

Delay-Tolerant Communication using Mobile Robotic Helper Nodes

(Invited Paper)

Daniel Henkel

Interdisciplinary Telecommunications Program
University of Colorado at Boulder
henk@colorado.edu

Timothy X Brown

Electrical and Computer Engineering
University of Colorado at Boulder
timxb@colorado.edu

Abstract—Delay tolerant networks can exhibit long phases of complete unconnectedness, where several nodes will not be able to communicate effectively. In those circumstances controlling the movement of helper nodes, that facilitate message exchanges instead of relying on direct communication, can improve network performance. We investigate when controlled movement is beneficial and how helper nodes need to move or position themselves to achieve maximum network performance. Three communication modes on a single link and in a hub model are analyzed: direct communication, communication through one or more relays, and communication via data ferries. Hereby we extend the traditional in-range, out-of-range communication model to include a distance dependence which gives a differentiated rate-distance profile. Further, we characterize controlled mobility of mobile data ferries to form static relay paths or dynamically ferry data between nodes. Achievable throughput, delay, and distance performance of the link modes change with channel and ferry characteristics. The union of these performance regions forms the achievable communication space for a given DTN scenario. Novel phase plots show when it is feasible to use one of the mobility schemes versus direct communication. The implications of these modes for practical link-layer designs are discussed using a real-world DTN example of sensor data collection with unmanned aircraft.

I. INTRODUCTION

Delay tolerant networks (DTN) typically try to cope with long delays and asymmetrical connections by passively waiting for the next opportunity to transmit data. Links, however, can also be established pro-actively by using *mobile helper* nodes which can be moved in a controlled fashion. We distinguish between two modes of controlled mobility which augment direct communication. In *relay communication* helper nodes are moved to intermediate spots to serve as relays and to establish temporary multihop links as in an ad hoc network. In *ferry communication* the helper nodes physically carry traffic between source and destination. Both modes are distinct from direct communication, where node pairs are typically close and communicate directly with each other.

Communication performance varies with distance. In principle, direct communication can always be used. However, at long distances direct communication has low performance and the other modes may be able to provide higher throughput and lower delay. The best mode depends on communication requirements, distance, and the mobility scheme used. We consider networks that operate in all three modes. Such a

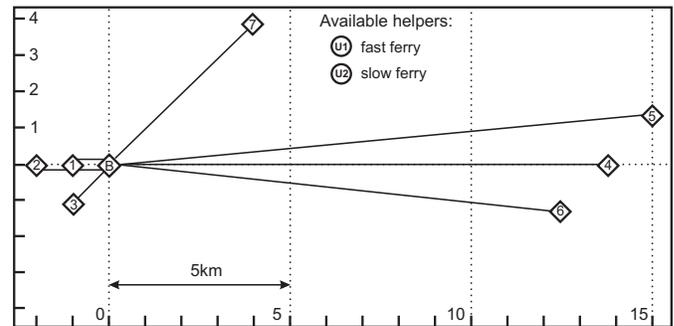


Fig. 1. A sensor data collection problem. Base, B, must collect 100kbps from each of 7 widespread sensors. The mobility of two unmanned aircraft (U1, U2) can be controlled to support the communication.

network can occur in sensor data collection, which we have investigated in prior research [1]. In sensor networks, sensors may end up in close proximity to the data sink or be far flung and require external agents such as an overhead unmanned aircraft to provide connectivity. Fig. 1 shows an example network where 7 widespread sensors must each send 100 kbps of data to a base B (indicated by the thin lines). Direct communication can not carry the traffic since distances are too long. Two unmanned aircraft, i.e., ferries of different airspeeds of 50 m/s (U1) and 10 m/s (U2), are available to support the communication but how to use them is not obvious. For instance, using parameters detailed in a later example, there is no single ferry tour of the nodes that can support the traffic flows. However, by controlling helper mobility and selecting appropriate links (i.e., direct, relay, or ferry) a DTN can proactively construct a network that can address this problem. The goal here is to understand the relationship between *distance dependent communication* and *controlled mobility at the link layer*. This understanding provides a basis for designing higher level DTN communication functions.

Mobility in wireless networks has been well studied in the literature. Most of this research, however, focuses on how to adapt traditional protocols to cope with mobility external to the communications. The goal is to detect topology changes and adjust communication parameters to improve system performance. There is prior research that aims to improve routing [2], make networks more secure [3], save energy [4],

or improve network throughput [5] by *reacting* to movement and without *proactively* controlling node mobility. A proposed “space-time routing framework” utilizes knowledge about future locations of nodes to make this carry-or-forward decision with the goal to minimize packet delay [6].

Others have started to consider the role of controlled mobility where nodes can deliberately change their trajectories to specifically aid the communication [7], [8]. The notion of the *support of the network* with special helper nodes has been introduced to decouple a node’s primary task from the communication support [9]. Controlled mobility together with some nodes physically carrying data can become an important networking paradigm since it helps overcome network partitioning [10]–[13] and it performs well in high-interference environments or jamming situations where traditional approaches to wireless networking fail [14]. This leads to a trade-off between better communication and higher mobility cost, which can be optimized [15]. Some papers analyze mobility support in the network layer and above while mobility support in lower layers is assumed [16], [17]. The scheduling of helper nodes strongly determines network performance. Some work has examined route designs for helper nodes to implement QoS schemes [17], [18], to lower the average packet delay [19], to prevent packet loss through buffer overflows [20], to extend the network lifetime by minimizing energy consumption [21], or to restore network connectivity after network element failures [14], [16].

The special case of controlled mobility of helper nodes to improve the performance of delay-tolerant networks using unmanned aircraft as communication helpers is seen by many as a promising technique [22], [23]. In addition, combining mobility and distance-dependent communication can lead to better-performing mobile networks.

Relay-based multi-rate communication has been studied in infrastructure wireless LANs. A proposed Relay-Based Adaptive Auto Rate protocol helps a fringe node find suitable relays for higher throughput routes [24]. In a similar way a relay-enabled DCF coordination scheme for wireless ad hoc networks has been studied. This protocol is shown to provide greatly improved system performance when channel quality is poor [25]. These solutions specifically focus on studying the 802.11 MAC access details to find improvements.

A common communication model is the *disc model* where two nodes are either out of range and can not communicate or they are in range and can communicate at some fixed rate. In this case, the correct mode is trivial, communicate directly at the fixed rate if the receiver is within the disc radius, relay at half the rate and twice the delay if within two radii, or ferry at up to half the rate otherwise. The disc model does not take into account the distance dependence of signal strength and the resulting distance dependence of the communication rate. Multiple data rates are common in many digital wireless interfaces such as IEEE 802.11 [26]. When communicators are close the communication rate is high (e.g. 54 Mbps in 802.11g). When communicators are far the communication rate is lower (802.11g has 14 different rates between 54 and 1

Mbps). More generally we can have an arbitrary monotonically decreasing rate versus distance function.

What is missing is a unifying view of controlled mobility’s potential assuming variable data rate links. This paper describes the performance gains in choosing the correct link layer communication mode and in proactively controlling mobility where possible. In Section II we introduce the analytic models for each mode. In Section III we contribute a fundamental understanding of the tradeoffs between different modes at the link layer. In particular, we derive performance upper bounds for the different modes with and without mobile nodes. From these we develop a notion of a phase plot that indicates the feasible distance, rate, and delay combinations. Section IV shows how the three modes could be deployed in a network example for sensor data collection. Finally, we describe our first results implementing a delay-tolerant data collection system in a testbed with unmanned aircraft (UA).

II. MODELS

Consider communication on a single link between a sender and receiver. The nodes may be fixed or mobile but their mobility is outside the control of the link layer and so it is not considered. The only significant geometry is their separation distance. The sender has some traffic flow it wants to send. Sender and receiver discovery and the general problem of link establishment are not trivial. Nevertheless we assume that this can be done and focus on the modes of communication that can be used. This section explores different modes that combine distance dependent communication with the use of additional helper nodes. The best mode according to some performance metric is chosen. To enable this comparison, this section provides performance models for the different communication modes.

A. Distance Dependent Direct Links

As the distance between two nodes increases, the expected signal to noise ratio (SNR) decreases. The SNR can be empirically modeled as [27]

$$S(d) = \frac{K}{d^\epsilon}, \quad (1)$$

where d is the transmitter-receiver separation, K is a constant related to radio parameters, and ϵ is the path-loss exponent, usually between 2 and 4¹. At this point we could consider a number of models that predict the SNR at the receiver based on site-specific information, such as terrain layout and obstacles [28]. Empirical models have been widely used in the cellular industry for propagation predictions. In order to make concrete statements we will use (1) to derive our results in the remainder of the paper.

For calculating the link throughput for a given SNR value we are assuming a function $R_D(\text{SNR})$ which is monotonically decreasing with decreasing SNR values. In practice, the

¹Very short distances have a different model than longer distances. The model used here generally over estimates the SNR at these short distances. However, the performance of the system is dominated by the characteristics at longer distances and so we will ignore this difference.

relationship between SNR and throughput is fairly complex. However, a number of models have been studied in the research literature. For instance, an exponential model of the form $R = R_{max} (1 - e^{-A_\epsilon(SNR - SNR_0)})$ can be derived from empirical measurements through curve fitting techniques [29].

For our studies we have chosen the well-known Shannon capacity [30] which relates the SNR to the maximum reliable communication capacity: $C = W \log_2(1 + \text{SNR})$, where W is the total channel bandwidth in Hertz. This provides an upper bound to what can be achieved in an actual system. This bound has been closely approached in a number of practical systems. It would typically be achieved in discrete rate steps as is the case with 802.11g. We will combine the Shannon capacity with the SNR model in (1) as a proxy for the data rate that can be achieved as a function of distance in a direct communication:²

$$R_D(d) = W \cdot \log_2 \left(1 + \frac{K}{d^\epsilon} \right). \quad (2)$$

Studies show that in practice a monotonically decreasing function of throughput over distance can be expected. This is the basis of our model, and the theoretical results in later sections stay qualitatively unchanged as long as any function $R_D(d)$ satisfies this requirement. Thus, it would be possible to derive detailed results for an empirically determined approximation to (2) for a given scenario.

In addition to the data rate, we will also consider the delay in delivering a message. The delay in delivering a message of length L over a distance d is assumed to be only a function of the insertion time, i.e., propagation delay and channel access delay are assumed to be small as compared to insertion delay. For direct communication, the time delay is given as:

$$\tau_D(d, L) = \frac{L}{R_D(d)}. \quad (3)$$

In summary, for the direct mode of communication, the available data rate and time delays are driven explicitly by the SNR, which is taken to be a function of distance only.

B. Relay Helper Model

The previous section introduced a model of direct communication with variable, distance-dependent data rates. We compare this to communication with mobility controlled helpers. We distinguish re-locatable helpers, called *relays*, and mobile helpers, called *ferries*. We first introduce the relay model in this section, then focus on the ferry model in the following section.

A relay can be placed in the middle of a direct link, halving the distance of the wireless transmission, thus increasing the nominal data rate according to (2). However, since the medium is accessed by the sender and then the relay, resulting in two insertion delays, the achievable end-to-end data rate is halved.

²We ignore effects such as shadow fading and Rayleigh fading so as to focus on the main ideas.

The actual data rate for one relay placed in the middle between the communicating nodes, R_R , is thus

$$R_R(d) = \frac{1}{2} R_D(d/2), \quad (4)$$

with d being the separation distance between sender and receiver.

Link performance may improve with multiple relays. For multiple relay communication, we distinguish between single transmission and parallel transmission cases.

1) *Single Transmitter Case:* We can easily show when relaying with one relay is better than direct communication, cf. (13), but beyond this point, how many relays is optimal to serve this link? We first consider the case that one message is forwarded completely from source to destination before the next message is forwarded.

We assume multiple mobile relays move so that they are evenly spaced between transmitter and receiver. With k relays, the total distance is divided into $k+1$ shorter links of distance $d_k = d/(k+1)$, resulting in a larger data rate at each hop. Each link requires one medium access, reducing overall throughput. In this case the data rate and delay with k relays is:

$$R_{RS}(d, k) = \frac{1}{k+1} R_D(d/(k+1)) \quad (5)$$

$$\tau_{RS}(d, L, k) = \frac{(k+1)L}{R_D(d/(k+1))} \quad (6)$$

For large enough d each additional relay initially helps because the rate increases more than the cost of another message store and forward. However, at some point additional relays become counterproductive. To find the optimal number we treat (5) as being continuous in k and minimize to:

$$k_{\min} = \frac{c_\epsilon d}{K^{1/\epsilon}} - 1 \quad (7)$$

where K and ϵ are from (1) and c_ϵ is the solution to the equation $\ln(1 + c^\epsilon) = \epsilon c^\epsilon / (1 + c^\epsilon)$. For $\epsilon \in \{1, 2, 4, \infty\}$; $c_1 = 0$, $c_2 = 1.98$, $c_4 = 2.65$, and $c_\infty = e = 2.72$. By definition the k that maximizes the throughput minimizes the delay.

Equation (7) implies an optimal spacing of relays for a given noise level and pathloss exponent. The relay spacing, d_k , is $d_k = \frac{d}{k+1} = K^{1/\epsilon} / c_\epsilon$. The optimal throughput is $K^{1/\epsilon} / (c_\epsilon d) R_D(K^{1/\epsilon} / c_\epsilon)$. For a single transmitter at a time, the throughput and delay decrease with distance. However at a rate proportional to $1/d$ compared to $1/d^\epsilon$ in the direct case³.

2) *Parallel Transmitter Case:* End-to-end throughput can be increased in certain cases when multiple relays forward traffic simultaneously along the relay path. However, since every actively transmitting relay poses as an interferer to its neighbors, there is a trade-off between the number of simultaneous transmissions and the additional interference they create.

Let ρ be the number of times a message is relayed along a path before the sender can transmit its next message. Now let

³Note: in the bilinear rate model in [29] the same conclusion can be drawn.

us consider multiple nodes in the relay link all transmitting at the same channel. Equation (5) changes to

$$R_{RP}(d, k) = \max_{\rho} \left\{ \frac{1}{\min\{k+1, \rho\}} R_I(d, k, \rho) \right\} \quad (8)$$

where $R_I(d, k, \rho) = W \log_2(1 + P_S/(P_I + P_N))$ is the relay rate with interference and noise. The signal to interference ratio $P_S/(P_I + P_N)$ consists of P_S , the received signal strength, P_I , the received interference power, and P_N , the received noise power. In a k hop relay, each hop is length $d/(k+1)$, and the received signal power is $P_S = \beta(k+1)^\epsilon/d^\epsilon$, where $\beta = KP_N$.

For the interference we sum over all possible interference sources for the worst case relay. The worst case relay is a node near the center of the relay chain since it will have nearby simultaneous interferers on each side. Assuming β as being equal for all nodes we get:

$$\begin{aligned} P_I &= \sum_{i=1}^{i \leq \frac{k-2+\rho}{2\rho}} \frac{\beta}{\left((i\rho-1)\frac{d}{k+1}\right)^\epsilon} + \sum_{i=1}^{i \leq \frac{k-2}{2\rho}} \frac{\beta}{\left((i\rho+1)\frac{d}{k+1}\right)^\epsilon} \\ &= \beta \left(\frac{k+1}{d}\right)^\epsilon s(\rho, \epsilon, k). \end{aligned}$$

where

$$s(\rho, \epsilon, k) = \sum_{i=1}^{i \leq \frac{k-2+\rho}{2\rho}} \frac{1}{(i\rho-1)^\epsilon} + \sum_{i=1}^{i \leq \frac{k-2}{2\rho}} \frac{1}{(i\rho+1)^\epsilon}$$

The achievable data rate with interference is therefore

$$R_I(d, k, \rho) = W \log_2 \left(1 + \frac{1}{s(\rho, \epsilon, k) + K\left(\frac{d}{k+1}\right)^\epsilon} \right). \quad (9)$$

As k increases, the dependence on distance, d , and radio factors, K , goes to zero in (9).

Across a variety of scenarios it was found in [31] that the optimal ρ is almost always 5.

This implies that in the limit of large k :

$$R_{RP}(d, \infty) = \frac{1}{5} W \log_2 \left(1 + \frac{1}{s(5, \epsilon, \infty)} \right).$$

Interestingly, this limit only depends on W and ϵ . Fig. 2 shows the maximal achievable throughput for $W = 1\text{MHz}$. For large ϵ , it can be shown that $R_{RP}(d, \infty) = W0.4\epsilon$. Thus, more attenuation yields greater limiting throughput.

Another noteworthy fact is the dependency of the optimal number of relays on the path loss exponent. For all ϵ below a threshold, ϵ_t , the peak data rate occurs at between seven and nine relays (determined empirically). Adding more relays reduces the data rate and lets it approach the limiting rate from above when $k \rightarrow \infty$. On the other hand, when $\epsilon > \epsilon_t$ the maximum data rate for a given distance can always be increased by adding more relays, and it approaches the limiting rate from below. This ϵ_t depends on the T-R separation distance and decreases with increasing distance. The characteristic ‘‘bend’’ in the graph in Fig. 2 around $\epsilon = 3.8$ marks ϵ_t for $d = 10\text{km}$.

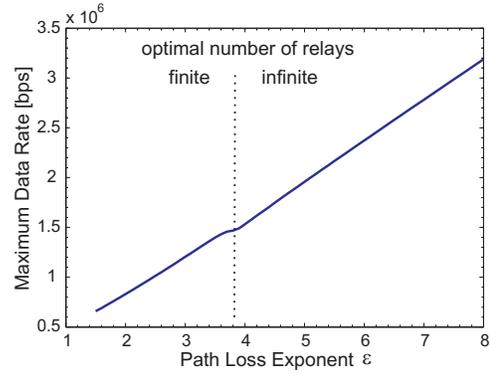


Fig. 2. Maximal achievable throughput over path loss exponent ϵ given optimal ρ for $d=10\text{km}$.

The lowest delay in the parallel transmission case is given by:

$$\tau_{RP}(d, L, k) = \min_{\rho} \frac{(k+1)L}{R_I(d, k, \rho)}$$

Since $R_I(d, k, \rho)$ is monotonically increasing in ρ , its maximum is achieved when $\rho = k+1$. In other words the delay is minimized when there is no interference. This is exactly the single transmitter case in which there is an optimal k that minimizes delay, 7.

C. Ferry Model

A ferry physically carries a buffer of size b between a sender and receiver. In the worst case the ferry must travel the entire distance between the two nodes. However, because it can transfer data while on the move and data rates increase as the distance decreases, the ferry can complete the exchange by traveling only a portion of the distance between the sender and receiver. The optimal strategy for a single ferry is for the ferry to start at the midpoint, move toward the source node until its buffer is half full, then move away from the node continuing to fill the buffer. Upon reaching the midpoint, its buffer is full. Moving onward toward the destination, it starts to unload its buffer, until it has unloaded half of its buffer, then travels away from the destination. Upon reaching the midpoint, the rest of the buffer will have been unloaded and the transfer will be complete. In this way, the ferry only moves as close to a node as is necessary to complete a transfer.

In this paper we are concerned with the region when the ferries and nodes are relatively widely spaced. In this regime, $\text{SNR} \ll 1$ and we can approximate the Shannon capacity by the upper bound:

$$R_D(d) \leq 1.44W \frac{K}{d^\epsilon} = R_D(d_0) \frac{d_0^\epsilon}{d^\epsilon}. \quad (10)$$

where d_0 is a chosen reference distance. Consider two nodes separated by distance d . When a ferry is at distance y from a node and traveling at velocity v , as it moves over a distance dy , it will relay $\frac{R(y)dy}{v}$ bits. So moving from the midpoint to within x of the node it will load:

$$\begin{aligned}
B(x) &= \int_x^{d/2} \frac{R_D(d_0)d_0^\epsilon}{vy^\epsilon} dy = \frac{R_D(d_0)d_0^\epsilon}{v(\epsilon-1)} \left(\frac{1}{x^{\epsilon-1}} - \frac{1}{(d/2)^{\epsilon-1}} \right) \\
&= \frac{R_D(d_0)d_0^\epsilon}{v(\epsilon-1)x^{\epsilon-1}} \left(1 - \left(\frac{2x}{d} \right)^{\epsilon-1} \right).
\end{aligned}$$

Setting $B(x) = b/2$, setting $d_0 = d/2$ and solving for x yields

$$x = \frac{d}{2} \left(\frac{1}{1 + \frac{b(\epsilon-1)v}{R_D(d/2)d}} \right)^{\frac{1}{\epsilon-1}}.$$

The time to complete one load-and-unload cycle is:

$$T_{cycle} = 4 \frac{(d/2 - x)}{v} = \frac{2d}{v} \left(1 - \left(1 + \frac{b(\epsilon-1)v}{R_D(d/2)d} \right)^{-\frac{1}{\epsilon-1}} \right).$$

The time $2d/v$ is the total time for the ferry to travel from one node to the other and back. Of this time, the quantity in the brackets is the fraction required to actually complete the cycle. If the ferry continues this cycle over time, the average data rate is:

$$R_F(d) = \frac{b}{T_{cycle}} = \frac{bv}{2d} \frac{1}{1 - \left(1 + \frac{b(\epsilon-1)v}{R_D(d/2)d} \right)^{-\frac{1}{\epsilon-1}}}. \quad (11)$$

To better see what is going on, we compute upper and lower bounds on this equation. Noting that $1/(\epsilon-1) \leq 1$ since $\epsilon \geq 2$, we can take the first two terms of the binomial expansion of the denominator and simplify to get:

$$R_F(d) \leq \frac{R_D(d/2)}{2} + \frac{vb}{2d}$$

with equality if $\epsilon = 2$ (so-called free-space pathloss). This is an upper bound further since the simplified rate model in (10) is an upper bound on the achievable rate; and the SNR model in (1) is an upper bound on the SNR at short ranges. As a lower bound, the ferry could choose one of two strategies, either to stay in the middle as a relay in which case the rate would be $\frac{R_D(d/2)}{2}$ or it could travel the complete distance back and forth between the two nodes, in which case the rate would be $\frac{vb}{2d}$.⁴ Thus, a lower bound is

$$R_F(d) \geq \max \left\{ \frac{R_D(d/2)}{2}, \frac{vb}{2d} \right\} \geq \frac{R_D(d/2)}{4} + \frac{vb}{4d} \quad (12)$$

where the last inequality uses the fact that the maximum of a set of values is an upper bound for their average. Comparing the upper and lower bounds, the performance of ferrying is determined within a factor of two by quantities which are

⁴We assume that the maximum achievable transfer rate when the ferry is in the vicinity of nodes, R_{\max} , is large enough so that the transfer time at each end can be ignored. If not, then a lower bound strategy would include the flight time to fly back and forth between the two nodes and the transfer time in and out of the buffer of size b in the vicinity of the two nodes. In this case, the right hand side of (12) becomes

$$\frac{R_D(d/2)}{4} + \frac{vb}{4d(1 + \frac{vb}{dR_{\max}})}.$$

readily measured: the data rate at the midpoint; the velocity of the ferry; the size of the ferry buffer; and the distance between the nodes. Note that with sufficiently large buffer size, b , any rate is possible with ferrying (up to $R_{\max}/2$ as described in footnote 4).

The delay with ferrying depends on the model. We will use a simple model in this paper. The sender is assumed to be sending a steady stream of data at rate $R_F(d)$ and the ferry sends the data in a FIFO order. In this case each message spends on average T_{cycle} waiting to be loaded onto the ferry, being carried by the ferry, and waiting to be unloaded by the ferry. Based on (11) and the bounds on the ferry rate, the delay is:

$$\frac{2}{\frac{R_D(d/2)}{b} + \frac{v}{d}} \leq \tau_F(d, L) = \frac{b}{R_F(d)} \leq \frac{4}{\frac{R_D(d/2)}{b} + \frac{v}{d}},$$

where we assume that the ferry buffer is at least the message size ($b \geq L$).

In the case of multiple ferries the main throughput analysis between ferry and end node stays the same. However, two different ferry deployment schemes determine end-to-end performance. In a *conveyor-belt mode* each ferry travels the whole way from source to destination, then returns to the source and repeats the movement in a circuit. In a *chain-relay mode* all ferries are distributed along a connecting path from source to destination forming a chain. The source passes available data on to the first ferry, which physically transports it some way along the path and hands it off to the next ferry. This hand-off procedure is repeated along the chain until the last ferry hands the data to the destination. Detailed description and performance analysis have been published earlier [32].

All of the results in this section assumed a unidirectional flow from one node to the other. Bidirectional flows can also be considered with only minor modifications to the results. We will not consider bidirectional flows here since it does not significantly change the main ideas in the paper. Instead we focus on the roles of the different communication modes in the following sections.

III. SINGLE-LINK ANALYSIS

The previous section derived the data rate and delay for the direct, relay, and ferry communication modes. Consider a pair of nodes which wish to communicate at some rate of at least r and delay less than t . The question is what modes, if any, can be used to carry this communication given the separation distance between the two nodes? We look at this from two perspectives. One is a rate-distance phase diagram and the other is delay-rate phase diagram.

A. Rate-Distance Phase Diagram

The rate-distance phase diagram graphically represents the achievable link layer modes in a continuous rate-distance space. A given distance-rate pair, (d, r) , can be mapped by a function, M , to one of the three link layer modes, either direct, relay, or ferrying; or unachievable if no mode can deliver data

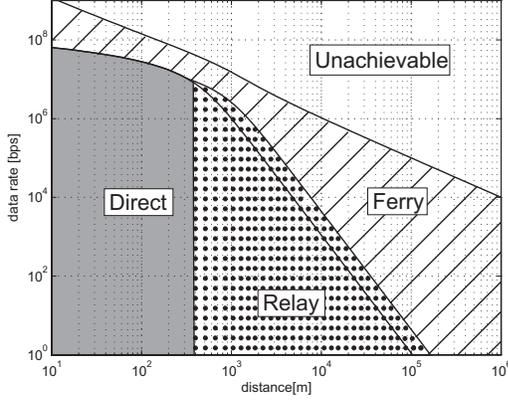


Fig. 3. Phase plot showing the communication mode that minimizes delay as a function of the link distance and the desired data rate.

rate r at distance d ;

$$M : (d, r) \rightarrow m \in \{D, R, F, U\}.$$

For a given distance, the achievable rates and the delays are compared to compute the mode that can achieve the desired data rate with the lowest delay:

$$M(d, r) = m \text{ if } R_m(d) \geq r \text{ and}$$

$$\tau_m(d, L) \leq \tau_{m', L}(d) \text{ for all } m' \text{ where } R_{m'}(d) \geq r,$$

where for the purposes of this definition $R_U(d) = \tau_U(d, L) = \infty$.

An example plot is shown in Fig. 3 based on the following parameters. Consider a radio link that has a rate of 1Mbps at 1km separation, 100Mbps at 1m separation, and the path loss exponent is 3 (i.e., $W = 3.6\text{MHz}$, $K = 2.1 \cdot 10^8$, and $\epsilon = 3$). The bottom curve in Fig. 3 (thin black line) shows the obtainable rates as a function of distance. The allowable rate decreases from 1Mbps at 1km to 1bps at 100km. If a relay node is added at the midpoint, then the obtainable rates are given by the second curve. Below about 400m, relaying does not add to the obtainable rate since the rate achieved by relaying through two shorter links is less than twice the direct link rate. Relaying is used when $R_R(d) > R_D(d)$, which, based on (2) and (4), is when

$$d > \left(\frac{K}{2^\epsilon - 2} \right)^{\frac{1}{\epsilon}}. \quad (13)$$

Instead, if we use helper ferry U1 that can move at 50m/s and has a 50MB buffer, the upper curve in Fig. 3 shows the throughputs that are achievable are further expanded. Decreasing the buffer size or ferry velocity causes the obtainable ferry rates to decrease. For small buffer sizes the ferry throughput curve follows the relay throughput curve until large distances are reached. Rates above the ferry curve are unachievable without increasing the buffer or velocity. In principle, any rate is achievable if the buffer size (or velocity) is not bounded. Examining 25km, the link can sustain a data rate of 70bps using a direct link, 280bps using a relay link and 400,000bps using a ferry link. The ferry in this example has much greater throughput. However, the delay is very high for the ferry. For

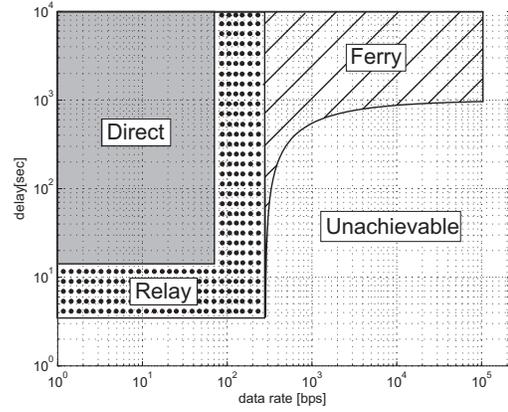


Fig. 4. Phase diagram showing the achievable delay vs. data rate combinations for each mode at a given distance, $d = 25\text{km}$.

a message of length $L = 1000$ bits the delay is about 14sec for direct, 3.5 sec for relay, and 1000 sec for the ferry. Relay has lower delay than direct when $\tau_R(d, L) < \tau_D(d, L)$ which based on the definition is satisfied for the d defined in (13). The ferry delay is greater than the direct and relay link delays over the range of this example. However, at extreme ranges (above 400km,) the ferry has lower delay. According to the definitions, and using the approximation in (10), ferrying has lower delay when

$$d > 2 \left(\frac{1.44WK}{vL} \right)^{\frac{1}{\epsilon-1}}.$$

B. Delay-Rate Phase Diagram

The achievable delay-rate graph shows the combinations of delay and rate (t, r) that are achievable at a given distance d and message size L . This is used when a transmitter and receiver are at a distance d and want to see if a QoS measure for an application (t, r) can be met with one of the communication modes. Given d and L , the set of achievable delay-rate pairs for the direct and relay communication modes is given by:

$$A_m(d, L) = \{(t, r) | t \geq \tau_m(d, L) \text{ and } r \leq R_m(d)\}.$$

where $m \in \{D, R\}$.

The achievable delay rate pairs for ferrying are slightly more complex. The ferry can lower its delay at the expense of lower rate by partially filling its buffer. As a minimum it needs to fill its buffer with a single message L . We implicitly define a function $b(r)$ which is the size of the buffer required to achieve rate r . We use $\tau_F(d, L, x)$ and $R_F(d, x)$ to indicate the delay or rate when the buffer size is x instead of b . With these definitions:

$$A_F(d, L) = \{(t, r) | t \geq \tau_F(d, L, L) \text{ and } r \leq R_D(d, L)\} \\ \cup \{(t, r) | t \geq \tau_F(d, L, b(r)) \text{ and } L \leq b(r) \leq b\}.$$

Fig. 4 shows the delay rate phase plot for $d = 25\text{km}$. A higher layer service could use this plot to choose a mode of communication and possibly adapt the service to the available rates and delays.

C. Disc Model versus Variable Rate Ferrying

It is instructive to compare ferrying with the distance dependent model versus the disc communication model. In the disc model the data rate is R_0 within a disc of radius d_0 and zero otherwise. Substituting these values into (10), $R_D(d/2) = R_0(2d/d_0)^\epsilon$. Substituting this into the single-ferry model yields:

$$R_F(d) \leq \frac{2^{\epsilon-1}R_0d_0^\epsilon}{d^\epsilon} + \frac{vb}{2d}.$$

According to [32] the rate for the disc model is

$$R_F^{dm}(d) = \frac{1}{\frac{2(d-2d_0)}{vb} + \frac{2}{R_0}}.$$

We note that $R_F^{dm} \leq R_0/2$, while R_F is unbounded as v or b increases. In the best case, when $d = 2d_0$, $R_F^{dm} = R_0/2$ which is less than $R_F = R_0/2 + vb/(4d_0)$. While R_F is generally superior, for large d they both have similar throughput of about $vb/(2d)$.

The superiority of the distance dependent rate model is more pronounced in the case of multiple ferries. In the chain relay mode, the effective d is shortened with each additional ferry and the communication is more into the regime where the distance dependent rate method dominates. In the conveyor belt mode, more ferries can always be added by passing closer to the source and destination nodes where the rates are higher. Thus, the variable rate model differs qualitatively from the disc model in that the rate can always be increased by adding more ferries.

IV. MIXED DIRECT, RELAY, AND FERRY HUB

The paper so far has examined the optimal communication mode for a single link. Another scenario for mode adaptation is an access point like setup where k stations transfer data to a single hub node. We assume full knowledge of network parameters like station flow rates f_i , distances d_i from the hub, ferry speed v , and ferry buffer size b . The hub talks to the stations one at a time, i.e., Round Robin fashion, and it can choose between either of the three modes. Nearby stations could use direct communication and further stations could use ferrying, while stations with low delay requirements are best served by relays. The decisions of which mode to use are guided by phase plots similar to Fig. 3 and Fig. 4, but specifically adapted to the scenario. Thus, a hub can use a combination of ferry and direct to achieve a higher throughput than either alone and deploy a relay to guarantee a delay constraint. However, the system must be carefully designed such that the flows can be carried by the specific mode employed. For brevity we are omitting the detailed analysis at this point.

A. Example Scenario

To show how the tools presented in this paper enable a solution to a challenged network, we return to the example in Fig. 1 with the parameters used in Fig. 3. The goal is not to present a general assignment algorithm, but to illustrate how a

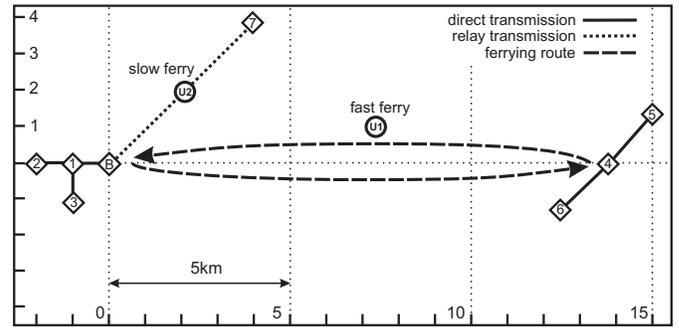


Fig. 5. Solution to Fig. 1. Sensor 1 is a local hub with direct links (thick single arrows). Sensor 4 is a hub with direct links and supported by unmanned aircraft U2 (thin double arrows). Sensor 4 is connected to the base via unmanned aircraft U1.

combination of modes can be used to solve a communication problem. Fig. 5 indicates how the two unmanned aircraft of different capabilities are assigned. Sensor 4 acts as a hub that collects data from sensors 5 and 6 in direct communication fashion. The slow moving ferry U2 is positioned as a relay between the base and node 7. This combination satisfies the hub network's stability criteria and the delay requirement of node 7 according to Fig. 4. The faster ferry U1 is used to carry the collected 300kbps of traffic the 14 km between sensor 4 and the base. Next, sensors 1–3 are in the vicinity of the base. A feasible solution is found if sensor 1 acts as a hub (with 100 kbps coming in from sensors 2 and 3 and 300 kbps sent out to the base). A potential alternate solution to the methods presented here is to have an unmanned aircraft make a ferry tour of the nodes. A minimum tour of the base and the 7 sensors is approximately 50km requiring 1000 sec (with U1 at 50m/s) in which 85 MB of data is generated by the sensors exceeding the ferry's buffer. A pure direct solution is also not possible because of the long distances (see the achievable direct throughputs for distances greater than 10 km in Fig. 3). Therefore, a solution requires mixed communication modes.

The delay to deliver a message from sensor 4 to the base using a ferry across the 14km is on average $2d/v = 560$ sec. If more helpers were available, how small could this be made? The lowest delay is with single-transmission multiple relays.⁵ If, according to (7), $k_{\min} = 58$ helpers are available, and the message size is 1000 bits, then the delay is 4msec. Thus, there is the potential for five orders of magnitude decrease in delay if more helpers become available.

B. Helper Signaling

Helpers are a scarce resource and thus can be seen as a shared medium which needs to be scheduled for task node access. Nodes contend for usage of the shared medium resources. A fundamental problem in delay tolerant networks is scheduling resources when the network is challenged. We

⁵At k_{\min} the throughput is 250kbps which is not enough to carry the 300 kbps of sensors' 4–6 data to the base. However, by introducing a few widely-spaced parallel transmissions (e.g. $\rho = 20$ in (8)) the rate can be made larger than 300 kbps while producing insignificant interference.

sketch a method here to suggest how it would work with controlled mobility.

A simple mechanism notifies the helper of medium access requests. A node that wants to send a message first broadcasts a request for establishment of a link. If it does not get a response, it lowers its data rate, effectively increasing its transmission power, and re-sends the request at decreasing rates until it gets a response from either the destination node or a helper. The request contains the node's location as well as the message's destination ID, and it is overheard by all nodes within the signaling range. If the destination is within this range, and the SINR permits effective direct transmission, the two nodes start communicating in direct mode and a helper is not utilized. If the destination is out of signaling range but a helper overhears the request, it acknowledges the request to the requesting node with an estimated service time. The helper then starts moving towards the requester to act as relay or ferry. If it is a short transfer it receives the transfer while flying towards the requester and when complete immediately flies toward the destination. If it is a longer transfer it evaluates the different modes and chooses the mode that satisfies the request.

V. IMPLEMENTATION IN THE AUGNET TESTBED

To verify some of the theoretical results with real-life networking equipment, we have created an outdoor ad hoc network comprising mobile and stationary nodes. Since the mobile helper nodes are small unmanned aircraft, the facility is called the Ad hoc Unmanned aircraft Ground NETwork (AUGNet). This section describes the initial phase of the project.

Our testbed represents a scenario where an area has intelligence, surveillance, and reconnaissance (ISR), environmental monitoring, or other sensors that can range from small and simple thermometers to large and highly functional RF scanners. A concept diagram of the outdoor test area and possible node constellation is shown in Fig. 6. These sensors need to report occasional data readings in near realtime to one or more sensor monitoring stations (SMS) located at remote locations. In general, these sensors are not assumed to have wireless radios, but could have short range radios in some cases. The sensor interfaces with a stationary MANET node (MN) which in turn delivers the data packets to a gateway node (GW) which connects to an external backhaul network to the SMS. The MANET can have stationary ground nodes to support this communication, but can as well be very sparsely distributed. A ground node typically has short range and can not support sensors over a wide area. Therefore, in order to provide a node capable of long-range communication and also having wide-ranging coverage, unmanned aircraft (UA) are outfitted with radio interfaces.

The goal of the network setup is to provide reliable backhaul of sensor data through the sparsely connected network to the multiple SMS and to enable remote sensor configuration from the SMS. The network should transparently carry traffic to and from the sensors such that sensors and SMS are isolated from

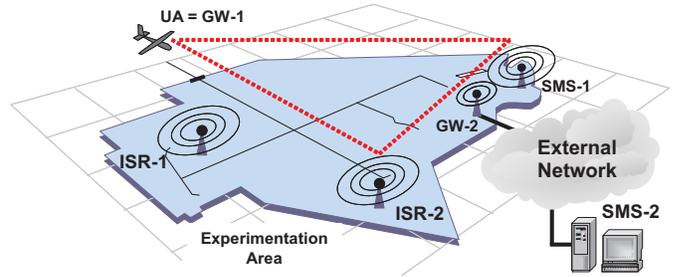


Fig. 6. Sensor Data Collection System Components

the network specifics. The UA can move so that it collects the sensor data according to waypoint plans. The waypoints can be updated through the network. The networking in this project extends availability through ad hoc protocols and uses delay tolerant network concepts when unavailable [33], [34]. Further it implements a full-scale test bed to expose the significant interactions between network protocols, radio propagation, and vehicle mobility constraints [35].

A. Communication Architecture

Two types of communication are supported: delivery of sensor data from the ISR sensors to all active SMS and control communication between a single SMS and ISR. To ensure reliable forwarding through the network, we implemented two complementary techniques, namely custody transfer, i.e., staged delivery, and a reliable forwarding mechanism similar to the R-UDP protocol [36], but stripped of unneeded functional overhead.

The staged delivery of data enables packet transmission in intermittently connected networks without contemporaneous end-to-end connections between ISR and SMS. This is important since the aircraft act as communication helpers and go in and out of a ground node's coverage area for prolonged time periods.

Traversal of different networks is achieved through Network Address Port Translation (NAPT) [37]. In order for the multicast delivery of sensor data to be successful, we implemented two node discovery mechanisms for gateways and SMS. Heartbeat packets are used by GW to keep track of which SMS are active. Gateways advertise their presence via a periodic flood in the MANET.

B. Practical Implementation Details

All ad hoc network nodes are running on Soekris net4511 single board computers with an AMD 100MHz CPU, 64MB RAM, 512MB flash memory cards, and two Ethernet ports. An Atheros-chipset-based miniPCI card provides 802.11abg wireless connectivity, while an RF power amplifier increases the coverage area. A GPS receiver provides location and altitude information to the nodes. The nodes run a custom-built embedded Gentoo Linux with the madwifi-ng driver⁶ operating the Wifi card.

All ad hoc routing and data collection logic was implemented using the Click Modular Router [38] framework. The

⁶Detailed capability description at <http://www.madwifi.org>.

internal network routing is based on our DSR implementation [39] with monitoring function. Gateway advertising, SMS discovery, reliable UDP forwarding, heartbeat protocols, and NAPT are written as modular elements that can be added as needed. Testing takes place in the lab and at our local 7km² test range.

C. Data Forwarding Experiments

Baseline network measurements have been performed in a lab setting. One data packet was generated every second at the ISR sensor, its timestamp and sequence number recorded, sent through the network to the SMS, and from there reflected back to the ISR.

The network test results show a robust forwarding of data packets through the network, with an average round trip time with fully connected network of 41ms from ISR to SMS and back. In a worst-case intermittent scenario the links go down and up in a cascading fashion, and the worst-case RTT can reach the packet retransmission timeout multiplied by number of hosts in the packet’s path plus each link’s disconnection time.

D. Aircraft-Specific Experiments

We have flown an experiment involving one unmanned aircraft, one monitor station, and several well-connected MANET ground nodes. Thus, communication to the aircraft should only temporarily be interrupted due to adverse propagation characteristics, e.g., banking of the plane or interference. The UA is a small (10kg) vehicle built at the University of Colorado and composed of a gateway node, autonomous navigation flight computer, and a scientific sensor payload for temperature, air pressure, and humidity.

The goal of the experiment was to demonstrate an aircraft’s ability to make mission level decisions based upon network communication status. If communication between aircraft and monitor station gets impaired beyond a tolerable level, the plane is supposed to autonomously switch from its current flight plan to an alternate flight plan which restores communication. Two simulated impairments were tested: a gradual drop in signal strength and a complete breakdown of communication. The autonomy of the plane to make independent waypoint control decisions is a critical element of switching between communication modes depending on the observed link characteristics.

Fig. 7 shows the primary results of the experiment. All of the graphs shown depict a value vs. time in minutes. The top graph shows the destination waypoints the plane is following. From start of experiment in minute 58, the aircraft flies on flight plan 1 between waypoints 10–15. As per the second graph, we simulated a gradual decrease in signal strength as measured by the plane. As the signal level drops below 20 dBm in minute 60, the plane switched to a different flight pattern until it reached a location with sufficient signal strength. The bottom graph shows round-trip ping times in milliseconds. At 64min 30sec we simulated a complete loss of communication (stop ping utility,) which prompts the aircraft

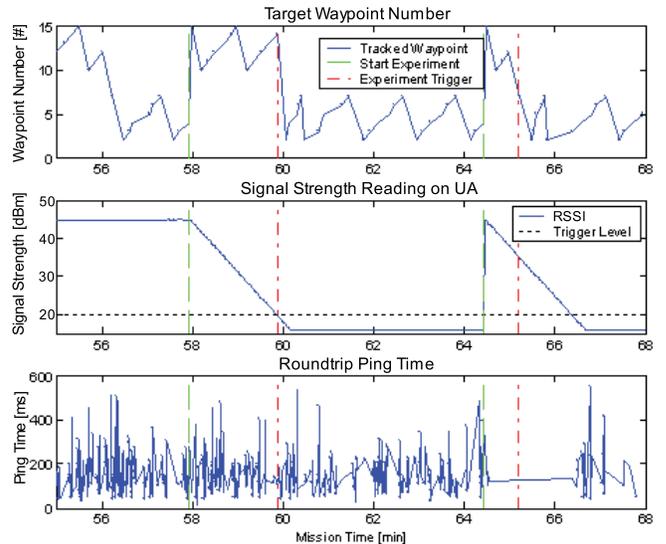


Fig. 7. Results: a. Destination waypoint number b. Recorded onboard temperature c. Recorded ping times from the plane to the monitoring station

to again autonomously switch to a different flight plan after an allowed timeout of 40sec.

VI. CONCLUSION

In this paper we have analyzed the impact of controlled mobility of helper nodes on the performance of delay tolerant networks. In addition to direct communication, we have taken into account two communication modes that rely on deliberate node movement: 1) relay communication using one or more intermediate helper nodes, and 2) controlled mobility of mobile data ferries which physically transport data from source to destination. We have introduced distance dependent data rates in all three modes and single and parallel transmission cases in the relay mode. We have derived performance results for these modes for a single link between source and destination and we have combined this analysis in a data collection scenarios which comprises multiple sensor and helper nodes.

Plotting the performance regions results is the notion of distance-data rate and data rate-delay phase diagrams, which graphically show the trade-offs and boundary conditions of the three communication modes. For a given (distance, data rate)-combination a node pair can choose the lowest-delay mode to optimize its communication. Detailed analysis of the communication modes in terms of throughput and average packet delay show a superior performance of the multi-rate approach compared to the commonly assumed disc model. In modern digital communication this multi-rate model becomes more and more feasible.

Our research indicates that the Ferrying mode can yield substantial throughput gains compared to Direct and Relay if the transmitted data can tolerate longer delays. However, the path taken by the ferry has a major influence on performance. In related work we compute an optimal ferry path using the agent-learning technique of reinforcement learning [40].

The feasibility of delay-tolerant networks in wireless communication has been shown in the literature, and with the possibility of choosing the correct communication mode our research suggests further performance increases. The approach taken in this paper was more theoretical so as to establish the broad performance characteristics. However, our unmanned aircraft testbed enables us to actually implement these algorithms and to conduct some exciting practical work in the future.

ACKNOWLEDGMENT

This work was supported by NSF award CNS 0428887 and AFOSR grant no. FA9550-06-1-0205.

REFERENCES

- [1] A. Jenkins, D. Henkel, and T. X. Brown, "Demo: A delay-tolerant sensor data collection architecture using unmanned aircraft," CHANTS'07, ACM Workshop on Challenged Networks, Montreal, CA, Sept. 14 2007.
- [2] A. Vahdat and D. Becker, "Epidemic routing for partially-connected ad hoc networks," Duke University, Tech. Rep., 2000.
- [3] S. Capkun, J.-P. Hubaux, and L. Buttyan, "Mobility helps security in ad hoc networks," in *Proc. of the 4th ACM International Symposium on Mobile Ad Hoc Networks and Computing (MobiHoc)*, June 2003.
- [4] S. Chakraborty, Y. Dong, D. K. Yau, and J. C. Lui, "On the effectiveness of movement prediction to reduce energy consumption in wireless communication," *IEEE Transactions on Mobile Computing*, vol. 5, no. 2, pp. 157–169, Februar 2006.
- [5] M. Grossglauser and D. N. C. Tse, "Mobility increases the capacity of ad hoc wireless networks," *IEEE/ACM Transactions on Networking*, vol. 10, no. 4, August 2002. [Online]. Available: citeseer.ist.psu.edu/article/grossglauser01mobility.html
- [6] S. Merugu, M. Ammar, and E. Zegura, "Routing in space and time in networks with predictable mobility," Georgia Institute of Technology, Technical Report GIT-CC-04-07, March 2004.
- [7] Q. Li and D. Rus, "Sending messages to mobile users in disconnected ad-hoc wireless networks," in *Mobile Computing and Networking (Mobicom)*, 2000, pp. 44–55. [Online]. Available: citeseer.ist.psu.edu/li00sending.html
- [8] D. Goldenberg, J. Lin, A. S. Morse, B. E. Rosen, and Y. R. Yang, "Towards mobility as a network control primitive," in *Proc. of Mobihoc'04*, ACM, Roppongi, Japan: ACM, May 2004.
- [9] I. Chatzigiannakis, S. E. Nikolettseas, N. Paspallis, P. G. Spirakis, and C. D. Zaroliagis, "An experimental study of basic communication protocols in ad-hoc mobile networks," in *Algorithmic Engineering*, ser. Lecture Notes in Computer Science. London: Springer-Verlag, 2001, vol. 2141, pp. 179–183.
- [10] W. Zhao and M. Ammar, "Message ferrying: Proactive routing in highly-partitioned wireless ad hoc networks," in *the 9th IEEE International Workshop on Future Trends of Distributed Computing Systems*, May 2003.
- [11] W. Zhao, M. Ammar, and E. Zegura, "A message ferrying approach for data delivery in sparse mobile ad hoc networks," in *MobiHoc'04*, May 24-26 2004.
- [12] —, "Controlling the mobility of multiple data transport ferries in a delay-tolerant network," in *Proc. of InfoCom 2005*, 2005.
- [13] R. C. Shah, S. Roy, S. Jain, and W. Brunette, "Data mules: Modelling and analysis of a three-tier architecture for sparse sensor networks," *Ad Hoc Networks Journal*, vol. 1, no. 2-3, pp. 215–233, Sept. 2003.
- [14] P. Basu and J. Redi, "Movement control algorithms for realization of fault-tolerant ad hoc robot networks," *IEEE Network*, pp. 36–44, July/August 2004.
- [15] C. Tang and P. K. McKinley, "iMobif: An informed mobility framework for energy optimization in wireless ad hoc networks," in *Proc. of Second International Workshop on Wireless Ad Hoc Networking (WWAN ICDCSW'05)*, 2005, pp. 939–945.
- [16] Z. Ye, S. V. Krishnamurthy, and S. K. Tripathi, "A framework for reliable routing in mobile ad hoc networks," in *Proc. of INFOCOM 2003*, vol. 1, March 2003, pp. 270–280.
- [17] K. Sanzgiri and E. M. Belding-Royer, "Leveraging mobility to improve quality of service in mobile networks," in *Proc. of MOBIQUITOUS'04*, August 22-26 2004, pp. 128–137.
- [18] M. C. Chuah and P. Yang, "A message ferrying scheme with differentiated services," in *Proc. of IEEE Military Communications Conference (MILCOM 2005)*, October 17-20, 2005.
- [19] Y. Gu, D. Bozdog, E. Ekici, F. Ozguner, and C.-G. Lee, "Partitioning-based mobile element scheduling in wireless sensor networks," in *Proc. 2nd Annual IEEE Conference on Sensor and Ad Hoc Communications and Networks*, Santa Clara, CA, USA, September 26-29, 2005.
- [20] A. A. Somasundara, A. Ramamoorthy, and M. B. Srivastava, "Mobile element scheduling for efficient data collection in wireless sensor networks with dynamic deadlines," in *Proc. of the 25th IEEE International Real-Time Systems Symposium*, 2004.
- [21] Y. Tirta, Z. Li, Y.-H. Lu, and S. Bagchi, "Efficient collection of sensor data in remote fields using mobile collectors," in *Proc. of the 13th International Conference on Computer Communications and Networks (ICCCN'04)*, Chicago, IL USA, October 11-13, 2004.
- [22] A. Ajorlou, A. Homaifar, and A. Esterline, "Robust multi-UAV data collection," in *Proc. of AIAA InfoTech 2007*, Rohnert Park, CA, May 7-10 2007.
- [23] C. Christmann and E. Johnson, "Design and implementation of a self-configuring ad-hoc network for unmanned aerial systems," in *Proc. of AIAA InfoTech 2007*, Rohnert Park, CA, May 7-10 2007.
- [24] J.-S. Liu and C.-H. R. Lin, "A relay-based MAC protocol for multi-rate and multi-range infrastructure wireless LANs," *Wireless Personal Communications*, vol. 34, no. 1-2, pp. 7–28, 2005.
- [25] H. Zhu and G. Cao, "rDCF: A relay-enabled medium access control protocol for wireless ad hoc networks," in *Proc. of INFOCOM 2005*.
- [26] IEEE, "802.11g wireless LAN medium access control (MAC) and physical layer (PHY) specifications—amendment 4: Further higher-speed physical layer extension in the 2.4 GHz band," 2003.
- [27] T. Rappaport, *Wireless Communications: Principles and Practice*. Prentice Hall, 1996.
- [28] G. Durgin, T. Rappaport, and H. Xu, "Measurements and models for radio path loss and penetration loss in and around homes and trees at 5.85 ghz," *Communications, IEEE Transactions on*, vol. 46, no. 11, pp. 1484–1496, Nov 1998.
- [29] C. Na, J. Chen, and T. Rappaport, "Measured traffic statistics and throughput of IEEE 802.11b public WLAN hotspots with three different applications," *IEEE Transactions on Wireless Communications*, vol. 5, no. 11, pp. 3296–3305, November 2006.
- [30] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, July 1948.
- [31] D. Henkel and T. X. Brown, "Optimizing the use of relays for link establishment in wireless networks," in *Submitted to IEEE Wireless Communications and Networking Conference (WCNC'07)*, Hong Kong, March 11–15, 2007.
- [32] —, "On controlled node mobility in delay tolerant networks of UAVs," in *Proc. of the International Symposium on Advanced Radio Technologies '06*, ser. NTIA Special Publications SP-06-438, Boulder, March 7-9, 2006.
- [33] V. Cerf and S. Burleigh, "Delay tolerant networking architecture," IETF RFC 4838, <http://www.ietf.org/rfc/rfc4838.txt>, April 2007.
- [34] K. Fall, "A delay tolerant network architecture for challenged networks," in *Proceedings of ACM SIGCOMM*. ACM, 2003, pp. 27–31.
- [35] T. X. Brown, S. Doshi, S. Jadhav, D. Henkel, and R.-G. Thekketkunnel, "A full scale wireless ad hoc network test bed," in *Proceedings of ISART'05*, ser. NTIA Special Publications SP-05-418, March 2005, pp. 51–60.
- [36] T. Bova and T. Krivoruchka, "Reliable UPD protocol," IETF Internet Draft, February 1999.
- [37] P. Srisuresh and M. Holdrege, "IP network address translator (NAT) terminology and considerations," RFC 2663, IETF, August 1999. [Online]. Available: [ftp://ftp.ietf.org/rfc/rfc2663.txt](http://ftp.ietf.org/rfc/rfc2663.txt)
- [38] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, "The click modular router," *ACM Transactions on Computer Systems*, vol. 18, no. 3, pp. 263–297, August 2000.
- [39] S. Doshi, S. Bhandare, and T. X. Brown, "An on-demand minimum energy routing protocol for a wireless ad hoc network," *Mobile Computing and Communications Review*, vol. 6, no. 3, pp. 50–66, July 2002.
- [40] D. Henkel and T. X. Brown, "Ferry route planning using reinforcement learning," in *submitted to IEEE Workshop on Autonomous Opportunistic Communications, AOC'08*, Newport Beach, CA, June 23, 2008.