Relays, Base Stations, and Meshes: Enhancing Mobile Networks with Infrastructure

Nilanjan Banerjee, Mark D. Corner, Don Towsley, Brian N. Levine
Department of Computer Science
University of Massachusetts
Amherst, MA
{nilanb, mcorner, towsley, brian}@cs.umass.edu

ABSTRACT

Networks composed of mobile nodes inherently suffer from intermittent connections and high delays. Performance can be improved by adding supporting infrastructure, including base stations, meshes, and relays, but the cost-performance trade-offs of different designs is poorly understood. To examine these trade-offs, we have deployed a large-scale vehicular network and three infrastructure enhancement alternatives. The results of these deployments demonstrate some of the advantages of each kind of infrastructure; however, these conclusions can be applied only to other networks of similar characteristics, including size, wireless technologies, and mobility patterns. Thus we complement our deployment with a demonstrably accurate analytical model of large-scale networks in the presence of infrastructure.

Based on our deployment and analysis, we make several fundamental observations about infrastructure-enhanced mobile networks. First, if the average packet delivery delay in a vehicular deployment can be reduced by a factor of two by adding x base stations, the same reduction requires $2x$ mesh nodes or $5x$ relays. Given the high cost of deploying base stations, relays or mesh nodes can be a more cost-effective enhancement. Second, we observe that adding small amount of infrastructure is vastly superior to even a large number of mobile nodes capable of routing to one another, obviating the need for mobile-to-mobile disruption tolerant routing schemes.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: [Modeling techniques]

General Terms

Experimentation, Measurement, Performance

Keywords

Hybrid Mobile Networks, Base stations, Relays, Mesh, Ordinary Differential Equations

1. INTRODUCTION

Mobile networks incur higher delays and more frequent disconnections than tethered networks. While a complete and robust infrastructure is necessary to support applications most sensitive to delay, such as VOIP, many applications from environmental monitoring [16] to software updates, email, and instant messaging can tolerate longer delays and intermittent connectivity. However, the lower the delay, the greater the number of applications the network can support.

Recent efforts have proposed adding infrastructure/stationary resources to mobile networks to improve connectivity [3, 7, 11, 4]. Infrastructure, such as open access points, can be used opportunistically [3, 7, 13]; however, the use of such infrastructure may not be available in many settings, such as remote monitoring. In contrast, one can build new infrastructure, but it can only be justified if the additional costs are manageable and the performance enhancement is significant.

The costs of adding different types of infrastructure is highly variable. For example, installing wired *base stations* connected to the Internet can lower delays; but they require costly installation of power and wired network connectivity—these costs can be as high as US$5,000 per base station [17]. Although opportunistic use of open base stations is free, finding volunteers that permit access represents a management cost and may not provide sufficient reliability. Moreover, their use may be legally problematic. An alternative is to deploy a wireless *mesh network* [5] from short-range, high-bandwidth technologies, e.g., WiFi, or from long-range, low-bitrate radios, e.g., those in the 900MHz band. This saves the cost of connecting the access points to the Internet, but requires a minimum density to maintain a connected topology. Lastly, and perhaps most inexpensively, one can place *relays* in the network that require no connections to electrical infrastructure or to the Internet and can be placed anywhere in the network [4, 14]. Relays are limited to routing information between mobile nodes in a disruption-tolerant fashion [15].

Unfortunately, the relative performance enhancement of each kind of infrastructure is poorly understood. For instance, if opportunistic access to base stations is available, is mobile-to-mobile routing necessary? If opportunistic access to base stations is unavailable, should one deploy relays or base stations to lower delays? How do these trade-offs scale as the network grows in size or density?

To answer these questions, we present results from a set of field experiments and analysis that compares the benefits of each kind of hybrid mobile network. While the deployment study demonstrates many of the practical issues involved—including node placement, real world propagation, and dynamic routing protocols, the results heavily depend on the particular system, underlying technology, and mobility patterns. To address this shortcoming, we use observa-

MobiCom'08, September 14–19, 2008, San Francisco, California, USA.

tions from the deployment study to develop analytical models for large-scale hybrid networks based on simple, ordinary differential equations (ODE). In the analysis, we present asymptotic results on the expected packet delivery delay and resource consumption for epidemic routing.

Research Contributions

Our deployment and analysis yield the following results.

- From our experiments in a deployed mobile network, we identify scenarios where a mesh or relay hybrid network is a better choice over a base station network. For example, if the average packet delivery delay in a vehicular deployment can be reduced by a factor of two by adding x base stations, the same reduction requires $2x$ mesh nodes or $5x$ relays. While a larger number of relays and mesh nodes are needed, the ease and cost of deployment make them a better choice in many circumstances.

- We observe that a small amount of added base station or mesh infrastructure can quickly obviate the need for mobile-to-mobile routing schemes. Simple two-hop forwarding algorithms, where a node delivers data to either a base station or mesh node followed by a pickup by the intended recipient, provide excellent performance.

- Based on the observations made in our deployment, we develop an analytical model of large-scale networks with infrastructure. The results from the model are within 10% of our deployment. For special cases, we derive closed-form expressions for the expected packet delivery delay of hybrid networks.

- Based on the deployment and analysis, we observe that an increase in the number of mobile nodes has little impact on expected packet delivery delay when there is infrastructure present. This differs dramatically from what is observed in an infrastructureless network where the addition of a small number of nodes does lead to a large reduction in expected delay.

Our study can be used as the basis of a cost-benefit analysis in deploying a mobile network or augmenting an under-provisioned one. For example, underwater networks can use wireless buoys as relays in the network [20], while other underwater networks may use fully interconnected mesh networks [23]. In a vehicular network, one may use relays [4] or wireless base stations [7]. Throughout this paper we have made every effort to develop a general set of tools to examine infrastructure across a large number of scenarios and technologies.

2. NETWORK MODEL

We begin with a network consisting exclusively of mobile nodes. As in a disruption tolerant network (DTN) [15], nodes transmit replicas of a packet to passing mobile nodes, which physically move and transmit the replicas to additional nodes—eventually a copy of the packet is delivered to its destination. The packet spreads through such a network in a controlled epidemic fashion and uses replication to reach its destination. It is important to note that such a DTN can incur packet delays from seconds to days, but it is not appropriate for very short delay applications, such as those in the millisecond range, where delays are determined primarily by congestion rather than mobility and disconnection.

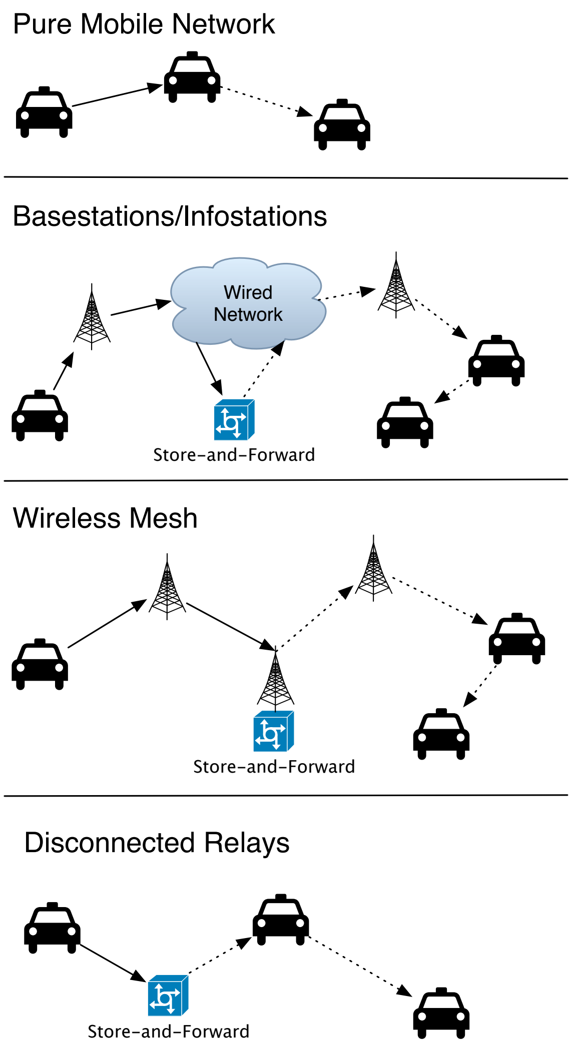

Figure 1: **A sparse mobile network and three methods of augmenting the network with resources to reduce delay. In the first case, base stations communicate with mobile nodes storing packets on a node in the network for eventual delivery to another node. The second case, a wireless mesh, operates similarly to the base station, but the mesh nodes must be physically proximate to transmit packets. In the third case, disconnected relays store packets, for delivery to other relays or mobile nodes, propagating information towards the final destination.**

There are at least three options for enhancing such a mobile network through additional stationary infrastructure, as illustrated in Figure 1. (1) Wireless base stations provide new opportunities to propagate packets, typically via an in-network proxy with storage so that mobile nodes can pass packets without contemporaneous connections through the base stations. (2) A wireless mesh works in the same way, but the wireless mesh nodes must be geographically placed such that a wireless backbone is connected. (3) Disconnected store-and-forward relays route packets to any passing mobile node without connection to a mesh or the Internet. With any of these methods, and sufficient resources, one can reduce the delay of the network to arbitrarily small amounts, bounded only by transmission delays and buffering.

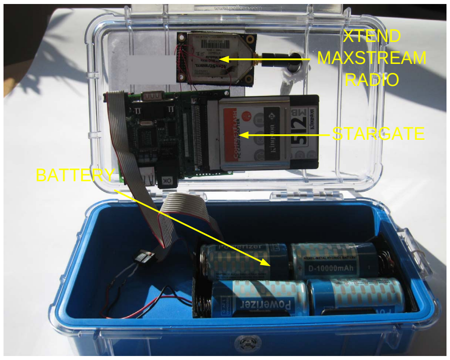

Figure 2: **Prototype for a relay/mesh node.**

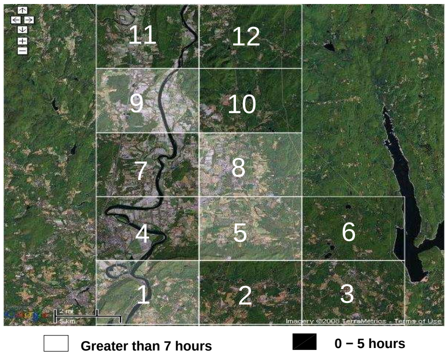

Figure 3: **Heatmap showing the amount of time vehicles spend in different regions of the network during a day.**

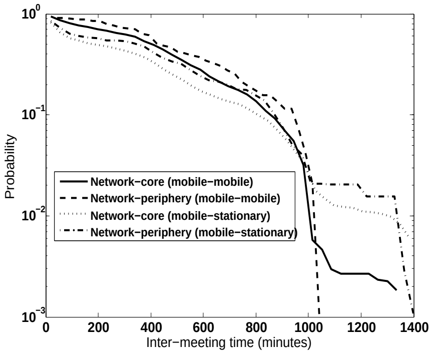

Figure 4: **CCDF of aggregate pairwise inter-meeting times between contacts for mobile nodes and mobile and stationary nodes.**

Intuitively, base stations can support much lower delays than an equivalent number of relays—the wired network effectively provides a wormhole across a geographic area. However, the key questions are *how many relays are needed to lower packet delivery delay to an acceptable level, and what are the relative costs of these options?* While base stations and meshes are necessarily more effective than relays, relays avoid the constraints of mesh wireless connectivity and the costs of base station wiring.

We analyze the relative benefit of adding varying numbers of base stations, mesh nodes, and relays using two methods:

- a physical deployment of base stations, meshes, and relays in a network of vehicular nodes;
- an ordinary differential equation (ODE) model based on observations made in the deployment, using methods from viral worm propagation [26].

3. VEHICULAR NETWORK DEPLOYMENT

The performance of hybrid mobile networks can be affected by a myriad of factors, including routing protocols, node mobility patterns, stationary node placement, and real world propagation. To better understand their performance, one first needs to understand the impact of each of these factors in a real mobile network. Hence, we perform a systematic study of base stations, relays, and mesh nodes in a large-scale vehicular network.

3.1 Testbed Overview

Our mobile network, UMass DieselNet [6], consists of an average of 20 transit buses in use on the road 18 hours/day everyday in an area of about 360 square kilometers. The buses are equipped with a Linux box, a WiFi radio, a GPS unit, and a long-range, low-bandwidth 900MHz XTend radio. The buses communicate with each other over TCP connections using WiFi radios. To examine different kinds of infrastructure, we deployed three different kinds of networks.

Relays: We placed six stationary relay nodes in the network for a period of 20 days. The nodes consist of a PDA-class Stargate device equipped with a WiFi CF card and 64 MB of storage (see Figure 2). When a bus comes within WiFi range, the two exchange data over a TCP connection until the contact opportunity ends. The relay node stores those packets, waits for another transit bus, and then exchanges packets with that bus, propagating packets from mobile node to mobile node.

Characteristics	Network-core	Network-periphery
average pairwise inter-contact time	434 min	506 min
average contact duration	8590 ms	2802 ms
number of contacts (per day)	145	7

Figure 5: **The above table shows the characteristics of the two regions in the network.**

Mesh Nodes: These same stationary nodes can be configured as a mesh network. Using an additional radio, in this case a 900MHz radio, the nodes form a mesh network. When one of the mesh nodes receives data from a bus over WiFi it propagates those packets to other nodes in the mesh network.

Base Stations: For base stations, we used a combination of APs that we deployed and a set of open-access points set up by third parties. When a bus passes an AP it exchanges data with a central server through the AP, reachable by all of the base stations. We collected the base station traces from October–November 2007.

We collected only data inherent to the network and type of infrastructure so that we could later apply the traces to a variety of scenarios. Each node, both mobile and stationary, collected the duration of time and when it was in contact with other nodes, as well as the amount of data it transferred. For the mesh nodes, we collected data on the WiFi-based bus-to-mesh and the XTend-based mesh-to-mesh connections. The mobile nodes collected detailed mobility traces from GPS units. We filled the network with dummy data to ensure complete traces—this also allows us to ignore the effects of routing protocols and apply those effects later in trace-driven simulations.

3.2 Network Characteristics

To understand how placement and mobility effects hybrid network performance, we first look at the geographic characteristics of our mobile network testbed. Figure 3 shows a heatmap of the amount of time vehicles spend in different parts of the network. The darker squares are areas where the vehicles spend the smallest amount of time while lighter regions are areas where the vehicles spend the largest amount of time. From the figure, we can divide the network into two disjoint regions: (1) a *network core*, where the mobile nodes reside for a large fraction of time (squares 1, 5, 8, and 9 in Figure 3); and (2) a *network periphery*, where the mobile nodes spend the rest of the time. Figure 5 shows the characteristics of these two regions. Note that the pairwise contact durations, the pairwise inter-contact times (the time between *particular* pairs of nodes meeting [1]) and the overall number of contacts is different for both regions.

As we show in Section 4, the dichotomy of regions is essential in developing an accurate mathematical model for mobile networks with infrastructure. Further, this feature is not unique to this kind of deployment—a dichotomy of regions is inherent in many other mobile networks. For example, in a network of humans [9], nodes are likely to spend more time in certain parts of the network as compared to others. The bipartition of the mobile network also provides insight on how stationary nodes should be placed in a mobile network to maximize performance.

While the two regions exhibit different absolute contact statistics, the pairwise inter-contact times between mobile nodes and mobile and stationary nodes for both regions have similarly *shaped* distributions. Figure 4 shows the CCDF of aggregated inter-contact time in each region. From the graph we see that 90% of all the contacts approximately follows an exponential distribution. This provides a strong hint that the mobile-to-mobile and mobile-to-stationary pairwise meeting time in each region can be modeled as Poisson processes. The noisy behavior in the tail of the distribution is due to end-of-day effects, called *partially observed contacts*—these are contacts that have a start time but not an end time recorded in the logs. We calculate the inter-meeting time for such contacts using the end of the day as the last contact between the node pair (similar to previous work in vehicular networks [25]).

3.3 Trace-driven Simulator

Using the traces of our deployment, we built a trace-driven simulator to evaluate the performance benefits of alternatives to enhancing a mobile network. The simulator can be used to examine three key factors that affect performance: the type of infrastructure enhancement, the placement of that infrastructure, and the choice of the routing protocol. Examining each type of infrastructure enhancement (relay, mesh, and base stations) is straightforward, we modify the simulator to simulate each communication path—relays can only communicate with mobile nodes, meshes and base stations with other infrastructure nodes and mobile nodes. Modifying the simulator for routing protocols is also possible: the packets exchanged in the simulator between nodes is determined by the particular algorithm we choose.

Placement requires restricting infrastructure to a set of feasible locations. Relays can be placed anywhere in the network. Using solar cells and batteries, relay nodes can be placed independent of power and network infrastructure, as well as independent of one another. Mesh nodes are more constrained, in that they must be placed within range of one another (in the case of our mesh nodes this range is 1650 meters). We have built these mesh nodes from the same solar-powered boxes as the relays, and the collection of connected nodes can be placed anywhere subject to the maximum distance constraint. The base stations can only be placed where there is wired connectivity. However, they do not need to be in proximity to one another. A pair of base stations can communicate across the full geographic diameter of the network. To approximate the placement constraint for base stations, we assume that the base stations can only be placed at current access point locations.

The remaining parameters come from the traces themselves. We use the vehicle GPS traces to calculate the time and duration of contacts between mobile and stationary nodes. We use measured values for the WiFi radio (100 meters) and the Digi-XTend radio (1650 meters) from our six node deployment. The simulator uses the bus-to-relay and bus-to-AP bandwidth distributions measured from our deployment for the amount of data transferred during each contact. For the mobile-mobile contacts, we use the bus-bus connections logged by the nodes in our testbed. The data includes the time, location, duration of contacts, and the amount of data transferred during connection events. The movement of vehicles is taken directly from the GPS traces.

3.4 Stationary Node Placement

Intuitively, there are two strategies for stationary node placement based on regions in the mobile network.

- **Uniform placement:** place the nodes uniformly across the entire network limited only by the placement constraints described above.
- **Non-uniform:** place more nodes in the network core, while still following the placement constraints. We use a simple heuristic for such a placement. The number of nodes in each square (see Figure 3) is proportional to the amount of time mobile nodes spend in that square. The placement within a square is uniform.

We evaluate the effect of each placement strategy using our trace-driven simulator. As a starting point, we use RAPID [1] as our routing protocol. We chose RAPID since it has been shown to perform close to optimal. RAPID is a routing protocol that can optimize for a given metric. It treats routing as a resource allocation problem that converts a routing metric to a per-packet utility that determines the degree of replication of a packet. In our experiments we use average packet delivery delay as our routing metric. We present results from experiments using other routing schemes in Section 3.5. We compare performance using two metrics (1) *average packet delivery delay*: the lifetime of a packet from its creation to the first time it is delivered to the destination; (2) *average number of transmissions per packet*: the number of copies made by the time the packet is first delivered to its destination. This can be used to estimate energy and bandwidth as both are proportional to the number of packet copies.

The simulation generates 1KB packets at a uniform rate of 5 packets per destination per hour. This load is intentionally small so that it does not exceed the network capacity. The small load also isolates the effect of delay from bandwidth. The destination for each packet is divided into two equal classes: with probability one half, the destination is a sink node; otherwise it is equally likely to be destined for any other mobile node. The sink node for the relay and mesh case is one random node, and in the base station case it is any base station node.

Figure 7 shows the average packet delivery delay for different infrastructures for both uniform and non-uniform placement across a range of numbers of infrastructure nodes. The error bars represent the 95% confidence interval around the sample mean. Figure 8 shows the number of transmissions per packet, including replication. From the figures, we make the following observations about our deployment. (1) We need only 5–7 times as many relays or 2–3 times as many mesh nodes as base stations to reduce delay by a factor of two. Given that the cost of deploying base station networks can be greater than 7 times the cost of relays or a mesh network (we provide some realistic cost estimates in Section 5.5), either relays or meshes are better choices. (2) The non-uniform placement of nodes leads to 10% lower delays as compared to uniform placement. Therefore, placing infrastructure in popular regions can lead to better performance. (3) The average number of transmissions per packet for different infrastructures is similar. In terms of resource consumption, all infrastructure networks are similar.

---

[1] Note that this is not the same as the inter-meeting time between a node and *any* other mobile node. This distinction is important for the remainder of the paper.

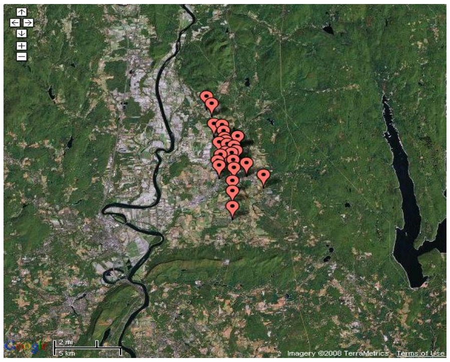

Figure 6: Placement of 25 base stations non-uniform across the mobile network.

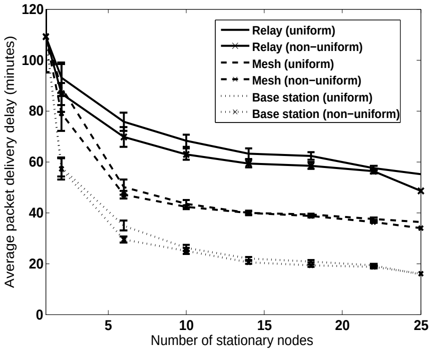

Figure 7: The average packet delivery delay with a varying number of stationary nodes for uniform and non-uniform placement.

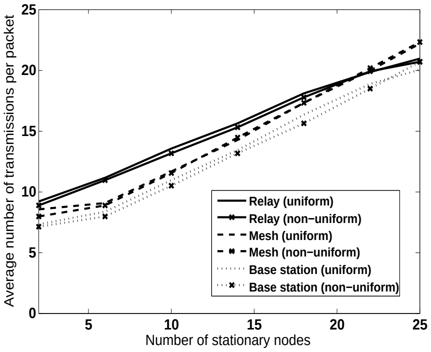

Figure 8: The average number of transmissions per packet with a varying number of stationary nodes for non-uniform and non-uniform placement.

In this section, we discussed uniform and non-uniform placement of stationary infrastructure based on *regions* in the network. However, we did not present any algorithm to determine the exact locations for placing base stations, relays or mesh nodes. Although Wenrui et al. [27] present schemes for placing relays in mobile disruption tolerant networks to maximize performance, deciding the optimal location for placing base stations and mesh nodes is an open problem.

3.5 Dynamic Routing Protocols

To determine the effect of different routing schemes on different kinds of infrastructure, we perform a similar experiment with three routing protocols: (1) RAPID, (2) two-hop: a protocol that does not use any mobile-to-mobile routing—instead the mobile source only passes the packet to the infrastructure, and the destination can only receive the packet from it as well, and (3) Random-epidemic: a base-line protocol that creates copies of a packet randomly, i.e., when two mobile or stationary nodes meet the packets they transfer are chosen randomly from their buffers.

Figures 9, 10, and 11 compare the average packet delivery delay for random, RAPID, and two-hop routing for relay, mesh, and base station networks. From these graphs, comparing two-hop routing with RAPID, we find that the improvement yielded by mobile-to-mobile routing is marginal for mesh and base stations. *Base stations and meshes substantially help reduce delays. However, the additional benefit of mobile-to-mobile routing over a two-copy approach yields little additional benefits.* This is because a large number of packets are transferred using the wired and wireless backbone. Somewhat intuitively, mobile-to-mobile routing helps in the relay case, as the destination must meet one of the same relay nodes that the source saw. Moreover, from the figures, we find that Random-epidemic performs close to RAPID. From this we can conclude that even for relay infrastructure, simple epidemic schemes like random have excellent performance. This is because the increase in capacity of the network through additional infrastructure overwhelms the benefits of priority schemes to select packets for replication, as used by RAPID, over Random-epidemic.

Since most of our results are for low traffic load, we also experimented with higher packet generation rates for RAPID and Random-epidemic. We found that increasing the packet generation load from 5 packets per hour per destination to 120 packets per hour per destination (a factor of 24) leads to a 29%, 12%, and 28% increase in average packet delivery delay for base stations, meshes, and relays, respectively. Moreover, even for higher loads the performance of Random-epidemic is close to RAPID. From the experiment, we conclude that the addition of infrastructure manages increasing network load well.

4. ANALYTICAL MODEL

Although our deployment study presents interesting trade-offs between the three type of infrastructure it suffers from two limitations. (1) The number of mobile nodes are fixed, hence we cannot validate our results for a large number of mobile and stationary nodes. (2) We cannot infer anything about the asymptotic behavior of hybrid networks. For example, how should the number of relay and base station nodes scale in the number of mobile nodes?

To address these limitations, we develop detailed analytical models for large-scale networks in the presence of infrastructure. We validate our model with results from the deployment and use it to study the performance of infrastructure-enhanced mobile networks when the number of mobile and stationary nodes are large. Finally, for a special case, we derive closed-form analytical expressions for delay and number of packet transmissions for base station and relay mobile networks. For N mobile nodes, we show that the network needs $\omega(N)$ relays and $\omega(\sqrt{N})$ base stations before the stationary nodes substantially affect the average packet delivery delay for epidemic routing.

4.1 Model and Network Parameters

We model the mobile network as $N+1$ mobile nodes and M stationary nodes. The area covered by the mobile network is divided into k disjoint *regions*. We observed that $k = 2$ for the deployment described in Section 2. The mobile nodes spend time in each of the k regions. The stationary nodes can either be placed uniformly— every region has the same number of stationary nodes or non-uniformly— some regions have more nodes than other regions. The stationary nodes are placed uniformly within a region.

An example of a two region network is shown in Figure 12. We assume that the *pairwise meeting times* between mobile nodes within region i are represented by exponentially distributed random variables with mean inter-meeting time $1/\beta_i$, and the *pairwise meeting times* between mobile and stationary nodes are represented by exponentially distributed random variables with mean inter-meeting time $1/\gamma_i$. The time at which mobile nodes move from region i to j are assumed to be represented by exponentially distributed random variables with mean rate μ_{ij}.

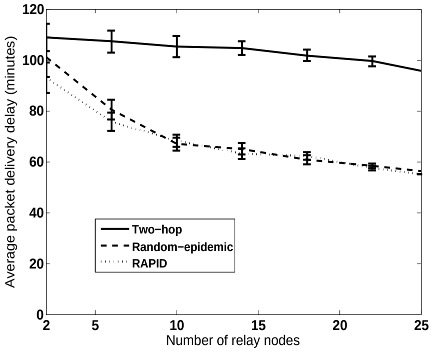

Figure 9: **The average packet delivery delay with a varying number of relay nodes for the three routing protocols.**

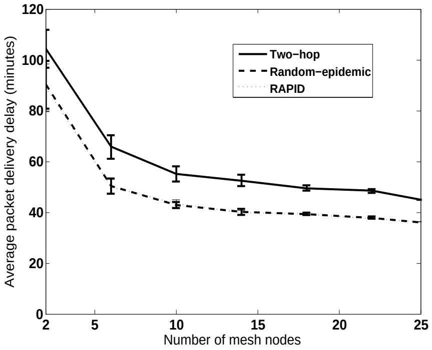

Figure 10: **The average packet delivery delay with a varying number of mesh nodes for the three routing protocols.**

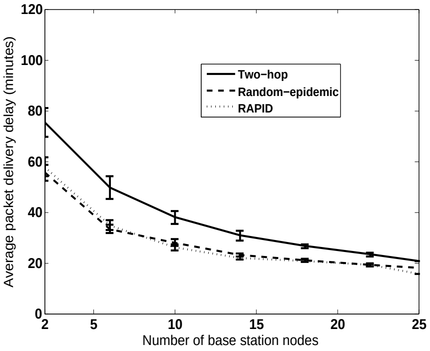

Figure 11: **The average packet delivery delay with a varying number of base station nodes for the three routing protocols.**

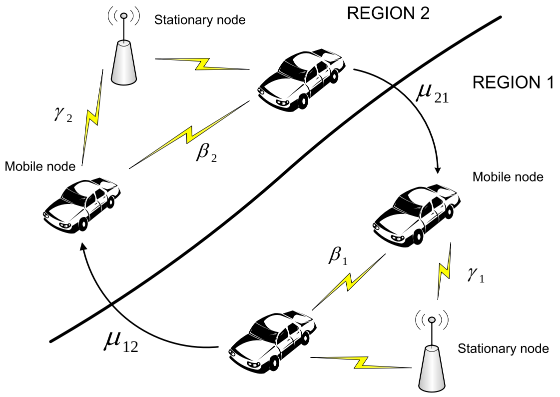

Figure 12: **The division of the mobile network into regions.**

The exponential assumption for pairwise meeting times between mobile nodes follow from a result by Groenevelt et al. [12], which showed that pairwise meeting times are nearly exponential if nodes move in area A according to common mobility models (such as random way point and random direction) and if their transmission range d is small compared to A. The exponential assumption for pairwise meeting times between mobile and stationary nodes follow from a result by Ibrahim et al. [14], which showed that the pairwise meeting times between mobile and stationary nodes are nearly exponential if mobile nodes move in area A according to common mobility models (such as RWP and RD). Our deployment results confirm that for a large number of meetings, the pairwise meeting times for mobile-mobile and mobile-stationary contacts (Figure 4) can be approximately represented as exponentially distributed random variables.

4.2 Routing Protocol and Traffic Model

We assume a very general traffic model: *(i)* traffic in the network is unicast; *(ii)* traffic sources and destinations are uniformly random; and *(iii)* stationary nodes route data and can serve as sinks and not sources. For traffic destined to a sink, all base stations, one of the mesh access points, and one of the relays are connected to the Internet.

Our analytical model evaluates a general epidemic routing protocol where every node forwards packets to every other node it meets. Two nodes meet when they are within transmission range of each other. Every node has a large buffer and every transfer opportunity is long enough such that all packets in a node's buffer can be transferred to its peer. Hence, every packet can be considered independent of all other packets.

4.3 Epidemic Spread

We model the spread of a packet and its replicas as an *epidemic infection* among nodes in the network. When nodes meet one another they exchange packets, infecting each other with the packets they possess until the packet infects the eventual destination. Similar models have been used in the analysis of worm propagation and purely mobile networks [26].

This packet infection model consists of a system of non-linear differential equations with two variables for every region i: $x_i(t)$ and $y_i(t)$. The number of mobile nodes infected with a packet at time t within region i is $x_i(t)$, and the number of infected stationary nodes at time t within region i is $y_i(t)$. The network delivers a packet once the destination is infected with the packet, and our goal is to determine $P(t) = Pr[T < t]$, the probability that the time to deliver a packet is less than t. The expected delivery delay of a packet is given by $\int_0^\infty (1 - P(t))dt$.

Similar to our deployment study, our other goal is to estimate the amount of resources used in the network, including bandwidth, storage, and energy costs. All are strongly correlated with the number of transmissions/copies per packet. The expected number of copies per packet, $E[C]$, is given by the following.

$$E[C] \quad = \quad \int_0^\infty (\sum_{i=1}^{i=k}(x_i(t) + y_i(t)))dP(t). \tag{1}$$

We assume in our analysis that all replicas of a packet are removed once the packet is delivered to its destination. We discuss how this assumption can be relaxed in Section 5.3. In this section, we describe the set of differential equations for each of the three cases: relays, base stations, and meshes.

4.4 Mobile Network with Untethered Relays

Recall that at any instant of time t, there are $x_i(t)$ infected mobile nodes and $y_i(t)$ infected relay nodes within region i. The total number of mobile nodes (infected + susceptible) in region i is given by $n_i(t)$ such that, $\sum_{1 \le i \le k} n_i(t) = N + 1$. The total number of

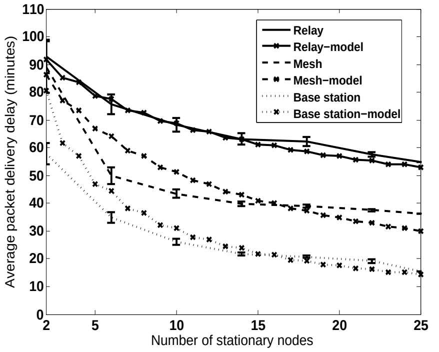

Figure 13: **The average packet delivery delay for uniform placement of stationary nodes from the model and the deployment.**

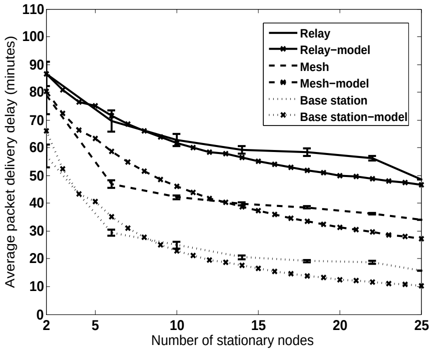

Figure 14: **The average packet delivery delay for non-uniform placement of stationary nodes from the model and the deployment.**

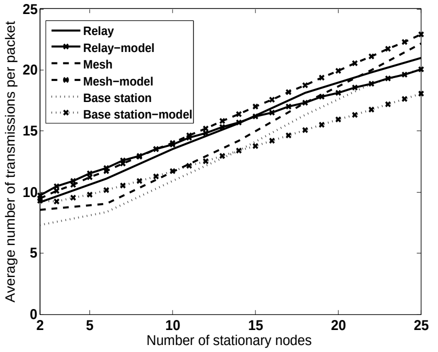

Figure 15: **The average number of transmissions per packet from the model and the deployment.**

relays in region i is M_i. We initially assume that all packets are destined towards mobile nodes. Later in the section, we modify the differential equations for packets destined to a randomly chosen relay which can act as a sink. The differential equations that govern the dynamics of the infection process in region i is given below:

$$\begin{aligned} x_i'(t) &= (\beta_i x_i(t) + \gamma_i y_i(t))(n_i(t) - x_i(t)) \\ &\quad - \sum_{\forall j \neq i} \mu_{ij} x_i(t) + \sum_{\forall j \neq i} \mu_{ji} x_j(t), \\ y_i'(t) &= \gamma_i x_i(t)(M_i - y_i(t)), \\ n_i'(t) &= \sum_{\forall j \neq i} \mu_{ji} n_j(t) - \sum_{\forall j \neq i} \mu_{ij} n_i(t), \\ x_1(0) = 1, &\quad \forall i \neq 1, x_i(0) = 0, \forall i, y_i(0) = 0, \\ &\quad n_1(0) = N + 1. \end{aligned} \tag{2}$$

The rate of change of $P(t)$, $P'(t)$, is given by the following:

$$\begin{aligned} P'(t) &= \frac{\sum_{i=1}^{i=k} x_i'(t)}{N}, \\ &P(0) = 0. \end{aligned} \tag{3}$$

Here $x_i'(t)$ is the rate at which mobile nodes in the network are infected by the packet in region i. It is the sum of the rate of mobile nodes meeting other infected mobile nodes in the region, the rate of mobile nodes meeting infected stationary nodes in the region, and the rate at which infected mobile nodes move into region i from other regions in the network. We also subtract the rate at which infected mobile nodes move out of region i. The rate of change of infected stationary nodes in region i, $y_i'(t)$, is equal to the rate at which infected mobile nodes in region i meet uninfected/susceptible stationary nodes. The rate at which total number of mobile nodes in region i changes, $n_i'(t)$, is the rate at which nodes from other regions enter region i minus the rate at which nodes from region i move to other regions. The rate of change of $P(t)$ is given by the rate of change of infected mobile nodes in all regions at time t divided by the total number of mobile nodes excluding the source node. This is because the destination of a packet is chosen uniformly at random from the set of mobile nodes excluding the source. The last set of equations describe the initial conditions for the system. Initially, only one mobile node in region 1 and no stationary node is assumed to be infected with the packet. Without loss of generality, initially, all the $N + 1$ mobile nodes are assumed to be in region 1.

The above set of differential equations assume that the destination is always a mobile node. However, in many cases data is destined towards a sink (e.g., the Internet). To account for node to sink traffic we assume that one of the relay nodes in region 1 is a sink. Such a node could be located close to an open access point and can act as a gateway to the Internet. We only need a minor modification to $P'(t)$ to account for the source-sink traffic. For packets destined to the sink, $P'(t)$ is the following.

$$\begin{aligned} P'(t) &= \frac{y_1'(t)}{M_1}, \\ P(0) &= 0. \end{aligned} \tag{4}$$

In the above differential equation, $P'(t)$ is equal to the rate at which relays are infected in region 1 divided by the total number of relays in the region. Now, assuming that a fraction f of the traffic is destined to a sink, the expected packet delivery delay is given by $(1-f) \cdot E[T_{pp}] + f \cdot E[T_{ps}]$, where $E[T_{pp}]$ is the expected delivery delay for peer-to-peer traffic and $E[T_{ps}]$ is the expected delivery delay for sink traffic.

4.5 Mobile Networks with Base Stations

The key advantage of relays is that mobile nodes do not need to be in the same place at the same time to exchange packets, they only have to meet the same relay at different times. However, if the network is augmented with wired base stations, packets can be propagated over large geographical distances with minimal latency. In terms of the infection model, this means that once a single base station is infected by a packet, all other base stations can be considered infected by the packet. Packets may reach their destinations through any combination of mobile nodes and base stations.

We model the M base stations as one node. We assume that the pairwise meeting times between mobile nodes and the single node in region i is represented by exponentially distributed random variables with mean inter-meeting time $1/(M_i\gamma_i)$. For peer-to-peer traffic, the differential equations governing the dynamics of the packet spread in the network is given in Equation 5. The rate of change of $P(t)$ is given in Equation 6.

There are two differences in the differential equations for relays and base stations. First, we consider a single node for all base stations as described above. Second, we consider a single instance of $y(t)$ since all base stations are infected when a single base station is infected by a packet. The rate of change of the number of infected base stations, $y'(t)$, is given by the sum of the rates of mobile nodes meeting base stations in each region i in the network.

$$\begin{aligned} x_i'(t) &= (\beta_i x_i(t) + \gamma_i M_i y(t))(n_i(t) - x_i(t)) \\ &\quad - \sum_{\forall j \neq i} \mu_{ij} x_i(t) + \sum_{\forall j \neq i} \mu_{ji} x_j(t), \\ y'(t) &= (\sum_{i=1}^{i=k} M_i \gamma_i x_i(t))(1 - y(t)), \\ n_i'(t) &= \sum_{\forall j \neq i} \mu_{ji} n_j(t) - \sum_{\forall j \neq i} \mu_{ij} n_i(t), \\ x_1(0) = 1, &\quad \forall i \neq 1, x_i(0) = 0, y(0) = 0 \\ &\quad n_1(0) = N + 1. \end{aligned} \tag{5}$$

$$\begin{aligned} P'(t) &= \frac{\sum_{i=1}^{i=k} x_i'(t)}{N}, \\ &\quad P(0) = 0. \end{aligned} \tag{6}$$

For mobile node to sink traffic any base station can act as a sink since all of them are connected to a wired infrastructure. We make a simple modification to $P'(t)$ to accommodate traffic to the sink.

$$\begin{aligned} P'(t) &= y'(t), \\ P(0) &= 0. \end{aligned} \tag{7}$$

Now, similar to the relay case, assuming that a fraction f of the traffic is destined to a sink, the expected packet delivery delay is given by $(1-f) \cdot E[T_{pp}] + f \cdot E[T_{ps}]$, where $E[T_{pp}]$ is the expected delivery delay for peer-to-peer traffic and $E[T_{ps}]$ is the expected delivery delay for sink traffic.

4.6 Mobile Network with a Mesh

An alternative to building the network using wired base stations is to build a mesh network. A mesh network has the advantage that it does not require wired connectivity at every node, although it often requires line-of-sight links between the nodes and incurs higher delays. The mesh can be built over a high-bandwidth, short-range radio, such as 802.11 or CC2420 [21] or a long-range, low-bandwidth radio, such as Digi-XTend [22] (as in our deployment). Note that the mesh nodes need to be placed in wireless range of each other to form a connected network. However, for simple mobility models, such as random way point and random direction, the mean pairwise meeting rate/intensity for mobile and stationary nodes can be approximated as $\gamma = \frac{2r}{E[V^{-1}]} \int_A f(x,y) g(x,y) dx dy$ [14]. The closely placed mesh nodes only affect $g(x,y)$, the density of the locations of the stationary nodes, and the value of γ.

We can model the wireless mesh as a special case of a relay network. The difference between a relay network and a mesh network is that once a mesh node is infected by the packet, all nodes in a mesh can be infected by the packet using the mesh radio. This rate of infection among mesh nodes, denoted by α, depends on the bit rate of the mesh radio and the number of packets transferred simultaneously over a mesh link. The ODEs governing the rate of change of $x_i(t)$, $P(t)$, and $n_i(t)$ are the same as the relay case. The difference lies in the differential equations governing the rate of infection among the stationary mesh nodes in region i ($y_i'(t)$).

Parameter	region 1 (mins)	region 2 (mins)
mobile-mobile	434	506
mobile-base station (uniform)	229	265
mobile-base station (non-uniform)	217	261
mobile-relay (uniform)	239	262
mobile-relay (non-uniform)	236	260

Figure 16: The average pairwise inter-meeting times for mobile-mobile contacts and mobile-stationary contacts.

$$y_i'(t) = \gamma_i x_i(t)(M_i - y_i(t)) + f(\alpha, y_1(t), ..., y_k(t), M_i) \tag{8}$$

The function f depends on the topology of the mesh. For example, if we consider a clique, f is the following.

$$f(\alpha, y_1(t), ..., y_k(t), M_i) = \alpha(\sum_{i=1}^{i=k} y_i(t))(M_i - y_i(t)) \tag{9}$$

The above definition of f means that the rate of infection among mesh nodes in region i is equal to the rate at which the infected mesh nodes (in the entire network) infect the nodes in region i.

4.7 Model Validation

We validate our analytical model by comparing it to the results of simulations using traces from the deployment. We fix the number of mobile nodes to 20, the average number of buses running per day in our deployment. We vary the number of stationary nodes from 2 to 25, the total number of stationary nodes in the deployment. To mirror our deployment scenario, we set the number of regions to $k = 2$. We ran experiments for uniform and non-uniform placement of stationary nodes. The parameters for the model are calculated from our deployment and are shown in Figure 16. The mean pairwise inter-meeting rate for mobile-mesh contacts is similar to that for mobile-relay contacts. Note again that this is the pairwise meeting rate between two nodes, which is different from the rate at which a node meets any other node. The mean rate at which mobile nodes move from region 1 to region 2 is once every 53 minutes and from region 2 to 1 is once every 46 minutes. The fraction of traffic to the sink is set to $f = 0.5$. For our simulations, we use RAPID as our routing protocol that has been shown to perform close to optimal and therefore is a close match for the ideal infinite bandwidth routing scheme assumed in our model.

The average packet delivery delay from the model and the deployment for uniform and non-uniform placement of nodes is shown in Figures 13 and 14. The number of packet copies calculated from the model and the deployment for uniform placement is shown in Figure 15. From the figures, we find that there is close confirmation between the model and the deployment results. The difference between the model and the deployment is less than 10% in most cases, quantitatively validating the model. However, the trace-driven simulation shows lower delays for smaller number of base stations and mesh nodes. This is because the model makes an implicit assumption that the pairwise inter-meeting times do not change with the number of stationary nodes. However, for small number of nodes placement constraints cause the model to deviate from this assumption.

5. RESULTS FROM THE MODEL

In this section, we use the analytical model to understand how the performance of a hybrid network *scales* with the number of mobile and stationary nodes.

5.1 Network size

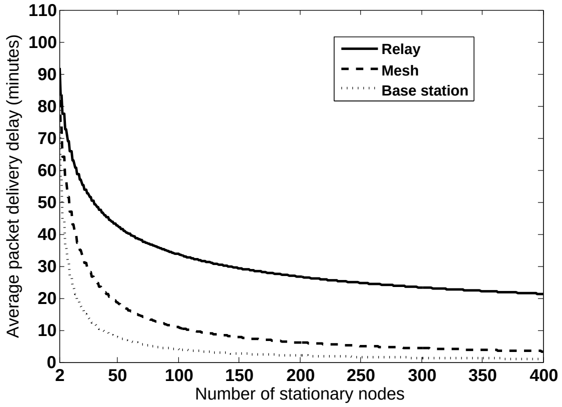

Figure 17: **The average packet delivery delay as a function of the number of stationary nodes. The number of mobile nodes is fixed at 20 and the network parameters are taken from our deployment.**

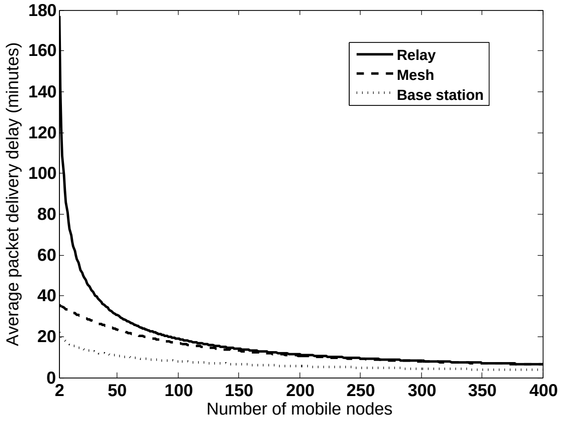

Figure 18: **The average packet delivery delay as a function of the number of mobile nodes. The number of stationary nodes are fixed at 25 and the network parameters are taken from our deployment.**

We use the model to evaluate the performance of hybrid networks with large numbers of stationary and mobile nodes. We perform two experiments with the model. We fix the number of mobile nodes in the first experiment to 20 and vary the number of stationary nodes. In the second experiment the number of stationary nodes is fixed at 25 and we vary the number of mobile nodes. The model parameters are taken from the uniform placement of stationary nodes in our deployment and $f = 0.5$.

Figures 17 and 18 show the average packet delivery delay for the first and second experiment. From the figures, we draw the following conclusions. (1) With a very large number of stationary nodes, we obtain very low delays. However, if we desire delays of the order of seconds (e.g., for an instant messaging application) we need a very large number of base station nodes. (2) To reduce packet delivery delay by a factor of two we need around 5–7 times as many relays and 2–3 times as many mesh nodes as base stations. (3) Adding mobile nodes to a network produces substantial improvement in performance of a relay network but only a small improvement in the performance of either a mesh or base station network. This result confirms that mobile-to-mobile routing leads to marginal improvement in the presence of base stations or a mesh.

5.2 Asymptotic Results

We are interested in understanding how the number of base stations and relays should scale with the number of mobile nodes in order to obtain a substantial improvement in performance. For example, do the stationary nodes need to grow super-linearly or sub-linearly in the number of mobile nodes for large performance improvements? To answer the above question, we consider a special case of our model. We assume that there is only one region in the network and all traffic is peer-to-peer. The pairwise meeting time between mobile nodes are represented by exponentially distributed random variables with mean inter-meeting time $1/\beta$. The pairwise meeting time between mobile and stationary nodes are represented by exponentially distributed random variables with mean inter-meeting time $1/\gamma$. The differential equations governing the spread of a packet in a relay network is shown below.

$$\begin{array}{rcl} x'(t) & = & (\beta x(t) + \gamma y(t))(N - x(t)), \\ y'(t) & = & \gamma x(t)(M - y(t)), \\ P'(t) & = & (1 - P(t))(\beta x(t) + \gamma y(t)), \\ & x(0) = 1, \quad y(0) = 0. & \end{array} \tag{10}$$

The expected packet delivery delay $E[T_d]$ is given by $\int_0^\infty (1 - P(t))dt$. Solving for $P'(t)$ we can show the following:

$$E[T_d] \quad = \quad \int_0^\infty e^{-\int_0^t (\beta x(t') + \gamma y(t'))dt'} dt \tag{11}$$

For large values of M and N, $f(t) = e^{-\int_0^t (\beta x(t') + \gamma y(t'))dt'}$ is small except around $t' = 0$ [14]. Hence, we approximate $x(t')$ and $y(t')$ by a Taylor expansion around $t' = 0$. Using an expansion up to degree three we can approximate $E[T_d]$ as follows.

$$\begin{array}{rcl} E[T_d] & = & \sqrt{\dfrac{\pi}{2 \cdot (\beta^2(N-1) + \gamma^2 M)}} \\ & & + \dfrac{\beta^3(N-1)(N-2) + \gamma^2 M(2\beta(N-1) - 1)}{3(\beta^2(N-1) + \gamma^2 M)^2} \end{array} \tag{12}$$

Now if $M = o(N)$, $E[T_d] = (1/\beta)(\sqrt{\frac{\pi}{2N}} + 1/3) = \Theta(1/\sqrt{N})$ and if $M = \omega(N)$, $E[T_d] = \sqrt{\frac{\pi}{2M\gamma^2}} = \Theta(1/\sqrt{M})$. *Hence, the number of relay nodes must grow super-linearly in the number of mobile nodes for relays to have substantial effect on the average packet delivery delay.*

For base stations, the differential equations governing the spread of packets among nodes in the network are given below.

$$\begin{array}{rcl} x'(t) & = & (\beta x(t) + \gamma M y(t))(N - x(t)), \\ y'(t) & = & \gamma M x(t)(1 - y(t)), \\ P'(t) & = & (1 - P(t))(\beta x(t) + \gamma M y(t)), \\ & x(0) = 1, \quad y(0) = 0. & \end{array} \tag{13}$$

Using the same technique as the relay case, we can derive the following expression for $E[T_d]$.

$$E[T_d] = \sqrt{\frac{\pi}{2 \cdot (\beta^2(N-1) + \gamma^2 M^2)}} + \frac{\beta^3(N-1)(N-2) + \gamma^2 M^2(2\beta(N-1) - \gamma M)}{3(\beta^2(N-1) + \gamma^2 M^2)^2} \tag{14}$$

If $M = o(\sqrt{N})$, then $E[T_d] = \sqrt{\frac{\pi}{N}} + 1 = \Theta(1/\sqrt{N})$. However, if $M = \omega(\sqrt{N})$, then $E[T_d] = (1/(\gamma M)) = \Theta(1/M)$. Therefore, the number of base stations should grow as $\omega(\sqrt{N})$ to have a substantial effect on the packet delivery delay. Stated another way, if we are trying to augment a base station network with mobile nodes, and N grows as $o(M^2)$, the mobile nodes do not effect the performance of the network. The use of mobile nodes to augment a base station network is futile unless we have a very large number of mobile nodes. This result confirms with our deployment results where we showed that adding mobile nodes does not improve the performance of base station hybrid networks.

These asymptotic results for base stations are similar to those by Liu et al. [19] on throughput capacity of *non-mobile* hybrid wireless networks. Comparing their results with ours, we can argue that node mobility only has a marginal effect on performance of base station networks.

The asymptotic behavior of mesh hybrid networks depend on the value of α and the topology of the mesh. If packets can quickly spread across the mesh (i.e., a large α and high average degree of nodes), they would be similar to base stations. However, if the spread of packets among mesh nodes is slow, the asymptotic behavior would be similar to relay networks.

5.3 Model Improvements

Our analytical model has two limitations. (1) It assumes that all packet replicas are removed once the packet is delivered to the destination. (2) It assumes that the contacts are long enough for all packets in a node's buffer to be transferred. We can overcome the first limitation by introducing two random variables R_i^M and R_i^N that tracks the number of recovered mobile nodes and recovered stationary nodes in region i. We can then model recovery schemes like VACCINE or IMMUNE_TX [26]. It is more involved to address the second shortcoming. We conjecture that we can model each node as a queue for packets and use simple priority schemes like FIFO (first-in-first-out) to determine which packets would be serviced at a node during a contact. However, we leave this complexity to future work.

5.4 Discussion

In this paper, we used average packet delivery delay to evaluate infrastructure-enhanced mobile networks. Though average delay is an important metric for building applications, other metrics like network throughput and loss probability could also be used to evaluate performance. Network throughput, and packet loss probability are especially important when the number of mobile and stationary nodes are very large. Though in this paper we do not study the effect on these metrics, we believe that our work would be a starting point for future research on analyzing these metrics in hybrid mobile networks.

5.5 Cost-Benefit Analysis

In this section, we discuss the cost and benefit of each type of infrastructure in our vehicular network. Unfortunately, their cost varies depending on the deployment environment. For example, a base station network in a vehicular deployment has a different cost than a base station network in an underwater deployment.

Our relay node (in Figure 2) costs less than US$800 as a custom-built device. With a Digi-XTend radio our mesh node costs less than US$1000. Note that the use of a gumstix platform instead of a Stargate can bring down the cost to US$300 and US$500 for the relay and mesh, respectively. Although an inexpensive indoor base station might cost as low as US$100, our experience shows that an outdoor weather-proof base station with wiring and installation typically costs as much as US$3,000. Other studies have shown that cost of wiring could be as high as US$5,000 per base station [17]. Recall that for similar performance improvements we require 5–7 times as many relays and 2–3 times as many mesh nodes as base stations. Therefore, in several situations using relay or mesh nodes is more cost effective than base stations.

6. RELATED WORK

There is a large body of research on analyzing hybrid networks. These include studies on the effect of placing a sparse set of well connected base stations in an ad hoc wireless network [19], improving performance of sparse mobile networks using autonomous agents [10], and adding relays in a purely mobile network [4]. Measurement studies on throughput capacity of vehicular networks show the feasibility of Internet-based applications from mobile nodes [7]. Similarly, systems of base stations, called Infostations, have been designed to provide intermittent coverage and connectivity in mobile networks [8]. Although, augmenting a mobile network with stationary nodes is a well studied area, there is little or no work that analyzes different hybrid network configurations under one unified framework. Our study offers a general analytical model and a more constrained real-world deployment study comparing the performance of different hybrid networks.

The analytical model in the paper derives from a body of work on using Markov chains to model mobile networks. The ODE model as a fluid limit of Markov chains was first introduced to study epidemic routing in sparse mobile networks [26]. Markovian models have been used to study various routing protocols: epidemic routing, 2-hop routing [12], and Spray and Wait [24]. More fundamental work on modeling inter-meeting time between nodes following common mobility models was performed by Kurtz [18].

A recent work by Ibrahim et al. [14] uses the Markov model to analyze a single region network with untethered relays and relays connected through a wired infrastructure. That paper derives asymptotic expressions using the fluid limit of Markov chains for a simple MTR routing protocol. We build on their work and present a more general model for epidemic routing for mobile networks with infrastructure based on ordinary differential equations. We argue that epidemic routing is more general than MTR routing since most routing protocols for sparse mobile networks are variants of epidemic routing. Moreover, their work lacks a deployment which is important for understanding the effect of practical issues like dynamic routing protocols, node placement, and real world propagation. More recent work by Balasubramanian et al. [2] confirms our results that a significantly sized base station infrastructure can quickly overwhelm the benefits of add mobile-to-mobile DTN routing in the context of web search and retrieval on DieselNet.

7. CONCLUSIONS

We have performed an experimental and analytical study of mobile networks enhanced with relays, meshes, and wired base stations. We first explore the trade-offs with each type of infrastructure experimentally in the context of a deployed vehicular testbed. However, since the number of mobile and stationary nodes in the deployment is small, we complement this effort by developing analytical models of large-scale networks in the presence of different infrastructures.

Based on the model and the deployment, our study draws three main conclusions. (1) We need less than 5–7 times as many relays and 2–3 times as many mesh nodes as base stations for a similar enhancement in performance. Considering the high cost of deploying base stations, it is often a better choice to deploy untethered relays or mesh nodes; however, when extremely small delays are sought, base stations are required. (2) The addition of infrastructure often obviates the need for mobile-to-mobile routing. For example, simple two-hop forwarding provides excellent performance for base station and mesh networks. (3) The number of mobile nodes required to influence performance given a fixed number of base stations is very large. From these results, we can conclude that a small amount of infrastructure in certain cases is vastly superior to even a large number of mobile nodes capable of routing data to one another.

Acknowledgments

We are thankful to the anonymous reviewers for their insightful comments and suggestions. This work was funded by National Science Foundation grants CNS-0519881, NSF-0133055, CNS-0615075, CNS-0520729, CNS-0447877, ARE award W911NF-07-1-0281, U.S. Army Research Laboratory and the U.K. Ministry of Defense under Agreement Number W911NF-06-3-0001. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation 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 Defense 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.

8. REFERENCES

[1] Balasubramanian, A., Levine, B. N., and Venkataramani, A. DTN Routing as a Resource Allocation Problem. In *Proceedings of ACM Sigcomm* (2007).

[2] Balasubramanian, A., Levine, B. N., and Venkataramani, A. Enabling Interactive Applications in Hybrid Networks. In *Proceedings of ACM Mobicom* (September 2008).

[3] Balasubramanian, A., Zhou, Y., Croft, W. B., Levine, B. N., and Venkataramani, A. Web Search From a Bus. In *Proceedings of ACM Workshop on Challenged Networks (CHANTS)* (September 2007).

[4] Banerjee, N., Corner, M. D., and Levine, B. N. An Energy-Efficient Architecture for DTN Throwboxes. In *Proceedings of Infocom* (May 2007).

[5] Bruno, R., Conti, M., and Gregori, E. Mesh Networks: Commodity Multihop Ad hoc Networks. *IEEE Communications Magazine 43*, 3 (March 2005).

[6] Burgess, J., Gallagher, B., Jensen, D., and Levine, B. N. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. In *Proceedings of Infocom* (April 2006).

[7] Bychkovsky, V., Hull, B., Miu, A., Balakrishnan, H., and Madden, S. A Measurement Study of Vehicular Internet Access Using in situ Wi-Fi Networks. In *Proceedings of Mobicom* (2006).

[8] Cavalcanti, D., Sadok, D., and Kelner, J. Mobile Infostations: A Paradigm for Wireless Data Communications. In *Wireless and Optical Communications* (2002).

[9] Chaintreau, A., Hui, P., Crowcroft, J., Diot, C., Gass, R., and Scott, J. Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms. In *Proceedings of Infocom* (2006).

[10] Dousse, P. O., and Hasler, M. Connectivity in Ad-hoc and Hybrid Networks. In *Proceedings of Infocom* (2002).

[11] Goodman, D. J., Borras, J., Mandayam, N. B., and Yates, R. D. INFOSTATIONS: A New System Model for Data and Messaging Services. In *Proceedings of Vehicular Technology Conference* (Phoenix, AZ, May 1997), pp. 969–973.

[12] Groenevelt, R., Nain, P., and Koole, G. The Message Delay in Mobile Ad hoc Networks. In *Proceedings of Performance* (2005).

[13] Hull, B., Bychkovsky, V., Chen, K., Goraczko, M., Miu, A., Shih, E., Zhang, Y., Balakrishnan, H., and Madden, S. CarTel: A Distributed Mobile Sensor Computing System. In *Proceedings of Sensys* (October 2006).

[14] Ibrahim, M., Hanbali, A. A., and Nain, P. Delay and Resource Analysis in MANETs in Presence of Throwboxes. In *Proceedings of International Symposium on Computer Performance, Modeling, Measurements, and Evaluation* (2007).

[15] Jain, S., Fall, K., and Patra, R. Routing in a Delay Tolerant Network. In *Proceedings of ACM Sigcomm* (August 2004).

[16] Juang, P., Oki, H., Wang, Y., Martonosi, M., Peh, L., and Rubenstein, D. Energy-Efficient Computing for Wildlife Tracking: Design Tradeoffs and Early Experiences with ZebraNet. *SIGOPS Oper. Syst. Rev. 36*, 5 (2002), 96–107.

[17] JupiterResearch. Municipal Wireless: Partner to Spread Risks and Costs While Maximizing Benefit Opportunities. Tech. rep., Jupitermedia Corporation, 2005.

[18] Kurtz, T. J. Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Process. In *Journal of Applied Probabilities* (1970).

[19] Liu, B., Liu, Z., and Towsley, D. On the Capacity of Hybrid Wireless Networks. In *Proceedings of Infocom* (2003).

[20] Marn, W., Rice, J., Fletcher, C., Creber, R., Babicz, R., and Rogers, K. The Evolution of Radio/Acoustic Communication Gateway Buoys. In *Proceedings of IEEE Oceans* (2005).

[21] Polastre, J., Szewczyk, R., and Culler, D. Telos: Enabling Ultra-low Power Wireless Research. In *Proceedings of IPSN/SPOTS* (April 2005).

[22] Radio. 9XTend^{TM} OEM RF Module, Product Manual.

[23] Rice, J. Seaweb Acoustic Communication and Navigation Networks. In *Proceedings of International Conference Underwater Acoustic Measurements: Technologies & Results* (July 2005).

[24] Spyropoulos, T., Psounis, K., and Raghavendra, C. S. Spray and Wait: An Efficient Routing Scheme for Intermittently Connected Mobile Networks. In *Proceedings of ACM Workshop on Delay Tolerant Networking* (2005).

[25] Zhang, X., Kurose, J., Levine, B. N., Towsley, D., and Zhang, H. Study of a Bus-Based Disruption Tolerant Network: Mobility Modeling and Impact on Routing. In *Proceedings of Mobicom* (September 2007).

[26] Zhang, X., Neglia, G., Kurose, J., and Towsley, D. Performance Modeling of Epidemic Routing. In *Proceedings of IFIP Networking* (2006).

[27] Zhao, W., Chen, Y., Ammar, M. H., Corner, M., Levine, B. N., and Zegura, E. Capacity Enhancement using Throwboxes in DTNs. In *Proceedings of IEEE International Conference on Mobile Ad hoc and Sensor Systems (MASS)* (Oct 2006).