

IBM Research Report

Supporting Materials for "Enabling Inter-domain DTN Communications by Networked Static Gateways"

Ting He, Kang-Won Lee
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598
USA

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



Supporting Materials for “Enabling Inter-domain DTN Communications by Networked Static Gateways”

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

I. INTRODUCTION

In this report, we provide supporting materials for [1], consisting of a complete survey of related DTN routing schemes, proofs of theoretical results, and extended simulations. We refer to [1] for a complete problem formulation and the proposed solutions.

II. TAXONOMY OF DTN ROUTING SCHEMES

A critical difference between inter-domain and intra-domain designs is the heterogeneity of routing scheme. There have been a number of routing schemes for DTN stemming from a wide range of diverse applications. In this section, we categorize the existing DTN routing schemes so as to extract the key parameters for later analysis.

We base our taxonomy on the previous work [2], where routing schemes are categorized according to their resource requirements (storage and contact volume) and relay operation (replication or forwarding). We notice that besides the above criteria, how many copies of a message can exist simultaneously in a domain is also a key factor for performance. Therefore, we further divide schemes under replication into those which constrain the total number of replicas and those which do not. We now apply our taxonomy to related routing schemes in the literature, as shown in Table I¹.

Resource: Earlier work [6] assumes perfect data link layer and unlimited storage at each relay for simplicity. In reality, such an assumption relies on several factors including node storage capacity, message size, mobility, radio range, and channel capacity. For contact volume to be unlimited, the applications need to have slowly-moving nodes ([5]), short messages ([11]), or high communication capacity ([9]). For

Research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. 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 U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence 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.

¹For clarity, we have combined schemes under either unlimited storage or unlimited contact volume into the “unlimited resource” category and refer to [2] for further distinction between the two.

TABLE I
CLASSIFICATION OF DTN ROUTING SCHEMES

	Resource	Number of replicas	Previous work
P1	Unlimited	Forwarding	[3]–[5]
P2	Unlimited	Unlimited replication	[6]–[9]
P3	Unlimited	Limited replication	[10]–[12]
P4	Finite	Forwarding	[13], [14]
P5	Finite	Unlimited replication	[2], [15]
P6	Finite	Limited replication	[16]

storage to be unlimited, the messages need to be short and sporadic ([11]) or the nodes need to have plenty of storage ([5]). In general, applications will impose certain constraints on both contact volume and storage, and the trend of DTN routing is to handle these constraints explicitly (P4–6).

Number of replicas: Judging by the number of replicas coexisting simultaneously, routing schemes can be classified into single-copy routing (forwarding) and multi-copy routing (replication) [4]. Moreover, among replication schemes, some explicitly control the total number of replicas in the network, such as the Spray-and-Wait scheme [10] and its variations [11], [12], [16], and others replicate as many as possible with the available resources. The number of replicas determines the propagation rate of messages inside the domain and thus affects the decision on gateway deployment.

Besides the above criteria, there are other dimensions to categorize DTN routing schemes. For example, based on different performance focuses, existing schemes can be partitioned into heuristic schemes and optimization-based schemes. Heuristic schemes use heuristics related to the likelihood of delivery or the delay of delivery that can be computed at local nodes based on their own encounter history, e.g., [4], [7], [8]. Optimization-based schemes, on the other hand, are explicitly designed to optimize given routing performance metrics such as average delay or delivery ratio [2], [14]. To our problem of gateway deployment, however, the detailed routing operations do not matter as long as we correctly abstract their influence on the performance benefit of each potential gateway deployment, which will be done through analysis of the number of contemporary replicas and their delivery rates

as discussed above. In addition, all the schemes above belong to ad hoc routing schemes where nodes are peers and operate only at the network layer. In some applications, it may be possible to have specialized nodes as message ferries [9], [17], or employ network coding to boost performance in certain cases [18]. Such schemes will require special devices and are thus outside the scope of the paper.

III. PROOF OF THEORETICAL RESULTS

In [1], we have presented several theoretical results but omitted the proofs due to space limit. We hereby provide the missing proofs (all notations are defined as in [1]).

A. Performance Guarantee for Greedy and Backward Greedy Algorithms

In Section III.B in [1], we have established sufficient conditions for the proposed greedy or backward greedy algorithms to be ϵ -close to the optimal solution. The conditions under equal cost are as follows.

Proposition 3.1: 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 $|\mathcal{S}_1| = |\mathcal{S}_2|$,

$$\frac{\tilde{U}(l|\mathcal{S}_1)}{\tilde{U}(l|\mathcal{S}_2)} \geq 1 - \epsilon, \quad (1)$$

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

$$\tilde{U}(\mathcal{L}^g) \geq (1 - \epsilon)\tilde{U}(\mathcal{L}^o), \quad (2)$$

$$\tilde{U}(\mathcal{L}) - \tilde{U}(\mathcal{L}^{bg}) \leq \frac{1}{1 - \epsilon} \left(\tilde{U}(\mathcal{L}) - \tilde{U}(\mathcal{L}^o) \right). \quad (3)$$

Proof: Under equal cost, all solutions have equal size $g \triangleq \lfloor C/c_i \rfloor$. For the greedy solution, we can decompose the total utility as

$$\tilde{U}(\mathcal{L}^i) = \tilde{U}(l_1^i) + \tilde{U}(l_2^i|l_1^i) + \dots + \tilde{U}(l_g^i|l_1^i, \dots, l_{g-1}^i) \quad (4)$$

for $i = o, g$, assuming that l_j^g is the location selected by the greedy solution in the j th iteration. Then by the greedy algorithm, we have that for $j = 1, \dots, g$,

$$\tilde{U}(l_j^g|l_1^g, \dots, l_{j-1}^g) \geq \tilde{U}(l_j^o|l_1^o, \dots, l_{j-1}^o),$$

and by (1), the right-hand side is in turn bounded as

$$\tilde{U}(l_j^g|l_1^g, \dots, l_{j-1}^g) \geq (1 - \epsilon)\tilde{U}(l_j^o|l_1^o, \dots, l_{j-1}^o).$$

Combining these inequalities and plugging the result into (4) yields that $\tilde{U}(\mathcal{L}^g) \geq (1 - \epsilon)\tilde{U}(\mathcal{L}^o)$.

For the backward greedy solution, we can similarly decompose the utility gap as

$$\begin{aligned} \tilde{U}(\mathcal{L}) - \tilde{U}(\mathcal{L}^i) = & \tilde{U}(l_1^i|\mathcal{L} \setminus \{l_1^i\}) + \tilde{U}(l_2^i|\mathcal{L} \setminus \{l_1^i, l_2^i\}) + \dots \\ & + \tilde{U}(l_{L-g}^i|\mathcal{L} \setminus \{l_1^i, \dots, l_{L-g}^i\}), \end{aligned} \quad (5)$$

where $i = o, bg$, and l_j^{bg} is the location removed by the backward greedy solution in the j th iteration. By the backward greedy solution, we have that

$$\tilde{U}(l_j^{bg}|\mathcal{L} \setminus \{l_1^{bg}, \dots, l_j^{bg}\}) \leq \tilde{U}(l_j^o|\mathcal{L} \setminus \{l_1^{bg}, \dots, l_{j-1}^{bg}, l_j^o\}),$$

and by (1),

$$\tilde{U}(l_j^o|\mathcal{L} \setminus \{l_1^{bg}, \dots, l_{j-1}^{bg}, l_j^o\}) \leq \frac{1}{1 - \epsilon} \tilde{U}(l_j^o|\mathcal{L} \setminus \{l_1^o, \dots, l_j^o\}).$$

Therefore, $\tilde{U}(\mathcal{L}) - \tilde{U}(\mathcal{L}^{bg}) \leq \left(\tilde{U}(\mathcal{L}) - \tilde{U}(\mathcal{L}^o) \right) / (1 - \epsilon)$. \blacksquare

For unequal costs, ϵ -approximation can be guaranteed if the conditional efficiencies have bounded variation, as stated below.

Proposition 3.2: For the greedy solution, if $\exists \epsilon_1 \in [c_{\max}/C, 1)$ for $c_{\max} \triangleq \max_j c_j$ such that

$$\frac{\tilde{U}(l_1|\mathcal{S}_1)}{c_1} \geq \frac{(1 - \epsilon_1)C}{(C - c_{\max})} \cdot \frac{\tilde{U}(l_2|\mathcal{S}_2)}{c_2} \quad (6)$$

for all $l_i \in \mathcal{L}$ and $\mathcal{S}_i \subseteq \mathcal{L} \setminus l_i$ ($i = 1, 2$), then

$$\tilde{U}(\mathcal{L}^g) \geq (1 - \epsilon_1)\tilde{U}(\mathcal{L}^o). \quad (7)$$

For the backward greedy solution, if $\exists \epsilon_2 \in [c_{\max}/(C_0 - C + c_{\max}), 1)$ for $C_0 \triangleq \sum_{j=1}^L c_j$ such that

$$\frac{\tilde{U}(l_1|\mathcal{S}_1)}{c_1} \leq \frac{(C_0 - C)}{(1 - \epsilon_2)(C_0 - C + c_{\max})} \cdot \frac{\tilde{U}(l_2|\mathcal{S}_2)}{c_2}, \quad (8)$$

then

$$\tilde{U}(\mathcal{L}) - \tilde{U}(\mathcal{L}^{bg}) \leq \frac{1}{1 - \epsilon_2} \left(\tilde{U}(\mathcal{L}) - \tilde{U}(\mathcal{L}^o) \right). \quad (9)$$

Proof: For the greedy solution, we can decompose the total utility as

$$\tilde{U}(\mathcal{L}^i) = c_1^i \frac{\tilde{U}(l_1^i)}{c_1^i} + c_2^i \frac{\tilde{U}(l_2^i|l_1^i)}{c_2^i} + \dots + c_{L_i}^i \frac{\tilde{U}(l_{L_i}^i|l_1^i, \dots, l_{L_i-1}^i)}{c_{L_i}^i},$$

where $i = o, g$ and $L_i \triangleq |\mathcal{L}^i|$. Let $x_0 = \max \tilde{U}(l|\mathcal{S})/c$ denote the maximum conditional efficiency over all locations and condition sets. Each conditional efficiency of the greedy solution is lower bounded by $x_0(1 - \epsilon_1)C/(C - c_{\max})$ by (6), and the sum cost of selected locations is no smaller than $C - c_{\max}$. Thus, the total utility of the greedy solution satisfies

$$\tilde{U}(\mathcal{L}^g) \geq (1 - \epsilon_1)C x_0.$$

On the other hand, it is easy to see that the optimal solution satisfies $\tilde{U}(\mathcal{L}^o) \leq C x_0$. Combining both results yields (7).

For the backward greedy solution, we rewrite the utility gap as

$$\begin{aligned} \tilde{U}(\mathcal{L}) - \tilde{U}(\mathcal{L}^i) = & c_1^i \frac{\tilde{U}(l_1^i|\mathcal{L} \setminus \{l_1^i\})}{c_1^i} + c_2^i \frac{\tilde{U}(l_2^i|\mathcal{L} \setminus \{l_1^i, l_2^i\})}{c_2^i} \\ & + \dots + c_{L-L_i}^i \frac{\tilde{U}(l_{L-L_i}^i|\mathcal{L} \setminus \{l_1^i, \dots, l_{L-L_i}^i\})}{c_{L-L_i}^i} \end{aligned} \quad (10)$$

for $i = o, bg$. Let $y_0 = \min \tilde{U}(l|\mathcal{S})/c$ denote the minimum conditional efficiency. For the backward greedy solution, all conditional efficiencies are no more than $(C_0 - C)y_0/(1 -$

²Without loss of generality, we assume $c_{\max} < C$.

$\epsilon_2)(C_0 - C + c_{\max})]$ by (8), and the total cost of removed locations is bounded by $C_0 - C + c_{\max}$, which implies

$$\tilde{U}(\mathcal{L}) - \tilde{U}(\mathcal{L}^{bg}) \leq (C_0 - C)y_0/(1 - \epsilon_2).$$

Meanwhile, note that $\tilde{U}(\mathcal{L}) - \tilde{U}(\mathcal{L}^o) \geq (C_0 - C)y_0$. Combining both yields (9). ■

Intuitively, 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 more than ϵ fraction smaller than the optimal solution, and the backward greedy solution will achieve a utility gap with the full deployment that is at most $1/(1 - \epsilon)$ times the gap for the optimal solution. Similar conclusions can also be made under unequal costs; see Appendix for details. In practice, the above conditions can often be induced by a minimum distance between candidate gateway locations.

B. Analytical Utility Calculation

In Section IV.B, we have provided explicit or even closed-form formulas to analytically calculate the utilities per domain (measured by mean delay or mean number of replicas per message) for several classes of routing schemes. We now present the omitted steps in our derivation.

1) *Source-Gateway Hop, Unlimited Replication*: Under unlimited replication, we have decomposed the expected delay $\mathbb{E}[D_{\text{UR}}]$ and the expected number of replicas $\mathbb{E}[R_{\text{UR}}]$ into:

$$\mathbb{E}[D_{\text{UR}}] = \sum_{j=0}^{N-1} \Pr\{T_j \leq D_{\text{UR}} < T_{j+1}\} \cdot \mathbb{E}[D_{\text{UR}} | T_j \leq D_{\text{UR}} < T_{j+1}], \quad (11)$$

$$\mathbb{E}[R_{\text{UR}}] = \sum_{j=0}^{N-1} \Pr\{T_j \leq D_{\text{UR}} < T_{j+1}\} (j+1), \quad (12)$$

and used the following lemmas to characterize their values.

Lemma 3.3: For $j = 0, \dots, N-1$,

$$\Pr\{T_j \leq D_{\text{UR}} < T_{j+1}\} = \frac{\lambda_l}{(N-j-1)\lambda_n + \lambda_l} \prod_{i=1}^j \frac{(N-i)\lambda_n}{(N-i)\lambda_n + \lambda_l}. \quad (13)$$

Proof: We use the property of exponential random variables that if X, Y are independent exponential random variables with parameters λ_X, λ_Y , then

$$\Pr\{X < Y\} = \frac{\lambda_X}{\lambda_X + \lambda_Y}. \quad (14)$$

We first show by induction that

$$\Pr\{D_{\text{UR}} \geq T_j\} = \prod_{i=1}^j \frac{(N-i)\lambda_n}{(N-i)\lambda_n + \lambda_l}, \quad 0 \leq j \leq N-1. \quad (15)$$

It trivially holds for $j = 0$. For $j = 1$, since T_1 has rate $(N-1)\lambda_n$ and D_{UR} has rate λ_l before T_1 , (14) implies

$$\Pr\{D \geq T_1\} = \frac{(N-1)\lambda_n}{(N-1)\lambda_n + \lambda_l}.$$

Suppose the formula holds up to $j-1$ ($2 \leq j \leq N-1$). We have

$$\Pr\{D_{\text{UR}} \geq T_j\} = \Pr\{D_{\text{UR}} \geq T_{j-1}\} \cdot \Pr\{T_j - T_{j-1} \leq D_{\text{UR}} - T_{j-1} | D_{\text{UR}} \geq T_{j-1}\}.$$

Since conditioned on $D_{\text{UR}} \geq T_{j-1}$, the residual delay $D_{\text{UR}} - T_{j-1} \sim \text{Exp}(j\lambda_l)$ (for $D_{\text{UR}} < T_j$), and $T_j - T_{j-1} \sim \text{Exp}(j(N-j)\lambda_n)$, we can apply (14) to obtain

$$\Pr\{D_{\text{UR}} \geq T_j\} = \Pr\{D_{\text{UR}} \geq T_{j-1}\} \cdot \frac{(N-j)\lambda_n}{(N-j)\lambda_n + \lambda_l},$$

which proves (15).

Then we note that

$$\Pr\{T_j \leq D_{\text{UR}} < T_{j+1}\} = \Pr\{T_j \leq D_{\text{UR}}\} \cdot \Pr\{D_{\text{UR}} - T_j < T_{j+1} - T_j | D_{\text{UR}} \geq T_j\},$$

and by similar arguments as before,

$$\Pr\{D_{\text{UR}} - T_j < T_{j+1} - T_j | D_{\text{UR}} \geq T_j\} = \frac{(j+1)\lambda_l}{(j+1)\lambda_l + (j+1)(N-j-1)\lambda_n}.$$

Therefore,

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

Lemma 3.4: For $j = 0, \dots, N-1$, $\mathbb{E}[D_{\text{UR}} | T_j \leq D_{\text{UR}} < T_{j+1}]$ is lower bounded by³

$$\left(\frac{1}{(j+1)\lambda_l} - \frac{1}{(j+1)(N-j-1)\lambda_n} \right)_+ + \sum_{i=1}^j \left(\frac{1}{i(N-i)\lambda_n} - \frac{1}{i\lambda_l} \right)_+ \quad (16)$$

and upper bounded by

$$\frac{1}{(j+1)\lambda_l} + \frac{1}{\lambda_n} \sum_{i=1}^j \frac{1}{i(N-i)} \quad (17)$$

Proof: We use the property of exponential random variables that if X, Y are independent exponential random variables with parameters λ_X, λ_Y , then

$$\left(\frac{1}{\lambda_X} - \frac{1}{\lambda_Y} \right)_+ \leq \mathbb{E}[X | X < Y] \leq \frac{1}{\lambda_X}. \quad (18)$$

³Define $(x)_+$ as $\max(x, 0)$.

We first show by induction that for $j = 0, \dots, N - 1$,

$$\mathbb{E}[T_j | T_j \leq D_{\text{UR}}] \in \left[\sum_{i=1}^j \left(\frac{1}{i(N-i)\lambda_n} - \frac{1}{i\lambda_l} \right)_+, \frac{1}{\lambda_n} \sum_{i=1}^j \frac{1}{i(N-i)} \right]. \quad (19)$$

It trivially holds for $j = 0$. For $j = 1$, since $T_1 \sim \text{Exp}((N-1)\lambda_n)$ and $D_{\text{UR}} \sim \text{Exp}(\lambda_l)$ before T_1 , applying (18) gives the bounds. For $j > 1$, note that

$$\begin{aligned} \mathbb{E}[T_j | T_j \leq D_{\text{UR}}] &= \mathbb{E}[T_{j-1} | T_j \leq D_{\text{UR}}] \\ &\quad + \mathbb{E}[T_j - T_{j-1} | T_j - T_{j-1} \leq D_{\text{UR}} - T_{j-1}], \end{aligned}$$

where the first term on the right-hand side is equal to $\mathbb{E}[T_{j-1} | T_{j-1} \leq D_{\text{UR}}]$, and the second term is bounded in the interval

$$\left[\left(\frac{1}{j(N-j)\lambda_n} - \frac{1}{j\lambda_l} \right)_+, \frac{1}{j(N-j)\lambda_n} \right]$$

according to (18). This proves (19).

We then note that

$$\begin{aligned} \mathbb{E}[D_{\text{UR}} | T_j \leq D_{\text{UR}} < T_{j+1}] &= \mathbb{E}[T_j | T_j \leq D_{\text{UR}} < T_{j+1}] \\ &\quad + \mathbb{E}[D_{\text{UR}} - T_j | 0 \leq D_{\text{UR}} - T_j < T_{j+1} - T_j], \end{aligned}$$

where the first term is just $\mathbb{E}[T_j | T_j \leq D_{\text{UR}}]$ given by (19), and the second term can be bounded in the interval

$$\left[\left(\frac{1}{(j+1)\lambda_l} - \frac{1}{(j+1)(N-j-1)\lambda_n} \right)_+, \frac{1}{(j+1)\lambda_l} \right]$$

by applying (18). Plugging in the above results proves the lemma. \blacksquare

In the case when λ_l and λ_n are comparable and N is large, we can reduce the results into the following closed-form characterization:

$$\mathbb{E}[D_{\text{UR}}] \approx \frac{\log N}{N} \left(\frac{1}{\lambda_l} + \frac{1}{\lambda_n} \right), \quad (20)$$

$$\mathbb{E}[R_{\text{UR}}] \approx \frac{1+N}{2}, \quad (21)$$

where we have used the approximation that $\Pr\{T_j \leq D_{\text{UR}} < T_{j+1}\} \approx 1/N$ and applied the Harmonic number approximation to the upper bound in Lemma⁴ 3.4.

⁴The upper bound can be simplified by rewriting $\sum_{i=1}^j 1/(i(N-i)) = \left(\sum_{i=1}^j 1/i + \sum_{i=N-j}^{N-1} 1/i \right) / N$ and plugging in the approximation of the Harmonic number $\sum_{i=1}^j 1/i \approx \log j$.

2) *Source-Gateway Hop, Limited Replication*: The analysis under limited replication schemes is similar to that under unlimited replication. Specifically, the average delay $\mathbb{E}[D_{\text{LR}}]$ and the average number of replicas $\mathbb{E}[R_{\text{LR}}]$ can be decomposed into:

$$\begin{aligned} \mathbb{E}[D_{\text{LR}}] &= \sum_{j=0}^{r-1} \Pr\{T_j \leq D_{\text{LR}} < T_{j+1}\} \mathbb{E}[D_{\text{LR}} | T_j \leq D_{\text{LR}} < T_{j+1}] \\ &\quad + \Pr\{D_{\text{LR}} \geq T_r\} \mathbb{E}[D_{\text{LR}} | D_{\text{LR}} \geq T_r], \end{aligned} \quad (22)$$

$$\begin{aligned} \mathbb{E}[R_{\text{LR}}] &= \sum_{j=0}^{r-1} \Pr\{T_j \leq D_{\text{LR}} < T_{j+1}\} (j+1) \\ &\quad + \Pr\{D_{\text{LR}} \geq T_r\} (r+1). \end{aligned} \quad (23)$$

A key observation here is that before the last replication at time T_r , the constraint on replication has no effect, and thus all the analysis of D_{UR} still holds for D_{LR} conditioned on $D_{\text{LR}} < T_r$. Therefore, (23) can be evaluated by Lemma 3.3 and

$$\Pr\{D_{\text{LR}} \geq T_r\} = \prod_{i=1}^r \frac{(N-i)\lambda_n}{(N-i)\lambda_n + \lambda_l}$$

derived from Lemma 3.3. For (22), the first summation on the right-hand side can be evaluated in the same way as in (11) by Lemma 3.3 and 3.4. The second term can be evaluated by rewriting

$$\mathbb{E}[D_{\text{LR}} | D_{\text{LR}} \geq T_r] = \mathbb{E}[T_r | D_{\text{LR}} \geq T_r] + \mathbb{E}[D_{\text{LR}} - T_r | D_{\text{LR}} \geq T_r],$$

where $\mathbb{E}[T_r | D_{\text{LR}} \geq T_r]$ is bounded from both above and below as in (19), and $\mathbb{E}[D_{\text{LR}} - T_r | D_{\text{LR}} \geq T_r] = 1/[(r+1)\lambda_l]$ (because given delivery has not occurred at T_r , the remaining time till delivery is exponentially distributed with rate $(r+1)\lambda_l$).

For comparable λ_l and λ_n , by similar simplification as for unlimited replication, it can be shown that for $0 < r < N - 1$ (nontrivial replication constraint),

$$\begin{aligned} \mathbb{E}[D_{\text{LR}}] &\approx \frac{1}{\lambda_l N} \left(\log r + \frac{N-r}{r+1} \right) + \frac{1}{\lambda_n N^2} \left((r-1) \log(N-1) \right. \\ &\quad \left. + (N-r) \log \frac{r(N-1)}{N-r-1} + \sum_{j=1}^{r-1} \log \frac{j}{N-j-1} \right), \end{aligned} \quad (24)$$

$$\mathbb{E}[R_{\text{LR}}] \approx \frac{(r+1)}{N} (N - \frac{r}{2}). \quad (25)$$

We point out that (24) may not be monotone decreasing in r , and additional monotonicity constraint can be imposed to produce a better approximation.

IV. EXTENDED SIMULATION RESULTS

Besides the selected results presented in Section V of [1], we have also conducted extended simulations under similar settings. We now present the results.

A. Evaluation on Synthetic Data

1) *Evaluation of Utility Computation*: For a network of N nodes, we generate contact traces according to *i.i.d.* exponential inter-contact times, where the contact rates are λ_n per node pair and λ_l per node-gateway pair (assume a single

gateway). New messages arrive according to a Poisson process of rate λ and are uniformly distributed among the nodes. We simulate two versions for each scheme: constrained and unconstrained versions. For the constrained version, assume the buffer size is B messages, the contact volume is one message (in either direction), and messages not delivered after TTL (Time-To-Live) time are dropped (their delays are considered as TTL). For the unconstrained version, we ignore all the above constraints. The unconstrained version is used to verify the analytical results, whereas the constrained version tests how well the results approximate the actual utilities.

In Fig. 1, we plot the average delays for direct delivery, forwarding, and unlimited replication schemes with and without resource constraints as functions of the node-gateway contact rate λ_l and compare them with the analytical results. The unconstrained schemes achieve smaller delays on the average than their constrained counterparts by eliminating queuing delays, even though their maximum allowable delays are unbounded. Moreover, although our analytical results underestimate the delays under resource constraints, the values are reasonably close, especially at relatively high contact rates.




Fig. 1. Average delay vs. contact rate ($N = 5$, $\lambda_n = 1$, $\lambda = 1$, $TTL = 10$, $B = 10$, $M = 100$, 100 Monte Carlo runs): constrained versions of direct delivery, forwarding, and unlimited replication; unconstrained version of unlimited replication; analytical results.

For limited replication⁵, the performance can vary between those of forwarding and unlimited replication depending on how many contemporary replicas are allowed. Therefore, we plot the results as functions of the maximum number of replicas per message; see Fig. 2–3. As the scheme becomes more aggressive by increasing r , the average delay decreases sharply while the average number of generated replicas increases, indicating a tradeoff between performance and resource consumption. In particular, a single replication can already reduce the delay by more than 40%. The unconstrained version again outperforms the constrained counterpart in delay at the cost of generating more replicas per message. As we compare the simulation results with our analytical approximation in (24,

⁵We 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).

25), we see that the approximation follows the actual values closely.




Fig. 2. Average delay of limited replication scheme vs. the maximum number of replicas per message ($N = 10$, $\lambda_n = 1$, $\lambda_l = 0.8$, $\lambda = 2$, $TTL = 10$, $B = 10$, $M = 100$, 100 Monte Carlo runs): constrained version, unconstrained version, and its analytical approximation as in (24) with enforced monotonicity.

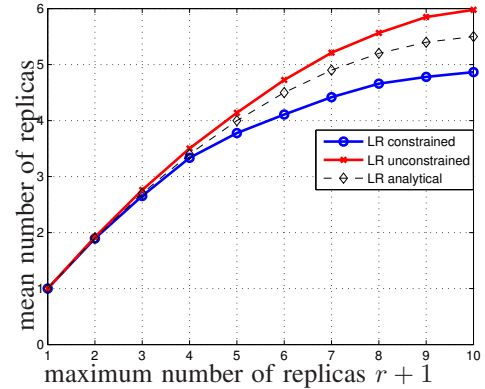


Fig. 3. Average number vs. the maximum allowable number of replicas per message (for the same parameters as in Fig. 2).

2) *Evaluation of Gateway Placement:* We simulate two domains of N_i ($i = 1, 2$) nodes each performing localized random walks on a $G \times G$ grid [19]. As illustrated in Fig. 4, there is a “home cell” in each domain, and nodes randomly move to one of their neighboring cells (including the current cell) with higher probabilities to move towards the home cell. Specifically, the transition probability to move to a neighboring 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 (the larger τ is, the more localized the mobility is). We randomly choose $L \leq G^2$ cells as candidate gateway locations and deploy $g < L$ gateways. Assume nodes and gateways can only communicate when they are in the same cell, with unlimited bandwidth and buffer size (the unconstrained case). New messages arrive in each slot with

⁶That is, $p_i = p'_i / \sum_{\text{neighbor } j} p'_j$ for $p'_i = e^{-\tau d_i}$.

probability λ ($0 < \lambda \leq 1$) and are uniformly distributed among the nodes in domain 1 with random destinations in domain 2. We implement two methods to compute the utility: simulations and analytical calculation⁷. For each method, we deploy gateways according to the optimal (*i.e.*, brute-force search) or the suboptimal solutions (*i.e.*, greedy and backward greedy algorithms). We deploy gateways according to one data set (contacts and traffic arrivals) and test its performance on another data set generated independently. For comparison, we also simulate random deployment as well as the oracle-assisted deployment by using the same data set for training and testing.

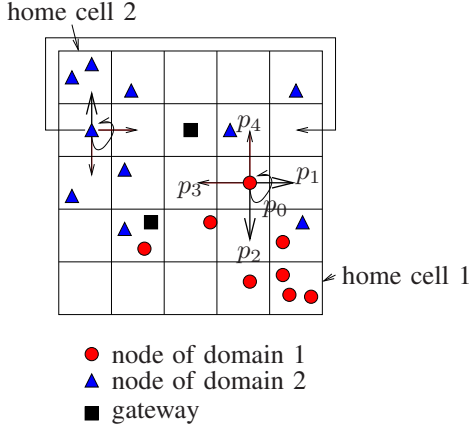


Fig. 4. Two domains performing localized random walks on a grid (with wrapping boundaries). Nodes in different domains can only communicate through gateways.

Fig. 5 compares the calculated end-to-end delays with the actual simulated delays under various combinations of routing schemes. We see that although the contact processes in this case are not Poisson, the calculated delays still approximate the trend of the actual delays reasonably well.




Fig. 5. 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). The calculated delay under direct delivery in both domains is scaled by 1.3818 to fit the simulated value.

⁷Single-domain average delays and numbers of replicas are calculated based on sample contact rates and then summed up to approximate the end-to-end utilities.

Using delay as the utility measure, we plot the performance of various deployment strategies as a function of the number of deployed gateways in Fig. 6 under unlimited replication in domain 1 and direct delivery in domain 2. From the results, we see that: (i) the proposed strategies (greedy/backward greedy algorithm based on calculated delay) perform almost as well as the optimal strategy (optimal algorithm based on simulated delay); (ii) the proposed strategies significantly outperform random deployment. The first observation is due to the facts that the calculated delays approximate the simulated ones closely, and the greedy and the backward greedy algorithms perform almost as well as brute-force search. The oracle strategy assumes the future contact and traffic traces to be completely known at the deployment phase, and thus its performance is absolutely the best among all deployment strategies. Since the contacts and traffic at run time may be different from those during deployment, the gap between the proposed strategies and the oracle strategy represents the inherent variation in the underlying networks rather than a deficiency of the strategies. Similar observations have been made under other routing schemes, as shown in Fig. 7–8 (backward greedy algorithm provides similar results as greedy algorithm and is thus omitted), as well as the utility measure of the number of replicas per message, as shown in Fig. 9.

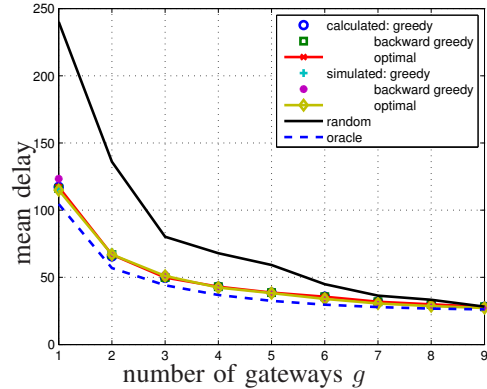


Fig. 6. Average delay vs. number of gateways ($N_1 = 5$, $N_2 = 10$, $G = 10$, $L = 10$, $\lambda = 0.05$, home cells for domain 1 and 2 at upper left and lower right corners, respectively, $\tau = 1$ for both domains, 1000 slots, 50 Monte Carlo runs).

B. Evaluation on Traces

In order to evaluate the performance of the framework proposed under realistic conditions, we have conducted a series of trace-driven simulations. UMass DieselNet [20] is a mobile network testbed consisting of wi-fi nodes mounted to buses in the area of Amherst, MA. Contacts of mobile nodes to various access points as well as contacts between mobile nodes have been logged. When using traces collected over such a scenario, there is a series of practical concerns that need to be taken into account. We observed for example, that mobile to mobile contacts were very sparse, especially when compared with the contacts of the mobile nodes with




Fig. 7. Same as Fig. 6 but under direct delivery in both domains.

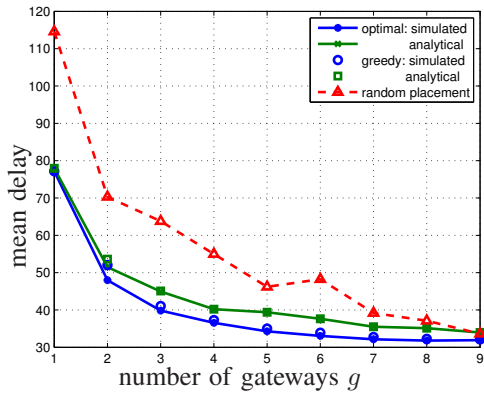


Fig. 8. Same as Fig. 6 but under unlimited replication in both domains.

the whole set of the Access Points. In addition, the contact traces are highly heterogeneous for different mobile nodes, as it obviously depends on the mobility pattern of each one. For this reason, we had to use only a subset of the fore mentioned trace, comprising the peak hours (7 am-7 pm) of 8 days. More specifically, the trace used describes the contacts among 3 mobile nodes per day and between these mobile nodes and 10 access points, which form our candidate solution set. The choice was made based on appearance consistency (i.e. the APs chosen should have come in contact with a mobile node on every day of the traces used), sufficiency of the number of contacts per day and elimination of nodes that demonstrated extremely large gaps in communication. All contacts between the same pair of nodes or node and gateway that were closer than 60 sec to each other were merged in one contact, as in [21]. At the same time we try to keep the time correlation of contacts between the chosen APs and the mobile nodes to a minimum. Finally, as we are evaluating the efficiency of approximation algorithms, the candidate solution set can not be very extensive, so that we can compute the optimal solution and use it as a benchmark.

1) *Contact Distribution*: To facilitate a better understanding of the trace used, we are plotting the Complimentary Cumulative Distribution Function of the pairwise inter-contact times




Fig. 9. Average number of replicas per message vs. number of gateways (for the same parameters as in Fig. 6) under unlimited replication in domain 1 and direct delivery in domain 2. The analytical calculation is based on (12) and Lemma 3.3 as the approximation in (21) does not depend on gateway placement.

between mobile nodes in Fig. 10, while the CCDF of the contacts between every mobile node and the APs (the set of the APs is viewed as a merged node in this context) is shown in Fig. 11. In both figures, we are comparing the empirically calculated probability with the analytical calculation under the assumption of exponentially distributed inter-contact times, where the mean of the distribution is estimated by the data available. For the case of mobile to merged gateway contacts, the empirical CCDF has a heavier tail than expected; this might be due to the fact that contacts between a mobile node and the different gateways are correlated in time, appearing in bursts, thus resulting in a big number of small inter-contact times when the node is in a “hot” area.

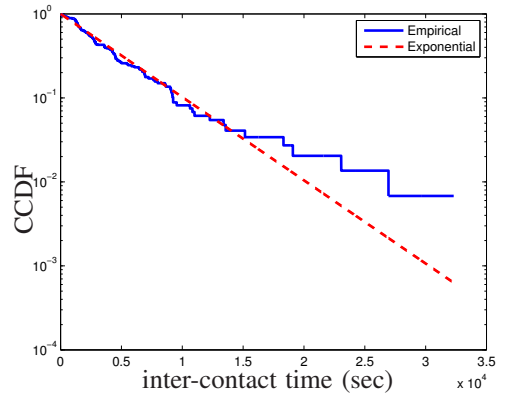


Fig. 10. CCDF of all the mobile-mobile contacts for the subset of the traces used (8 days of 3 mobile nodes per day).

2) *Accuracy of Utility Calculation*: In the following, we are considering all nodes of one day to belong to the same domain. We are then matching the days into 4 inter-domain pairs, and are using the contact traces to simulate direct delivery or unlimited replication routing schemes. Aiming to evaluate the merit of the analytical calculations of utility for the case




Fig. 11. CCDF of all the mobile-merged AP contacts for the whole subset of the traces that is used (8 days, 3 mobiles per day, 10 APs).

of overall delay and number of packets, we are simulating all possible gateway combinations when the number of gateways deployed varies from 2 to 8 (out of 10), and plotting the mean calculated and simulated utilities as these vary for different number of deployed gateways. We have performed 4 Monte Carlo runs for every possible gateway combination based on the different contact traces; in addition, for the case of simulated utility, 10 different Monte Carlo runs on traffic realizations have been run (every node of domain 1 sends messages to nodes of the other domain, at the rate of 5 packets per hour; the destinations are picked randomly among the nodes of domain 2). Fig. 12 shows the mean end-to-end delay when both domains use direct delivery or unlimited replication routing schemes. For simulation, messages that have not been delivered until the end of the simulation are considered delivered at the minimum of (i) the last contact of the node carrying the message or (ii) the last contact time observed for the destination node; this is necessary as not all of the nodes' trace spans across the exact same time period. Fig. 13 compares the calculated with the simulated utility when the metric concerned is the mean number of replicas per packet.

3) *Performance of Gateway Placement Strategies:* In order to evaluate the gateway placement algorithms, we are using every pair of domains to acquire solution sets based on three different methods, namely the greedy and the backward greedy algorithms as well as the optimal algorithm (brute-force search). As two different metrics can be used, i.e. the calculated or the simulated value of utility, every pair of domains will give 6 different, but not necessarily distinct solution sets. The evaluation consists of simulating the deployment of the suggested solution sets on the other three pairs of domains (namely the training set is not part of the testing set). Every method will be tested against 120 different Monte Carlo runs (4 deployment solutions \times 3 pairs of contact traces \times 10 different traffic realizations). Simulation settings are the same as above. We are plotting the simulated end-to-end utilities acquired when evaluating the 6 different methods with respect to the number of gateways deployed. As a benchmark, we are

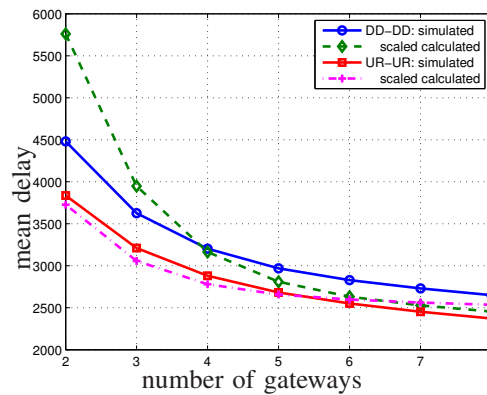


Fig. 12. Comparison of simulated and calculated end-to-end delay for direct delivery (“DD”) and unlimited replication (“UR”) schemes; we have scaled the calculated delays by a constant factor of 0.7171 for direct delivery and 0.4859 for unlimited replication to fit the simulated values.

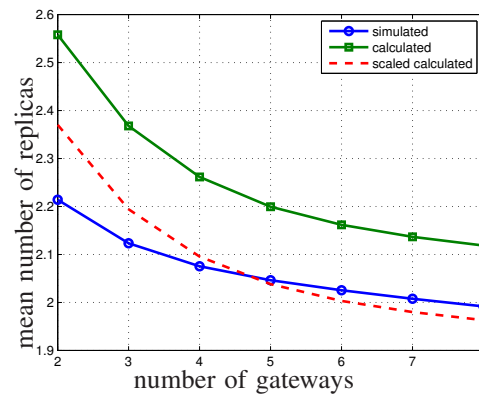


Fig. 13. Comparison of simulated and calculated number of replicas per packet, together with a scaled version of the calculated values by a constant factor 0.9267. Both domains use unlimited replication.

also plotting the mean end-to-end utilities averaged over all the possible combinations of gateways (which defines the average performance of random placement), as well as the maximum utility (i.e., the minimum delay) that can be achieved under oracle deployment, defined as the value of the utility when the deployment is evaluated on the training data set.

Fig. 14 shows the mean delay when both domains use direct delivery schemes. All the methods based on the calculated utilities yield the same results (more specifically, the backward greedy algorithm yields exactly the same results as the optimal algorithm, while the greedy algorithm has very slight deviations). The methods based on the simulated utilities yield different results most of the times. Although there is no clear pattern in the relative performance of the two groups of methods, it is encouraging to see that the solutions of the methods based on analytical utility calculation perform equally well as the ones based on simulation. We also observe that the proposed methods yield better results than random placement, especially when the number of deployed gateways is small. As the number of deployed gateways increases,




Fig. 14. Evaluation of proposed methods for direct delivery schemes, when minimizing overall delay.

the deployment solutions converge, all converging to a full deployment. Similar results are obtained for the unlimited replication scheme (see Fig. 15) (for unlimited replication, calculation-based methods work better than simulation-based ones for the majority of the cases) and for the minimization of the mean number of replicas per packet (see Fig. 16).




Fig. 15. Evaluation of proposed methods for unlimited replication schemes when minimizing overall delay.

REFERENCES

- [1] T. He, N. Sofra, K.-W. Lee, and K. Leung, "Enabling Inter-domain DTN Communications by Networked Static Gateways," submitted to the 29th IEEE Conference on Computer Communications (INFOCOM'10), July 2009.
- [2] A. Balasubramanian, B. N. Levine, and A. Venkataramani, "DTN Routing as a Resource Allocation Problem," in *Proceedings of ACM SIGCOMM*, (Kyoto, Japan), August 2007.
- [3] 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.
- [4] T. Spyropoulos, K. Psounis, and C. Raghavendra, "Single-Copy Routing in Intermittently Connected Mobile Networks," in *Proceedings of IEEE SECON*, October 2004.
- [5] 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.
- [6] A. Vahdat and D. Becker, "Epidemic Routing for Partially Connected Ad Hoc Networks," in *Duke Technical Report CS-2000-06*, July 2000.

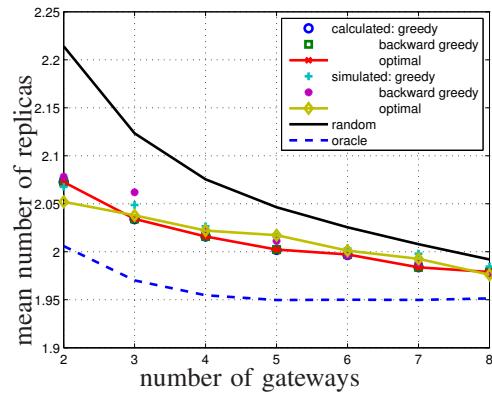


Fig. 16. Evaluation of proposed methods for unlimited replication schemes, when minimizing mean number of replicas per packet.

- [7] J. Davis, A. Fagg, and B. N. Levine, "Wearable Computers and Packet Transport Mechanisms in Highly Partitioned Ad Hoc Networks," in *Proceedings of IEEE ISWC*, pp. 141–148, October 2001.
- [8] A. Lindgren, A. Doria, and O. Schelen, "Probabilistic Routing in Intermittently Connected Networks," in *Proceedings of SAPIR Workshop*, pp. 239–254, 2004.
- [9] 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.
- [10] 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.
- [11] T. Small and Z. Haas, "Resource and Performance Tradeoffs in Delay-Tolerant Wireless Networks," in *Proceedings of ACM WDTN*, pp. 260–267, August 2005.
- [12] 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.
- [13] E. Jones, L. Li, and P. Ward, "Practical Routing in Delay-Tolerant Networks," in *Proceedings of ACM Chants Workshop*, pp. 237–243, August 2005.
- [14] S. Jain, K. Fall, and R. Patra, "Routing in a Delay-Tolerant Network," in *Proceedings of ACM SIGCOMM*, pp. 145–158, 2004.
- [15] 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.
- [16] G. Sandulescu and S. Nadjm-Tehrani, "Opportunistic DTN Routing with Window-Aware Adaptive Replication," in *Proceedings of ACM AINTEC*, (Bangkok, Thailand), November 2008.
- [17] R. C. Shah, S. Roy, S. Jain, and W. Brunette, "Data MULEs: Modeling a Three-tier Architecture for Sparse Sensor Networks," in *IEEE Workshop on Sensor Network Protocols and Applications (SNPA)*, pp. 30–41, May 2003.
- [18] J. Widmer and J. Y. L. Boudec, "Network Coding for Efficient Communication in Extreme Networks," in *Proceedings of ACM WDTN*, pp. 284–291, August 2005.
- [19] 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.
- [20] 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.
- [21] 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.