

An Energy-Efficient Architecture for DTN Throwboxes

Nilanjan Banerjee Mark D. Corner Brian Neil Levine
 Department of Computer Science, Univ. of Massachusetts, Amherst
 {nilanb,mcorner,brian}@cs.umass.edu

Abstract—*Disruption Tolerant Networks* rely on intermittent contacts between mobile nodes to deliver packets using store-carry-and-forward paradigm. The key to improving performance in DTNs is to engineer a greater number of transfer opportunities. We earlier proposed the use of *throwbox* nodes, which are stationary, battery powered nodes with storage and processing, to enhance the capacity of DTNs. However, the use of throwboxes without efficient power management is minimally effective. If the nodes are too liberal with their energy consumption, they will fail prematurely. However if they are too conservative, they may miss important transfer opportunities, hence increasing lifetime without improving performance.

In this paper, we present a hardware and software architecture for energy efficient throwboxes in DTNs. We propose a hardware platform that uses a multi-tiered, multi-radio, scalable, solar powered platform. The throwbox employs an approximate heuristic for solving the NP-Hard problem of meeting an average power constraint while maximizing the number of bytes transferred by the throwbox. We built and deployed prototype throwboxes in UMassDieselNet – a 40 bus DTN testbed. Through extensive trace-driven simulations and prototype deployment we show that a single throwbox with a 270 cm² solar panel can run perpetually while improving packet delivery by 37% and reducing message delivery latency by at least 10% in the network.

I. INTRODUCTION

Recent efforts in Disruption Tolerant Networking (DTN) promise networking infrastructures that are robust to widespread failures and operate in highly challenged environments [7, 9, 2]. Such networks can be used to route information in completely unconnected and decentralized scenarios, including natural disasters, sparse sensor deployments, underwater networking, and highly mobile systems.

DTNs rely on intermittent contacts between mobile nodes to deliver packets using a store-carry-and-forward paradigm. The routing performance in a DTN, including deliverability and delay, is critically dependent on the node mobility patterns that drive the frequency, duration, and sequence of contact opportunities. One method to improve these metrics is to engineer a greater number of opportunities. For instance, several papers have proposed altering the mobility of nodes to enhance network performance [3, 4, 26, 8]. Unfortunately, movement patterns are often inherent to the nodes and cannot be modified.

Another solution for improving DTN performance is to place additional *stationary* nodes in the network, which increases the number and frequency of contact opportunities. In previous work, we proposed the use of *throwboxes* within a DTN for this purpose [27]. Throwboxes are inexpensive, battery-powered, stationary nodes with radios and storage. When two nodes pass by the same location at different times, the throwbox acts as a router, creating a new contact opportunity. In our previous work, we focused on algorithms that place throwboxes in loca-

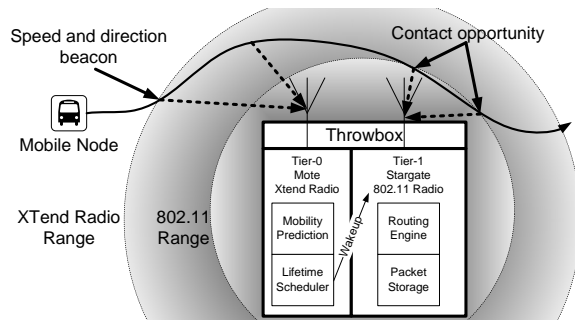


Fig. 1. Overview of our throwbox architecture.

tions that maximize network performance. We found that even a small number of throwboxes can significantly increase the delivery rate found in a DTN. However, this addressed only part of the problem.

Throwbox deployment without the use of power management is minimally effective. The high power consumption of always-on nodes results in a short lifetime; conversely, nodes powering on and off intermittently without regard to node mobility may miss numerous connection opportunities, increasing lifetime without necessarily improving performance. The key goal is to provide energy efficiency without sacrificing network performance. Moreover, throwboxes are constrained to consume power at a rate dictated by a battery lifetime requirement or the availability of scavenged energy, such as solar power. In this paper, we address these issues in the context of the following challenge: *how can a throwbox in a DTN meet an average power constraint and simultaneously maximize the number of packets delivered?*

We show that the primary source of overhead is the energy cost of neighbor discovery: idling or turning the platform and radio on and off to search for contacts in a sparse network wastes the vast majority of node energy. Due to the sparseness of a DTN, waking up for non-existent, or brief, contact opportunities is not worth the cost of turning on the platform and radio.

In this paper, we propose a novel paradigm for power management in DTNs that provides more efficient neighbor discovery by detecting the mobility of other nodes at a minimum cost and predicting the cost and opportunity of each possible contact. Using these predictions, the throwbox can intelligently choose the most fruitful contact opportunities, and it can limit the number of opportunities to meet energy constraints. In our architecture, we pair a *tier-1* platform — a PDA-like Stargate and 802.11 radio — with a low-power *tier-0* platform — a Mote and XTend radio. The XTend radio is a *long-range, low-bitrate* radio for

neighbor and mobility discovery. It has a range of more than a kilometer, which is more than five times the range of WiFi radios. By coupling the XTend radio with a low-power Mote, we reduce the total energy cost of the throwbox platform.

Figure 1 presents an overview of our approach and the problems we address in this paper. Mobile nodes (shown as a bus in the diagram) beacon their position, direction, and speed using the long-distance radio. While listening for other nodes, the throwbox needs very little computational power and memory, and the tier-0 platform is sufficient to process the beacons. If a throwbox hears a beacon, tier-0 predicts if, when, and for how long the mobile node will come in contact with the tier-1, 802.11 radio. Using a constrained single-objective optimization to meet a power consumption target, the node decides whether to take the transfer opportunity. If it does, tier-0 wakes the Stargate and WiFi radio in advance of the mobile node entering radio range, which then transfers DTN data. After the contact opportunity, the tier-1 platform returns to a sleep state. For this scenario, we develop a method for predicting the contact opportunities based on the observed mobility of remote nodes and an algorithm for choosing which contacts to take given power constraints.

We meet our objectives through a mix of theoretical and practical results. In Section III, we present a mobility prediction algorithm that is sufficiently lightweight to operate on a Mote. Section IV presents a duty-cycling algorithm for the long range radio. In Section V, we detail our scheduling algorithms that help meet an average power constraint for the throwbox while maximizing the number of packets delivered in the network. We present theoretical bounds on the complexity of the optimization problem and the performance of an online algorithm based on token-buckets. We show the problem is NP-Hard and present competitive bounds on the performance of the algorithms for a simple node mobility pattern. In Section VI, we detail our design and implementation of a prototype throwbox and its deployment and integration into the UMassDieselNet DTN [2]. And finally, we show the performance of the system through an extensive set of trace-driven simulations and experiments on the deployed platform. Our methods reduce the power consumption of the throwbox from 2500 mW to 80 mW while delivering almost as many packets. Stated another way, we have deployed a throwbox that runs perpetually on solar cells similar to the size of the box itself.

II. BACKGROUND

Peers in a mobile network alternate between two basic operations: *neighbor discovery* and *opportunistically transferring data*. In DTNs the latter operation has received much attention as part of routing protocols [2, 7]. However, in a sparse DTN network, searching for other nodes consumes a large percentage of time in comparison to transferring data. Consequently, searching for other nodes becomes the dominant drain on the energy of a battery-powered DTN node. For instance, we show in this paper that using an 802.11 radio to search for contacts in a DTN devotes 99.5% of the total energy just to find other nodes to exchange data with. Simple power management schemes are insufficient—a radio in idle mode has little savings as compared to its listening mode.

Previous work has addressed mobile system power man-

agement by using 802.11 radios for data transfers and *low-power, short-range* radios (e.g., 802.15.4 [14], Bluetooth [15], or CC1000 [16]) for neighbor discovery tasks. While this solution is appropriate for dense mobile networks, our recent work has shown it is inefficient for sparsely populated DTNs [11]—we strengthen this claim in this paper. This is because short-range radios miss too many connection opportunities, which tend to be short-lived compared to the delay and cost of switching 802.11 radios from sleep to wake modes. Moreover, these architectures do not reduce the cost of other idle hardware components, predict the mobility of the node, nor choose which opportunities to take.

A key to realizing the gains of a second radio is to be able to predict neighbor mobility *before* the node is within contact of the throwbox. Mobility prediction algorithms in the past have been used to predict future network topology to perform route reconstruction proactively [23], accurately predicting hand-offs and bandwidth provisioning in cellular networks [5, 20], and location tracking in wireless ATM networks [13]. Such algorithms either require coordinated information from more than one base station [20, 21], they are computationally too expensive to implement on an extremely low power tier-0 subsystem [21], or they are built assuming very specific movement patterns of nodes [20]. Most mobility predictors are assumed to be implemented on powerful base stations, which are unconstrained platforms. Mobility prediction has also been used in DTNs, taking advantage of a known schedule for mobile nodes in the network [12]. In our work, we develop prediction algorithms that are lightweight enough to run on a small, extremely efficient platform, such as a Mote. This helps minimize the cost for deciding which contact opportunities to take. Additionally, our technique does not depend on any centralized infrastructure, or scheduled node mobility.

III. MOBILITY PREDICTION ENGINE

Because DTNs are sparsely populated, throwboxes must expend significant energy searching for contact opportunities. Because of node mobility, throwboxes have a limited amount of time to exchange data with peers. Searching for contacts efficiently requires waking and sleeping rapidly (called *duty cycling*), and conversely, data transfer requires high-bandwidth connections. Our design philosophy is to separate neighbor discovery from data transfer, and to divide these tasks across the hardware that is best suited to each task. The long-distance radio and tier-0 platform can duty-cycle, listen, and idle efficiently, while the high-bandwidth radio and tier-1 platform can transfer large amounts of data during limited connection opportunities. While the tier-0 platform is listening for contacts, the tier-1 platform remains asleep.

The decision to wake the tier-1 platform and radio depends on an engine that detects, and predicts the mobility of nodes in the network, based on information gathered by the tier-0 platform and radio. By predicting the trajectory of an on-coming mobile node, the tier-0 platform can determine if, when, and for how long it will enter the range of the data transfer radio. If the throwbox determines that the mobile node will not come in contact with its data transfer radio, it can safely ignore that node. If it determines that the mobile node will enter at a certain time, it

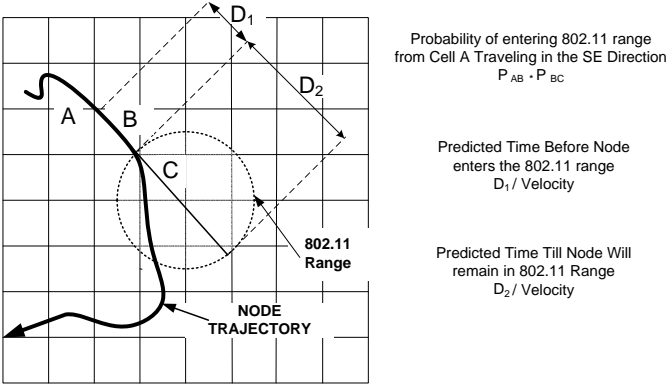


Fig. 2. The figure depicts the working of the Mobility Prediction Engine

must make sure that the throwbox can make the most out of the contact opportunity. When the throwbox estimates the time that the mobile node will enter 802.11 range, it wakes the tier-1 node in advance of the mobile node arriving—a process that can take several seconds. The predicted length of the contact opportunity is then used to determine which contact opportunities to take, a process we explain in Section V.

To determine trajectories, we require that mobile nodes periodically transmit location, speed, and direction over the long-distance radio as a beaconing message. By requiring mobile nodes to beacon their positions, we have shifted some of the energy burden from the stationary throwboxes to the mobile nodes. In this paper, we study the use of throwboxes in a vehicular network, where the mobile nodes have plentiful energy sources. Nonetheless, the techniques will aid systems where all nodes are energy constrained as it lowers the cost of discovery; however, we leave this as future work.

A. Prediction Algorithm

We model the movement pattern of the mobile nodes as a Markovian process. Hence, we assume the location of a node at time T_i is dependent only on its location at time T_{i-1} . Such a model is accurate for movement patterns that are both predictable and semi-predictable.

The algorithm models the problem as a virtual square of width $2r$, where r is the radius of the long-range radio. The throwbox is located at the center of the square, as illustrated in Figure 2. The large square is divided into square cells numbered from 0 to k . The algorithm makes use of a probability transition matrix T for mobile nodes in the network. Entry T_{ij} of the matrix is the probability that a node will transition to cell j given that it is in cell i .

Mobile nodes beacon data in tuples in the form $\langle p, v, t, d \rangle$, where p is the present GPS location of the mobile node, v is its speed, t is the time at which the data was sent, and d is the direction of motion of the node. When a node n approaches, the throwboxes collect a set of tuples $\{\langle p_1, v_1, t_1, d_1 \rangle, \dots, \langle p_k, v_k, t_k, d_k \rangle\}$. Using this set, as well as historic information, the algorithm estimates:

- The probability that the mobile node n would be in any cell within the data transfer radio range after time ΔT , where ΔT is the time required to wakeup tier-1 and start transferring ap-

plication data—it is typically on the order of seconds

- The predicted time when the node will be in range of the data transfer radio and, subsequently, the amount of data it is likely to transfer

Figure 2 illustrates this process. At time $t = 0$, the mobile node's beacon from cell A is received. By dead-reckoning along a straight line of motion, the prediction engine estimates D_1 , the distance until the node is within 802.11 range, and D_2 , the distance for which the node would be within range. Accordingly, the predicted length of time until the mobile node reaches the range of the data radio is D_1/v ; if this time is greater than the transition time to wakeup tier-1, then the mobile node is ignored. Otherwise, the throwbox calculates Pr , the probability that the node will enter the range of the data radio. (In the simple example presented in the figure, Pr is the product of the probability of transition from cell A to B and from B to C; i.e., $Pr = T_{AB} \cdot T_{BC}$.) The expected duration of time the mobile node will stay in data radio range is $Pr \cdot (D_2/v)$. Below we show how we calculate Pr under Markovian assumptions.

B. Estimating the Probability of Entering Data Radio Range

The prediction engine assumes that the node would move in the direction that it was last seen. Therefore, it constructs a straight line in the current direction of motion and finds the intersection with the range of the data-transfer radio range (idealized as a circle). Let p be the cell at which the node enters the data radio range; let D_{p,p_k} be the Euclidean distance between the present location of the node and the location of the point where it enters the data transfer radio range (assuming the present direction of motion of the node); let v_k be the velocity of node n ; and let c_i, \dots, c_{j+i} be a sequence of cells that are in the predicted path of motion of the node. Let $\{d_1, \dots, d_i\}$ be the set of i possible directions of motion. For example, a direction could be north, south, east or west. The probability that a node enters the data transfer radio range at position p after time ΔT is given by

$$Pr[X_{t_k+\Delta T} = p | (X_{t_k} = p_k, d_{t_k} = d_k), \dots, (X_1 = p_1, d_{t_1} = d_1)], \quad (1)$$

where X_t is the position of the node at time t . The above probability can be approximated as

$$P_{appr} = Pr[X_{t_k+\Delta T} = p | X_{t_k} = p_k, d_{t_k} = d_k] \quad (2)$$

assuming that the node movement is modeled as a Markovian process. P_{appr} is evaluated by the algorithm as

$$Pr[X_{t_k+\Delta T} = p | (X_{t_k} = p_k, d_{t_k} = d_k)] = \begin{cases} T_{c_i, c_{i+1}} \cdot \dots \cdot T_{c_{i+j-1}, c_{i+j}}, & \frac{D_{p,p_k}}{v_k} \leq \Delta T \\ 0, & \frac{D_{p,p_k}}{v_k} > \Delta T \end{cases} \quad (3)$$

Equation 3 shows that the probability of a node entering a cell of the data transfer radio after time ΔT is equal to the transition probability of the node from its present cell to cell c_{i+j} if the node is traversing fast enough to cover the shortest distance between the two cells in time ΔT .

IV. DISCOVERY RADIO DUTY-CYCLE CONTROLLER

If multiple contacts are simultaneously discovered by the throwbox, it takes the size of the connection event as the sum of the predicted connection size of each contact.

We also note that while searching for mobile nodes, the long-distance radio does not need to be constantly powered; it needs to only wake-up often enough to predict if, when, and for how long the mobile node will be within range of the data transfer radio. This is consistent with past approaches to lowering wireless power consumption such as PSM [1], which duty-cycle the radio to save energy.

Our throwbox design uses a MaxStream XTend radio that supports several very efficient duty-cycling modes at the hardware and MAC layer. The MAC protocol built into the long-range radios is very similar to the B-MAC protocol built for CC1000 Mote radios [17].

The radios used in this work are built with k software configurable cyclic sleep modes $\{S_1, S_2, \dots, S_k\}$. In sleep mode S_k , the radio wakes up every t_k seconds for time τ ($\tau \ll t_k$) looking for RF activity. If any RF activity is discovered, the tier-0 Mote collects the data it received.

All mobile DTN nodes are equipped with the long range radio and beacon their location, speed and direction of motion periodically.

If the radio is always on, the power consumed by the radio would be high and if the radio is in sleep mode most of the time, the throwbox might miss important information regarding location and mobility of DTN nodes and consequently miss important contact opportunities. Therefore, there is an inherent tension between power consumed by the radio and the amount of information that it can gather. Moreover, the duty-cycling algorithm should be able to adapt to changing mobility patterns of nodes in the network.

Therefore, we propose an adaptive duty-cycling heuristic for the long range radio. If the mobility prediction engine requires m data points to accurately predict the movement pattern of a node, the aim of the heuristic is to gather these m points at minimal energy cost. The number of points the mobility predictor needs must be determined through experimentation, and through experience we have found that a small number of points, such as four, is sufficient to make accurate predictions in our network.

The algorithm constructs m virtual concentric circles of radii $r_n = R + \frac{n \cdot (r - R)}{m}$ ($1 \leq n \leq m$), where r and R are the ranges of the discovery radio and the data transfer radio. The number of circles m , constructed depends on two factors. 1) the number of data points required by the mobility prediction engine to accurately predict when the node will enter the data transfer radio range. 2) the distance $r_1 - R$ is taken such that the time for a node to enter circle R is less than the transition time to wake up the data transfer radio and tier-1 platform. The virtual concentric circles helps the throwbox determine how close it is to a mobile node. We assume that the radios range which is roughly circular, it is intuitive to select circles to dictate how close a node is from data transfer range. The sleep mode S_k corresponding to the shortest sleep interval is taken as the default sleep mode. Therefore, when there are no nodes in range the radio is duty-cycled in sleep mode S_k . When the throwbox

receives a beacon from a mobile node it determines the circles of radii r_i and r_{i-1} within which the node is presently located. The algorithm also knows the velocity \bar{v} with which the node is moving. The duty-cycle controller calculates the interval of time $T = [(r_i - r_{i-1})/\bar{v}, \sqrt{(r_i^2 - r_{i-1}^2)/\bar{v}}]$ within which the node is expected to be within the circle of radius r_{i-1} ($r_{i-1} < r_i$). The algorithm finds the sleep mode S_j with the longest t_j such that $t_j \leq (r_i - r_{i-1})/\bar{v}$. Consequently, the duty-cycle controller puts the discovery radio to sleep in sleep mode S_j . If the discovery radio does not find any data after time t_k , the node is put to sleep again in the default sleep mode. Assuming that the variance in the speed of the mobile nodes is not very large, the algorithm would not miss many beacons on the long range radio and would remain asleep as long as possible. If the throwbox records beacons from multiple nodes, the heuristic estimates the shortest time after which a node will enter the adjacent concentric circle.

The choice of appropriate radio as the discovery radio is important for maximal discovery of connection opportunities at minimal power. Following our conjecture made in [11], we found it more energy efficient and useful to use a long range low bit-rate radio as a discovery radio in place of an short range low power radio. We further strengthen the claim through the proposition below.

We show in the following proposition, that is is more efficient to use a long range low bitrate radio as the discovery radio than a small range low power radio. We assume a simple mobility model where nodes can move in chords once they enter the data transfer radio range.

Proposition 1: Consuming the same amount of energy, the XTend long range radio helps a throwbox detect all connection events with a high probability while the Zigbee radio can detect only $\frac{1}{9}$ of the connection events.

Proof: We prove the proposition through the following lemmas. ■

Lemma 1: If the range of the discovery radio, $r \leq R \cdot \sin(\frac{\pi}{2k})$, where R is the range of the data transfer radio and $k > 1$, the probability that the percentage of connection events discovered is more than $\frac{1}{k}$ is bounded by $e^{-\Omega(n)}$, where n is the number of contact opportunities.

Proof: The scenario for the proof is illustrated in Figure 6. We assume that the throwbox is located at the center of the circle and a node enters the range of the large range radio at point A. The probability that the node would enter the range of the low power radio is at most $\frac{2 \cdot \sin^{-1}(\frac{r}{R})}{\pi}$. We assume that the connection events occur independently of each other and follow a binomial distribution with parameters $(n, \frac{2 \cdot \sin^{-1}(\frac{r}{R})}{\pi})$. By Chernoff's bound, $P[B(n, p) > (1 + \delta)np] \leq e^{-\frac{\delta^2 np}{3}}$. Therefore, $P[B(n, p) \geq n/k] \leq e^{-\Omega(n)}$ if $1 - \frac{1}{2p} > 0$. Replacing p by $\frac{2 \cdot \sin^{-1}(\frac{r}{R})}{\pi}$, yields the result in the proposition. Also note that the size of the connection event detected by the short range radio is at most half of the total size of the connection event. ■

Corollary 1: If we use the Zigbee radio (range = 50 m) as the 802.11 radio (range = 300 m) $k = 9$. This implies that the combination of a Zigbee and a 802.11 radio $\frac{8}{9}$ of the connection events would not be detected with a very high probability.

Lemma 2: If $r \geq 3R$, and the long range radio is awake with a probability p , then the probability that the throwbox would detect every contact opportunity is at least $1 - e^{-[1 - \frac{v}{2Rp}]^2 \cdot \frac{Rp}{v}}$, where v is the maximum speed of the node moving towards the throwbox.

Proof: We assume that a node moves towards the throwbox with a maximum speed of v and beacons at a rate of 1 beacon/sec. The event of the throwbox receiving a beacon is assumed to follow a binomial distribution. In the worst case the node would move directly towards the throwbox at a speed v . The probability that the throwbox would receive a beacon is at least p and the number of beacons transmitted by the node is $\frac{2R}{v}$. Applying chernoff's inequality, yields the bound in the lemma. ■

Corollary 2: If the long range radio is a XTend radio (range = 1000m) which is 16% duty-cycled, the radio will detect a connection event with a probability at least 0.7 if the maximum speed of the approaching node is 80 km/hr.

Corollary 1 and Corollary 2 yields the result in Proposition 1.

Comment 1: If a node entering the data transfer radio range is assured with a probability 1, to enter the short range radio it is still more energy efficient to use a long range radio for discovery.

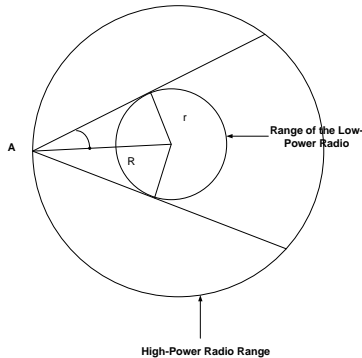


Fig. 3. The figure illustrates the scenario for Proposition 1

Proof: If a node is assured to enter the short range radio (with a range r), once it is within the 802.11 range (R), the throwbox would lose $1 - \frac{r+R}{2R} \approx \frac{1}{2}$ (assuming a small r) bytes of a connection event. Therefore, at the same energy cost, using the short range radio transfers data corresponding to approximately half of the connection events while the long range discovery radio transfers all the data with a high probability. ■

The above proposition further validates the conjecture made by Hyewon et. al [11].

V. TOKEN BUCKET LIFETIME SCHEDULER

Mobile nodes may present a throwbox with many transfer opportunities. However, because they have limited energy, they cannot participate in and draw power for every opportunity. A limited energy supply is better viewed as a constraint on the *average power* that the system consumes in two different scenarios. First, the throwbox may be designed to last for a certain period of time — computing the average power from the capacity of the batteries is straightforward. Second, throwboxes may capable of scavenging energy, such as the solar panels used in our prototype system. In that case, we assume the average

power constraint to be the average power produced by the solar cells. This does assume that the solar cells charge a battery large enough to smooth out variations in solar energy production; however, we omit this complexity in this paper.

In this section, we first show that computing the optimal subset of opportunities to participate in is NP-Hard. We then present our sub-optimal solution, which is linear in the number of connections and requires constant memory. In Section VII, we show that our algorithm performs within 80 % of the optimal.

We assume that energy is the only optimization criteria; however, other optimization criteria exist and they can be considered concurrently. For example, mobile nodes may transmit the priority of packets that they are carrying, or the throwbox can decide on which opportunities to take based on the likelihood of being able to route the packet to its final destination. We leave these possibilities as future work.

A. Complexity of the Optimization Problem

The most efficient strategy is for the throwbox to amortize the transition energy over the largest transfer opportunities. For example, if a mobile node skirts the outer range of the throwbox, it is relatively expensive to wake the tier-1 system for such a short opportunity. However, determining the optimal set of transfer opportunities to wake the tier-1 system is NP-Hard, even as an offline problem, as we show below.

The inputs to the above optimization problem, O , are a set of connection events $C = \{c_1, c_2, \dots, c_n\}$ that occur within time interval $[0, t]$ and the throwbox's energy constraint $P \cdot t$. Each connection event c_i is associated with an energy cost e_i that includes the transition energy to wake up tier-1 and the data transfer radio and the amount of energy spent while transferring data during the connection event. Let d_i be the number of bytes likely to be delivered as a result of the data transfer during c_i . The solution to the optimization problem O is a subset of events from C , such that energy goal $P \cdot t$ is not exceeded and maximum number of bytes is delivered by the throwbox.

The decision version of this problem, O_d , takes as input an additional positive number k . A solution to O_d does not spend more than $P \cdot t$ energy and delivers at least k bytes of data. We prove below that the above optimization problem is NP-Hard.

Theorem 1: O_d is NP-Complete.

Proof: We first prove that O_d is in NP. We nondeterministically guess a subset S of connection event. We need to check the following two items. (1) If $\sum_{c_i \in S} e_i \leq P \cdot t - I$, where I is the idle energy consumption of the system in time t . The idle energy includes the energy consumed by the tier-0 system and the discovery radio. (2) If $\sum_{c_i \in S} d_i \geq k$. These checks can be done in time polynomial in the input size. Therefore, $O_d \in NP$.

We next prove that $O_d \in NP$ -Hard through the following reduction from the knapsack problem. The input to the knapsack problem is the set of items $I = \{I_1, I_2, \dots, I_n\}$ where each item is associated with a weight w and a value v . The goal is to choose a subset of the items such that the sum of their weights is less than a capacity C and the sum of their values is greater than a fixed number V . The reduction works as follows.

An item in the knapsack is mapped to a connection event in C . The value of the items are mapped to number of packets delivered and the weight of the items are mapped to the energy

cost associated with a connection event. C , the capacity of the knapsack is mapped to the energy constraint $P \cdot t - I$. The value k is taken the same for both problems. It is clear that a subset of items in the knapsack problem which satisfy the weight and value constraints exists if and only if there is a subset of the connection events that meet the energy constraint and delivers more than k bytes. ■

B. A Token-Bucket Approach

Since O_d is NP-Complete, it is clear that the online version of the above optimization problem does not have any efficient solution. However, an approximate online solution would follow the following intuition: since an online algorithm has no knowledge of the future, at any time t it should take a contact opportunity only if the energy cost does not lead to a violation of the average power constraint for time t . In other words, the heuristic should regulate energy flow based on the average power constraint. Moreover, since contact opportunities may occur in bursts, the scheduler should also be capable of handling bursts of energy consumption.

Token buckets have been used in a similar scenario in networking to regulate network traffic [24]. Token buckets allow bursty traffic to continue transmitting while there are tokens in the bucket, up to a user-configurable threshold thereby accommodating traffic flows with bursty characteristics [10].

Algorithm 1 Token Bucket Scheduler

```

Average power constraint =  $P$ 
Generate energy tokens at the rate of  $P/\text{sec}$ 
Energy cost of present connection event =  $E_p$ 
Size of present connection event =  $S_p$ 
Number of tokens accumulated till present =  $n$ 
if  $n \cdot P < E_p$  then
    Do not wakeup Tier-1
else
    Mean size of the last  $k$  connection events =  $S_k$ 
    Mean energy of the last  $k$  connection events =  $E_k$ 
    Mean inter-arrival time of connection events =  $T_m$ 
    if  $n \cdot P - E_p + T_m \cdot P < E_k$  and  $S_p < S_k$  then
        Do not wakeup Tier-1
    end if
end if

```

The scheme, shown in Algorithm 1, generates energy tokens at the rate of the average power constraint. For any given connection event if the number of tokens accumulated is less than the amount of energy required, the event is ignored. However, if the number of tokens accumulated is greater than the energy required for a connection event the system makes a simple choice: should it take this connection event, the next connection event, or both. First, the algorithm estimates the size and energy cost of the current connection event based on the mobility prediction engine. It then estimates the size and energy cost of the next connection event based on the mean of the last k connection events and assumes that the connection event arrives at the mean inter-arrival time of the last k connections. Taking into account the number of tokens that are accumulated between now and the next connection event, if it estimates it can take both

connection events it takes the current one. Otherwise it chooses the larger of the two. This process repeats for the next contact opportunity.

The algorithm runs in time linear in the number of connection events and requires $O(1)$ amount of space. Therefore, the algorithm is space and time efficient and is easily implementable on the low power tier-0.

C. Performance of Token-bucket Scheduler

We perform a competitive analysis on the performance of the token bucket scheduler for a simple mobility model. We assume that there is only one mobile node which moves passes the radio range of the throwbox following a fixed schedule with inter-contact time t_r . The amount of data that is transferred during one contact is assumed to follow a normal distribution. We bound the maximum energy for each contact by $2 \cdot t_r \cdot P$, where P is the average power constraint on the system. Such an assumption is on a average true for any DTN where the nodes follow a schedule about a fixed point (eg. UMassDiselNet [2]).

We prove (proof given in the appendix Lemma 3), that the probability that the size of a sample connection event is larger than the mean of the last m connection events is bounded by a constant c . Since the expected energy for each is bounded by $2 \cdot t_r \cdot P$, the token-bucket scheduler would have accumulated enough energy for one out of two consecutive connection events. Therefore, the probability that the scheduler would wakeup tier-1 for the i^{th} connection event is at least $(1 - Pr_{i-1}) \cdot c$, where Pr_{i-1} is the probability that the $(i-1)^{\text{th}}$ connection event was chosen. Therefore, the average number of bytes transferred by the token-bucket scheduler is given by $\sum_{i=1}^{i=n} (S_i \cdot Pr_i)$, where n is the total number of connection events. Assuming large n , the above sum is equal to $\sum_{i=1}^{i=n} ((1 - Pr_{i-1}) \cdot c \cdot S_i) \geq (S \cdot c \cdot (1 - c))$, where S is the sum of the sizes of all connection events. Note that the algorithm always meets its average power constraint since it never takes a connection event if it does not have enough tokens for it. The optimal algorithm can transfer at most S bytes. Therefore, the token-bucket scheduler is $(\frac{1}{c \cdot (1 - c)})$ -competitive.

We prove in the appendix (Lemma 4) that if the sizes of the transfer opportunities are normally distributed (approximately true for UMassDieselNet), the token-bucket scheme is $\frac{64}{15} \approx 4$ -competitive (Note that in Lemma 4 if the probability of failure is taken as $1/2$ ($3/8 < 1/2$), then the algorithm is at most 4-competitive).

In section VII, we show that the algorithm performs better than the above bound in practice.

VI. THROWBOX PROTOTYPE IMPLEMENTATION AND DEPLOYMENT

We constructed prototype throwboxes, shown in Figure 4, using a Crossbow Stargate (tier-1) [25] and a TelosB Mote (tier-0) [18]. We chose these hardware platforms because they handle the two DTN activities, data transfer and neighbor discovery, efficiently. The Stargate platform runs Linux, allowing us to run the same Java-based routing software used in UMassDieselNet.

The Stargate contains a 32-bit, 400 MHz PXA255 XScale processor, 64 MB of RAM, 32 MB of internal flash, and a




Fig. 4. Prototype throwbox system

Characteristic	XTend	Dlink-Air
Ranges	1000 <i>m</i>	150 <i>m</i>
Receive Power	360 <i>mW</i>	1000 <i>mW</i>
Transmit Power (max)	1 <i>W</i>	1.2 <i>W</i>
Sleep Power	10 <i>mW</i>	50 <i>mW</i>

Fig. 5. Characteristics of the two radios

DLink-Air 802.11b interface. The TelosB Mote contains an 8-bit, 8MHz microcontroller, 10 KB of RAM, and 1 MB of external flash. The Mote is attached to a Maxstream XTend 900 MHz OEM module. The prototype also employs two 5V PowerFilm solar panel as additional energy source. The characteristics of the two radios are summarized in Table 5.

We fabricated a power supply board for attaching both platforms to a single battery and support charging from solar energy. This board provides an additional hardware element necessary to throwboxes, a Maxim DS2770 fuel gauge chip. The fuel gauge chip gives accurate readings of the energy stored in the battery, similar to those found in laptops. This accounts for energy consumed by the platforms and radio, as well as energy produced by the solar panels. The TelosB Mote reads this value periodically and corrects the token bucket to account for the actual energy used and produced.

To support a real-world test of the throwbox, we used our DTN testbed, the UMassDieselNet [2]. The testbed normally consists of 40 buses covering an area of more than 150 square miles. However, when the experiments were performed, during a reduced summer bus schedule, only 10 buses were running on three routes. Each bus is a highly mobile DTN node using a small computer with an attached access point and WiFi interface. Buses constantly scan for other nodes and transfer DTN data whenever a connection can be made.

We augmented the equipment on the buses with an XTend radio and added scripts to beacon the position, speed, and direction of motion of the buses once each second. We deployed three always-on throwbox prototypes in fixed locations for three weeks on the UMassDieselNet bus routes (shown in Figure 6). The throwboxes were placed according to our deployment algorithm in previous work [27].

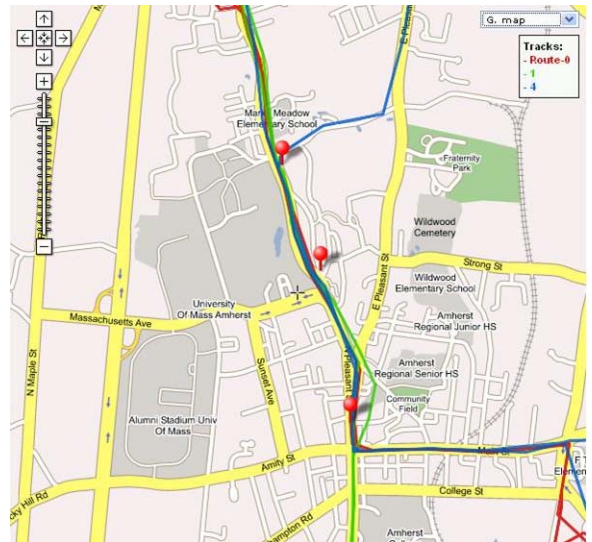


Fig. 6. The Map shows the locations where the throwboxes were placed

VII. EVALUATION

We evaluated the throwbox system through two techniques: trace-driven simulations and prototype experimentation. The simulations use traces collected by placing three always-on throwboxes in UMassDieselNet for three weeks changing batteries manually. These prototypes logged connection events on the 802.11b radio and the XTend radio and we fed the data into a Java-based simulator.

In the simulations, we compare our system with the following systems.

- **Optimal (Dual platform):** The optimal system uses the two-radio, two-platform system, but with a perfect mobility prediction algorithm and an optimal scheduler. It has an oracle that knows the exact time and length of every connection opportunity. The system uses an optimal dynamic programming algorithm for the knapsack problem to select the exact set of connection of events that maximizes throughput while meeting an energy constraint [6]. This system represents the best our dual hardware design can do with an oracle of future events.
- **PSM*:** This system is a single tiered single radio system that periodically powers off its wireless interface to save energy. In some ways this is similar to the PSM system found in WiFi cards. PSM* wakes up its wireless interface and scans for connection events and goes back to sleep if it finds none. When the WiFi card is switched off, the platform is in its idle state. We exhaustively searched the state space to set parameters that use the minimum energy with equivalent data transferred. This provides a comparison to the best system that does not use extra hardware.
- **WoW*:** This system is adapted from Wake-on-Wireless [19] which uses a second radio as a discovery radio. The published WoW system always wakes up when it sees data addressed to it on the discovery radio and it is assumed that the discovery radio has the same range as the data transfer radio. To provide a fair comparison, we adopt WoW to a DTN environment. The WoW* system uses our mobility prediction engine to intelligently decide when to wakeup the data transfer radio. Without




Fig. 7. The PDF of the amount of data transferred per contact opportunity with a throwbox




Fig. 8. The PDF of the inter-contact time between a throwbox and a mobile node

this modification, the WoW system would wakeup on every spurious contact over the long-range radio. However, it does not use the scheduling algorithm to decide which opportunities to take, nor does it duty-cycle the long-range radio.

- **Always-on:** This throwbox remains on all of the time, and thus does not use the second radio, and it is not penalized by the additional energy needed for the long-distance radio and Mote.
- **No throwbox:** Without a throwbox the DTN delivers data normally.

First, we show the statistics of the collected data to provide some insights into the rest of the results. We also compare the data collected by the throwbox to the statistics from a purely mobile DTN. Second, we show how well the mobility prediction and token bucket schedulers operate in comparison to the optimal and WoW* schemes. We show that the mobility prediction algorithm has a zero percent false-positive rate, a 10% false-negative rate, and generally predicts a node's arrival within 4 seconds. While the scheduling algorithm is inferior to the optimal algorithm (1.25-competitive), it delivers significantly more packets than WoW* and meets the average power constraint. Third, we show the trade off between the network-wide throwbox benefits and its average power constraint. The results demonstrate that almost all of the benefits can be obtained with a very modest power budget of 80 mW, 3.2% of the always-on throwbox and 8% of the PSM* system. Fourth, we demonstrate that this power constraint requires small-sized solar cells that fit on our prototype's case. Lastly we show the results of deploying a real throwbox and demonstrate that small-sized solar cells can meet our goal of a perpetually operating throwbox while delivering almost as many packets as an always on system.

A. Data Statistics

Figures 7 and 8 show the PDF of the amount of data transferred per contact and the inter-contact time between a throwbox and a mobile node. The amount of data transferred during one contact is seen to be approximately a log-normal distribution. This is similar to the distribution observed for bus-to-bus data transfer [2]. However, the mean of the data transferred between a throwbox and a bus is around two times that of a bus to bus




Fig. 9. The PDF of the error introduced by the Mobility Prediction Engine

transfer. This increase in the throughput is due to two factors. First, the throwboxes are stationary increasing the time they are in contact with a bus. Second, the throwboxes have been strategically placed to increase the delivery rate [27]. The locations are in relatively dense areas of connectivity in the DTN. The inter-contact time is shown to be a bimodal distribution.

B. Trace-Driven Simulation Results

B.1 Mobility Prediction Engine

We next test the accuracy of the mobility prediction engine. The traces collected from the XTend radio of the always-on throwbox prototypes are taken as an input to this experiment. The mobility prediction engine is run on this trace to predict probable connection events, their time of occurrence, and sizes of the connection events. The 1km-by-1km square around a throwbox was divided into 25 200m-by-200m cells. The prediction engine was warmed-up with only one contact opportunity from each route. Figure 9 shows the PDF of the error of the mobility prediction engine.

The results show that the prediction engine was able to predict probable connection events with an average error of less than 4 seconds. From Figure 9, it is clear that the system was woken




Fig. 10. The number of bytes transferred by the throwbox for different average power constraints

up earlier than the actual connection events most of the time and hence it does not miss many transfer opportunities, though it consumes slightly more energy than the ideal case. Two additional results not shown in the graph are that the system never predicts false-positives (buses that it predicts will enter 802.11 range, but don't) and a 10% false negative rate (connection opportunities that it misses). The false negative rate comes from the duty-cycling algorithm used in the long-distance radio—described in Section IV.

B.2 Token Bucket Scheduler

In the next experiment, we evaluate the token bucket scheduler. The experiment uses the predicted length of the connection events and the predicted time of occurrence from the mobility prediction engine. We evaluate whether the token bucket scheduler is able to meet an average power constraint while delivering a large number of packets. We have set the average power constraint to 80 mW and compare the results with a WoW* system and an optimal system. Recall that an optimal system uses the same hardware as a throwbox but has a perfect mobility prediction engine and prior knowledge of all connection events. Figure 10 shows the amount of data delivered by each of the systems, and Figure 11 depicts the depletion of battery energy with time.

The results show that the WoW* system delivers less than 20% of the amount of data delivered by the token bucket scheme. The token bucket scheduler transfers more than 80% of the amount of data transferred by the optimal scheme, and hence is approximately 1.25-competitive. Figure 11 shows that the throwbox scheduler meets the average power constraint (shown as a straight line), while the WoW* system quickly depletes its energy. This is due to two effects. First, WoW* does not discriminate based on the size of the connection events and takes short events, providing little benefit and wasting energy. Second, the WoW* system does not duty-cycle the long distance radio, an additional benefit of our throwbox design.

To show these results in more detail, we compare the throwbox system with the WoW* as well as the PSM* system. We did not show the PSM* system in the previous graphs as it per-

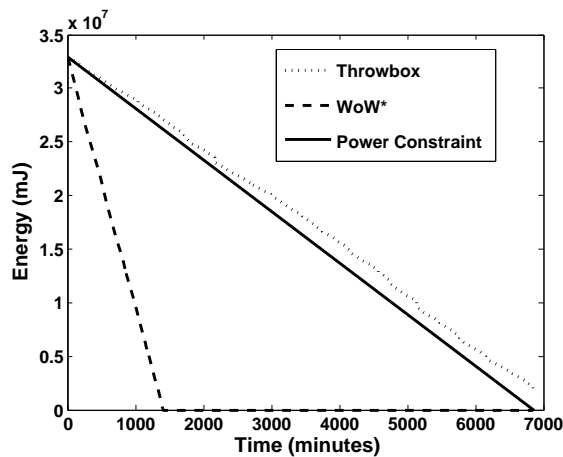


Fig. 11. How energy is used by the different schedulers as a function of time




Fig. 12. The breakup of energy consumed by components

forms orders of magnitude worse than the other systems. We ran a trace-based simulation of the throwbox system with a power constraint of 80 mW. We then ran simulations of the WoW* and PSM* systems to find the amount of power used by the systems to transfer the same number of bytes. Figure 12 shows the average power of these systems, broken down by how the power is spent. We have divided the bars into: energy for useful work (Data Transfer), the number of tokens left over by the scheduler at the end of the simulation (Tokens Left), the energy used by the long distance radio (Discovery Radio), the cost of turning the tier-1 system on (Transition Cost), the cost of the idling tier-1 system while it is searching for contacts (Idle Cost), and the power consumed by the tier-0, Mote system (TelosB).

The results of the PSM* system plainly illustrates the value in using the tier-0 platform and radio. PSM* devotes 99.5% of its energy turning the platform on and off and idling while searching for contacts. The WoW* system virtually eliminates the idle cost as it only turns on the tier-1 system when there is a contact present. However, it does devote a large amount of energy to the long-distance radio, as it does not use the duty-cycling of employed by our throwbox design—described in Section IV. Additionally, the WoW* system has higher costs for transitions as it amortizes that cost over smaller connection events as compared to the throwbox system.




Fig. 13. The increase in delivery rate with an 80 mW power constraint

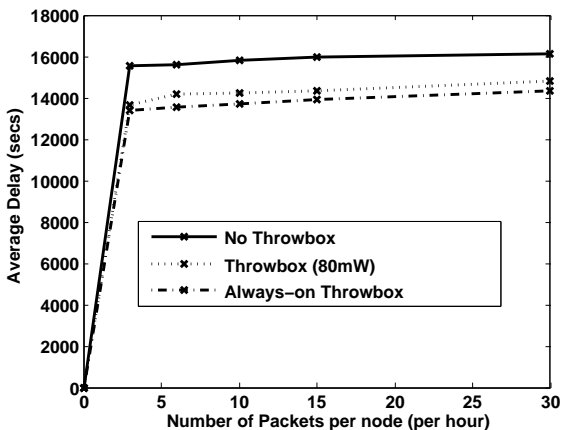


Fig. 14. The decrease in latency with an 80 mW power constraint

B.3 Average Power Constraint Versus Network Performance Boost

We next study the performance boost provided by placing a throwbox in the UMassDieselNet [2] for a given power constraint. We use the MaxProp [2] routing protocol to transfer data among nodes in the network. Each node used a 1 GB buffer and the packets transferred were 1 KB in size. We varied the number of packets per hour and calculated the delivery rate and average latency of transferred packets. We ran the experiments for three weeks of simulated time. Figures 13 and 14 show the increase in packet delivery and decrease in packet delivery time when deploying a throwbox with an average power constraint of 80 mW.

We find from the results presented in the figures that the throwbox with an average power constraint of 80 mW performs as well as an always-on throwbox and leads to an increase in delivery percentage of more than 37% and decrease of at least 10% in the packet delivery time. When compared to an always-on throwbox, the power constraint of 80 mW yields a system that can last 31 times longer on the same battery while delivering almost as many packets.




Fig. 15. The consumption of energy for the throwbox prototype over a period 24 hours

C. Prototype Evaluation

We deployed a two-platform, two-radio prototype with a power constraint of 80 mW, using the mobility prediction, scheduling, and duty-cycling algorithms described in the paper. The system was equipped with solar panels 220 cm^2 in area and a 1 Ah battery that was 20% full. The box was deployed in the UMassDieselNet for a day and it logged the battery capacity remaining in the battery as a function of time. The results are shown in Figure 15. The results show that during the day the throwbox stores excess energy and during nighttime it spends the excess accumulated energy. We determined through this experiment that the solar panels produce on an average power of 65 mW over 24 hours. This is about 15 mW less than that the amount necessary to make the throwbox last perpetually. The amount of energy produced will vary from day to day depending on the intensity of the sun. Therefore, throwboxes can be made to run perpetually through the use of a slightly larger solar cell area (270 cm^2) and accurate modeling of the solar power (e.g., using eFlux [22]).

VIII. CONCLUSION

The paper presents a novel paradigm for power management in DTNs that provides efficient neighbor discovery and prediction of cost and opportunity of each possible contact. Throwboxes, through these predictions can select the most useful contact opportunities such that energy constraints can be met while maximizing the number of packets delivered. Our methods reduce the power consumption of the throwbox from 2500 mW to 80 mW, while delivering almost as many packets. We show, through extensive trace-driven simulations and prototype experimentation that a throwbox can run perpetually on solar cells slightly larger than the size of the box itself.

REFERENCES

- [1] M. Anand, E. B. Nightingale, and J. Flinn. Self-Tuning Wireless Network Power Management. In *Proc. ACM Intl Conf on Mobile Computing and Networking (MobiCom)*, 2003.
- [2] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. In *Proc. IEEE INFOCOM*, April 2006.
- [3] B. Burns, O. Brock, and B. N. Levine. MV routing and capacity building in disruption tolerant networks. In *Proc. IEEE INFOCOM*, pages 398–408, March 2005.

- [4] B. Burns, O. Brock, and B. N. Levine. Autonomous Enhancement of Disruption Tolerant Networks. In *Proc. IEEE International Conference on Robotics and Automation*, May 2006.
- [5] S. Choi and K. G. Shin. Predictive and adaptive bandwidth reservation for hand-offs in qos-sensitive cellular networks. In *ACM Sigcomm*, 1998.
- [6] T. H. Cormen, C. E. Lieserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2001.
- [7] J. Davis, A. Fagg, and B. N. Levine. Wearable Computers and Packet Transport Mechanisms in Highly Partitioned Ad hoc Networks. In *Proc. IEEE Intl. Symp on Wearable Computers (ISWC)*, pages 141–148, October 2001.
- [8] M. Dunbabin, P. Corke, I. Vasilescu, and D. Rus. Data muling over underwater wireless sensor networks using an autonomous underwater vehicle. In *Proc. 2006 IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 2091–2098, May 2006.
- [9] K. Fall. A delay-tolerant network architecture for challenged internets. In *ACM Sigcomm*, 2003.
- [10] P. Ferguson and G. Huston. *Quality of Service: Delivering QoS on the Internet and in Corporate Networks*. Wiley and Sons, Inc, 1998.
- [11] H. Jun, M. H. Ammar, M. D. Corner, and E. Zegura. Hierarchical Power Management in Disruption Tolerant Networks with Traffic-Aware Optimization. In *Proc. ACM SIGCOMM Workshop on Challenged Networks (CHANTS)*, September 2006.
- [12] H. Jun, W. Zhao, M. Ammar, E. Zegura, and C. Lee. Trading latency for energy in densely deployed wireless ad hoc networks using message ferrying. *Elsevier Journal of Ad Hoc Networks*, 2006. To Appear.
- [13] T. Liu, P. Bahl, and I. Chlamtac. Mobility modeling, location tracking, and trajectory prediction in wireless ATM networks. *IEEE JSAC*, 16(6):922–936, 1998.
- [14] N. Mishra, K. Chebrolu, B. Raman, and A. Pathak. Wake-on-WLAN. In *Proc. Intl Conf on the World Wide Web (WWW)*, pages 761–769, 2006.
- [15] T. Pering, Y. Agarwal, R. Gupta, and R. Want. CoolSpots: Reducing the Power Consumption of Wireless Mobile Devices with Multiple Radio Interfaces. In *Proc. ACM MobiSys*, pages 220–232, June 2006.
- [16] T. Pering, V. Raghunathan, and R. Want. Exploiting Radio Hierarchies for Power-Efficient Wireless Device Discovery and Connection Setup. In *VLSI Design*, pages 774–779, 2005.
- [17] J. Polastre, J. Hill, and D. Culler. Versatile low power media access for wireless sensor networks. In *ACM Sensys*, 2004.
- [18] J. Polastre, R. Szewczyk, and D. Culler. Telos: Enabling Ultra-Low Power Wireless Research. In *Proc. Intl Conf on Information Processing in Sensor Networks (IPSN/SPOTS)*, April 2005.
- [19] E. Shih, P. Bahl, and M. J. Sinclair. Wake on Wireless: An Event Driven Energy Saving Strategy for Battery Operated Devices. In *Proc. ACM Mobicom*, pages 160–171, September 2002.
- [20] W.-S. Soh and H. S. Kim. Dynamic Bandwidth Reservation in Cellular Networks Using Road Topology Based Mobility Predictions. In *Proc. IEEE Infocom*, March 2004.
- [21] L. Song, U. Deshpande, U. C. Kozat, D. Kotz, and R. Jain. Predictability of WLAN Mobility and its Effects on Bandwidth Provisioning. In *Proc. IEEE INFOCOM*, 2006.
- [22] J. Sorber, A. Kostadinov, M. Brennan, M. Corner, and E. Berger. eFlux: Simple Automatic Adaptation for Environmentally Powered Devices. (Poster/Demo). In *Proc. IEEE workshop on Mobile Computing Systems and Applications (HotMobile/WMCSA)*, April 2006.
- [23] W. Su, S. Lee, and M. Gerla. Mobility prediction and routing in ad hoc wireless networks. *International Journal of Network Management*, 2001.
- [24] A. S. Tanenbaum. *Computer Networks*. Prentice-Hall, 3rd edition, 2003.
- [25] R. Want, T. Pering, G. Danneels, M. Kumar, M. Sundar, and J. Light. The Personal Server - Changing the Way We Think about Ubiquitous Computing. In *Proc. ACM UbiComp*, Sept. 2002.
- [26] W. Zhao, M. Ammar, and E. Zegura. Controlling the mobility of multiple data transport ferries in a delay-tolerant network. In *IEEE INFOCOM*, 2005.
- [27] W. Zhao, Y. Chen, M. H. Ammar, M. Corner, B. N. Levine, and E. Zegura. Capacity Enhancement using Throwboxes in DTNs. In *Proc. IEEE Intl Conf on Mobile Ad hoc and Sensor Systems (MASS)*, Oct 2006.

APPENDIX

Lemma 3: Given any distribution D , for a large k , the probability that the mean of the k sample point is smaller than a random point is bounded by a constant c .

Proof: Given sample points in the distribution D , with a mean μ and standard deviation σ , the mean of any subset of the sample set of size k approximately follows a normal distribution with parameters $N(\mu, \frac{\sigma}{\sqrt{k}})$. This follows

from the Central Limit Theorem. Therefore the probability $P\{\frac{\sum_{i=1}^{i=k} D(C_i)}{k} >$

$$D(C_j)] = \int_0^\infty P[X = y] \cdot \int_y^\infty \frac{\sqrt{k}}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2 k}{(\sigma)^2}} dx.$$

$$\begin{aligned} P &= \int_0^\infty P[X = y] \int_y^\infty \frac{\sqrt{k}}{\sigma \cdot \sqrt{2\pi}} e^{-\frac{(x-\mu)^2 k}{\sigma^2}} dx dy \\ &= \int_0^\infty P[X = y] \int_{\frac{(y-\mu)\sqrt{k}}{\sigma}}^\infty \frac{1}{\sqrt{2\pi}} e^{-z^2} dz dy \\ &= \int_0^\infty P[X = y] \frac{1}{2\sqrt{2}} [1 - \text{erf}(\frac{(y-\mu)\sqrt{k}}{\sigma})] dy \\ &\leq \int_0^\infty P[X = y] \frac{1}{\sqrt{2}} dy = 1 - c \end{aligned} \quad (4)$$

Note that the error function erf is greater than zero for all values greater than zero. One minus the above probability proves the Lemma

Lemma 4: If the distribution above is normal, then the above constant is $\frac{5}{8}$.

Proof: Substituting $P[X = y]$ with $\frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(y-\mu)^2}{\sigma^2}}$, we get

$$\begin{aligned} P &\leq \frac{1}{\sqrt{2}} \int_0^\infty P[X = y] [1 - \text{erf}(\frac{(y-\mu) \cdot \sqrt{k}}{\sigma})] dy \\ &\leq \frac{1}{4\sqrt{\pi}} \int_{-\frac{\mu}{\sigma}}^0 e^{-z^2} dz + \frac{1}{4 \cdot \sqrt{\pi}} \int_0^\infty e^{-z^2} dz \\ &\leq \frac{1}{4} + \frac{1}{8} = \frac{3}{8} \end{aligned} \quad (5)$$

The derivation uses the fact that when $k < 0$, the minimum value of $\text{erf}(k) = -1$ and when $k > 0$, the minimum value of $\text{erf}(k) = 0$. One minus the above probability yields the Lemma. ■