Adaptive Contact Probing Mechanisms for Delay Tolerant Applications

Wei Wang, Vikram Srinivasan and Mehul Motani

Department of Electrical and Computer Engineering
National University of Singapore, Singapore
{wang.wei,elevs,motani}@nus.edu.sg

ABSTRACT

In many delay tolerant applications, information is opportunistically exchanged between mobile devices who encounter each other. In order to effect such information exchange, mobile devices must have knowledge of other devices in their vicinity. We consider scenarios in which there is no infrastructure and devices must probe their environment to discover other devices. This can be an extremely energy consuming process and highlights the need for energy conscious contact probing mechanisms. If devices probe very infrequently, they might miss many of their contacts. On the other hand, frequent contact probing might be energy inefficient. In this paper, we investigate the trade-off between the probability of missing a contact and the contact probing frequency. First, via theoretical analysis, we characterize the trade-off between the probability of a missed contact and the contact probing interval for stationary processes. Next, for time varying contact arrival rates, we provide an optimization framework to compute the optimal contact probing interval as a function of the arrival rate. We characterize real world contact patterns via Bluetooth phone contact logging experiments and show that the contact arrival process is self-similar. We design STAR, a contact probing algorithm which adapts to the contact arrival process. Via trace driven simulations on our experimental data, we show that STAR consumes three times less energy when compared to a constant contact probing interval scheme.

Categories and Subject Descriptors: C.2.1 [Network Architecture and Design]: Wireless communication

General Terms: Measurement, Design, Algorithm.

Keywords: Delay Tolerant Networking, Bluetooth, Energy efficiency.

1. INTRODUCTION

Since its inception, the goal of networking research has been to provide instant, anytime, anywhere access to information. However, in recent times, research interest has been piqued by a new class of applications which are tolerant to delay. In several of these applications, information is exchanged opportunistically between devices when they are within communication range of each other. In other words, information transport is governed by the mobility of information carriers and their underlying contact patterns.

The notion of delay tolerance is useful in a variety of scenarios. One compelling application is providing connectivity and network services during disasters and in rural environments, where network infrastructure is minimal or nonexistent. Another example comes from software developers who are developing dating applications for mobile phones. The profile of an ideal partner is entered into a Bluetooth based mobile phone, which alerts the user whenever a matching profile is detected in the vicinity (e.g., www.bedd.com).

Current research and development efforts for delay tolerant applications fall broadly into two categories, Delay Tolerant Networking (DTN) [1] and Delay Tolerant Databases (DTD). For DTN applications, the goal is to enable communication between specific source-destination pairs in the network. Research in this area has involved studying algorithmic issues such as routing in networks [2], fundamental issues such as scaling laws [3] and performance bounds of routing algorithms [4] based on real world contact patterns.

Delay Tolerant Database applications have been driven by the observation that mobile devices are becoming increasingly powerful in terms of computation and storage and have multiple radio interfaces such as Bluetooth, 3G, WiFi etc [5]. An effort is also being made by phone manufacturers to embed sensors in these phones to acquire and store personal information (health related) and for environmental monitoring [6]. As a consequence, these devices store large volumes of digital information such as songs, photographs and sensory data and constitute a distributed geographic database. The dating application stated earlier is an example of a DTD application. The research community has investigated opportunistic query propagation and data aggregation algorithms, based on device proximity, in [5], [7], [8], and [9].

For both DTN and DTD applications, the common fundamental primitive is the opportunistic exchange of information between mobile devices when they are in communication range of each other. In order to enable such exchanges, devices will have to constantly probe the environment to discover other devices in the vicinity. Not surprisingly, device discovery[1] is an extremely energy consuming process. To

MobiCom'07, September 9–14, 2007, Montréal, Québec, Canada.

[1] We use device discovery and contact probing interchangeably in this paper.

understand this better, we made measurements on a Nokia 6600 mobile phone. The current drawn was (*i*) 38.61mA for Bluetooth device discovery, (*ii*) 9.33mA for response to a device discovery, (*iii*) 51.47mA for Bluetooth data transfer and (*iv*) 38.68mA for making a phone call. In other words, the device discovery process is as energy intensive as making a phone call!

Our measurements clearly motivate the need for energy conscious device discovery algorithms. One strategy to conserve energy is to increase the time between subsequent device discoveries. The consequence of this is that devices in the vicinity may go undiscovered and opportunities to exchange data are lost. This points to a trade-off between energy and missed opportunities. For strategies which use a constant device discovery interval, the larger the probing interval, the larger the missed opportunities and vice-versa. However, in stochastic environments, device discovery should be done adaptively by choosing the probing interval based on the state of the environment. For example, late at night at home, device discovery can occur at large intervals without missing many contacts, while on the subway to work, device discovery should be done more frequently to catch the myriad of new contacts.

In this paper, we investigate the design of energy conscious, adaptive contact probing algorithms which trade-off energy consumption and the probability of missing a contact. Specifically, we make the following contributions:

- *Theoretical Framework:* We first lay the theoretical base to characterize the trade-off between the average contact probing interval T and the contact missing probability P_{miss}. We show that if the contact duration distribution is *i.i.d.* (independent and identically distributed) and the contact arrival rate is constant, then for a given missing probability constraint, the optimal contact probing interval is (*i*) constant and (*ii*) depends only on contact duration distribution, independent of the contact arrival distribution. When the contact arrival rate is time varying, the optimal contact probing interval is a function of the contact arrival rate. We provide an optimization framework to compute the optimal contact probing interval as a function of the contact arrival rate. This theoretical base provides us with bounds on performance and also aids in the investigation and design of adaptive contact probing algorithms. See Sections 2, 3.
- *Real World Contact Pattern Experiments and Analysis:* To understand real world contact patterns, we conducted a large scale data logging experiment [10]. Nine volunteers were given Bluetooth devices equipped with a software program which probed for contacts every 30 seconds and logged information about other Bluetooth devices which came within range. Our database contains the largest number of unique devices discovered, compared to existing work [11], [12], [13]. We conduct rigorous analysis on the data. We confirm that the contact duration is Pareto distributed. Moreover, our data analysis indicates weak correlations in contact patterns at 24 hour time lags. Finally, we show that the contact process is self-similar with a Hurst parameter of 0.76. See Section 4.
- *Algorithm design and validation:* Finally, using insights gleaned from the theory and the data analysis, we investigate adaptive contact probing algorithms. We compare the performance of our algorithms against two benchmark algorithms. The first, a non-adaptive scheme provides a lower bound to performance, while the second is an idealized genie aided algorithm which utilizes non-causal information to determine the contact probing interval. We evaluate the performance of our algorithms via trace driven simulations on the data which we have collected. We then design STAR, an algorithm which strives to estimate the arrival rate and adapt the contact probing interval based on this estimate. Our results indicate that for a given contact missing probability constraint, STAR's energy consumption is one third of the constant non-adaptive contact probing algorithm. We also compare STAR with an AIMD (Additive Increase/Multiplicative Decrease) algorithm. STAR can conserve up to 20% more energy than the AIMD algorithm. See Section 5.

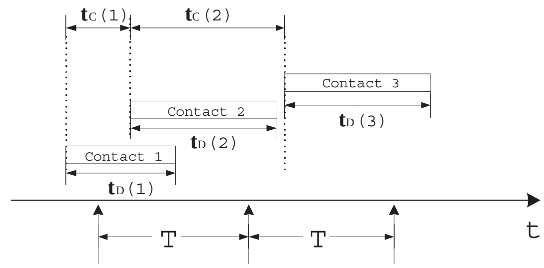

Figure 1: Illustrating the contacts for a specific user probing at a constant interval T.

2. MODELING THE CONTACT PROCESS

2.1 System Model and Assumptions

Assume that every device probes the environment governed by some contact probing algorithm. When a device A, probes its environment, all devices which hear the probe respond to device A with some information (e.g., identity, services available etc.). Based on this information, A may choose to exchange data with some of its neighbors. We define two devices A and B, to be in *contact*, if they are within communication range of each other[2]. The duration over which devices A and B are continuously in contact is called the *contact duration* for that contact. If neither device probes its environment during the contact duration, then we have a *missed contact*. We further assume for the sake of analysis that each probe is an impulse and does not consume any time.

We assume that each probe consumes equal energy so that the energy consumption rate of the device can be converted to the average probing frequency. With the same number of missed contacts, the algorithm which uses fewer probes (longer average probing interval) will be more energy efficient.

For a given device, we assume that the contact durations $\mathbf{t}_D(i)$ are *i.i.d. stationary* random variables with CDF (Cumulative Distribution Function) of $F_D(x)$ and mean $\mathbb{E}\{\mathbf{t}_D\} = 1/\mu$. We assume that the time between subsequent contacts, defined as *inter-contact time* $\mathbf{t}_C(i)$ are *stationary* random variables with CDF of $F_C(x)$, and mean $\mathbb{E}\{\mathbf{t}_C\} = 1/\lambda$. See Fig. 1 for an illustration of contact duration $\mathbf{t}_D(i)$ and inter-contact time $\mathbf{t}_C(i)$.

[2] We do not assume a perfect communication region here. The imperfectness in communication is embedded in our contact data gathering process which uses real devices.

2.2 Missing Probability

The missing probability P_{miss} is the probability that a contact which occurs is not detected. For the following analysis, we assume that for a device A, a contact B is missed, if B is not discovered by A's probes. We will relax this later to compute the missing probability when neither A nor B's probes discover each other. Let us first consider the simplest possible contact strategy, where each device probes for contacts at constant intervals of T (See Fig. 1).

If $\mathbf{t}_D > T$, the contact will never be missed. Consider a contact i which is initiated at time $[nT - x, nT - x + dx]$, where dx is an arbitrarily small value so that the time interval can be treated as a time point. This contact will not be detected by the mobile at time nT if $\mathbf{t}_D(i) < x$, which happens with probability of $F_D(x)$. By Blackwell's Theorem in renewal theory [14], when the contact process is a general renewal process with a nonlattice distribution function, we can calculate the long-term average missing probability given the probing interval of T as:

$$\lim_{n\to\infty} P_{miss}(T) = \frac{1}{T}\int_0^T F_D(x)dx \qquad (1)$$

Note that we do not need to restrict the inter-contact time distribution to be memoryless to perform the averaging procedure in Eq.(1).

Consider the mean value of $\mathbf{t}_D$ given that $\mathbf{t}_D < T$, which is $\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D < T\} = \int_0^T [1 - F_D(x)/F_D(T)]dx$, then $P_{miss}(T)$ can be expressed as:

$$P_{miss}(T) = \frac{T - \mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D < T\}}{T} \times F_D(T) \qquad (2)$$

For a fixed contact probing interval of T, any contact with duration larger than T will always be detected. Eq.(2) shows that for a contact with duration shorter than T, the expected probability that the contact will be missed is $\frac{T - \mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\}}{T}$. So, only contacts with duration smaller than T will be missed. If all the contacts with $\mathbf{t}_D < T$ have zero contact duration, then $P_{miss}(T)$ will be exactly $F_D(T)$.

The key observation here is that for a constant probing scheme, if the contact duration is stationary and *i.i.d* and the inter-contact times are stationary, then the missing probability depends only on the distribution of the contact duration and the contact probing interval T. It is independent of the inter-contact time distribution.

2.3 Optimal contact probing Scheme

We will now prove the following theorem about the fixed contact probing scheme described earlier.

THEOREM 1. *Consider a contact process for which the contact duration distribution is stationary and* i.i.d *and the inter-contact time distribution is stationary, with an expected inter-contact-time of* $\frac{1}{\lambda}$. *Consider the class of contact probing strategies which do not exploit knowledge of the contact process. Then, among all contact probing strategies with the same average contact probing interval, the strategy which probes at constant intervals achieves the minimum missing probability.*

PROOF. Consider a large interval of time L. Let us consider all strategies which probe the environment n times in this interval. As shown previously, for the strategy which probes at constant intervals $T = \frac{L}{n}$, the missing probability over duration L is $P_{miss}(T) = \frac{1}{T}\int_0^T F_D(x)dx$. Assume that an arbitrary scheme probes n times at $t_1, t_2, \ldots, t_n$. Define $t_0 = 0$ and $t_{n+1} = nT$, then we have $n+1$ intervals of $I_1 = t_1 - t_0, I_2 = t_2 - t_1, \ldots, I_{n+1} = t_{n+1} - t_n$. Since the scheme selects probe time t_k without knowledge of the contact process, the expected number of missed contacts in an interval of $[t_{k-1}, t_k]$ is $\lambda\int_0^{I_k} F_D(x)dx$, for $k \neq n+1$. All the contacts which occur in $[t_n, nT]$ will be missed. The expected missing probability is:

$$\hat{P}_{miss} = \frac{1}{\lambda L}\left[\sum_{k=1}^{n}\lambda\int_0^{I_k} F_D(x)dx + \lambda I_{n+1}\right] \qquad (3)$$

since the expected number of contacts arriving in duration L is λL. For $I_k \geq T$, we have

$$\begin{aligned}\lambda\int_0^{I_k} F_D(x)dx &= \lambda\left[\int_0^T F_D(x)dx + \int_T^{I_k} F_D(x)dx\right] \\ &\geq \lambda\int_0^T F_D(x)dx + \lambda\int_T^{I_k} F_D(T)dx \\ &= \lambda\int_0^T F_D(x)dx + \lambda(I_k - T)F_D(T) \quad (4)\end{aligned}$$

Similarly, when $T_k < T$ we have:

$$\begin{aligned}\lambda\int_0^{I_k} F_D(x)dx &= \lambda\left[\int_0^T F_D(x)dx - \int_{I_k}^{T} F_D(x)dx\right] \\ &\geq \lambda\int_0^T F_D(x)dx - \lambda\int_{I_k}^{T} F_D(T)dx \\ &= \lambda\int_0^T F_D(x)dx + \lambda(I_k - T)F_D(T) \quad (5)\end{aligned}$$

We also have:

$$\lambda I_{n+1} \geq \lambda I_{n+1} F_D(T) \qquad (6)$$

Putting all these into Eq.(3), we have:

$$\begin{aligned}\hat{P}_{miss} &\geq \frac{1}{\lambda nT}\Big[\sum_{k=1}^{n}\lambda\left(\int_0^T F_D(x)dx + (I_k - T)F_D(T)\right) \\ &\qquad + \lambda I_{n+1}F_D(T)\Big] \\ &= \frac{1}{nT}\left[\sum_{k=1}^{n}\int_0^T F_D(x)dx + F_D(T)(\sum_{k=1}^{n+1} I_k - nT)\right] \\ &= P_{miss}(T) \quad (7)\end{aligned}$$

□

2.4 Trade-offs in Energy Efficiency and P_{miss}

Having established that a constant contact probing interval is optimal under certain assumptions, we now explore the trade-off between energy efficiency and the missing probability. When contact durations are distributed according to a given distribution, we can analytically determine the relationship between T and $P_{miss}(T)$ according to Eq.(1). Here, we demonstrate the relationship between energy efficiency and missing probability for several distributions. In section 4 we will focus on distributions obtained from real world Bluetooth contact logs.

a). Poisson Process

For Poisson process, we have $F_D(x) = 1 - e^{-\mu x}$. Using Eq.(1), we have:

$$P_{miss}(T) = \frac{\mu T - 1 + e^{-\mu T}}{\mu T} \qquad (8)$$

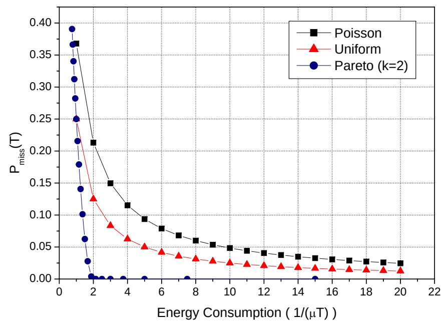

Figure 2: Trade-off between energy consumption and missing probability on different distributions.

b). Uniform Distribution
The uniform distribution is:

$$F_D(x) = \begin{cases} 0 & x < 0 \\ \frac{\mu x}{2} & 0 \leq x \leq 2/\mu \\ 1 & x > 2/\mu \end{cases} \quad (9)$$

And, we have:

$$P_{miss}(T) = \begin{cases} \frac{\mu T}{4} & T < \frac{2}{\mu} \\ \frac{\mu T - 1}{\mu T} & T \geq \frac{2}{\mu} \end{cases} \quad (10)$$

c). Pareto Distribution
We have:

$$F_D(x) = \begin{cases} 0 & x < \tau \\ 1 - (x/\tau)^{-k} & x \geq \tau \end{cases} \quad (11)$$

In this case, we have $1/\mu = k\tau/(k-1)$ when $k > 1$. The mean is unbounded when $k \leq 1$. Using Eq.(1), we have:

$$P_{miss}(T) = 1 + \frac{k\tau}{T(1-k)} - \frac{\tau^k}{T^k(1-k)}, T > \tau \quad (12)$$

Fig. 2 shows the trade-off between energy consumption and missing probability for these distributions. The energy consumption is computed as $\frac{1}{\mu T}$, which is the contact probing rate normalized by the average contact duration (i.e., the number of probes taken during the average contact duration). We see that for Poisson and uniform distributions, the missing probability of $5 - 10\%$ is near the knee of the curve which is a good trade-off point between energy consumption and missing probability. This means the contact probing interval should be around $1/6 - 1/3$ of the average contact duration. However, for Pareto distribution, the contact probing interval should be around τ to achieve a near zero missing probability. In other words, for a constant arrival rate and a Pareto contact duration distribution, it is difficult to trade-off between P_{miss} and T.

2.5 Double Detection

As we stated earlier, a contact between device A and B is missed only if neither device probes the environment during the contact. Consider the case when two users A and B are independently probing the environment. Assume that both users are using the same fixed contact probing interval of T. Then, the probability that A does not discover B is $P_{miss}(T)$. However, the probability that neither A nor B discovers the other during a contact is much higher than $P^2_{miss}(T)$, even though their contact probing processes are independent. Suppose one user probes at times of $T, 2T, \ldots, nT$, and the other probes at $y, y+T, \ldots, y+(n-1)T$. Without loss of generality, we can assume that $y < T/2$. Then, the probability that during a contact, neither user discovers the other is given by:

$$\tilde{P}_{miss}(T, y) = \frac{1}{T}\left[\int_0^y F_D(x)dx + \int_0^{T-y} F_D(x)dx\right]$$

When $y = T/2$, $\tilde{P}_{miss}(T, y)$ has a minimum value of $P_{miss}(T/2)$, and $\tilde{P}_{miss}(T, y)$ has maximum value of $P_{miss}(T)$ when $y = 0$. Since the two users are probing independently, y is uniformly distributed in $[0, T/2]$, and the average missing probability is:

$$\begin{aligned}\tilde{P}_{miss}(T) &= \frac{2}{T}\int_0^{\frac{T}{2}} \frac{1}{T}\left[\int_0^y F_D(x)dx + \int_0^{T-y} F_D(x)dx\right]dy \\ &= \frac{2}{T^2}\left[\int_0^{\frac{T}{2}}\int_0^y F_D(x)dxdy + \int_0^{\frac{T}{2}}\int_0^{T-y} F_D(x)dxdy\right] \\ &= \frac{2}{T^2}\left[\int_0^{\frac{T}{2}}\int_0^y F_D(x)dxdy + \int_{\frac{T}{2}}^{T}\int_0^{y} F_D(x)dxdy\right] \\ &= \frac{2}{T^2}\int_0^T\int_0^y F_D(x)dxdy \end{aligned} \quad (13)$$

For example, when the contact duration is uniformly distributed as shown in Eq.(9), we have $\tilde{P}_{miss}(T) = \frac{\mu T}{6} = \frac{2}{3}P_{miss}(T)$, which is smaller than the single user missing probability only by a constant factor.

2.6 Bounding the Missing Probability

In the real world, the contact duration distributions as seen by a given mobile device may change over time. This could occur for a variety of reasons, e.g., relocation to a new city or increasing Bluetooth penetration. This implies that the mobile device will have to adapt its probing interval over large time scales. In order to do this adaptation effectively, it needs to estimate the missing probability. However the mobile has no knowledge of contacts it has missed. Here we show how a mobile device can estimate a bound on the missing probability purely from local information. We define a contact to be a short contact if it is detected in only one probe, but not its previous or subsequent probes. We define $P_{short}(T)$ as the probability that a device sees a short contact when the contact probing interval is T.

Consider the probability that a contact i occurs between $[nT - x, nT - x + dx]$, which is detected by the mobile at nT, but not at $(n+1)T$. Thus, the contact duration should be $x \leq \mathbf{t}_D(i) \leq x + T$. Averaging over $[0, T]$, the probability for a contact to be a short contact is:

$$\frac{1}{T}\int_0^T [F_D(x+T) - F_D(x)]dx \quad (14)$$

Note that the probability a contact will be detected is $1 - P_{miss}(T)$. Thus, $P_{short}(T)$ measured by the mobile will be the number of short contacts (note that short contacts are always detected) divided by the total number of detected contacts:

$$\begin{aligned}P_{short}(T) &= \frac{\frac{1}{T}\int_0^T [F_D(x+T) - F_D(x)]dx}{1 - P_{miss}(T)} \\ &\geq \frac{\frac{1}{T}\int_0^T [F_D(T) - F_D(x)]dx}{1 - P_{miss}(T)}\end{aligned}$$

$$
\begin{aligned}
&= \frac{\frac{F_D(T)}{T}\int_0^T [1-\frac{F_D(x)}{F_D(T)}]dx}{1-P_{miss}(T)} \\
&= \frac{F_D(T)\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\}}{T(1-P_{miss}(T))} \\
&= \frac{TP_{miss}(T)}{T-\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\}}\frac{\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\}}{T(1-P_{miss}(T))} \\
&= \frac{\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\}}{T-\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\}}\frac{P_{miss}(T)}{1-P_{miss}(T)} \qquad (15)
\end{aligned}
$$

We have:

$$P_{miss}(T) \le \frac{P_{short}(T)}{P_{short}(T)+\frac{\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\}}{T-E\{\mathbf{t}_D|\mathbf{t}_D<T\}}} \qquad (16)$$

When $\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\}$ is small, we will have $\frac{\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\}}{T-\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\}} \approx 0$, and the bound will be close 1. To avoid trivial cases when $\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\}$ goes to zero, we assume that contacts smaller than a duration τ can be ignored. For example, we only consider contacts with contact duration longer than $\tau > 30$ seconds in practice. This is reasonable in practical examples such as Bluetooth exchange, where a device discovery, followed by a service search in the presence of several Bluetooth devices, is of the order of a few seconds. In this case, we have $\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\} \ge \tau$ and $\frac{\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\}}{T-\mathbb{E}\{\mathbf{t}_D|\mathbf{t}_D<T\}} > \frac{\tau}{T}$. Thus:

$$P_{miss}(T) < \frac{P_{short}(T)}{P_{short}(T)+\frac{\tau}{T}} \qquad (17)$$

Note that this bound may be loose when $\tau \ll T$.

2.7 P_{short} for Different Distributions

We derive $P_{short}(T)$ for several useful distributions in this section. From the relationship between T and $P_{short}(T)$, we show $\frac{1}{2}P_{short}(T)$ can be used as an accurate estimate of P_{miss} for certain distributions.

a). Poisson Process

Using Eq.(15), we have:

$$P_{short}(T) = \frac{(1-e^{-\mu T})^2}{\mu T(1-P_{miss}(T))} \qquad (18)$$

b). Uniform Distribution

Using Eq.(15) and (9), we have:

$$P_{short}(T) = \frac{2\mu T}{4-\mu T} = \frac{2P_{miss}(T)}{1-P_{miss}(T)} \qquad (19)$$

when $T < 1/\mu$.

c). Pareto Distribution

Using Eq.(15) and (11), we have:

$$P_{short}(T) = \begin{cases} \frac{2\tau^k T^{1-k} - k\tau - \tau^k(2T)^{1-k}}{T(1-k)(1-P_{miss}(T))} & T \ge \tau \\ \frac{k\tau + 2T(1-k) - \tau^k(2T)^{1-k}}{T(1-k)(1-P_{miss}(T))} & \tau/2 \ge T < \tau \\ 0 & T < \tau/2 \end{cases}$$

The relationship between $P_{miss}(T)$ and $P_{short}(T)$ is shown in Fig. 3. We see that for both Poisson and uniform distribution, $\frac{1}{2}P_{short}(T)$ is quite close to $P_{miss}(T)$. With some algebra, we can also show that $P_{short}(T)/P_{miss}(T)$ converges to 2 as P_{miss} goes to zero for these distributions. This shows that $\frac{1}{2}P_{short}(T)$ can serve as an estimate of $P_{miss}(T)$ for these distributions.

However, the estimation error can be quite large for the Pareto distribution. The $\frac{1}{2}P_{short}(T)$ and $P_{miss}(T)$ for Pareto distribution with different values of k are plotted in Fig. 4.

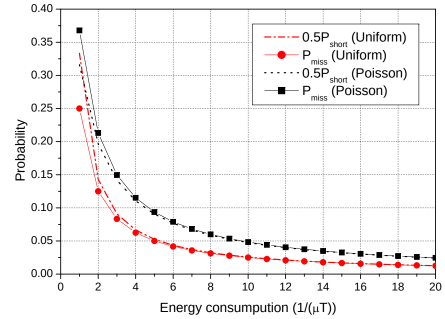

Figure 3: Relationship between $P_{miss}(T)$ and $P_{short}(T)$ for Poisson and uniform distribution.

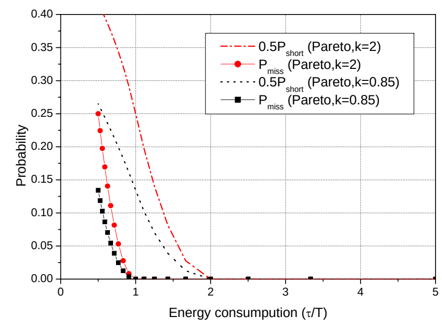

Figure 4: Relationship between $P_{miss}(T)$ and $P_{short}(T)$ for Pareto distribution. We use $\frac{\tau}{T}$ instead of $\frac{1}{\mu T}$ for energy consumption since $\frac{1}{\mu}$ is unbounded for $k = 0.85$

In this case, $P_{miss}(T)$ can be much smaller than $\frac{1}{2}P_{short}(T)$ so that we may unnecessarily drive $P_{miss}(T)$ to a small value when using $\frac{1}{2}P_{short}(T)$ to estimate it. This may result in the device choosing a small T and wasting energy. However, we can still safely upper bound P_{miss} by $\frac{1}{2}P_{short}(T)$ for the Pareto distribution.

3. TIME-VARYING INTER-CONTACT TIME DISTRIBUTIONS

Until now, we have made the simplifying assumption that the inter-contact-time distribution is stationary. This assumption is clearly not true. For example, one would expect the inter-contact-time patterns, late at night to differ significantly from the inter-contact-time distribution during lunch hours or peak traffic hours. This implies that the optimal contact probing interval will vary with time. In the following, we assume that the average contact arrival rate varies with time.

We assume that the distribution is constant over a duration L. We assume that the expected inter-contact-time in the i^{th} interval of duration L is $\frac{1}{\lambda_i}$.

From Theorem 1, for the i^{th} interval of duration L, there exists a fixed contact probing interval T_i, which is optimal. Then, the expected number of contacts captured during time period i will be $\lambda_i(1-P_{miss}(T_i))$. Assume that we have a constraint N, on the number of probes taken over a certain time period (e.g. you should probe the environment not more than 300 times in a 24 hour period). We can then formulate the following optimization problem:

$$\begin{aligned} Maximize \ & \textstyle\sum_i \lambda_i(1 - P_{miss}(T_i)) \\ s.t. \quad & \textstyle\sum_i \frac{L}{T_i} \leq N \\ & T_i \geq 0 \quad \forall i \end{aligned} \tag{20}$$

Defining the variable x_i as $1/T_i$ and $U_i(x_i) = \lambda_i(1-P_{miss}(1/x_i))$, the optimization in Eq.(20) can be recast as:

$$\begin{aligned} Maximize \quad & \textstyle\sum_i U_i(x_i) \\ s.t. \ & \textstyle\sum_i Lx_i \leq N \\ & x_i \geq 0 \quad \forall i \end{aligned} \tag{21}$$

It can be verified that $U(x_i)$ is always a concave and increasing function for the distributions we discussed above. Using the Karush-Kuhn-Tucker (KKT) condition [15], we have:

$$U'(x_i^*) = c \qquad \forall i \tag{22}$$

for the optimal solution of x_i^*, where c is a function of N, contact duration distribution and the set λ_i. For a given set of λ_i, c is determined by N. The larger the value of N, the smaller the value of c. In other words c determines the energy efficiency. The optimal T_i^* can be solved from:

$$T_i^{*2} P'_{miss}(T_i^*) = \frac{c}{\lambda_i} \tag{23}$$

for any given λ_i. Thus, we can just choose a fixed constant c and use Eq.(23) to calculate the optimal T_i^* based on the estimate of λ_i for a time period i. This solution minimizes the missing probability over all time, subject to energy constraints, which are determined by the constant c.

Example: Let the contact duration distribution be uniform, $L = 1$ hour, N is the number of times the device can probe in 24 hours. Then, the optimal contact probing interval T_i^* for the arrival rate λ_i is:

$$T_i^* = \sqrt{\frac{4c}{\mu\lambda_i}} \tag{24}$$

by Eq.(23). If the λ_i are all equal, then, clearly $T_i^* = \frac{24}{N}, \forall i$. Therefore $c = \frac{\mu\lambda}{4}(\frac{24}{N})^2$ and $P_{miss}^*(T_i^*) = \frac{6\mu}{N}$.

4. REAL-WORLD CONTACT PROCESSES

Before addressing the issue of designing adaptive contact probing algorithms, we first take a look at some real contact pattern data which we collected.

4.1 Data Collection Experiments

In order to characterize contact distributions, we handed Bluetooth phones to nine volunteers. We also installed static Bluetooth probes in high traffic areas. The phones had a J2ME program running on them which initiated a Bluetooth device discovery every 30 seconds. If other Bluetooth devices are discovered in the vicinity, then the time of contact and Bluetooth address were captured. Since the probing software could terminate due to lack of energy, crashes etc., we captured the start time and end time of each probing session. This allowed us to capture contact duration and inter-contact-time distributions accurately. Over all, we did 424 man days of data collection. 12,332 unique devices were discovered in our experiment. To the best of our knowledge, this is the largest volume of unique devices discovered compared to other comparable studies[3]. For the details of the experiment, please refer to [10].

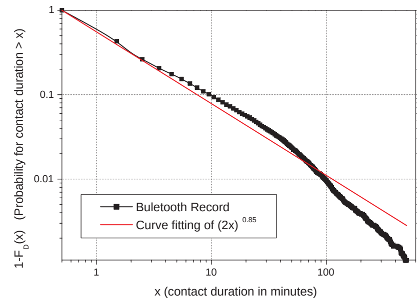

Figure 5: Contact duration distribution.

Each contiguous set of scans during which a device is discovered is counted as a *contact*. Assume that a device (say D) is discovered in m contiguous scans. Then the duration of the contact with D is the difference between the time of D's discovery in the m^{th} scan and the first scan. If a device is only detected in one scan, we treat the duration of contact as 30 seconds.

4.2 Distribution of Contact Duration

We first look at contact duration distribution of the Bluetooth contact data. Fig. 5 plots the $1 - F_D(x)$ curve (probability of $\mathbf{t}_D > x$) in log-log scale. We see that the contact duration distribution follows a power law (Pareto distribution). By curve fitting, we can estimate $F_D(x) = 1-(x/\tau)^{-k}$ with $\tau = 30$ seconds and $k = 0.85$. Although the mean for Pareto distribution with $k \leq 1$ is unbounded, the average contact duration is around 370 seconds in our case. This is because the probability of long contacts decreases rapidly when contact duration is longer than 2 hours, and in our experiments, there were no contacts lasting longer than 5 hours. The fact that the contact duration distribution follows a power law has also been verified by other studies (see [13] and references therein).

The fact that the contact duration distribution is Pareto distributed is a pessimistic result. For the Pareto distribution, beyond a small threshold probing interval, the missing probability rises sharply from close to zero as shown in Fig. 2. Fortunately, as a consequence of human mobility, contact arrival rates are time varying. As we show in later sections, this can be exploited to achieve significant energy savings.

4.3 Correlation Analysis

In order to see if there is any predictability in the contact patterns, which will aid us in designing good adaptive contact probing mechanisms, we investigate the autocorrelation of the contact processes.

Here, we use the autocorrelation function [16]: Define an aggregated measurement series $f(i)$ as the average of certain measurements, over a time window of $[(i-1)\tau', i\tau']$. We calculate the autocorrelation of the measurement as $A(n) = \sum_{i=1}^{\infty} f(i) \times f(i+n)$. In this paper, we consider the measurement of three types:

[3]In the Haggle [13] studies a maximum of 41 volunteers, probing at 120 second intervals, discovered 200 unique devices over 5 days. In the Serendipity [12] study, 100 volunteers, probing at 5 minute intervals, discovered 2798 unique devices over a 9 month period.

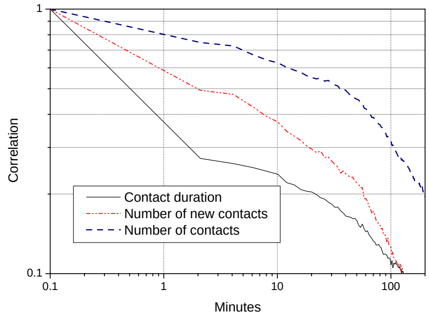

Figure 6: Normalized correlation of contact duration and contact number over short time period ($\tau' = 2$ minutes).

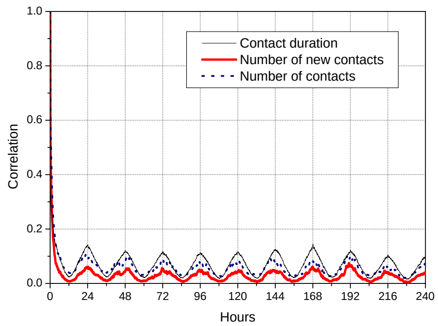

Figure 7: Normalized correlation of contact duration and contact number over long time periods ($\tau' = 5$ minutes).

• *Contact duration*: The average contact duration for the contacts occurring in a time window.

• *Number of new contacts*: The number of newly arrived contacts during the time window of $[(i-1)\tau', i\tau']$. This measurement determines the correlation of inter-contact-time.

• *Number of contacts*: The number of contacts the device observes during the time window, including those which arrived in previous windows.

As shown in Fig. 6, the number of contacts and number of new contacts were highly correlated over short periods and the correlation dropped faster when the correlation period is longer than 60 minutes. The contact duration was less correlated. This indicates that the number of contacts or new contacts are useful metrics which can be used to design adaptive contact probing algorithms.

4.4 Correlation at Large Time Scales

The contact processes are highly related to human behavior. Therefore, we expect to see diurnal variations in the contact patterns. Also, since most people are creatures of habit, one might expect to see periodicity in the contact processes over daily or weekly time scales. Fig. 7 shows the correlation over long time periods. We clearly see the diurnal variations. The autocorrelation dropped to zero over 12 hours and had a period of 24 hours. Finally, we do not see any significant correlation increase over 168 hours (one week) period.

4.5 Self-Similar Nature

The slowly decreasing slope of the autocorrelation for the number of new contacts hints that the contact arrival process

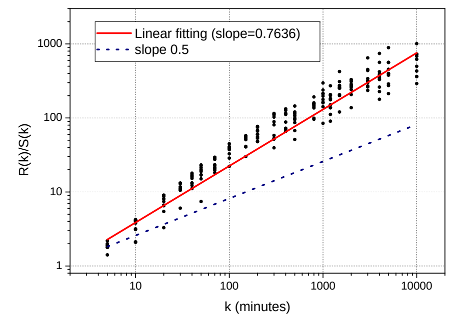

Figure 8: The R/S statistic of the contact arrival process.

is self-similar [17], [18]. We test for self-similarity over four time scales of 10, 100, 1000 and 10,000 minutes. We use the R/S measurement to verify this conjecture. Let X_i denote the number of new contacts seen in one minute and $Y_j = \sum_{i=1}^{j} X_i$ be the aggregated number of new contacts in j minutes. Define,

$$R(t,k) = \max_{0\le i\le k}[Y_{t+i} - Y_t - \frac{i}{k}(Y_{t+k} - Y_t)] - \min_{0\le i\le k}[Y_{t+i} - Y_t - \frac{i}{k}(Y_{t+k} - Y_t)]$$

and

$$S(t,k) = \sqrt{\frac{1}{k}\sum_{i=t+1}^{t+k}(X_i - \overline{X_{t,k}})^2}, \overline{X_{t,k}} = \frac{\sum_{i=t+1}^{t+k} X_i}{k}$$

If $E[R(k)/S(k)] = Ck^H$, with *Hurst* parameter $H \in (0.5, 1)$, then the process is self-similar [18]. From Fig. 8, we see that our process has a Hurst parameter close to 0.76.

This implies that the contact arrivals are not only long range dependent but also bursty. An intuitive model for this is that arrivals follow an ON-OFF process where the ON and OFF durations have memory. This can potentially be exploited to achieve energy savings.

4.6 Contact Arrival Rate vs. Time of Day

The contact arrival rate distribution varies with time of day. In Fig. 9, we plot the average rate at which new contacts are seen at different hours of the day. The contact arrival rate during the early morning was quite small. This implies that we can use longer contact probing intervals during the early morning to save energy as described in section 3. The variance of the contact arrival rate over time of day is also plotted in Fig. 9. We see that the variance was quite large. This shows the contact arrival rate at the same time period in different days may vary drastically. Therefore, we can not simply use the time of day to infer the contact arrival rate. As we show in section 5, estimation of arrival rate will be the key to the contact probing interval selection problem.

5. ADAPTIVE PROBING ALGORITHMS

5.1 Assumptions and Methodology

In this section we investigate the design of adaptive contact probing algorithms which achieve a good trade-off be-

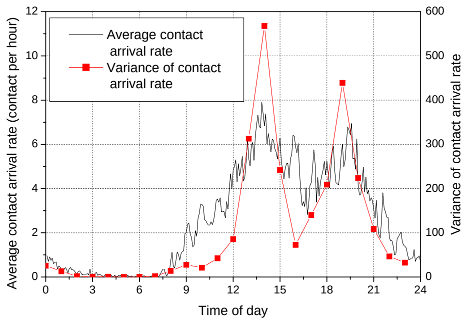

Figure 9: Contact arrival rate and its variance over time of day.

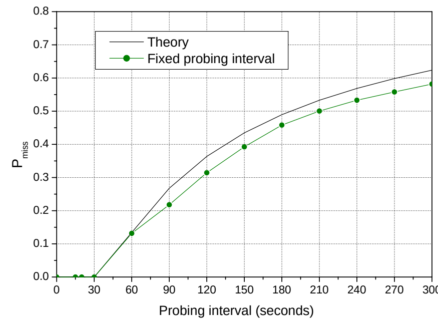

Figure 10: Comparing the missing probability of the Bluetooth log to theory results.

tween the probability of a missed contact and energy consumption. As described in section 2.1, the energy consumption rate is inversely proportional to the average probing interval, e.g., a device with average probing interval of 30 seconds will consume two times energy than a device with average probing interval of 60 seconds. From our earlier analysis in section 3, we see that in order to compute the optimal contact probing interval, we require k, the power law exponent for the contact duration distribution and the arrival rate. From the analysis of our logging data in section 4, we see that arrival rates vary over very short time scales. On the other hand, we contend that variations in the contact duration distribution, if any, will occur over large time scales. It is reasonable to hypothesize that contact duration distribution will always follow a power law, but with different decay coefficients k over large time scales. This is validated by the fact that experiments at widely differing geographical locations have all displayed a power law contact duration distribution, but with different decay coefficients k [13], [11], [19]. For the rest of this section, we assume that the contact duration distribution is known, and focus on the problem of adapting the contact probing interval to variations in arrival rates.

We compare the performance of different adaptive contact probing algorithms by running trace driven simulations on the data that we have collected. Our data set consists of contacts logged at a time granularity of 30 seconds.

When calculating the missing probability, we use the Bluetooth contact logs, which are based on a 30 second contact probing interval as the baseline. We apply different adaptive contact probing algorithms to filter the contact log data. Assume that a device D has been logged in our experiments from time t to $t + \tau$. Then, for a specific adaptive contact probing algorithm, we say that the contact has been missed, if the algorithm does not initiate any probe in the interval $[t, t + \tau]$. The contact missing probability is computed as the ratio of the number of contacts missed by the adaptive contact probing algorithm to the total number of contacts made in our data logging experiments.

5.2 Performance Bounds

First, for a constant contact probing interval T, we verify that the relationship between T and P_{miss} derived in equation (12) is correct. From our logging data, we have computed the decay exponent of the contact distribution $k = 0.85$. Since the granularity of our logging experiments was over 30 second intervals, we set $\tau = 30$. The comparison between the trace driven simulations and the analytical results are shown in Fig. 10. We see that the results of the simulation are quite close to the theoretical results. Since the constant contact probing interval algorithm does not adapt to the changes in the contact arrival rate, we will use this as the lower bound for the average contact probing interval to achieve a given missing probability.

We now turn our attention to algorithms which can adapt to the arrival rate. As shown in section 5, if the arrival rate is known, we can find the optimal value of the contact probing interval by solving (23). Since our contact duration is Pareto distributed, we have:

$$T_i^* = \tau \left(\frac{c(1-k)}{\lambda_i k \tau} + 1 \right)^{\frac{1}{1-k}}, \qquad (25)$$

where $k = 0.85$ and $\tau = 30$ seconds for our logging data. By changing the constant c, we can achieve different trade-off points between P_{miss} and T.

We first consider an idealized genie aided algorithm. In this algorithm, the adaptation is done over one hour blocks, i.e., the contact probing interval remains constant for one hour (e.g., $T = 5$ minutes from 1000 - 1059hrs, $T = 7$ minutes from 1100 - 1159hrs etc.). Furthermore, we assume that a genie provides the mobile device with an accurate forecast of the contact arrival rate λ_i for the subsequent hour i. The mobile device then sets the arrival rate to T_i^* computed from equation (25) for the next hour. The result of this ideal algorithm is compared to the fixed contact probing interval algorithm in Fig. 11. For the same missing probability of 10%, the ideal scheme uses an average contact probing interval of 300 seconds while the fixed contact probing interval scheme uses 50 seconds. This indicates that adaptation to the arrival rate can result in significant energy savings for a given missing probability constraint.

The improvement in energy consumption comes from the fact that human behavior is highly dynamic so that the contact arrival rate varies substantially over small time scales. Therefore, schemes which adapt their contact probing interval to the contact arrival rate can save significant amounts of energy. Unfortunately, the ideal algorithm is noncausal since it uses the future arrival rate in selecting the contact probing interval. This motivates us to design algorithms which adapt their contact probing interval based on estimates of the contact arrival rate.

5.3 Heuristic Algorithms

In this section, we discuss heuristic algorithms to estimate the optimal contact probing interval based on historical contact records. We first study the relationship between the

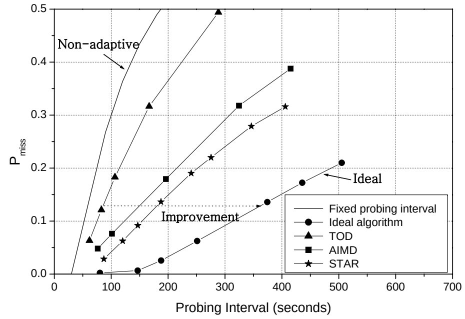

Figure 11: Improving energy efficiency by adapting to the contact arrival rate.

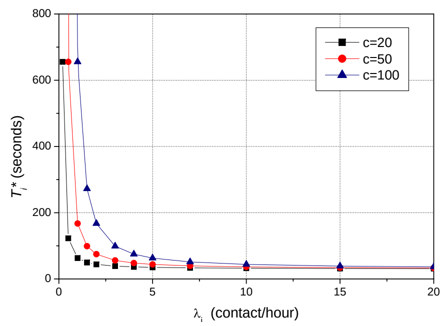

Figure 12: Relationship between optimal contact probing interval and the arrival rate.

optimal T_i^* and λ_i to guide our design of heuristics. Fig. 12 shows the optimal contact probing interval for different arrival rates. We see that when the arrival rate is high the optimal T_i^* quickly converges to 30 seconds. This means we can tolerate a certain amount of estimation error in λ_i when the contact arrival rate is high. However, when the contact rate is small, we need to estimate it accurately so that we will not underestimate the rate and use extremely long contact probing intervals.

We compare the performance of three heuristics:

• *TOD* (Estimation based on time of day): This heuristic is motivated by two facts, first people are creatures of habit and Fig. 7, which shows that the auto-correlation peaks at 24 hour lags. The goals is to see whether maintaining long term history of arrival statistics, will be useful in adaptation. For this heuristic, we adapt the contact probing interval over one hour blocks just as in the ideal genie aided algorithm. We assume that perfect knowledge of the average arrival rate at different hours in the past is available to the mobile device (e.g., average arrival rate between 1000hrs – 1059hrs in the past). The contact probing interval for any hour is then determined by Eq.(25).

• *AIMD* (Additive Increase / Multiplicative Decrease): We test this algorithm for two reasons. First, it is a robust adaptive algorithm used commonly in networking protocols such as TCP. Second, our observation that the contact arrival process is self-similar and consequently bursty, implies that the contact probing interval should decrease sharply when a new contact is seen. The algorithm adjusts the contact probing interval based on the presence or absence of new contacts seen in previous probes. If no new contacts

---

Calculating the contact probing interval T_i in probe i at time t_i

if n_i new contacts detected in this probe
    Set estimated new contacts in $(t_i - T_{i-1}, t_i]$ as
    $\hat{n}_i = n_i/(1 - P_{miss}(T_{i-1}))$,
    /* $P_{miss}(T_{i-1})$ *computed from Eq.*(12) */
endif
if At least one new contact detected in last m minutes
    Set $\lambda_i = \frac{60}{m} \sum_{t_j \in [t_i - m, t_i]} \hat{n}_j$
    Set $w = 1$
else
    $w = w + 1$.
    $\lambda_i = \lambda_{i-1} \times (w-1)/w$.
endif
if Time of day is 8 o'clock and $\lambda_i < 1$;
    Increase λ_i by 1.
endif
Set T_i based on Eq.(25).

Figure 13: STAR: Short Term Arrival Rate estimation algorithm

appear, it linearly increases the contact probing interval. Otherwise, it decreases the contact probing interval by a factor of 2. Therefore, the contact probing interval quickly decreases when new contacts arrive, while it slowly increases when the device sees no new contacts. We also use time of day information explicitly by setting $T_i = 90$ seconds at 8 AM in the morning. This is to adapt effectively to the sharp change in arrival rate from night to day. The actual time for resetting T_i can be either 8 AM or 9 AM, which has minor effects on the algorithm performance.

• *STAR* (Short Term Arrival Rate): STAR uses what we have learned earlier from the theoretical and data analysis. If the arrival rate can be estimated accurately over a short time interval, then the optimal contact probing interval can be chosen based on Eq.(25). The design of the algorithm is motivated by the self-similar observation. This implies that when a new contact is seen, then the estimate of the arrival rate must rise sharply. Similarly since the OFF duration (duration when there are no arrivals) also has memory, one might argue that when no arrivals are seen, one should decrease the estimate of the arrival rate sharply. However from Fig. 12, we see that when λ decreases, the contact probing interval T increases dramatically. Therefore, decreasing the estimate of the arrival rate too quickly, will result in large contact probing intervals and possibly several missed contacts. Therefore, we argue that the arrival estimate should decay slowly. The description of the STAR algorithm is given in Fig. 13. Note that, as in AIMD, we increase the arrival rate estimate at 8 AM. We estimate the arrival rate over a short time window of m minutes. If no arrivals are seen in the last m minutes, then the estimated arrival rate decays as a power law with exponent 1. If a new contact is seen during a probe, the number of potential new contacts since the last probe is estimated and used to estimate the arrival rate. For the results shown in this paper, we use $m = 5$.

For all algorithms, we only allow T_i to vary between 30 – 1800 seconds. Fig. 11 shows the relationship of the missing probability versus the average probing intervals. Recall that a algorithm is more energy efficient when it can achieve the same missing probability with longer probing interval. We

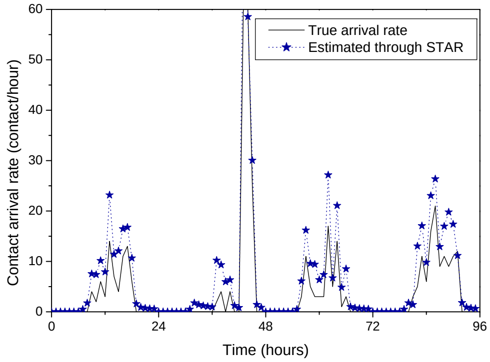

Figure 14: The figure shows the arrival rate estimate of STAR averaged over one hour blocks vs. the true arrival rate over the same one hour block

see that the TOD heuristic performs much poorer than the STAR and AIMD algorithms. For the same missing probability, the average contact probing interval for the TOD heuristic is at most 1.5 times larger than the non-adaptive scheme. This implies that there is no benefit to maintaining long term history about arrival rates.

The other two heuristics, AIMD and STAR, use the fact that the contact arrival rate is highly correlated in the short term (Fig. 6). The performance is much better since the estimation of arrival rates is more accurate. For example, at a missing probability of 20%, the average contact probing interval used by AIMD can be nearly 3 times larger than non-adaptive scheme and 2 times larger than TOD. In other words, the AIMD can use one third contact probes to achieve the same missing probability. The STAR algorithm is the best among the three heuristics and its average contact probing interval is around 20% larger than the AIMD heuristic for the same missing probability.

5.4 Analysis of the Heuristics

To better understand the performance of the STAR heuristic, we plot the arrival rate estimates of STAR in Fig. 14. We see that the estimated arrival rate is quite close to the real value except that it sometimes overestimates the arrival rate. As in our previous discussion, overestimation will not greatly impact the contact probing interval when λ_i is large. However, when λ_i is small, we will use a much smaller contact probing interval than the ideal algorithm. As shown in Fig. 15, although we only slightly overestimate the arrival rate during the night periods (e.g., time periods from 48-51 hours), the contact probing interval is much smaller than the ideal algorithm. This explains the gap between our heuristics and the ideal algorithm. Compared to the AIMD heuristic, the STAR algorithm increases the contact probing interval faster when the arrival rate is low and gives quicker response when the arrival rate increases, as shown in Fig. 15(b).

Fig. 16 shows the short term missing probability of STAR. We see that the short term missing probability can some times reach 100%. This is due to the fact that the STAR algorithm uses long contact probing intervals when the estimated arrival rate is low. However, as we see in Fig. 16, the missing probability is high only in time periods where contact arrival rate increases rapidly. The number of con-

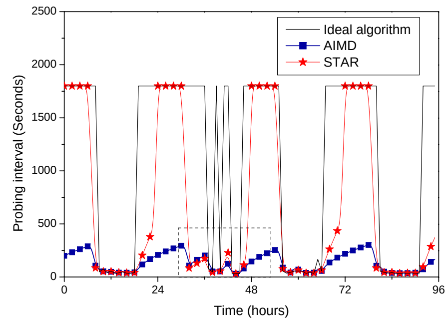

(a) Average contact probing interval

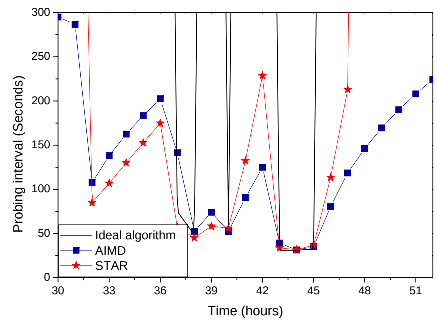

(b) Enlarged figure in the dashed box

Figure 15: Comparison of the estimated contact probing interval of STAR and AIMD versus the ideal algorithm

tacts in time periods when there is a big difference between the actual arrival rate and estimated arrival rate is small. Therefore, even though the missing probabilities in these time intervals are large, the average missing probability over time is small. As shown in section 2, the constant probing interval scheme is optimal if the application requires a hard guarantee on the missing probability over every time period. However, this scheme consumes more energy than adaptive schemes as shown in Fig. 11.

Finally, we discuss the relative performance of STAR versus AIMD. Note that although STAR tracks the ideal algorithm much more closely, its average contact probing intervals is only 20% better than AIMD. This is because when the contact arrival rate is large, the contact probing interval for both STAR and AIMD is small. However, the number of occasions when STAR's contact probing interval is much larger than AIMD (i.e. 1800 seconds for STAR), is relatively small. Consequently, upon averaging, STAR's average contact probing interval is only 20% larger than AIMD. An important fact to note that AIMD is much simpler than STAR and requires no prior knowledge of the contact distribution and yet uses only around 20% more energy than STAR. However, we see that the performance of STAR is not very sensitive to the parameters of contact duration distribution. Fig. 17 shows STAR running with different k values from 0.7–1.2. We see that the performance with inaccurate distribution parameters are very close to the one with $k = 0.85$. This significantly simplifies implementation of STAR, since exact knowledge of the parameters of the contact duration distributions is not required.

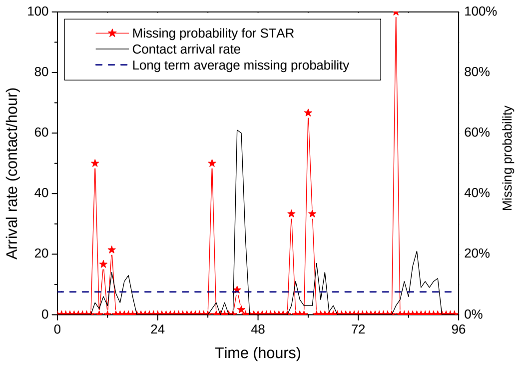

Figure 16: Short term missing probability of the STAR algorithm.

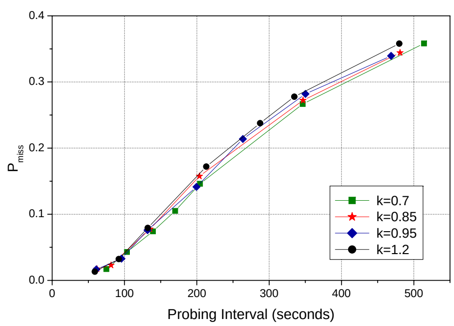

Figure 17: Performance of STAR with inaccurate k estimations.

5.5 Proof of Concept

To demonstrate that STAR is implementable and feasible in real world situations, we implemented the STAR algorithm in Symbian C++ and installed on Bluetooth mobile phones. The parameters were chosen based on the trace driven simulations for a missing probability of 0.1. The complied SIS file is only 12kB in size, illustrating that it can easily be deployed on a variety of wireless devices. Fig. 18 shows the probing interval for STAR for one mobile phone carried by a volunteer over a 24 hour period and the number of contacts detected during that period. It is difficult to do comparative studies of the algorithms in the real world. This is because multiple devices running the different probing algorithms will have to be carried by the same person over long durations. Moreover, the device discoveries from the different devices can interfere with each other, resulting in inaccurate statistics.

6. RELATED WORK

User mobility has profound impact on the performance of wireless networks. It has been shown that the capacity of wireless networks can be improved by random [3] or controlled [20] mobility. The impact of user mobility in delay performance and forwarding algorithm design has also been widely studied in DTNs which use Bluetooth or WLAN [4, 21, 22, 19]. In this paper, we study how user mobility will effect energy efficiency in device discovery for Delay Tolerant Applications. Our protocol provides a energy conscious device discovery service, which can be used by many delay tolerant applications.

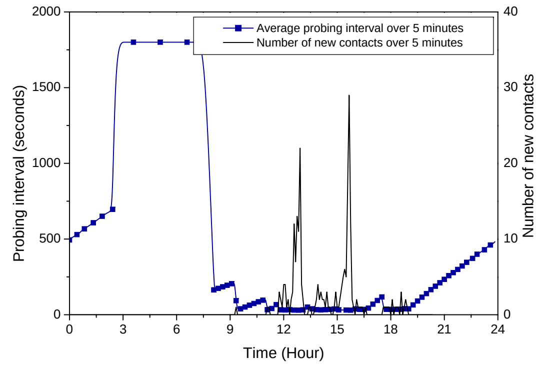

Figure 18: Number of new contacts and the probing interval for a given day.

In [23], the authors observe that existing mobility models are too simple and do not accurately reflect user mobility. Based on WLAN traces from various universities, they observe that people have location preferences, which they visit periodically. They then propose an empirical model to simulate user mobility and validate it with WLAN data traces. The goal of this paper is modeling user mobility and contact patterns and not to derive contact probing algorithms.

In [24], the authors again consider WLAN traces from UCSD and Dartmouth. They analyze the statistics of inter-pair contact time distribution, i.e., the time between subsequent meetings of a specific pair of nodes. This metric is especially useful for DTNs. They show that the inter-pair contact time distribution is self-similar. In our work, we consider traces of a different process, namely the set of contacts made as a user moves around. This contact process is significantly different from that obtained from WLAN traces. In WLAN traces, contacts between users can be inferred only if two users associate with the same access point at the same time. Contacts between users at locations where there are no access points cannot be inferred. Also, we characterize the inter-contact time, i.e., the time between the discovery of two new contacts. This metric is useful for both DTN and DTD applications.

Energy efficient Bluetooth device discovery has been studied in [25, 26]. However, they focus on the Bluetooth protocol stack. Our algorithms are independent of the communication protocol and adapt to user mobility to conserve energy for device discovery, making them applicable to many communication systems.

Stochastic event capturing schemes are studied in wireless sensor networks [27]. In [27], optimal visiting routes are devised for mobile sensors to capture events which randomly happen in different Points of Interest. The event process is assumed to be memoryless and the parameters of the process are known in advance. In our study, we investigate real-world contact process which are more complicated than memoryless arrival and departure process. The analysis in this paper can be applied to a general renewal process and our algorithm can dynamically adapt to unknown parameters in the process.

7. REFLECTIONS

In this paper, we have identified that contact probing mechanisms play a critical role in certain mobile delay tolerant applications. In these applications, mobile devices periodically probe their environment for the presence of new contacts. We investigated the design of energy conscious, adaptive contact probing algorithms which trade-off energy consumption and the probability of missing a contact. Our key contributions were (i) a theoretical foundation which aids in the design of adaptive contact probing algorithms, (ii) real-world experiments and characterization of empirical contact patterns, (iii) design and validation of an adaptive probing algorithm (called STAR) via trace driven simulations. We demonstrate that STAR is *three* times more energy efficient than a naive constant probing algorithm. We now reflect on what we have done:

• *Exploiting contact bursts:* Our empirical data shows that contact duration is Pareto distributed and the new contact arrivals are self-similar, meaning they are bursty. Not surprisingly, it is advantageous to exploit the bursty nature to find a good trade-off between missing probability and energy. This is what STAR implicitly does.

• *Sensitivity to contact duration distribution:* We note that STAR, which uses knowledge of the contact duration distribution, has an average contact probing interval that is only 20% larger than AIMD, which requires no such prior knowledge. However, we observe that the performance of STAR is not sensitive to the exact parameters of the contact duration distribution. This implies that individuals do not have to learn their personal contact duration statistics and greatly simplifies implementation.

• *Maintaining short term P_{miss}:* From Fig. 16, we see that even though STAR maintains a low missing probability over large time scales, it can miss a relatively large fraction of contacts in the short term. This is because we have so far assumed that when a device is not probing, it has no way of knowing if new contacts have arrived even if those new contacts are probing. However, if devices are cognizant of the fact that they have been probed, this information can clearly be used to better adapt the probing interval and thereby reduce the short term contact missing probability. We point out that our analysis and simulations were motivated by current Bluetooth implementations, which do not expose the fact that a device has been probed, to the application. However, this constraint is not fundamental and can easily be relaxed.

• *Synthetic Contact Models:* Our empirical data set for contact patterns is much larger than any comparable data set. We have learned quite a bit about mobile device contact patterns and are certain that there is much more that can be gleaned. In addition, we believe that our data set and analysis can be used to develop synthetic models for real world contact patterns, as done in [28] (for modeling Internet traffic) and [23] (for modeling association with WiFi access point).

Acknowledgments

We thank Fahad Kohwaja for the implementation on mobile phones. We also thank the anonymous reviewers and our shepherd, Elizabeth Belding, for their valuable feedback.

8. REFERENCES

[1] K. Fall, "A delay-tolerant network architecture for challenged internets," in *Proceedings of ACM SIGCOMM*, 2003.

[2] S. Jain, K. Fall, and R. Patra, "Routing in a delay tolerant network," *SIGCOMM Computer Communication Review*, vol. 34, no. 4, pp. 145–158, 2004.

[3] M. Grossglauser and D. Tse, "Mobility increases the capacity of ad hoc wireless networks," *IEEE/ACM Transactions on Networking*, 2002.

[4] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of human mobility on the design of opportunistic forwarding algorithms," in *Proceedings of IEEE INFOCOM*, 2006.

[5] M. Motani, V. Srinivasan, and P. Nuggenhalli, "Peoplenet: Engineering a wireless virtual social network," in *Proceedings of ACM MobiCom*, 2005.

[6] A. J. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, and R. A. Peterson, "People-centric urban sensing," in *Proceedings of WICON*, 2006.

[7] H. Y. O. Wolfson, B. Xu and H. Cao, "Search-and-discover in mobile p2p network databases," in *Proceedings of IEEE ICDCS*, 2006.

[8] M. Kuhn and R. Wattenhofer, "Community aware mobile networking," in *Proceedings of 1st Workshop on Mobile Services and Personalized Environments*, 2006.

[9] V. Srinivasan, M. Motani, and W. T. Ooi, "Analysis and implications of student contact patterns derived from campus schedules," in *Proceedings of ACM MobiCom*, 2006.

[10] A. Natarajan, M. Motani, and V. Srinivasan, "Understanding urban interactions from bluetooth phone contact traces," in *Proceedings of PAM*, 2007.

[11] J. Su, A. Chin, A. Popivanova, and E. d. L. Ashvin Goel, "User mobility for opportunistic ad-hoc networking," in *Proceedings of IEEE WMCSA*, 2004.

[12] N. Eagle and A. Pentland, "Social serendipity: Mobilizing social software," *IEEE Pervasive Computing*, vol. 4, no. 2, pp. 28–34, 2005.

[13] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Pocket switched networks: Real world mobility and its consequences for opportunistic forwarding," Technical Report No. 617, University of Cambridge, 2005.

[14] S. M. Ross, *Stochastic processes.* Wiley, 1983.

[15] S. Boyd and L. Vandenberghe, *Convex optimization.* Cambridge University Press, 2005.

[16] G. He and J. Hou, "An In-Depth, Analytical Study of Sampling Techniques for Self-Similar Internet Traffic," in *Proceedings of IEEE ICDCS*, 2005.

[17] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of Ethernet traffic," *IEEE/ACM Transactions on Networking*, vol. 2, no. 1, pp. 1–15, 1994.

[18] J. Beran, *Statistics for Long-Memory Processes.* Chapman and Hall, 1994.

[19] M. McNett and G. M. Voelker, "Access and mobility of wireless PDA users," *Mobile Computing and Communications Review*, vol. 9, no. 2, pp. 40–55, 2005.

[20] R. Shah, S. Roy, S. Jain, and W. Brunette, "Data mules: Modeling a three-tier architecture for sparse sensor networks," in *Proceedings of IEEE SNPA*, 2003.

[21] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot, "Pocket switched networks and human mobility in conference environments," in *Proceeding of the ACM SIGCOMM workshop on DTN*, 2005.

[22] T. Gross and C. Tuduce, "A mobility model based on WLAN traces and its validation," in *Proceedings of IEEE INFOCOM*, 2005.

[23] W.-J. Hsu, T. Spyropoulos, K. Psounis, and A. Helmy, "Modeling time-variant user mobility in wireless mobile networks," in *Proceedings IEEE INFOCOM*, 2007.

[24] L.-J. Chen, Y.-C. Chen, T. Sen, P. Sreedevi, K.-T. Chen, C.-H. Yu, and H.-H. Chu, "Finding self-similarities in opportunistic people networks," in *Proceedings of IEEE INFOCOM*, 2007.

[25] C. Drula, "Fast and energy efficient neighbour discovery for opportunistic networking with bluetooth," Master Thesis, University of Toronto, 2005.

[26] G. Zaruba and I. Chlamtac, "Accelerating Bluetooth inquiry for personal area networks," in *Proceedings of IEEE GLOBECOM*, 2003.

[27] N. Bisnik, A. Abouzeid, and V. Isler, "Stochastic event capture using mobile sensors subject to a quality metric," in *Proceedings of ACM MobiCom*, 2006.

[28] P. Barford and M. Crovella, "Generating representative web workloads for network and server performance evaluation," in *Proceedings of ACM SIGMETRICS/PERFORMANCE*, 1998.