

# Utility-based Gateway Deployment for Supporting Multi-domain DTNs

Ting He and Kang-Won Lee  
IBM T.J. Watson Research Center  
Hawthorne, NY, USA

Email: {the, kangwon}@us.ibm.com

Nikoletta Sofra and Kin K. Leung  
Department of Electrical and Electronic Engineering  
Imperial College, London, UK

Email: {nikoletta.sofra06, kin.leung}@imperial.ac.uk

**Abstract**—Due to technology or policy constraints, communications across network domains usually require the intervention of gateways, and their proper deployment is crucial to the overall performance. In this paper, we study the problem of placing static gateways in mobile DTNs consisting of multiple domains. Given a limited gateway budget, the problem is to select deployment locations to optimize certain performance. The challenge is that different domains may possess heterogeneous properties. To ensure general applicability of solution, we propose a unified framework based on utility optimization, and solve utility computation and placement optimization separately. To handle heterogeneity, we decompose utility computation into individual domains and derive closed-form solutions based on key domain characteristics with focus on the routing scheme. Moreover, we develop quadratic-complexity algorithms to solve the optimization efficiently, which has guaranteed performance under certain uniformity conditions. Although certain assumptions have been made in developing the solutions, evaluations based on synthetic data and real DTN traces both show that the proposed solutions can achieve near-optimal performance at much lower complexities, and the results are robust with respect to the routing schemes and the mobility patterns. Compared with utility-agnostic deployments, our solutions significantly improve the end-to-end performance.

**Keywords:** Resource allocation, Performance analysis, Inter-domain communications, Disruption-tolerant networking.

## I. INTRODUCTION

Disruption-tolerant networks (DTNs), which can provide practical means for communication in an environment where the conventional networking paradigms do not apply due to sporadic connectivity, have been a hot topic in the networking community in recent years. Due to the diverse nature of wireless technologies and application scenarios, it has been suggested that one solution does not fit all, and sometimes multiple DTNs need to work together. For instance, in an urban sensing project, called MESSAGE [1], whose goal is to monitor the pollution level in the London metro area, we have employed various sensing and wireless technologies: sensor

nodes mounted on buses and cars use WiFi or WiMax technologies, while sensors mounted on pedestrians and cyclists communicate via Bluetooth and Zigbee. In this scenario, all DTN nodes can communicate with access points deployed in the city, *e.g.*, to upload collected data and pass messages to each other; different groups of nodes implement different forwarding and replication strategies to best suit their purposes and mobility patterns (*e.g.*, buses moving along fixed routes require a different strategy from pedestrians exhibiting more random mobility). In a slightly different scenario, such as an emergency situation after a disaster, DTNs can play an important role in covering areas that lack ubiquitous connectivity. In this case, multiple units of civil and military forces (*e.g.*, fire fighters, police, medical crew, red cross, and military forces) also need to quickly set up an *ad hoc* communication infrastructure to coordinate their operations.

To support such on-demand networking needs across multiple groups, where each group may be equipped with its own wireless technology and data dissemination mechanism, we must develop a solution to enable inter-communications across different DTN domains<sup>1</sup>. One way to address this inter-domain communication issue is to deploy special devices, called gateways, that can bridge the technology gaps between domains. The main focus of this paper is to study how to optimally deploy gateways to best support inter-domain communications across multiple DTNs under certain budget constraints.

*Models of networks and gateways:* To illustrate the problem more concretely, consider an example in Fig. 1, where several DTN domains have (partially) overlapping operation regions and need to communicate with each other, possibly through intermediate domains. Given a limited deployment budget (*e.g.*, a total number of gateways and the associated cost of deployment), we want to strategically place the gateways to optimize certain metrics of inter-domain communications. We model each domain by several characteristics including their mobility patterns, radio range, channel conditions, network size (*i.e.*, number of nodes), and intra-domain routing schemes. We note that the first three characteristics jointly determine the node contact behavior and the opportunity to transfer data.

Research was sponsored in part by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-06-3-0001, and the U.K. Engineering and Physical Sciences Research Council and the U.K. Department for Transport under the MESSAGE project. The views in this document are those of the authors and do not represent the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence, the U.K. Department for Transport, or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

<sup>1</sup>Here we define a domain to be a logical grouping of mobile nodes that employ the same DTN protocol (*e.g.*, all cyclists connected via Bluetooth near South Bank).

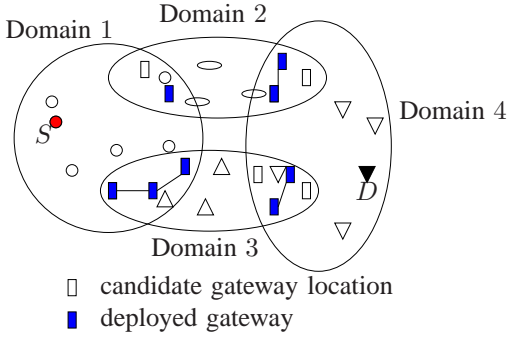


Fig. 1. Example: source  $S$  in Domain 1 needs to communicate with destination  $D$  in Domain 4, and networked gateways are deployed at selected locations in overlapped regions between neighboring domains to support such inter-domain communications.

Together with the latter two parameters, they determine the local communication performance inside a domain per specific gateway deployment. In addition, we model the global network characteristics by the domain-level topology (two domains are neighbors if their operation regions overlap) and the inter-domain routing scheme. It is also possible that domains can impose certain policies on inter-domain traffic flows, *e.g.*, a domain may route flows with different source and/or destination domains differently according to the policies. Together, these global characteristics determine how inter-domain traffic is routed among the domains. For the gateways, we consider static gateways with multiple communication interfaces and persistent storage that allow them to receive, buffer, and forward messages between supported domains. Depending on the installed interfaces, a gateway can support all or some of the domains, where it is often the case that the more domains to support, the higher the cost. We further assume that gateways can also communicate with each other through separate links, such as satellite links or a backbone network; in particular, we assume the gateways serving the same neighboring domains are fully connected (as shown in Fig. 1) so that they can share forwarding messages.

#### A. Our Approach & Contributions

We formulate the gateway deployment problem as a resource allocation problem, where the goal is to optimize the end-to-end performance of inter-domain communications, represented by utilities, under budget constraints. Different from infrastructure deployment in a single domain (*e.g.*, [2], [3]), the utilities for potential gateway placements depend on the characteristics (*e.g.*, the aforementioned five characteristics) of multiple constituent DTNs, which can vary drastically. In contrast to single-domain problems where the deployment solution can also alter other network characteristics (*e.g.*, routing) to suit its need [3], a solution for gateway deployment is often not allowed to change domain configurations but has to adjust itself to suit various types of domains. To maximize the applicability of our solution, we take a unified approach that jointly allocates gateway resources among neighboring domains while allowing models for specific types of DTNs to be plugged in. Our specific contributions are three-fold:

*Unified gateway deployment framework (UGDF):* We divide the problem into the subproblems of utility computation and gateway placement to decouple the performance prediction (*i.e.*, utility computation) that is network-dependent from the optimization (*i.e.*, gateway placement) that is independent of the networks. To handle domain heterogeneity, we decompose the global utility into local performance metrics in individual domains, which are then analyzed independently. We formulate the gateway placement problem as a generalized knapsack problem and develop two quadratic-complexity algorithms with guaranteed performance under certain uniformity conditions.

*Utility computation in heterogeneous domains:* To capture the impact of domain properties on the gateway deployment, we analyze two commonly-used utility metrics (delay and number of replicas per message) under the assumption of Poisson contacts. Based on their replication strategies, we divide DTN routing schemes into three general classes, namely forwarding, unlimited replication, and limited replication, and derive closed-form expressions of the above metrics for each class as functions of several key domain parameters including the contact rates and the number of nodes. Our results extend previous analysis to the context of multi-domain DTNs.

*Performance evaluation:* We test the proposed solutions on both synthetic mobility models and real DTN traces. While the actual contact traces are not exactly Poisson, our simulations show that our performance expressions correctly capture the trend of the actual performance, and the final deployment solutions are near-optimal. Specifically, trace-driven simulations for a bus-based DTN show that the calculated utilities are reasonably close to the actual values after a constant scaling (with a difference less than 10%), and the proposed solutions perform as well as the optimal solution from a brute-force search and up to 30% better than a random deployment.

#### B. Related Work

The problem of inter-domain communications has only been recently studied in the context of mobile ad hoc networks (MANETs). In the regime of mostly connected MANETs, Lee et al. [4] have studied the problem of selecting mobile nodes as gateways to optimize the performance of inter-domain routing protocols. Their solution, however, is for real-time communications and is not applicable to our problem.

To the best of our knowledge, there is no previous work on the optimal infrastructure deployment in multi-domain DTNs. In the single-domain context, the most relevant works are [2], [3]. Banerjee et al. [2] studied enhancing DTN performance through various types of infrastructures including base stations, wireless meshes, and isolated relays. Their focus is on comparing different types of infrastructures rather than designing the deployment of a specific infrastructure. Zhao et al. [3] studied the placement of isolated relays called throwboxes in conjunction with routing design for performance optimization. Our problem differs from this work in that we need to handle multiple domains with heterogeneous properties, where the routing scheme of each domain is given and cannot be changed

for inter-domain communications. Such heterogeneity imposes a challenge to our solution. An analogous problem in cellular networks is the deployment of cellular gateways to support communications across service vendors. The solution there, however, is reduced to finding a geographic coverage since only direct communications are allowed between users and gateways.

The rest of the paper is organized as follows. Section II reviews existing DTN routing schemes in preparation for later analysis. Section III presents a general framework for gateway deployment, followed by a detailed study of the utility computation element in Section IV. Section V evaluates the proposed solutions on synthetic data as well as traces, and Section VI concludes the paper.

## II. TAXONOMY OF DTN ROUTING SCHEMES

Given the large number of DTN routing schemes stemming from diverse applications, we start by categorizing the existing schemes so as to extract the key parameters for later analysis.

We base our taxonomy on the previous work [5], where routing schemes are categorized according to their resource assumption (unlimited or finite buffer/bandwidth)<sup>2</sup> and relay operation (replication or forwarding). We note that besides the above criteria, how many copies of a message can exist simultaneously in the domain is another key performance factor. Therefore, we further divide replication schemes into unlimited replication schemes (*e.g.*, [5], [6]), which replicate messages as many times as possible within the resource constraints, and limited replication schemes (*e.g.*, [7] and its variations [8], [9]), which explicitly control the total number of replicas to conserve resource. Table I gives a taxonomy of selected routing schemes; a more detailed description can be found in [10].

TABLE I  
CLASSIFICATION OF DTN ROUTING SCHEMES

	Resource	Number of replicas	Previous work
P1	Unlimited	Forwarding	[11], [12]
P2	Unlimited	Unlimited replication	[13], [14]
P3	Unlimited	Limited replication	[7]–[9]
P4	Finite	Forwarding	[15], [16]
P5	Finite	Unlimited replication	[5], [6]
P6	Finite	Limited replication	[17]

## III. UNIFIED GATEWAY DEPLOYMENT FRAMEWORK

In this section, we introduce a framework to plan the deployment of gateways, called the Unified Gateway Deployment Framework (UGDF). The UGDF framework consists of two elements: (1) utility computation element and (2) gateway placement element, explained in detail below.

<sup>2</sup>For clarity, we have combined schemes under either unlimited buffer size or unlimited bandwidth into the “unlimited resource” category and refer to [5] for further distinction between the two.

### A. Utility Computation

The utility computation element aims at accurately predicting the performance of potential deployments with respect to application requirements. Specifically, we represent the overall inter-domain communication performance as a global utility  $U_{\text{global}}$  which is a function of performance metrics, denoted by

$$U_{\text{global}} = U(X_1, X_2, \dots), \quad (1)$$

where  $X_i$ 's are the end-to-end performance metrics of interest. The specific utility functions will be application-dependent, but the involved metrics are generally in common, such as delay, delivery ratio, and number of replicas/transmissions per message. To handle heterogeneous domains, we further decompose these end-to-end metrics into single-domain metrics

$$X_i = f_i(x_i^{(1)}, x_i^{(2)}, \dots), \quad (2)$$

where  $x_i^{(j)}$  denotes type- $i$  metric for (intra-domain) communication at the  $j$ th *domain-level hop*, defined as the process of delivering a message from a node to the next-hop gateways or from gateways to nodes inside the next-hop domain, and  $f_i$  an aggregation operator. For example,  $f_i$  may be summation for additive metrics such as delivery delay and number of replicas/transmissions, multiplication for multiplicative metrics such as delivery ratio, or other operators such as maximization, minimization, etc. Note that (2) depends on the inter-domain path; if there are multiple pairs of source/destination domains and/or multiple paths between each pair, we need to take the average metric, *e.g.*, by computing a weighted sum of (2), weighted by the fraction of inter-domain traffic going through that path. After the decomposition, we can analyze the per-hop metrics  $x_i^{(j)}$ 's individually for each pair of neighboring domains. The core step of utility computation is to estimate these per-hop metrics as this is the step where various domain properties affect gateway deployments. We will analyze it in detail in Section IV.

*Remarks:* We note that the decomposition in (2) implicitly requires the propagations of a message in different domains to be independent and disjoint in time. While this may not hold in general, *e.g.*, whether an upstream domain uses single-copy or multi-copy routing affects the number of gateways possessing the message and thus the propagation rate in the downstream domain, it is satisfied if gateways between neighboring domains can share messages with small delays (*i.e.*, networked gateways). In practice, this can be enabled through satellite links or backbone connections between gateways.

### B. Gateway Placement

This element selects locations to deploy gateways based on their deployment costs and the utilities given by the utility computation element. Let the set of feasible deployment locations be  $\mathcal{L} = \{l_1, \dots, l_L\}$  and the deployment cost per location  $c_i$  ( $i = 1, \dots, L$ ). Since the performance metrics  $x_i^{(j)}$ 's are functions of the gateway placement, we can rewrite the overall utility as a function of gateway locations  $\mathcal{L}'$  ( $\mathcal{L}' \subseteq \mathcal{L}$ ) as

$$U_{\text{global}} = U(f_1(x_1^{(1)}(\mathcal{L}'), x_1^{(2)}(\mathcal{L}'), \dots), \dots) \triangleq \tilde{U}(\mathcal{L}'). \quad (3)$$



Given a deployment budget  $C$ , the problem is an optimization over candidate sets of deployment to maximize the overall utility, *i.e.*,

$$\begin{aligned} & \max \tilde{U}(\mathcal{L}') \\ \text{s.t.} \quad & \mathcal{L}' \subseteq \mathcal{L}, \quad \sum_{l_i \in \mathcal{L}'} c_i \leq C. \end{aligned} \quad (4)$$

Note that we use the joint utility  $\tilde{U}(\mathcal{L}')$  instead of a sum of utilities for individual locations  $\tilde{U}(l)$ 's because in general,  $\tilde{U}(\mathcal{L}') \neq \sum_{l \in \mathcal{L}'} \tilde{U}(l)$ . For example, two locations may be nearby and both have high contact rates with nodes, but it may not be worthwhile to deploy gateways at both of them because once one location has a gateway, the utility of the other will be lower.

Because of such correlation, the above optimization is harder than the 0-1 Knapsack Problem (and thus NP-hard) since the latter is a special case of (4). Although there are several algorithms that can solve the Knapsack Problem efficiently [18], they rely on the independence of utilities and thus are not applicable to (4). Even for the simpler case of equal cost (when the Knapsack Problem becomes trivial), we still have to enumerate over the  $O(L^g)$  possible solutions (assume  $g \triangleq \lfloor C/c_i \rfloor \ll L$ ), which is expensive for large  $L$ . To reduce the complexity, we investigate the following alternatives.

1) *Greedy and backward greedy solutions*: Given a set of instrumented locations (*i.e.*, locations deployed with gateways)  $\mathcal{L}'$ , define the *conditional utility* of instrumenting another location  $l$  at cost  $c$  by  $\tilde{U}(l|\mathcal{L}') \triangleq \tilde{U}(\mathcal{L}' \cup l) - \tilde{U}(\mathcal{L}')$  and the *conditional efficiency* by  $\tilde{U}(l|\mathcal{L}')/c$ . The *greedy solution* tries to find a good set of locations by sequentially selecting locations with the maximum conditional efficiencies without altering the existing selection, which works as follows:

1) Initially,  $\mathcal{L}' \leftarrow \{l^{(1)}\}$  such that

$$l^{(1)} = \arg \max_{l_i \in \mathcal{L}} \frac{\tilde{U}(l_i)}{c_i};$$

2) For  $j = 2, 3, \dots$ , do

$$\begin{aligned} l^{(j)} &= \arg \max_{l_i \in \mathcal{L} \setminus \mathcal{L}'} \frac{\tilde{U}(l_i|\mathcal{L}')}{c_i}, \\ \mathcal{L}' &\leftarrow \mathcal{L}' \cup l^{(j)}, \end{aligned}$$

until the total cost reaches  $C$ .

The greedy solution only has a complexity of  $O(gL)$ . A nice property of the greedy solution is that the sets of gateway locations are always nested, as stated below.

*Claim 3.1*: The solutions  $\mathcal{L}'_{C_1}$  and  $\mathcal{L}'_{C_2}$  of the greedy solution under budgets  $C_1$  and  $C_2$  satisfy  $\mathcal{L}'_{C_1} \subseteq \mathcal{L}'_{C_2}$  if  $C_1 \leq C_2$ .

This property ensures that the computed deployment plan is extendable when new budget becomes available. In other words, if the deployment occurs sequentially over time, then the greedy solution gives an order of deployments that optimizes the current utility without backtracking previous deployments.

When it comes to extendable deployment strategies, the greedy solution is not the only method. In fact, one can argue that if more budget is expected later, we can try to optimize future utilities instead of the current one. The *backward greedy solution* is based on this idea. It starts with full deployment and gradually removes the locations with the least impact on utility until the budget is satisfied:

- 1) Initially,  $\mathcal{L}' = \mathcal{L}$ ;
- 2) While the cost of  $\mathcal{L}'$  is greater than  $C$ , do

$$\begin{aligned} l^{(j)} &= \arg \min_{l_i \in \mathcal{L}'} \frac{\tilde{U}(l_i|\mathcal{L}' \setminus \{l_i\})}{c_i}, \\ \mathcal{L}' &\leftarrow \mathcal{L}' \setminus \{l^{(j)}\}. \end{aligned}$$

It has a complexity of  $O((L-g)L)$ .

2) *Performance guarantee*: Although reducing the complexity to  $O(L^2)$ , the greedy and the backward greedy solutions result in a loss of optimality as the optimal deployment solution may not be nested. Under certain conditions, however, their performance can be guaranteed as follows.

*Proposition 3.2*: Under equal cost, if the conditional utility has bounded variation, *i.e.*,  $\exists \epsilon \in (0, 1)$  such that for all  $l \in \mathcal{L}$  and  $\mathcal{S}_1, \mathcal{S}_2 \subseteq \mathcal{L} \setminus l$  with  $|\mathcal{S}_1| = |\mathcal{S}_2|$ ,  $\tilde{U}(l|\mathcal{S}_1)/\tilde{U}(l|\mathcal{S}_2) \geq 1 - \epsilon$ , then the utilities of the greedy solution  $\mathcal{L}^g$ , the backward greedy solution  $\mathcal{L}^{bg}$ , and the optimal solution  $\mathcal{L}^o$  satisfy

$$\tilde{U}(\mathcal{L}^g) \geq (1 - \epsilon)\tilde{U}(\mathcal{L}^o), \quad (5)$$

$$\tilde{U}(\mathcal{L}) - \tilde{U}(\mathcal{L}^{bg}) \leq \frac{1}{1 - \epsilon} \left( \tilde{U}(\mathcal{L}) - \tilde{U}(\mathcal{L}^o) \right). \quad (6)$$

*Proof*: See [10]. ■

The proposition says that if the performance benefit of deploying a gateway at a location does not depend too much on the locations of the other gateways, then the greedy solution will achieve a utility no smaller than  $1 - \epsilon$  fraction and the backward greedy solution will achieve a utility gap (to full deployment) at most  $1/(1-\epsilon)$  times compared with the optimal solution. Similar conclusions can also be made under unequal costs; see [10]. Intuitively, the condition in the proposition can be induced by a minimum distance between candidate gateway locations, as it guarantees certain performance gain for each new deployment even when the other gateways are deployed nearby. Simple as they are, we will show that these algorithms can achieve near-optimal performance (see Section V).

#### IV. UTILITY COMPUTATION IN HETEROGENEOUS DOMAINS

In this section, we will investigate utility computation in more detail. Recalling that each domain may have different properties, we aim to capture their impact on the gateway placement solution by deriving analytical expressions for utility computation based on these properties, with the focus on the routing scheme and the position on the inter-domain path.

##### A. Modeling Contact Processes

For each domain-level hop under consideration, we model the contact processes as Poisson processes with mean inter-contact times  $1/\lambda_n$  for each pair of nodes and  $1/\lambda_l$  for each

node and the gateway<sup>3</sup> ( $\lambda_n, \lambda_l$  may vary at different hops). Poisson contact processes have been widely used to analyze the performance of DTN routing schemes (e.g., [8], [19] and followups), and such contact behaviors have also been reported in empirical studies of real DTN traces [20].

### B. Utility Analysis

We focus on the metrics of delay and number of replicas, i.e., the utility function (1) is of the form  $U(\mathbb{E}[D], \mathbb{E}[R])$ , where  $\mathbb{E}[D]$ ,  $\mathbb{E}[R]$  denote the expected end-to-end delay and total number of replicas, although the approach is also applicable to other metrics. Under a gateway placement  $\mathcal{L}'$ , the utility computation is decomposed into

$$U_{\text{global}}|\mathcal{L}' = U\left(\sum_j \mathbb{E}[D_j]|\lambda_l^{(j)}(\mathcal{L}'), \sum_j \mathbb{E}[R_j]|\lambda_l^{(j)}(\mathcal{L}'),\right), \quad (7)$$

where  $\mathbb{E}[D_j]$ ,  $\mathbb{E}[R_j]$  are the metrics at the  $j$ th (domain-level) hop. These per-hop metrics depend on gateway placement only through the node-gateway contact rate  $\lambda_l^{(j)}$  at the  $j$ th hop, which is in turn a function of  $\mathcal{L}'$ .

We now analyze  $\mathbb{E}[D_j]$  and  $\mathbb{E}[R_j]$  at each hop for a given  $\lambda_l^{(j)}$  ( $j$  will be omitted for simplicity). Based on how routing schemes control message replication, we divide them into forwarding schemes (P1, P4 in Table I), unlimited replication schemes (P2, P5), and limited replication schemes (P3, P6) to analyze separately. Throughout this section, we assume there are  $N$  nodes and one (aggregated) gateway in the domain under consideration. As in previous studies (e.g., [2]), we assume unlimited contact volume and buffer size to focus on the impact of replication strategies (see Section V-A1 for an evaluation of the influence of this assumption).

1) *Source-Gateway Hop, Forwarding Schemes:* Forwarding routing schemes only keep a single copy of the message in the network, i.e., the average number of replicas per message is  $\mathbb{E}[R_F] = 1$ . The delivery rate of the message is determined by the delivery rate of the node carrying that copy at each time instant. Under the homogeneous contact assumption, the delivery rate is fixed at  $\lambda_l$ , and thus the average delay is  $\mathbb{E}[D_F] = 1/\lambda_l$ .

2) *Source-Gateway Hop, Unlimited Replication Schemes:* Define the domain-wide replication rate at time  $t$  as  $\lambda_R(t)$  and the domain-wide delivery rate as  $\lambda_D(t)$ . These rates evolve over time as follows. Let  $T_j$  denote the time of the  $j$ th contact between any node carrying the message and a new node; define  $T_0 \triangleq 0$ , and  $T_N \triangleq \infty$ . Upon  $T_{j-1}$  ( $1 \leq j \leq N$ ), there are  $j$  nodes with the message and  $N - j$  nodes without the message. The delay until the next replication (i.e.,  $T_j - T_{j-1}$ ) is the minimum of the node-node contacts for all the  $j(N - j)$  node pairs, which forms an exponential random variable with rate parameter  $j(N - j)\lambda_n$ . Meanwhile, the message will be delivered if any of the  $j$  nodes meets the gateway, the delay of

which is again exponentially distributed with rate  $j\lambda_l$ . Therefore, we have that for  $t \in [T_{j-1}, T_j]$ ,  $\lambda_R(t) = j(N - j)\lambda_n$ , and  $\lambda_D(t) = j\lambda_l$ .

We now use these rates to characterize performance. Let random variable  $D_{\text{UR}}$  denote the delivery delay under unlimited replication and  $R_{\text{UR}}$  the number of replicas (including the source copy) generated until delivery, assuming all replicas are purged afterwards. We can decompose the expectation of  $D_{\text{UR}}$  and  $R_{\text{UR}}$  into:

$$\mathbb{E}[D_{\text{UR}}] = \sum_{j=0}^{N-1} \Pr\{T_j \leq D_{\text{UR}} < T_{j+1}\} \cdot \mathbb{E}[D_{\text{UR}}|T_j \leq D_{\text{UR}} < T_{j+1}], \quad (8)$$

$$\mathbb{E}[R_{\text{UR}}] = \sum_{j=0}^{N-1} \Pr\{T_j \leq D_{\text{UR}} < T_{j+1}\}(j+1). \quad (9)$$

Observing that  $D_{\text{UR}} - T_j$  conditioned on  $D_{\text{UR}} \in [T_j, T_{j+1})$  is exponentially distributed with the corresponding delivery rate  $(j+1)\lambda_l$ , we have the following results.

*Lemma 4.1:* For  $j = 0, \dots, N-1$ ,

$$\Pr\{T_j \leq D_{\text{UR}} < T_{j+1}\} = \frac{\lambda_l}{(N-j-1)\lambda_n + \lambda_l} \prod_{i=1}^j \frac{(N-i)\lambda_n}{(N-i)\lambda_n + \lambda_l}. \quad (10)$$

Moreover,  $\mathbb{E}[D_{\text{UR}}|T_j \leq D_{\text{UR}} < T_{j+1}]$  is lower bounded by

$$\left(\frac{1}{(j+1)\lambda_l} - \frac{1}{(j+1)(N-j-1)\lambda_n}\right)_+ + \sum_{i=1}^j \left(\frac{1}{i(N-i)\lambda_n} - \frac{1}{i\lambda_l}\right)_+, \quad (11)$$

where  $(x)_+ \triangleq \max(x, 0)$ , and upper bounded by

$$\frac{1}{(j+1)\lambda_l} + \frac{1}{\lambda_n} \sum_{i=1}^j \frac{1}{i(N-i)}. \quad (12)$$

*Proof:* See [10]. ■

Substituting these results into (8, 9) gives  $\mathbb{E}[R_{\text{UR}}]$  and a pair of upper and lower bounds on  $\mathbb{E}[D_{\text{UR}}]$ . In the case when  $\lambda_l$  and  $\lambda_n$  are comparable and  $N$  is large, we have the following closed-form approximations (see [10] for details):

$$\mathbb{E}[D_{\text{UR}}] \approx \frac{\log N}{N} \left(\frac{1}{\lambda_l} + \frac{1}{\lambda_n}\right), \quad (13)$$

$$\mathbb{E}[R_{\text{UR}}] \approx \frac{1+N}{2}. \quad (14)$$

These approximations are consistent with previous results in [19] for homogeneous networks (up to a constant-factor difference for delay). A byproduct of the above analysis is a scaling law that says as the node density increases (i.e.,  $N \rightarrow \infty$ ), the number of replicas for unlimited replication increases at  $\Theta(N)$  and the delay decreases at<sup>4</sup>  $\Theta(\log N/N)$ .

<sup>3</sup>Since gateways between each pair of neighboring domains are networked, we can group them into an aggregated gateway and denote the sum of their contact rates with nodes by  $\lambda_l$ .

<sup>4</sup>A different scaling law is given in [2] but for  $\lambda_l \gg \lambda_n$ .

### 3) Source-Gateway Hop, Limited Replication Schemes:

The analysis under limited replication schemes is similar to that under unlimited replication, except that the scheme will impose a maximum number of replications  $r \geq 0$ . Under this constraint, the replication rate  $\lambda_R(t)$  and the delivery rate  $\lambda_D(t)$  will evolve the same as before, but they will be fixed after  $T_r$ . We assume centralized spray in analysis; distributed spray schemes tend to have longer delays and fewer replicas, but the differences are small as observed in Section V-A1.

Following similar arguments as before (see [10]), the average delay  $\mathbb{E}[D_{\text{LR}}]$  and the average number of replicas  $\mathbb{E}[R_{\text{LR}}]$  can be approximated by:

$$\mathbb{E}[D_{\text{LR}}] \approx \frac{1}{\lambda_l N} \left( \log r + \frac{N-r}{r+1} \right) + \frac{1}{\lambda_n N^2} \left( (r-1) \log(N-1) + (N-r) \log \frac{r(N-1)}{N-r-1} + \sum_{j=1}^{r-1} \log \frac{j}{N-j-1} \right), \quad (15)$$

$$\mathbb{E}[R_{\text{LR}}] \approx \frac{(r+1)}{N} \left( N - \frac{r}{2} \right) \quad (16)$$

for comparable  $\lambda_l$  and  $\lambda_n$  and any  $0 < r < N-1$ . The result improves and generalizes the closed-form bound on delay in [21], which is only for homogeneous networks and loose for large<sup>5</sup>  $r$ . Note that (15) may not be monotonically decreasing in  $r$ , and monotonicity constraint can be imposed to produce a better approximation.

4) *Analysis of the Other Hops:* We now briefly present the results for the other domain-level hops. Assume that the gateways purge a message upon the first delivery.

For an intermediate domain, the ingress hop has an average delay of  $1/(N\lambda_l)$  and does not generate additional message replicas (we only count replicas at nodes). Then the receiving node will act as a new source, and its performance in delivering the message to the downstream gateways can be analyzed as in Section IV-B1 to IV-B3.

For the destination domain, the analysis for the ingress hop is the same as before, but that for the final delivery is slightly different. With probability  $1/N$ , the node receiving the message from gateways is the destination itself. Otherwise, the same analysis as before applies, with  $N$  replaced by  $N-1$  and  $\lambda_l$  by  $\lambda_n$ . We summarize the results below. For forwarding schemes,

$$D'_F = \frac{N-1}{N\lambda_n}, \quad R'_F = \frac{N-1}{N}. \quad (17)$$

For unlimited replication,

$$D'_{\text{UR}} \approx \frac{2 \log(N-1)}{N\lambda_n}, \quad R'_{\text{UR}} = \frac{N-1}{2}. \quad (18)$$

<sup>5</sup> [21] also gives a tighter approximation for a specific routing scheme (Binary Spray and Wait) but not in closed form.

For limited replication ( $0 < r < N-2$ ),

$$D'_{\text{LR}} \approx \frac{1}{N\lambda_n} \left( \log r + \frac{N-r-1}{r+1} + \frac{(N-2)}{N-1} \log(N-2) \right) + \frac{1}{N(N-1)\lambda_n} \left( (N-r-1) \log \frac{r}{N-r-2} + \sum_{j=1}^{r-1} \log \frac{j}{N-j-2} \right), \quad (19)$$

$$R'_{\text{LR}} = \frac{(r+1)}{N} \left( N-1 - \frac{r}{2} \right). \quad (20)$$

Note that these results do not depend on  $\lambda_l$  since the last hop is entirely inside the destination domain and thus independent of gateway placement.

### C. Comparison with Previous Analytical Studies

The original analysis of DTN routing schemes is only for homogeneous MANETs (*e.g.*, see [8], [19]). Recent work [2], [22] extends the analysis to hybrid networks with infrastructure nodes, but their results are not directly applicable to our problem as they consider a single domain where an end-to-end path may not involve an infrastructure node. To our knowledge, this is the first performance analysis for multi-domain DTNs.

Technically, all existing work uses Markovian models (*e.g.*, [19], [22]) or its variations (*e.g.*, Ordinary Differential Equations (ODEs) [2], [23]). While ODEs generally give more succinct mathematical representations, they sacrifice certain accuracy by taking fluid approximations. In this paper, we follow the exact Markovian approach as in [19], but only focus on the evolution of delivery rate and discard the other details to simplify the derivation. Moreover, we focus on closed-form solutions for easy computation and insights. Although derived for the asymptotic case, our results have been verified to have reasonable accuracy even for small numbers of nodes and gateways (see Section V-A1).

## V. PERFORMANCE EVALUATION

We now evaluate the performance of the proposed deployment strategies through simulations. Consider the utility function  $U_{\text{global}} = -\mathbb{E}[D]$  which corresponds to minimizing the end-to-end delay (see [10] for parallel results for  $U_{\text{global}} = -\mathbb{E}[R]$ ). For routing schemes, we simulate direct delivery, forwarding, limited replication, and unlimited replication.

### A. Evaluation on Synthetic Data

1) *Evaluation of Utility Computation:* We first simulate one-(domain-level)-hop communications to evaluate the accuracy of the analytical results in Section IV-B. New messages arrive according to a Poisson process of rate  $\lambda$  and are uniformly distributed among the nodes. For each routing scheme, we simulate two versions: a constrained version where we impose a buffer size  $B$ , a contact volume of one message per contact, and a Time-To-Live (TTL) constraint, and an unconstrained version where these constraints are removed; see [10] for a detailed specification.

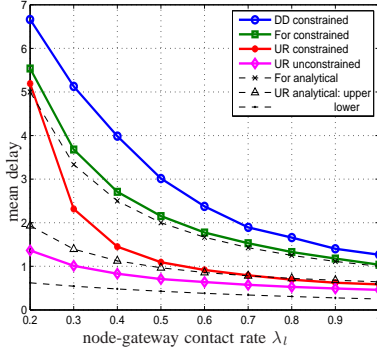


Fig. 2. Average delay of direct delivery, forwarding, and unlimited replication ( $N = 5$ ,  $\lambda_n = 1$ ,  $\lambda = 1$ ,  $TTL = 10$ ,  $B = 10$ ,  $M = 100$ , 100 Monte Carlo runs).

In Fig. 2, we plot the average delays for direct delivery (“DD”), forwarding (“For”), and unlimited replication (“UR”) schemes as functions of the node-gateway contact rate  $\lambda_l$  and compare them with the analytical results in Section IV-B1 and IV-B2 (unconstrained direct delivery and forwarding schemes have the same delay as the analytical result in Section IV-B1 and are thus omitted). The unconstrained schemes achieve smaller delays than their constrained counterparts by eliminating queueing delays. Moreover, our analytical results closely approximate the actual delays at relatively high contact rates. For limited replication<sup>6</sup> (“LR”), the performance depends on the parameter  $r$ , and thus we plot the average delay as a function of  $r$  in Fig. 3. As  $r$  increases, the scheme becomes more aggressive, and the delay decreases sharply at first but levels off after a few replications, indicating a saturation phenomenon. Again, our analytical approximation in (15) follows the actual value closely.

2) *Evaluation of Gateway Placement*: Next, we test the overall inter-domain communication performance under the proposed deployment strategies. We implement 6 deployment strategies, which are the combinations of 2 utility computation methods (simulations and analytical calculation based on estimated parameters) and 3 optimization algorithms (brute-force search, greedy, backward greedy). For comparison, we also simulate a bottom-line strategy that randomly selects deployment locations and an oracle deployment strategy that uses the same data for training and testing (in contrast, all the other strategies use one data set for computing utilities and placement solutions, and an independent data set for testing performance; see [10]). To simulate heterogeneous mobility for different domains, we use a localized random walk mobility model [24], where each domain is represented by a  $G \times G$  grid with a “home cell” that attracts nodes

<sup>6</sup>We have simulated binary spray for the constrained version of limited replication and centralized spray for the unconstrained version. Binary spray without resource constraints yields performance between the two (not shown).

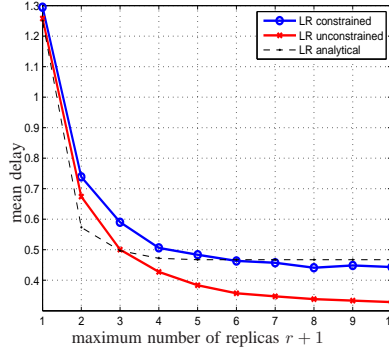


Fig. 3. Average delay of limited replication ( $N = 10$ ,  $\lambda_n = 1$ ,  $\lambda_l = 0.8$ ,  $\lambda = 2$ ,  $TTL = 10$ ,  $B = 10$ ,  $M = 100$ , 100 Monte Carlo runs).

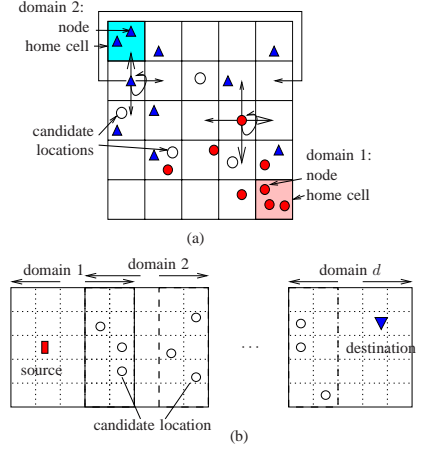


Fig. 4. Simulation scenario: (a) two overlapped domains; (b) multiple partially overlapped domains.

according to a tightness parameter<sup>7</sup>  $\tau$ . Assume contacts only occur within the same cell. We generate new messages with probability  $\lambda$  ( $0 < \lambda \leq 1$ ) per slot with random sources and destinations.

We simulate two scenarios: (i) fixed domains and increasing number of gateways, and (ii) fixed gateways and increasing number of domains. In scenario (i), there are two domains occupying the same area, but with different number of nodes ( $N_i$  for domain  $i$ ), home cells, and routing schemes (unlimited replication in domain 1 and direct delivery in domain 2), as illustrated in Fig. 4 (a). We randomly choose  $L$  cells as candidate gateway locations. Fig. 5 compares the calculated end-to-end delays with the actual simulated delays. We see that although the contact processes in this case are not Poisson, the calculated delays still approximate the trend of the actual delays well (up to a constant scaling, e.g., the calculated delay under direct delivery has been multiplied by 1.3818). Fig. 6 compares the performance under different strategies as more gateways are deployed. From the results, we see that the proposed strategies (greedy/backward greedy algorithm based on calculated delay) perform almost as well as the optimal strategy (brute-force search based on simulated delay). Moreover, the proposed strategies significantly outperform random placement and even closely approximate the oracle strategy. Similar observations have been made under other routing schemes (see [10]).

In scenario (ii), we extend the simulation to multiple (domain-level) hops by separating the source and the destination domains with pairwise partially overlapped intermediate domains, as illustrated in Fig. 4 (b). Each domain has a number of nodes randomly selected from  $[N_{\min}, N_{\max}]$  following a localized random walk with a randomly selected home cell and a tightness parameter uniformly distributed in  $[\tau_{\min}, \tau_{\max}]$ . Each pair of neighboring domains has  $L$  candidate gateway locations randomly selected from cells in the overlapped area. Given a total of  $g$  gateways, we need to place at least one

<sup>7</sup>For each cell  $i$ , the transition probability into cell  $i$  is proportional to  $e^{-\tau d_i}$ , where  $d_i$  is the taxicab distance between cell  $i$  and the home cell, and  $\tau > 0$  is a tightness parameter; see [24]. All boundaries are wrapping boundaries.



gateway per neighbor pair. Fig. 7 shows the results under different strategies as functions of the number of domains. Similar to scenario (i), our strategies yield similar performance to the brute-force strategy at a much lower complexity. Compared with random placement, our solution still shows clear improvement, and the gap grows with the number of hops as the freedom of placement increases. We observe a slightly larger gap with the oracle strategy than scenario (i), which is likely to be caused by more randomness in contact traces in the multi-domain simulations.

### B. Evaluation on Traces

We have also conducted a set of trace-driven simulations to test the robustness of our solutions. We use the contact traces from DieselNet as in<sup>8</sup> [2], which is a bus-based mobile network in an area with scattered access points (APs). To construct an inter-domain scenario, we use the contact traces in one day to represent the contacts within one domain and the APs to represent the candidate gateway locations (assume different domains to have the same coverage area). We preprocess the original traces by filtering out low-activity periods as well as inactive nodes and APs. The result is a subset of the traces during 7 am-7 pm of 8 days between 3 active nodes per day and 10 APs with sufficient contacts in all the selected days.

1) *Contact Distribution*: First, we examine the contact properties of the traces. Fig. 8 shows the empirical Complementary Cumulative Distribution Functions (CCDFs) for the inter-contact times between pairs of mobile nodes as well as between each mobile node and the set of APs. We then compare the empirical CCDFs with the CCDFs of the exponential distributions with parameters estimated from the traces. We see that the mobile-mobile contacts approximate the exponential distribution closely, whereas the mobile-AP contacts have a slightly heavier tail. The latter is due to the spatial adjacency of some of the APs, which causes their contacts to appear in bursts.

2) *Accuracy of Utility Calculation*: Next, we consider the traces from the same day as one domain and group the 8 days into 4 domain pairs. Aiming to evaluate the accuracy of the utility calculation in Section IV-B, we simulate direct delivery and unlimited replication under uniformly distributed traffic at 5 packets per hour per source node (with random destinations). We plot the simulated and the calculated delays in<sup>9</sup> Fig. 9, from which we see that although the absolute values of the calculated and the simulated delays are different, our calculation closely approximates the actual delay after scaling (*i.e.*, multiplication by a constant), with an error within 10% for direct delivery and 5% for unlimited replication. Such approximation is sufficient for our purpose because constant

<sup>8</sup>The traces, labeled “DieselNet Fall 2007”, can be downloaded from <http://traces.cs.umass.edu/index.php/Network/Network>. All contacts between the same nodes that are within 60 seconds of each other are merged as in [20].

<sup>9</sup>The values are averaged over 10 traffic realizations per trace pair, 4 pairs of traces, and all possible gateway deployments (*i.e.*, for 2 gateways, there are  $10 \times 4 \times \binom{10}{2}$  Monte Carlo runs).

scaling does not affect the optimization result of (4). Similar observation has been made for the number of replicas [10].

3) *Performance of Gateway Placement Strategies*: Finally, we conduct end-to-end simulations similar to those in scenario (i) of Section V-A2 to compare different deployment strategies. Except for the oracle strategy, we solve for deployments on a pair of days and test its performance on the other days<sup>10</sup>. As shown in Fig. 10, the overall comparison is similar to that in Fig. 6. In particular, for 70% of the time, the proposed algorithms based on calculated utilities even slightly outperform the brute-force search based on simulated utilities. This is because the strategies based on simulated utilities may overfit the training data from one pair of days and thus not work so well on testing data from different days. This result together with the accuracy result in Section V-B2 justifies the robustness of our utility calculation. As in Fig. 6, the proposed strategies again significantly outperform the random deployment, cutting down the delay by up to 30% (at 2 gateways), which confirms the benefit of strategically placing gateways, especially with very limited resources. Compared with the oracle strategy, the delays under the proposed strategies are within 10%. Similar results have been obtained under other routing schemes and for the utility measure of the number of replicas (see [10]).

## VI. CONCLUSION

We have developed a rigorous solution for placing gateways in heterogeneous, multi-domain DTNs via the approach of utility optimization. The proposed solution balances performance and complexity by employing decomposition and closed-form utility calculation at local domains together with efficient optimization algorithms network-wide. Although certain assumptions and simplifications have been made to obtain the solution, the results have been shown to be near-optimal and robust in both synthetic and real DTN scenarios.

## REFERENCES

- [1] “The MESSAGE Project.” <http://bioinf.ncl.ac.uk/message/>.
- [2] N. Banerjee, M. D. Corner, D. Towsley, and B. N. Levine, “Relays, Base Stations, and Meshes: Enhancing Mobile Networks with Infrastructure,” in *Proceedings of ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, (San Francisco, CA), September 2008.
- [3] W. Zhao, Y. Chen, M. Ammar, M. Corner, B. Levine, and E. Zegura, “Capacity enhancement using throwboxes in dtns,” in *IEEE International Conference on Mobile Adhoc and Sensor Systems (MASS)*, (Vancouver, Canada), pp. 31–40, October 2006.
- [4] S.-H. Lee, S. H. Wong, C.-K. Chau, S. Varadarajan, K.-W. Lee, J. Crowcroft, and M. Gerla, “Self-organizing Inter-Domain Routing for Heterogeneous MANETs,” tech. rep., UCLA, 2009. <http://www.cs.ucla.edu/shlee/papers/idrm.pdf>.
- [5] A. Balasubramanian, B. N. Levine, and A. Venkataramani, “DTN Routing as a Resource Allocation Problem,” in *Proceedings of ACM SIGCOMM*, (Kyoto, Japan), August 2007.
- [6] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine, “MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks,” in *Proceedings of IEEE INFOCOM*, April 2006.

<sup>10</sup>Each strategy is tested through 120 Monte Carlo runs, consisting of 4 deployment solutions, each tested on 3 pairs of contact traces for 10 traffic realizations per pair.



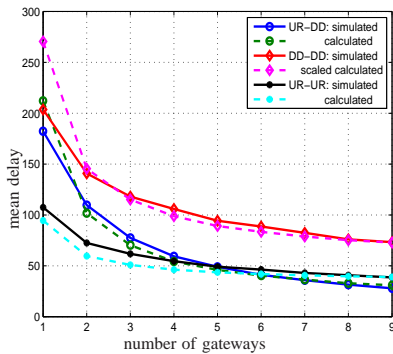


Fig. 5. Two-domain simulation: calculated vs. simulated delays ( $N_1 = 5$ ,  $N_2 = 10$ ,  $G = 10$ ,  $L = 10$ ,  $\lambda = 0.05$ ,  $\tau = 1$ , 1000 slots, 50 Monte Carlo runs).

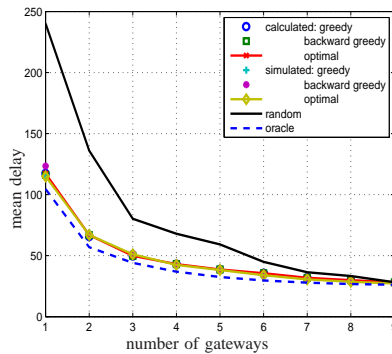


Fig. 6. Two-domain simulation: comparison of deployment strategies ( $N_1 = 5$ ,  $N_2 = 10$ ,  $G = 10$ ,  $L = 10$ ,  $\lambda = 0.05$ ,  $\tau = 1$ , 1000 slots, 50 Monte Carlo runs).

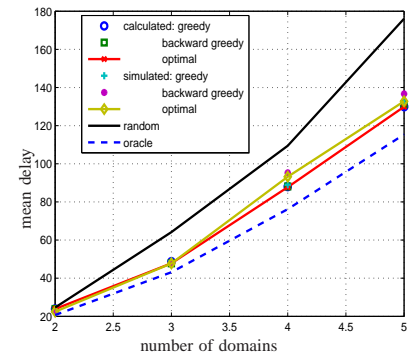


Fig. 7. Multi-domain simulation: comparison of strategies ( $[N_{\min}, N_{\max}] = [2, 5]$ ,  $[\tau_{\min}, \tau_{\max}] = [0, 0.5]$ ,  $G = 5$ , width of overlap = 2,  $L = 8$ ,  $g = 4$ ,  $\lambda = 1$ , 1000 slots, 20 Monte Carlo runs).

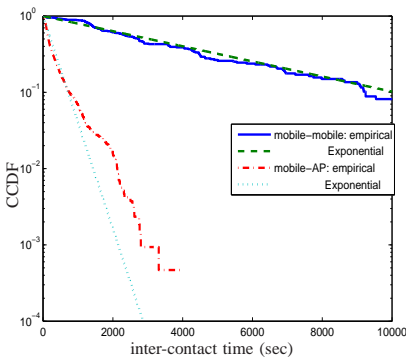


Fig. 8. CCDF of the mobile-mobile and the mobile-AP inter-contact times as compared with the exponential distribution.




Fig. 9. Simulated vs. calculated delays. We have scaled the calculated delays by 0.7171 for direct delivery and 0.4859 for unlimited replication to fit the simulated values.

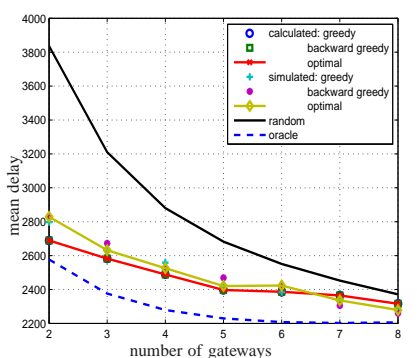


Fig. 10. Performance (measured by delay) of various deployment strategies under unlimited replication in both domains.

- [7] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Spray and Wait: An Efficient Routing Scheme for Intermittently Connected Mobile Networks," in *Proceedings of ACM WDTN*, pp. 252–259, August 2005.
- [8] T. Small and Z. Haas, "Resource and Performance Tradeoffs in Delay-Tolerant Wireless Networks," in *Proceedings of ACM WDTN*, pp. 260–267, August 2005.
- [9] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Spray and Focus: Efficient Mobility-Assisted Routing for Heterogeneous and Correlated Mobility," in *Proceedings of the Fifth Annual IEEE International Conference on Pervasive Computing and Communications Workshops (PerComW'07)*, 2007.
- [10] T. He, N. Sofra, K.-W. Lee, and K. K. Leung, "Supporting Materials for 'Utility-based Gateway Deployment in Multi-domain DTNs,'" Tech. Rep. RC24826, IBM, July 2009. [http://domino.research.ibm.com/comm/research\\_people.nsf/pages/ting\\_pubs.html/\\$FILE/rc24826.pdf](http://domino.research.ibm.com/comm/research_people.nsf/pages/ting_pubs.html/$FILE/rc24826.pdf).
- [11] J. Leguay, T. Friedman, and V. Conan, "DTN Routing in a Mobility Pattern Space," in *Proceedings of ACM Chants Workshop*, pp. 276–283, August 2005.
- [12] I. Cardei, C. Liu, J. Wu, and Q. Yuan, "DTN Routing with Probabilistic Trajectory Prediction," in *Proceedings of the International Conference on Wireless Algorithms, Systems, and Applications (WASA)*, 2008.
- [13] A. Lindgren, A. Doria, and O. Schelen, "Probabilistic Routing in Intermittently Connected Networks," in *Proceedings of SAPIR Workshop*, pp. 239–254, 2004.
- [14] B. Burns, O. Brock, and B. N. Levine, "MV Routing and Capacity Building in Disruption Tolerant Networks," in *Proceedings of IEEE INFOCOM*, pp. 398–408, March 2005.
- [15] E. Jones, L. Li, and P. Ward, "Practical Routing in Delay-Tolerant Networks," in *Proceedings of ACM Chants Workshop*, pp. 237–243, August 2005.
- [16] S. Jain, K. Fall, and R. Patra, "Routing in a Delay-Tolerant Network," in *Proceedings of ACM SIGCOMM*, pp. 145–158, 2004.
- [17] G. Sandulescu and S. Nadjm-Tehrani, "Opportunistic DTN Routing with Window-Aware Adaptive Replication," in *Proceedings of ACM AINTEC*, (Bangkok, Thailand), November 2008.
- [18] H. Kellerer, U. Pferschy, and D. Pisinger, *Knapsack Problems*. Springer, 2004.
- [19] R. Groenevelt, P. Nain, and G. Koole, "The Message Delay in Mobile Ad Hoc Networks," *Performance Evaluation*, vol. 62, pp. 210–228, Oct. 2005.
- [20] 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. 13th Annual ACM International Conference on Mobile Computing and Networking*, (Montreal, Canada), pp. 195–206, 2007.
- [21] T. Spyropoulos, K. Psounis, and C. Raghavendra, "Efficient Routing in Intermittently Connected Mobile Networks: The Multiple-copy Case," *ACM/IEEE Transactions on Networking*, vol. 16, pp. 77–90, Feb. 2008.
- [22] M. Ibrahim, A. Hanbali, and P. Nain, "Delay and Resource Analysis in MANETs in Presence of Throwboxes," in *Proceedings of International Symposium on Computer Performance, Modeling, Measurements, and Evaluation*, 2007.
- [23] X. Zhang, G. Neglia, J. Kurose, and D. Towsley, "Performance Modeling of Epidemic Routing," in *Proceedings of IFIP Networking*, 2006.
- [24] B. Walker, T. Clancy, and J. Glenn, "Using localized random walks to model delay-tolerant networks," in *Proc. of IEEE Military Communications Conference (MILCOM)*, (San Diego, CA), pp. 1–7, November 2008.