We also consider a discrete set \mathcal{T} of time instants that correspond to transitions or deadlines of data needs. This restriction does not compromise the result. In fact, given an optimal solution for demand cover, it is always possible to modify this solution in such a way that each transmission between two nodes is delayed until the contact between them ends (right before the link breaks) or a data need that involves one of the nodes expires, without increasing the cost. Since communication is assumed to be instantaneous, the contact length is not important.

We employ two classes of boolean variables. The first class contains variables of the kind $x_{i,j,t,r}$, where i and j represent nodes, t represents a time instant and $r = (i_r, t_r, \delta_r) \in I$ represents a data need, that models the flow of data objects. The variable $x_{i,j,t,r}$ has value 1 if i send a message to j at time t to satisfy the data need r . Variable of this kind are considered for $i \rightarrow_t j$ and $t_r - \delta_r \leq t \leq t_r$. The second class of variables, of the kind $y_{i,t}$, counts the number of remote transmissions. Each variable says whether a remote transmission between the central server and a particular node occurs at a certain time or not. The complete formulation follows.

$$\begin{aligned}
\min \quad & \sum_{t \in \mathcal{T}} \sum_{i=1}^n y_{i,t} \\
\text{s.t.} \quad & \sum_{\substack{t \in \mathcal{T} \\ t \geq t_r - \delta_r \\ t \leq t_r}} \sum_{\substack{i=0 \dots n \\ i \rightarrow_t i_r}} x_{i,i_r,t,r} \geq 1 \quad \forall r = (i_r, t_r, \delta_r) \in I \quad (1) \\
& \sum_{\substack{t' \in \mathcal{T} \\ t_r - \delta_r \leq t' \leq t}} \sum_{\substack{i=0 \dots n \\ i \rightarrow_{t'} j_1}} x_{i,j_1,t',r} - x_{j_1,j_2,t,r} \geq 0 \quad (2) \\
& \forall r = (i_r, t_r, \delta_r) \in I, \forall j_1, j_2, t \mid j_1 \rightarrow_t j_2 \\
& y_{i,t} \geq x_{0,i,t,r} \quad \forall r \in I, i = 1 \dots n, t \in \mathcal{T} \quad (3) \\
& x_{i,j,t,r}, y_{i,t} \in \{0, 1\}
\end{aligned}$$

Constraint 1 models the fact that for each data need, the data object must be sent to the destination at a time instant between the release time and the deadline. Constraint 2 models the propagation of data objects. It says that if a data object is transmitted from a source j_1 to j_2 at time t for satisfying a data need r , then j_1 must receive the object before t and after the release time. Finally, constraint 3 assigns value 1 to each variable of the kind $y_{i,t}$ if a message is transmitted from the central server to node i at time t in order to satisfy some data need.

Solving this formulation with standard solvers is infeasible on large instances. The main problem concerns the number of variables and constraints. Even considering a sparse network with 100 nodes, 100 encounters per node and 100 data needs, we have hundreds of millions of variables and constraints. In our experiments we obtained hundreds of billions of variables and constraints, therefore executing it was not possible.

B. A naive approach for the demand cover problem

An improvement on executing the ILP program can be obtained by reducing the problem to Set-cover. Each candidate remote transmission can cover a set of data needs. The minimum set of remote transmissions that covers all the data needs corresponds to the minimum Set-cover in the family of associated sets. Since remote transmissions can occur at any time, the number of sets for the Set-cover family is huge. However, not all time instants need to be considered. To guarantee that all the data needs are covered, one can consider only time instants that correspond to the release time of a data need. We note that the release time is the earliest time instant in which the data object needs to be sent for a data need to be satisfied. Delaying a remote transmission after the release time does not help with serving the data needs, unless some other release times are overtaken.

Given a candidate remote transmission, the set of covered data needs can be computed by exploring the space-time graph (through depth-first search or breadth-first search).

C. A compact graph representation

The naive approach requires exploring a space-time graph, whose size can be huge. However, this graph can be compacted thanks to two observations. First, in each snapshot graph, all the vertices that are in the same connected component have the same reachability properties, so one vertex can be taken as a representative of all the others. Second, when a transition occurs, connected components of the snapshot graph that do not contain any nodes involved in the transition are not influential.

To generate the *compressed graph*, we focus on two class of transitions. A *split transition* causes a connected component to divide in two connected components. A *merge transition* causes two components to merge in a unique connected component. We generate a space-time graph considering only this two kind of transitions. Then, for all snapshot graphs, each connected component is collapsed in one single vertex. At this point, all the edges of the graph are directed and can be classified as follows: (i) *split edges*, that connect splitting components with their partitions; (ii) *merge edges*, that connect merging components with the resulting components and (iii) *non-influential edges*, that connect components that do not change. Finally, each non-influential edge is removed by collapsing its endpoints and each vertex is labeled with its lifetime (note that a vertex can span several snapshots), that we call *component lifetime*.

One example of a compressed graph is shown in Fig. 2, that refers to the example in Fig. 1. Boxes represent vertices of the compressed graph. Edges of the connected graph are represented by solid lines. The extent of a box in time represents the component lifetime. For instance, the extent of the component c_1 is $[t_0, t_{MAX})$ (the whole time horizon) since this component is never involved in any split or merge. The naive approach can be executed on the compressed graph in place of the cumbersome space-time graph. The family for Set-cover is obtained by building a set for each vertex of the

compressed graph whose lifetime contains the release time of some data needs. Each set can be computed by exploring the compressed graph.

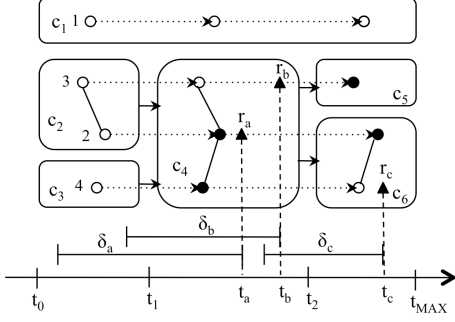


Fig. 2. The compressed graph representation of the example in Fig. 1. A compressed graph is depicted over the space-time graph. Boxes and solid lines represent vertices and edges of the compressed graph respectively. The extent of a box in time represents the component lifetime. Three data needs are represented (by filled triangles) with their extent in time. From left to right: $r_a = (2, t_a, \delta_a)$, $r_b = (3, t_b, \delta_b)$, $r_c = (4, t_c, \delta_c)$.

V. AN INDEXING SYSTEM FOR THE DEMAND COVER PROBLEM

Solving the demand cover problem efficiently raises several challenges. First, for each vertex v of the compressed graph, the set of data needs that can be covered by v needs to be retrieved. This operation may be very expensive when the size of the graph is large. Second, Set-cover is NP-hard, therefore no polynomial-time solutions exist (unless $P=NP$) in the general case. For small instances, Set-cover can be solved optimally in acceptable time by pruning techniques as *branch-and-bound*. In our case, however, the number of sets generated is usually very high, since a data need can potentially be covered by many vertices. Many of these sets are redundant, i.e. they are fully contained in other sets. For example in Fig. 2 the set of data needs covered by c_6 contains only r_c . The vertex c_4 covers the set $\{r_a, r_b, r_c\}$, that contains the data need covered by c_6 . Therefore, c_6 can be excluded by the computation since all the data needs that can be covered by it can also be covered by c_4 . Removing redundant sets leads to a considerable reduction of the Set-cover instance. However, removing the redundancy by traditional methods is expensive, since it requires one to find *maximal sets* [24]. Additionally, in a typical application, a large number of data objects are requested and each data object has its own set of data needs. Solving the demand cover problem for each data object can be extremely expensive.

We propose a novel approach, *Path-wise indexing* (PIE, for short), that builds an index of the set of trajectories with the purpose of efficiently performing queries of the form: *given a set of data needs, return the minimum set of remote transmissions that covers all the data needs*. We use a preprocessing-filtering-optimization scheme to solve the demand cover problem. Given a database of trajectories, a *preprocessing* phase generates a compact data index. When a query (represented by a set of data needs) has to be performed,

we use the data index to generate a lightweight instance of Set-cover (*filtering* phase). Set-cover is then solved optimally (*optimization* phase) and the solution is returned.

The proposed indexing system has several advantages. First, the index is much more compact than the compressed graph, and hence requires less memory and is much more efficiently manageable. Second, the set of vertices in the compressed graph from which the data needs are reachable can be identified fast. Note that current reachability indexes cannot be efficiently applied to our problem since many reachability tests need to be performed. Finally, we can efficiently prune nodes of the compressed graph that are not promising and generate a small instance for Set-cover. Next, we introduce the proposed index and describe the three phases of our indexing system: preprocessing, filtering and optimization.

A. Index structure

The key idea is that the set of data needs covered by a node in a path p of the compressed graph includes the set of data needs covered by other subsequent nodes p . Therefore, a node can be taken as a representative of a portion of the path. Moreover, a node of the compressed graph can be uniquely determined by a path and a time instant. This implies that we can use the coverage of a pair (p, t) in place of the coverage of the corresponding compressed node. We denote the coverage of (p_i, t) as $C(p_i, t)$. Based on these considerations, we partition the compressed graph into a set of disjoint paths and build a compact graph, named a *PIE graph*, whose vertices represent disjoint paths and edges preserve the connectivity across paths. Each vertex of the PIE graph is labeled with a time interval (named *lifetime*) that is the union of the lifetimes of its composing vertices. Instead of exploring all the nodes of a path, we can determine a set of time instants that is representative of the whole path by exploring the compact PIE graph.

Figure 3(a) shows an example. The small circles and thin edges form the compressed graph, while the big ovals represent disjoint paths. Consider the path p_3 . The set of data needs covered by p_3 at time t_a is $C(p_3, t_a) = \{r_1, r_2\}$. Since no other data needs (i, t, δ) have $(t - \delta) \in [t_a, t_b)$, (p_3, t_a) is representative of the interval $[t_a, t_b)$. t_b coincides with the *release time* of r_3 (i.e. $(t_{r_3} - \delta_3)$). Therefore, its coverage $(C(p_3, t_b) = \{r_3\})$ cannot be contained in $C(p_3, t_a)$. The pair (p_3, t_b) is instead representative of the remaining part of the path. The path p_3 produces only two sets $(C(p_3, t_a)$ and $C(p_3, t_b))$ for Set-cover. In general, up to 4 sets would be produced without indexing, since we may have many other non reachable data needs whose release times falls within all vertices of p_3 . Next we describe in details the three steps of our method: preprocessing, filtering and optimization.

B. Preprocessing

Given the set of trajectories, first a compressed graph (G_C) is generated. The graph is then decomposed into a disjoint set of paths. There is a large number of possible ways to partition the graph into disjoint paths. A suitable partition




Fig. 3. (a) An example of PIE graph. The small circles and thin arrows form the compressed graph. Each path is circumscribed by an oval and its lifetime is reported. Links between paths are represented by thick arrows. They are labeled by the ends of the lifetimes of their source vertices. Solid triangles within circles represent data needs. (b) Validity intervals of a set of data needs in a path p_3 . Bars represent the extent of validity intervals of data needs. The minimal family of sets for this path is $\{C(p_3, t_a), C(p_3, t_b)\}$.

strategy should satisfy the following properties: (i) the number of disjoint paths is minimal and (ii) the number of edges across two paths is minimal. In general, finding the minimum set of disjoint paths that covers a graph is a non-trivial problem [7]. However, since the compressed graph is a DAG and is generated by a simple split-merge model, we can use optimally the following simple strategy: pick one vertex a time (proceeding in time order) and elongate it by random walk until a vertex without outgoing edges is reached.

We can prove that across two paths no more than one edge exists in each direction. Indeed, each edge of the compressed graph comes from a merge or a split between two components. In the case of merge, the source vertex cannot have other outgoing edges, while in the case of split, the target vertex cannot have other incoming edges. This implies that no edges can exist between internal nodes of two different paths, and hence each edge connecting two paths can be either outgoing from the last vertex of the source path or incoming to the first vertex of the target path. Since the compressed graph is a DAG, and each path is elongated as much as possible, at most two edges can connect two paths, one in each direction.

We associate each edge (p_i, p_j) of the PIE graph with the end of the lifetime of the source vertex in p_i . We denote this time instant as $ft(p_i, p_j)$. It represents the time in which a data object can traverse the edge (p_i, p_j) . Figure 3 depicts an example of PIE graph. The small vertices and thin edges form the compressed graph, while the big vertices and thick edges represent the PIE graph.

C. Filtering

For each vertex p of the PIE graph, our filtering algorithm finds a set of time instants TI_p that are representative of the whole path p , and the family \mathcal{S} of corresponding sets. Our strategy guarantees that the coverage of each vertex of the compressed graph is fully contained in at least one set in \mathcal{S} . Since the PIE graph is much smaller than the compressed graph, exploring the former is much more advantageous in terms of elaboration time and memory consumption.

The filtering procedure considers two steps: *backflow* and *prune*. *Backflow* propagates the data needs in reverse order from the destination paths to all the possible source paths. For each path, we compute the *validity interval* of a data need, that defines the time interval in which the data object must reach the path for the data need to be covered. At the end, each path is associated with a set of data needs that it can cover with their validity intervals. The coverage of a pair (p, t) can be identified by the set of data needs such as their validity intervals in p include t . After the validity intervals are generated, the *prune* procedure computes the family of sets for Set-cover. It collects the family of maximal coverage sets of each path, i.e. the set of time instants whose coverage is not strictly contained in the coverage of any other time instant in the path.

Before describing these two procedures in detail, we give an example. Fig. 3(b) shows the path p_3 of the example in Fig. 3(a) and the validity interval of each data need in it. The validity intervals of r_1 and r_2 start at the beginning of the path, since their release time precede it. These intervals end at times t_1 and t_2 respectively, times associated to outgoing edges (see Fig. 3). Each of them represents the last time instant in which the data object must leave the path to be able to reach the respective data need. For the data need r_1 (r_2 resp.), if the data object leaves the path after t_1 (t_2 resp.), the destination cannot be reached. The validity interval of r_3 starts at time $t_b = t_{r_3} - \delta_3$, corresponding to the release time of r_3 , and ends at time t_3 , time associated to the unique outgoing edge that can reach r_3 . The representative time instants for this path are t_a and t_b , corresponding to maximal sets of data needs. Therefore, the minimal family of sets for this path is $\mathcal{S} = \{C(p_3, t_a), C(p_3, t_b)\}$. Note that no other time instants have a coverage that is not included in at least one set of the family.

1) *Backflow*: We define the *validity interval* of a data need $r = (i, t, \delta)$ in a path p (named $valid_int(r, p)$) recursively in the following way:

If p has lifetime $[b, e)$ and is the destination path of r (i.e. $t \in [b, e)$), we have: $valid_int(r, p) = [\max(b, t - \delta), t)$.

If p is a non-destination path with lifetime $[b, e)$ that links to a set of paths p_1, p_2, \dots, p_k with validity intervals $[b_1, e_1)$, $[b_2, e_2)$, \dots , $[b_k, e_k)$ respectively:

$$valid_int(r, p) = \begin{cases} \Phi & \text{if } ft(p, p_i) \notin [b_i, e_i) \forall i = 1 \dots k \\ [t_1, t_2) & \text{otherwise} \end{cases}$$

where $t_1 = \max(b, t - \delta)$ and t_2 is the maximum t' such as $t' = ft(p, p_i)$ for some $i = 1 \dots k$ and $t' \in [b_i, e_i)$.

Intuitively the end of a validity interval in a path is given by the last time instant in which the data object can flow in another path that has a compatible validity interval, while the beginning of a validity interval is limited by $t - \delta$ and the starting time of the path.

The coverage of a pair (p, t) can be identified by the set of data needs whose validity intervals include t . Intuitively, if validity intervals are represented by horizontal bars (as in Fig. 3(b)), the coverage of a pair (p, t) can be easily identified by drawing a vertical line and taking all the data needs whose validity intervals are intersected. For instance, in Fig. 3(b) a vertical line drawn at time t_a intersects the validity intervals of r_1 and r_2 . Therefore, the coverage of (p, t_a) is $\{r_1, r_2\}$. This property is formally stated by the following lemma:

Lemma 1: Let (T, I) be an instance of demand cover, where T is the set of trajectories and I is the set of data needs, and G_P be the corresponding PIE graph. Given a vertex p of G_P and a time instant t , the coverage of p at time t is:

$$C(p, t) = \{r \in I \mid t \in \text{valid_int}(r, p)\}$$

$\text{valid_int}(r, p)$ can be computed for all paths in a breadth-first search fashion, by starting from the path containing r and exploring the PIE edges in reverse time order until the release time is reached. When a new vertex is visited, the validity interval of r in it is updated. The complexity is $O(|E_P|)$, where E_P is the set of edges in the PIE graph.

2) *Prune:* For each path p , we identify the minimum-size set TI_p of time instants that is representative of the whole path, i.e. such that for all $t \in \text{lifetime}(p)$ we have $C(p, t)$ contained in at least one set $C(p, t')$ with $t' \in TI_p$. This problem corresponds to the problem of finding the maximal sets in the family of all possible coverage sets of p (i.e. $\{C(p, t) \mid t \in \text{lifetime}(p)\}$).

Figure 4 shows an example of path with the validity intervals of five data needs. The coverage of a time instant can be easily identified by drawing a vertical line and taking all the validity intervals that it intersects. The representative time instants for this path are t_1, t_2 and t_3 , corresponding to the maximal sets of data needs. Note that no other time instants have a coverage that is not included in the coverage of at least one of the time instants t_1, t_2 or t_3 .

In general, the maximal sets can be found in time $O(mn)$, where m is the number of maximal sets and n is the size of the input [24]. In our case, since each element corresponds to a contiguous interval, we can find the maximal sets in linear time. Our procedure slides a vertical line across the path in reverse time order, and takes all the time instants that correspond to maximal sets. Each position t of the line corresponds to a coverage $C(p, t)$. As the line is slid, the coverage is modified, by either adding or deleting data needs. Whenever a deletion follows an addition, the current coverage is taken as a maximal set. Note that additions correspond to the end of validity intervals, while deletions correspond to the beginning of validity intervals. In Fig. 4, the coverage associated with the sliding line is initially empty. When the




Fig. 4. An example of maximal coverage sets in a path. Bars represent the extent of validity intervals of data needs. The coverage of the time instants t_1, t_2 and t_3 are maximal sets among all the coverage sets in the path. The family of maximal sets can be found by sliding a vertical line in reverse time order and taking each time instant that corresponds to the beginning of a validity interval (indicated by the symbol “-” at the top) that occurs right after the end of the same or another validity interval (indicated by the symbol “+”). This family has minimum size.

line intersects the validity interval of r_4 , r_4 is added to the coverage (additions are indicated by the symbol “+” at the top). The interval of r_5 is then encountered and r_5 is also added to the coverage. When the beginning of the validity interval of r_4 is encountered (at time t_3), the current coverage is taken as maximal set and r_4 is deleted (indicated by the symbol “-”). Other two additions are then encountered (r_2 and r_3) followed by a deletion (r_5). The coverage at time t_2 (before deleting r_5) is then taken as another maximal set. The last maximal set is taken at time t_1 , after another addition and another deletion are encountered. The following lemma states that this procedure finds all and only the maximal sets in the family of coverage sets.

Lemma 2: Let (T, I) be an instance of demand cover and G_P be the PIE graph built from (T, I) . Given a path p of G_P , consider the sequence of time instants t_1, t_2, \dots, t_k corresponding to extremes (beginnings or ends) of validity intervals in reverse time order and the set $TI_p = \{t_i \mid t_i \text{ is a beginning time and } t_{i-1} \text{ is an ending time}\}$.

- 1) For each time instant $t \in \text{lifetime}(p)$ we have: $\exists t' \in TI_p \mid C(p, t) \subseteq C(p, t')$;
- 2) For each time instant $t' \in TI_p$ we have: $\nexists t \in \text{lifetime}(p) \mid C(p, t') \subset C(p, t)$;
- 3) For each pair of distinct time instants $t', t'' \in TI_p$ we have $C(p, t') \neq C(p, t'')$.

A clear consequence of this lemma is that the family of maximal sets generated by our procedure has minimum size. TI_p can be built in time $O(|I| \cdot \log(|I|))$.

D. Optimization

After the filtering process, a *post-pruning* (in short PP) phase is applied in order to remove sets that are fully contained in other sets. Note that although the purpose of the filtering procedure is to remove these sets, this procedure is not guaranteed to be exhaustive, since redundant sets can occur across different paths. The post-pruning phase can be applied to the naive approach as well.

We use an Integer Linear Program to solve Set-cover optimally. Finally, the optimal set of remote transmissions is extracted from the optimal subfamily returned by Set-cover.

E. Adaptive extension

In real world, it is difficult for many applications to guarantee that moving objects travel with known trajectories over a long time interval. However, it is reasonable to assume that moving objects stick to known traveling plans in near future. In this case, the time dimension is partitioned into discrete time slots, where trajectories of moving objects are updated after each time slot. PIE can adapt to this variation without much modification. In the following, we briefly introduce two possibilities: *null-initial-state* and *adaptive* extensions.

The most straightforward way is to build an independent index for each time slot. We call it null-initial-state extension because this method simply ignores previous knowledge and treats each time slot as a new start. One weakness of this method is that it neglects information objects transmitted during the prior time slots, producing more remote transmissions. An alternative is to apply the adaptive extension. To reuse data objects transmitted before, we keep track of the distribution of the objects over the nodes, together with the remote transmission time of each object, and use them as initial state for the new time slot. Some data needs can be satisfied without any additional remote transmissions and will not be considered in the computation.

VI. EXPERIMENTAL ANALYSIS

A. Dataset

Cabs Mobility [20] (CAB, for short) contains mobility traces of taxi cabs in San Francisco, USA. It consists of GPS coordinates of 536 taxis collected over 23 days in the San Francisco Bay Area. The average time interval between two consecutive location updates is less than 10 seconds.

GeoLife GPS Trajectories [1] (GeoLife, for short) is a GPS trajectory dataset collected in (Microsoft Research Asia) GeoLife project by 165 users in a period of over two years.

Synthetic trajectories (SYN, for short) consists of 10K nodes that move randomly on a 2-D plane with size 3600 km^2 over 10 days. Starting from a uniformly random position, the speed of each node is updated periodically with normal distribution ($\mu = 1.2$ m/s and $\sigma = 1$) as well as its direction ($\mu =$ current direction and $\sigma = 1$ radian). The updating rate is generated with exponential distribution ($\mu = 60$ sec).

For all datasets, the radio range is set to 100 meters. The data needs are generated by the following process. First, for each moving node, the number of data needs is generated with a Poisson distribution. Then, each data need is generated with a deadline uniformly distributed and a latency normally distributed ($\mu = 15$ min and $\sigma = 1$).

B. Implementation

We implemented the naive approach described in Sect. IV-B on the space-time graph. We also implemented the naive approach on the compressed graph (called *naive-c* for short)

and the PIE indexing system (Sect. V). All methods include the post-pruning phase described in Sect. V-D. We tried a version without the post-pruning phase, obtaining a slight degradation of performances in each method. We also tried to run the ILP program described in Sect. IV-A, but it did not terminate due to the huge number of variables and constraints (hundreds of billions) considered. All the approaches were implemented in C++ (Dev C++ IDE ver. 4.9.9.2). The experiments were performed on a DELL Intel core I7 CPU with 2 Gb of memory. For the ILP solver we use *lp_solver* 5.5.2.0 [4], an open source tool based on branch-and-bound.

C. Results

Each dataset is first preprocessed and its PIE index is generated. Fig. 5(a) reports the preprocessing time on CAB, concerning a number of datasets spanning from 1 to 13 days. Depending on the dataset size, the preprocessing phase takes through thousands of seconds. Although the preprocessing phase is sometimes expensive, it is executed only once. The rate of compression of the PIE graph and the compressed graph with respect to the space-time graph is shown in Fig. 5(b). The compressed graph is about four folds smaller than the space-time graph and PIE further reduces the size of about three times.

On CAB, the execution time for demand cover queries is shown in Fig. 5(c) and 5(d). Fig. 5(c) shows the execution time for a number of datasets spanning from 1 to 13 days. The average number of data needs per cab per day is set to 2. The reported times represent an average over 10 queries. PIE performs about one fold faster than naive-c and two folds faster than naive in almost all cases. In order to evaluate the scalability over the size of the query, we generate queries by varying the expected number of data needs per cab per day from 1 to 4. The results over 1 day are reported in Fig. 5(d). For more than 4 expected data needs, the naive method is unable to answer queries in acceptable time.

We also execute the adaptive extension (Sect. V-E) on CAB, for one day with time slot 15 minutes. Over a total number of 1089 data needs, null-initial-states method returns 675 remote transmissions, while the adaptive method returns 617 ones, with approximately a 10% improvement. For reference, the number of transmissions suggested by using full knowledge is 480. All the results of the adaptive extension refer to an average over 10 executions.

Fig. 5(e) show the execution time for demand cover queries on GeoLife. The expected number of data needs per person per day is set to 10. The results refer to a set of datasets, each of them spanning a time interval ranging from 1 to 30 days. As for CAB, the reported times represent an average over 10 queries. In this dataset, PIE scales better than naive and naive-c with length of the spanning interval. For SYN, the results are reported in Fig. 5(f). The naive approach here is not able to terminate in acceptable time even for one day, therefore we report only PIE and naive-c. PIE performs about three folds faster than naive-c. Additional results are provided in our extended technical report [18].

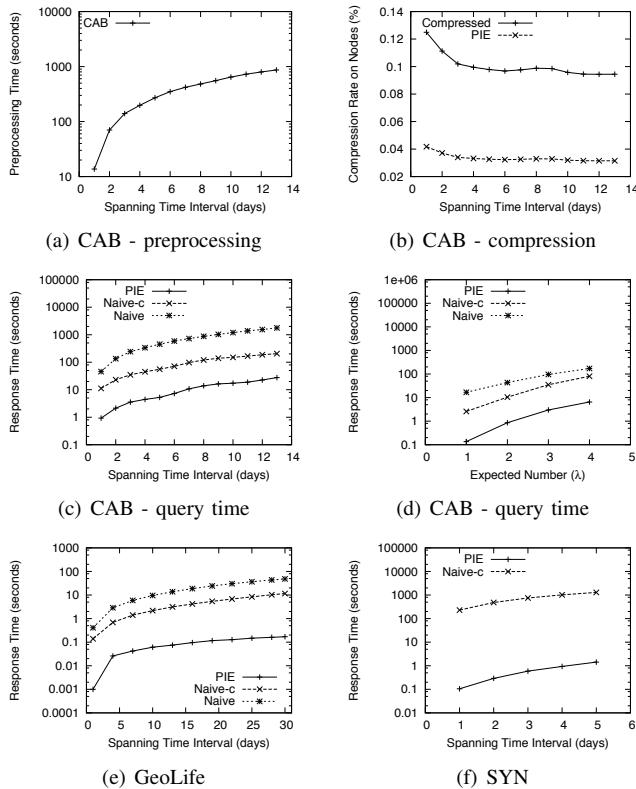


Fig. 5. Evaluating the response time as a function of the size of the dataset (number of days) and the number of data needs

VII. CONCLUSION

We presented a new approach that optimizes the long-range communication cost for multicast in DTN. After formalizing the demand cover problem and showing that it is NP-hard, we provided a graph-indexing-based solution for it. Our system can solve the demand cover problem optimally on large real instances (dataset with million of events and queries with thousands of nodes) in less than 10 seconds in most cases. We plan to extend our work in two ways. First, we aim to take into account the uncertainty in mobility and data needs. For this, we need to fit stochastic mobility models in our framework and optimize the expected communication cost. Finally, we plan to consider the problem of scheduling new trajectories with the purpose of guaranteeing the connectivity, in the case when the communication with a central data source is not always available.

ACKNOWLEDGEMENTS

Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-09-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

REFERENCES

- [1] GeoLife GPS Trajectories. <http://research.microsoft.com/>.
- [2] I. F. Akyildiz, O. B. Akan, C. Chen, J. Fang, and W. Su. Interplanetary internet: state-of-the-art and research challenges. *Comput. Netw.*, 43:75–112, October 2003.
- [3] E. Altman, G. Neglia, F. De Pellegrini, and D. Miorandi. Decentralized stochastic control of delay tolerant networks. In *Proc. of INFOCOM*, pages 1134–1142, 2009.
- [4] M. Berkelaar et. al. Mixed Integer Linear Programming (MILP) solver. <http://sourceforge.net/projects/lpsolve>.
- [5] V. Borrel, M. H. Ammar, and E. W. Zegura. Understanding the wireless and mobile network space: a routing-centered classification. In *Proc. of ACM CHANTS*, pages 11–18, 2007.
- [6] S. Chen et. al. ST2B-tree: a self-tunable spatio-temporal B+ tree index for moving objects. In *SIGMOD*, pages 29–42, 2008.
- [7] R. Diestel. *Graph theory*. Graduate texts in mathematics. Springer, 2006.
- [8] W. Gao, Q. Li, B. Zhao, and G. Cao. Multicasting in delay tolerant networks: a social network perspective. In *Proc. of MobiHoc*, pages 299–308, 2009.
- [9] R. Groenevelt, P. Nain, and G. Koole. Message delay in manet. *SIGMETRICS Perform. Eval. Rev.*, 33:412–413, June 2005.
- [10] H. V. Jagadish. A compression technique to materialize transitive closure. *ACM Trans. Database Syst.*, 15:558–598, December 1990.
- [11] S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. volume 34, pages 145–158, August 2004.
- [12] R. Jin et. al. Efficiently answering reachability queries on very large directed graphs. In *SIGMOD*, pages 595–608, 2008.
- [13] U. Lee, S. Y. O., L. K., and M. Gerla. Relaycast: Scalable multicast routing in delay tolerant networks. In *IEEE ICNP*, pages 218–227, 2008.
- [14] J. Leguay, T. Friedman, and V. Conan. Dtn routing in a mobility pattern space. In *Proc. of WDTN*, pages 276–283, 2005.
- [15] Q. Li and D. Rus. Sending messages to mobile users in disconnected ad-hoc wireless networks. In *Proc. of MobiCom*, pages 44–55, 2000.
- [16] C. Liu and J. Wu. An optimal probabilistic forwarding protocol in delay tolerant networks. In *Proc. of MobiHoc*, pages 105–114, 2009.
- [17] S. Merugu, M. Ammar, and E. Zegura. Routing in space and time in networks with predictable mobility. Technical report, 2004.
- [18] M. Mongioli, A. K. Singh, X. Yan, B. Zong, and K. Psounis. Efficient multicasting for delay tolerant networks using graph indexing - Technical report. <http://www.box.net/shared/7xm4ptobu4>, 2011.
- [19] M. Mongioli et. al. SIGMA: a set-cover-based inexact graph matching algorithm. *J Bioinform Comput Biol*, 8(2):199–218, 2010.
- [20] M. Piorkowski, N. Sarafjanovic-Djukic, and M. Grossglauser. CRAWDAD data set epfl/mobility (v. 2009-02-24). <http://crawdad.cs.dartmouth.edu/epfl/mobility>, Feb. 2009.
- [21] B. K. Polat, P. Sachdeva, M. H. Ammar, and E. W. Zegura. Message ferries as generalized dominating sets in intermittently connected mobile networks. In *Proc. of the MobiOpp*, pages 22–31, 2010.
- [22] T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Efficient routing in intermittently connected mobile networks: the multiple-copy case. *IEEE/ACM Trans. Netw.*, 16:77–90, February 2008.
- [23] T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Efficient routing in intermittently connected mobile networks: the single-copy case. *IEEE/ACM Trans. Netw.*, 16:63–76, February 2008.
- [24] D. M. Yellin. Algorithms for subset testing and finding maximal sets. In *SODA*, pages 386–392, 1992.
- [25] S. Zhang, J. Yang, and W. Jin. SAPPER: subgraph indexing and approximate matching in large graphs. *PVLDB*, 3:1185–1194, September 2010.
- [26] X. Zhang, J. Kurose, B. N. Levine, D. Towsley, and H. Zhang. Study of a bus-based disruption-tolerant network: mobility modeling and impact on routing. In *Proc. of MobiCom*, pages 195–206, 2007.
- [27] W. Zhao, M. Ammar, and E. Zegura. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *Proc. of MobiHoc*, pages 187–198, New York, NY, USA, 2004. ACM.
- [28] W. Zhao, M. Ammar, and E. Zegura. Multicasting in delay tolerant networks: semantic models and routing algorithms. In *Proc. of WDTN*, pages 268–275, 2005.