

Elsevier Editorial System(tm) for Pervasive and Mobile Computing

Manuscript Draft

Manuscript Number: PMC-D-07-00017R1

Title: BlueTorrent: Cooperative Content Sharing for Bluetooth Users

Article Type: Special Issue PerCom2007

Corresponding Author: Sewook Jung,

Corresponding Author's Institution: UCLA

First Author: Sewook Jung

Order of Authors: Sewook Jung; Uichin Lee; Alexander Chang; Dae-Ki Cho; Mario Gerla

BlueTorrent: Cooperative Content Sharing for Bluetooth Users

Sewook Jung ^a, Uichin Lee ^a, Alexander Chang ^a, Dae-Ki Cho ^a, and Mario Gerla ^a

^a *Department of Computer Science, University of California, Los Angeles, Los Angeles, CA, USA*

Abstract

People wish to enjoy their everyday lives in various ways, among which entertainment plays a major role. In order to improve lifestyle with more ready access to entertainment content, we propose BlueTorrent, a P2P file sharing application based on ubiquitous Bluetooth-enabled devices such as PDAs, cellphones and smart phones. Using BlueTorrent, people can share audio/video contents as they move about shopping malls, airports, subway stations etc. BlueTorrent poses new challenges caused by limited bandwidth, short communications range, mobile users and variable population density. A key ingredient is efficient peer discovery. This paper approaches the problem by analyzing the Bluetooth periodic inquiry mode and by finding the optimum inquiry/connection time settings. At the application layer, the BlueTorrent index/block dissemination protocol is then designed and analyzed. The entire system is integrated and implemented both in simulation and in an experimental testbed. Simulation and measurement results are used to evaluate and validate the performance of BlueTorrent in content sharing scenarios.

Key words: BlueTorrent, Bluetooth, P2P

1 Introduction

Bluetooth is an always-on, low-power, and short-range hookup for implementing Wireless Personal Area Networks (WPANs). It became the most popular PAN communication device. The Economists [11] reported that sales of Bluetooth devices are more than doubled in 2005 to reach 320M units, and the figure is expected to exceed 520M this year.

The concept behind Bluetooth is to support universal short-range wireless capabilities using the 2.4GHz globally unlicensed low power band. Bluetooth allows users to connect up to eight devices by forming a star-shaped cluster, called piconet.

The cluster head is called master and the other nodes are called slaves. Two Bluetooth devices within 10 m range can exchange data up to 723kbps. To minimize interference among different Piconets, Bluetooth uses frequency hopping (FH) with pseudo-random ordering of 79 frequencies in the same band.

Bluetooth is intended to various applications, including audio and data. As widely observed these days, Bluetooth gradually replaces cables to link computers to keyboards, printers, and mouse. It is often treated as a multimedia enabler such that users can listen to and download music from other machines using wireless headsets and MP3 players. More importantly, new legislation banning the use of mobile phones without a hands-free kit while driving boosted the use of Bluetooth-based mobile headsets in public. This adoption has in turn cleared the way for the inclusion of Bluetooth in all kinds of new products. For example, automotive manufacturers are looking to link Bluetooth devices with the on-board system to provide in-car stereo/hands-free solution.

With Bluetooth's continuing proliferation, Bluetooth-based applications are well positioned to deliver new opportunities to all facets of the industry. This paper deals with one fast growing application – namely, proximity advertising and marketing. In 2005, handset maker Nokia and music label EMI started a project to let coffee shop customers listen to music via Bluetooth [3]. Similarly, San Francisco's WideRay has placed Bluetooth-enabled kiosks in selected music stores, video game stores, and theaters across the country to send messages asking if users are interested in getting more information about various items the store is selling, e.g., music, ring tones, video etc. More interestingly, in August 2006 CBS announced that it will make clips from its prime-time fall program line-up available for free via CBS Outdoor billboards [7]. The digital billboards, initially located in NYC's Grand Central Station, will allow cellphone or PDA users to watch clips such as CSI. There are already off-the-shelf commercial products available for proximity marketing such as BlueBlitz and BlueCasting.

However, proximity marketing solutions so far are based on a client and server model. A service provider spreads advertisements using a Bluetooth Access Point (BT-AP) to mobile hosts within Bluetooth communication range (i.e., 10m). However, since users are mobile and the Bluetooth channel is error-prone, the total amount of data that a mobile user can download from the server is limited to a few hundreds kilobytes – i.e., only images or small video clips. Moreover, it is not practical nor economically feasible to install BT-APs every 10m. Under such circumstances, distributing contents larger than several hundred KB requires customers to stop in front of the BT-AP, unless we adopt a P2P technology.

In this paper, we propose BlueTorrent, a cooperative content sharing protocol for Bluetooth that exploits sparsely distributed BT-APs allowing mobile users to cooperatively download relatively large files (e.g., 1-10Mbytes). The challenge now becomes to support P2P connections among mobile Bluetooth users. In Bluetooth,

one node must be a master and the other node must be a slave: the role must be dynamically changed with a frequency that minimizes the discovery latency, yet allowing enough time for useful peer to peer data transfers. To this end, we analyze an existing standard function and find the optimal parameter configurations. To effectively share content in spite of short link duration, BlueTorrent uses BitTorrent-like file swarming. Content is divided into a number of small pieces, and mobile users can exchange whatever pieces are available. BlueTorrent uses a cooperative carry and forward strategy: pieces are forwarded whenever a connection is available in order to minimize download latency. Note that meta-data (e.g., unique file ID, title, media type) of a file is also spread opportunistically via mobility. BlueTorrent users can actively query other peers to search for a file of interests. The following is the key results of our study:

- We find that our simple periodical inquiry scheme at the application layer results in better performance than the standard inquiry mode. The latter only supports parameter control at the time scale of seconds, thus resulting in performance degradation.
- We identify key parameters and their configurations to minimize the peer discovery latency. Among them, we show that the inquiry scan period plays a key role in determining the optimal configuration of the other parameters.
- We validate the performance of BlueTorrent via simulations and testbed experiments. For some tested scenario, our cooperative data sharing results in more than 400% improvement in terms of download latency. We show that the size of a piece is important in mobile content sharing. For AP mode, we find that for a given user density, there exists the optimal speed, leading the best performance.

This paper is organized as follows. In Section 2, we explain Bluetooth. In Section 3, we analyze the periodic inquiry mode in Bluetooth for P2P connections and find optimal configurations. In Section 4, we propose and analyze our index and content sharing protocol. In Section 5, we evaluate our protocols through extensive simulations. In Section 6, we show the preliminary performance results of BlueTorrent in our testbed. In Section 7, we analyze power consumption and discuss incentive plans to increase cooperation. In Section 8, we review the related work and then conclude the paper in Section 9.

2 Bluetooth Overview

The total bandwidth in Bluetooth is divided into 79 channels (each 1 MHz). Frequency hopping (FH) occurs by jumping from one physical channel to another in a pseudorandom sequence. The same hopping sequence referred as an FH channel must be shared by two devices that communicate each other (i.e., by forming a piconet). The hop rate is 1600 hops per second, so that each physical channel is occupied for $625\mu s$. This interval is referred to as a slot and is numbered sequen-

tially. An FH channel is shared between a master and slaves using a time division scheme in which data are transmitted in one direction at a time, with transmissions alternating between two directions. Because more than two devices share the piconet medium (an FH channel), the access technique is Time Division Multiple Access (TDMA). Transmission of a packet starts at the beginning of a slot and can span multiple slots (1, 3, or 5 slots accordingly). The FH sequence is determined pseudorandomly based on piconet master Bluetooth ID. Therefore, several piconets can coexist with minimal interference. Occasionally, two piconets will use the same physical channel during the same time slot, causing a collision and data loss, but this happens rarely, and it is recovered by forward error correction (FEC) and error detection/ARQ techniques. Bluetooth standard defines an asynchronous connectionless (ACL) link for a point-to-multipoint link between the master and all the slaves in the piconet. Achievable data rates on the ACL link vary depending on the number of slots per packet and on the FEC strategy. FEC packets formats are DM1, DM3, and DM5 (with the digits indicating the number of slots used). The non-error coded formats are DH1, DH3, and DH5. Readers can find the details of Bluetooth in the published specification [5] or in the textbook [27]

3 Bluetooth for P2P

A peer in BlueTorrent must be able to discover other peers in order to share content. For two nodes to connect using Bluetooth, one node should be in the Inquiry state and the other in the Inquiry Scan state. However, since BlueTorrent users are randomly moving, their roles as Inquirers (masters) or Inquirees (slaves) cannot be predetermined, rather, they must be randomly alternated. Bluetooth supports such a random role selection through the periodic inquiry mode. As shown later, the performance of the periodic inquiry mode is dependent on various inquiry parameters. Careful analysis is required to minimize the connection setup latency. In this section, we begin with the overview of the Bluetooth inquiry procedure. We then analyze the mode and empirically find the optimal parameter setting via extensive simulations.

3.1 Bluetooth inquiry procedure

A master peer in the Inquiry state sends inquiry packets and waits for response packets from the potential slave peers. The number of physical channels used for the inquiry procedure is reduced from 79 to 32 for increased discovery efficiency. Moreover, the master peer uses the inquiry hopping sequence for inquiry (as Tx slots), and the potential slave peers use the inquiry response hopping sequence (different from the former) for response (as Tx slots). Therefore, the master peer must switch to inquiry response hopping sequence to listen to response packets (as

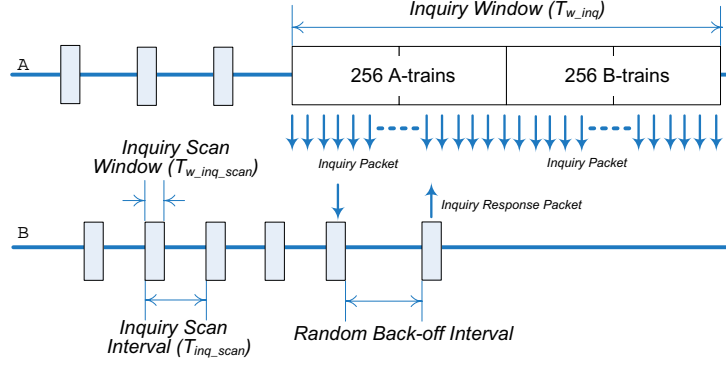


Fig. 1. Inquiry procedure example

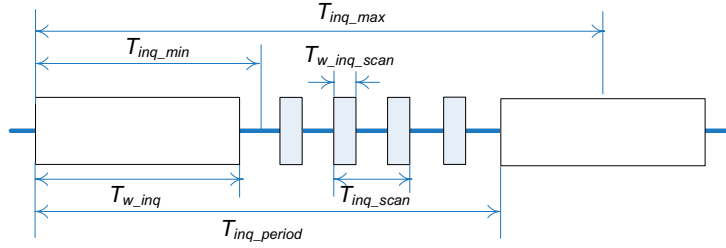


Fig. 2. Periodic inquiry mode

Rx slots). Conversely, the potential slave peers must switch to the inquiry hopping sequence to listen to inquiry packets (as Rx slots). The inquiry hopping sequence is divided into two distinct sequences called A- and B-train. Each train contains 16 physical channels and the total duration is 10ms ($=16 \cdot 0.625ms$). The Bluetooth specification mandates that each train must be repeated at least 256 times (i.e., 2.56s) before switching to another train. The hopping rate for inquiry is increased to 3200 hops/second (i.e., half-slots) so that the master peer can transmit very short ($68\mu s$) inquiry packets every $312.5\mu s$. In each Tx slot, the master can send two inquiry packets, and in next Rx slot, it listens to response packets from other devices. This alternation repeats during the period of the inquiry window $T_{w_inquiry}$. Note that the host controller interface allows us to set the interval as a *multiple* of 1.28s through *HCI_Inquiry* host controller interface (HCI) command [5]. Figure 1 shows an example of the inquiry procedure. The size of the inquiry scan window is 5.12s ($=4 \cdot 1.28s$).

Other peers that want to make connections to the master peer should be in the Inquiry Scan state. Each peer enters to the Inquiry scan interval for every inquiry scan interval (T_{inq_scan}) and stays there for inquiry scan window ($T_{w_inq_scan}$). $T_{w_inq_scan}$ should be greater than the length of a train (i.e., $>10ms$) in order to ensure that a frequency synchronization between inquiring and scanning peers can happen when the scanning frequency is in the currently active train of the inquirer, and it cannot be larger than T_{w_inq} . The ranges of T_{inq_scan} and $T_{w_inq_scan}$ are given as $[10.625 - 2560ms]$ and $[11.25 - 2560ms]$ respectively. These parameters can be set using *HCI_Write_Inquiry_Scan_Activity* HCI command. After receiving

the inquiry packet, the peer changes its state to the Inquiry Response state and sends back an inquiry response packet containing its Bluetooth device address and clock information. It is important to note that the peer backs off for a random number of slots to reduce the probability of colliding with other inquiry responses. This random number is drawn uniformly out of a range $[0, 1023]$ ($<640\text{ms}$) if the inquiry scan interval is larger than or equal to 1.28s; otherwise, a range $[0, 127]$ ($<80\text{ms}$) is used. The master peer can receive responses if they arrive within the inquiry window. There is no acknowledgment of a response packet.

3.2 Periodic inquiry mode analysis

The conventional neighbor discovery is asymmetric in the sense that for a given node, the role is fixed to be either as a master or a slave. In mobile peer-to-peer networks, however, peers should be able to randomly switch their roles in order to find each other, since a node can be either a client or a server. Bluetooth defines a symmetric neighbor discovery mode in the specification (as of v1.0), called *Periodic_Inquiry_Mode* where inquiry procedures are periodically executed. This mode is set by configuring the range of period, $[T_{inq_min}, T_{inq_max}]$, and inquiry length (T_{w_inq}) (see Figure 2). In each round, a node alternates an inquiry state immediately followed by an inquiry scan state. The length of an inquiry is fixed to T_{w_inq} , and the length of an inquiry scan state is uniform randomly chosen over $[T_{inq_min} - T_{w_inq}, T_{inq_max} - T_{w_inq}]$ in order to avoid synchronization. Note that this interval has nothing to do with the inquiry scan window ($T_{w_inq_scan}$) and interval (T_{inq_scan}).

Although the periodic inquiry mode is widely used, its optimization is not explored. For instance, BlueZ, an official Bluetooth protocol stack for Linux, includes Host Controller Interface Daemon (HCID).¹ HCID provides a wrapper function by internally setting the parameters, i.e., $T_{w_inq} = 8$ and $[T_{inq_min} = 16, T_{inq_max} = 24]$.² Since the unit is 1.28s, each round takes on average 35.84s. Therefore, the parameter selection is not proper for mobile P2P applications. Thus, the goal of this section is to optimize the P2P mode by tuning three key parameters, namely *inquiry window size* (T_{w_inq}), *the length of inquiry scan state* (via $[T_{inq_min}, T_{inq_max}]$), and *inquiry scan interval* (T_{inq_scan}).

The inquiry window size may vary based on which Bluetooth version we use. As of Bluetooth v1.2, the interlaced scan mode is used by default. Since the interlaced mode scans two trains in a row, for a given inquiry the probability of missing an inquiry packet (i.e., inquiry failure) is negligible [21]. Bluetooth v1.1, however, does not support it, requiring the inquiry window size T_{w_inq} at least 5 ($5 \times 1.28\text{s} = 6.4\text{s}$) in order to discover neighbors with high probability [21]. Bluetooth v1.1 devices

¹ <http://www.bluez.org>

² HCID exports functions via D-Bus, a system for *interprocess communication* (IPC)

gradually disappear and currently, v1.2 and v2.0 devices become dominant [15]. Thus, we focus only on Bluetooth v1.2 and v2.0.

The inquiry procedure is blocking. The function returns the list of inquiry responses received during the inquiry period (i.e., $1.28s * T_{w_inq}$). The obvious disadvantage is that we have to wait for the entire period. For mobile applications, non-blocking inquiry is preferable because the received inquiry responses can be read immediately via event polling; consequently, we can reduce the average discovery latency. The events of interests are `EVTINQUIRYRESULT` and `EVTINQUIRYCOMPLETE`. The inquiry procedure can be cancelled anytime by sending an HCI inquiry cancel command. The non-blocking inquiry procedure allows us to have an arbitrary length of inquiry window size, instead of a multiple of 1.28s.

The length of inquiry scan state is determined by setting $[T_{inq_min}, T_{inq_max}]$. The minimum value must be carefully chosen such that it should be larger than inquiry scan interval (T_{inq_scan}) since the first scan starts after T_{inq_scan} instead of starting immediately. In addition, the difference between those two values should be greater than zero. Otherwise, the process becomes deterministic – two nodes may never be able to discover each other. It is important to note that although the standard periodic inquiry mode uses the unit of 1.28s for the minimum and maximum values, we can set arbitrary numbers at the granularity of 0.625ms by implementing a periodic inquiry mode in the *application layer*.

The inquiry scan interval (T_{inq_scan}) is also important because it determines the usefulness of the inquiry scan state. The usefulness depends on how many “scans” actually happens during the inquiry scan period and can be measured by dividing the average length of inquiry scan period by T_{inq_scan} . The more the number of scans, the higher is the usefulness. As shown later, for a given T_{inq_scan} , there exists an optimal configuration of $[T_{inq_min}, T_{inq_max}]$. However, this comes at the cost of more energy consumption. This implication is discussed at the end of the section.

Along with the aforementioned parameters, the maximum back-off interval for the inquiry response has a critical impact, especially for Bluetooth v1.2 and v2.0. Even though a “single” inquiry is enough to receive an inquiry packet, an inquiry fails if the node backs off before returning an inquiry packet, and in the mean time the inquiring node switches its role; i.e., the failure probability depends on the maximum back-off interval. Bluetooth specification (v1.2 and v2.0) requires that the back-off interval be drawn uniformly from the range $[0,1023]$ if the inquiry scan interval is larger or equal to 1.28s, or from the range $[0,127]$ otherwise. But the actual implementation is dependent on chipmakers. For instance, in the extended version of this paper we show that Bluetooth v1.2 (Silicon Wave) uses $[0,1023]$, and Bluetooth v2.0 (Broadcom) does not implement random back-off.

For the sake of illustration, we provide a simple mathematical analysis to better understand the correlation among these parameters. The probability that one node

is in the inquiry scan mode is given as: $P_{scan}=1 - \frac{T_{w.inq}}{T_{w.inq}+T_{diff}/2}$ where $T_{diff} = T_{inq-max} - T_{inq-min}$. Let us assume that $T_{inq-min} = T_{w.inq}$ (since $T_{inq-min} \geq T_{w.inq}$). Given that the node is in the inquiry scan mode, the probability of successful inquiry, denoted as P_{si} , is determined based on the starting point of the inquiry period. When it is in $[0, T_{inq-scan} - T_{w.inq}]$, the inquiry window does not overlap with the inquiry scan window; thus, the probability of successful discovery is 0. When $[T_{inq-scan} - T_{w.inq}, T_{inq-scan} - T_{w.inq} + T_{max.bo}]$, the inquiry window overlaps with the inquiry scan window, but the probability of successful discovery is purely determined by the random back-off: i.e., $1/T_{max.bo}$. When $[T_{inq-scan} - T_{w.inq} + T_{max.bo}, T_{inq-scan}]$, the overlapped period is longer than $T_{max.bo}$; consequently, the probability is 1. Let us simply assume that the distribution of the starting point of the inquiry period is uniform.³ The probability of successful inquiry is given as: $P_{si} = \frac{T_{w.inq}-T_{max.bo}+1}{T_{inq-scan}+T_{diff}/2}$. Since discovery happens only when one is in the inquiry scan mode and the other is in the inquiry scan mode, the probability of successful discovery is proportional to these two events: $P_s = \alpha P_{si} P_{scan}$ where α is a constant normalization factor. Let us assume that each discovery trial⁴ is independent of one another. The process follows geometric distribution with success probability P_s . The average number of trials is simply $1/P_s$. Since the average length of a trial is $T_{w.inq} + T_{diff}/2$, the latency is given as $\frac{T_{w.inq}+T_{diff}/2}{P_s}$. One can reduce the average length of a trial by decreasing the inquiry window size, but this will decrease P_s as well. As a result, the average number of trials will increase (i.e., $1/P_s$). Similarly, one can maximize P_s with infinitely large inquiry window, but this makes the average length of a trial infinite. These examples show that we cannot make the inquiry window arbitrary small or long. Note that it is non-trivial to accurately model the overall process since each trial is dependant on the previous trial, which can be represented using a Markov chain. This analysis is part of our future work.

3.3 Periodic inquiry mode evaluation

We developed *InqSim* to empirically find an optimal parameter configuration. This simulator performs a slot-level discrete time, event-driven simulation and accurately models the P2P mode. Every time an event expires, it schedules the next event such as inquiry, scan, and random back-off. During the simulation, actual discovery is checked when a scan expires. A node then examines whether the other party is in the inquiry state and whether the state lasts longer than its random back-off value. If so, the other party will be notified and the inquiry stops immediately (i.e., non-blocking inquiry). To simulate a random encounter of two nodes, the actual measurement starts after passing 160,000 slots (i.e., 100s). Since the overall simulation is simple and requires slot level control and various parameter tuning,

³ This distribution is closely related to the periods of both inquiry and inquiry scan. If one is longer than the other, it is more likely that we observe the longer one.

⁴ Inquiry and inquiry scan comprise a trial.

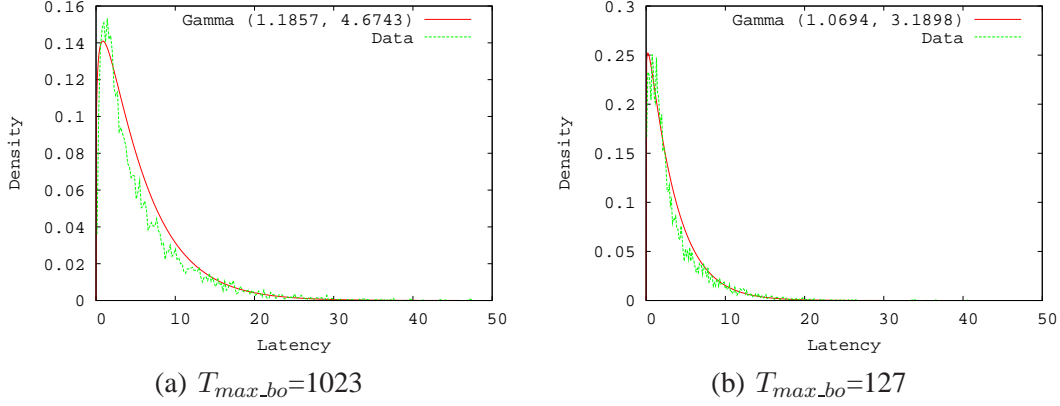


Fig. 3. Discovery latency distribution ($T_{w_inq} = 1.28s$, $T_{diff} = 3.5s$, $T_{min} = 1.28s$)

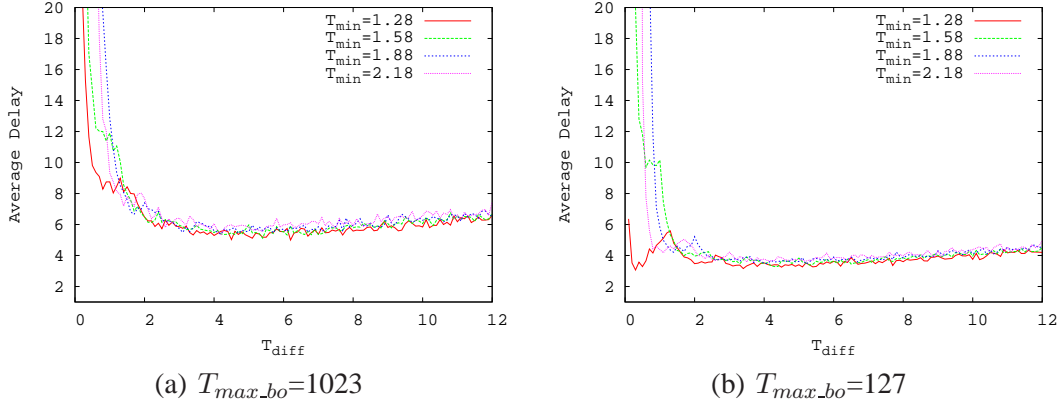


Fig. 4. Discovery latency with various T_{min}

we use *InqSim* instead of using UCBT, a simulator used in Section 5.

We measure the discovery latency of the periodic inquiry mode that is the elapsed time until either node successfully discovers the other party. Since we implement non-blocking inquiry, we use inquiry period (T_{w_inq}) in range of 0.38s to 2.48s with a gap of 0.3s. For a given range $[T_{inq_min}, T_{inq_max}]$, we define $T_{min} = T_{inq_min} - T_{w_inq}$ and $T_{diff} = T_{inq_max} - T_{inq_min}$; i.e., the inquiry scan period is in range of $[T_{min}, T_{min} + T_{diff}]$. We assume that each inquiry scan period contains at least one scan; thus, $T_{min} \geq T_{inq_scan}$. For given maximum back-off (T_{max_bo}) and inquiry scan intervals, we measure the discovery latency by varying T_{w_inq} , T_{min} and T_{diff} . For each configuration, the average of 3000 runs is presented.

Figure 3 shows the latency distribution with inquiry period of 1.28s, $T_{diff} = 3.5s$, $T_{min} = 1.28s$, and $T_{max_bo} = 1023/127$. It shows that the curves can be accurately fitted using *gamma distribution* with parameters (1.1857, 4.6743) for $T_{max_bo} = 1023$ and (1.0694, 3.1898) for $T_{max_bo} = 127$. In general, the latency distribution can be fitted using gamma distribution.

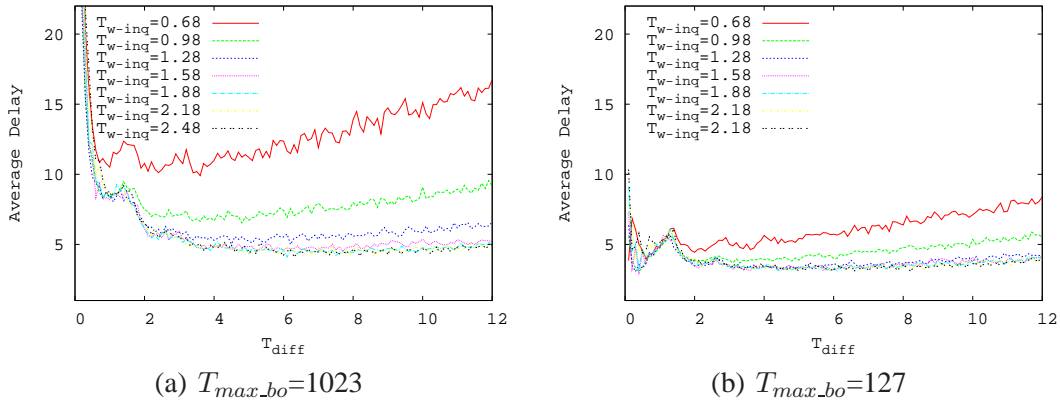


Fig. 5. Discovery latency with various lengths of inquiry period

Figure 4 shows the results of the average latency with the inquiry scan interval of 1.28s, which is the default scan interval, and with the inquiry period of 1.28s. To show the impact of the maximum back-off size, we use $T_{max_bo} = 1023$ slots (639.375ms) and 127 slots (79.375ms). The figure shows that too small T_{diff} results in high average delay: given that both nodes are in the same state, they need to spend more time to unlock the synchronization. As T_{diff} increases, the delay decreases till it reaches a certain threshold. As the inquiry scan period gets longer, randomly encountered nodes are likely in the inquiry scan state. Mathematically speaking, given a collection of random intervals, the length bias or inspection paradox tells us that longer intervals are more likely to be sampled than shorter intervals [22]. Both nodes tend to stay longer in the scan state, thus resulting larger latency. The figure also shows that the back-off value has a great impact on the average latency. When the maximum back-off interval is decreased to 127 slots, the average latency drops more than 30%. Interestingly, Figure 4(b) shows that the average latency has its minimum at $T_{diff} \approx 0.3s$ and then it increases again till it reaches $T_{diff} \approx 1.5s$. If $T_{diff} < 1.28s$, it does not improve the usefulness of the inquiry scan period; only a single scan per round is feasible. T_{diff} is mainly used to avoid synchronization. Figure 4(b) clearly shows that $T_{diff} \approx 0.3s$ is the optimal value, and unnecessarily large T_{diff} adversely affects the performance. In the case of $T_{max_bo} = 1023$ slots (639.375ms), it requires large T_{diff} to handle random back-off. Let us say that $T_{diff} = 0.1s$. For a given scan period of $[0, 1.38s]$, scan only happens during $[1.28, 1.38s]$. Although an inquiry packet can be successfully received during this period, backing off larger than the residual life time of the inquiry period results in failure. The results show that we conclude that the optimal value can be found when “ $T_{inq_scan} = T_{min}$ ” and this trend is consistent with different inquiry periods.

Figure 5 shows the results of various inquiry periods ranging from 0.38s to 2.48s with a gap of 0.3s. Both $T_{max_bo} = 1023$ and $T_{max_bo} = 127$ are used for simulations. The length of inquiry scan is 1.28s. The graph shows that the latency is heavily dependent on the length of inquiry period and T_{max_bo} , but it is less sensitive to the T_{diff} . This trend is even more evident as the length of inquiry period increases.

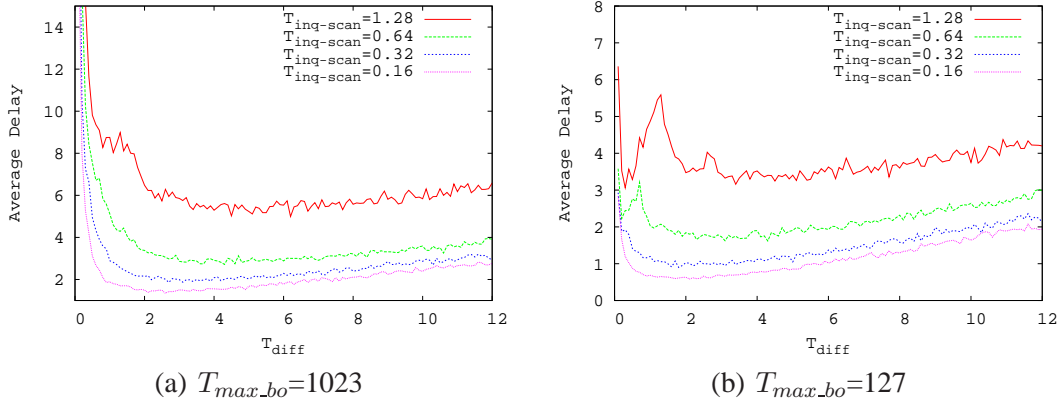


Fig. 6. Discovery latency with various T_{inq_scan}

The graph also shows that enlarging inquiry period after a certain threshold gives us a mere improvement: i.e., 1.58s and 1.28s for $T_{max_bo} = 1024$ and $T_{max_bo} = 128$ respectively. Given this, the length of T_{diff} is less sensitive: $T_{diff} = [4, 6]$ s for $T_{max_bo} = 1023$ (avg. latency 4.3s) and $T_{diff} = [3.5, 5.5]$ s for $T_{max_bo} = 127$ (avg. latency 3.3s).

Figure 6 shows the results as a function of inquiry scan interval with $T_{w_inq} = 1.28$. The figure shows the case when $T_{inq_scan} = T_{min}$, which minimizes the latency. As inquiry scan interval decreases, the average delay also decreases. Smaller inquiry scan interval results in more randomness, thus smoothing the curves. The figure also shows that as T_{inq_scan} decreases, the minimum latency tends to converge. For example, for $T_{min} = 0.16$ s, $T_{diff} \approx 2.1$ s minimizes the latency to 1.39s; for $T_{min} = 0.16$ s, $T_{diff} \approx 1.7$ s minimizes the latency to 0.59s. Therefore, the inquiry scan interval is a *critical* factor for the discovery latency.

In summary, for a given T_{inq_scan} , we show that the optimal T_{min} is equal to T_{inq_scan} , and the optimal T_{diff} can be found through simulations. The results show that optimal values are not multiple of 1.28s, and thus, the standard periodic inquiry mode is suboptimal. We also show that the inquiry scan interval is one of the critical factors in determining the discovery latency. In general, by reducing T_{inq_scan} , we can minimize the discovery latency; e.g., $T_{inq_scan} = 0.32$ s has 2.53s. One caveat is that this comes at the cost of more energy consumption. In fact, minimizing both discovery latency and energy consumption is two conflicting goals. One of the convincing solutions is to exploit the idea of sociological orbits, a probability mobility model where people move between a set of hubs or information bazaars [12]. Drula et al. propose adaptive energy conserving algorithms based on recent activity level in hubs [10]. For given a set of discovery modes, from aggressive (fast discovery, energy intensive) to lazy (slow discovery, low power consumption), nodes change one's mode based on the contact frequency. P2P content sharing shares the same idea, and thus, for a given set of constraints, our results can be used to optimize the algorithm. Note that readers can find the details of power consumption of each Bluetooth operation in [6].

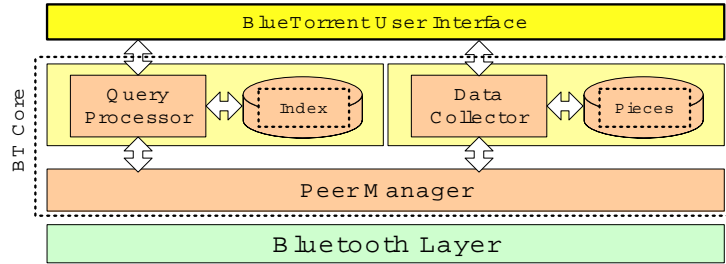


Fig. 7. BlueTorrent architecture

4 Bluetooth based content sharing protocol

The core components of BlueTorrent are presented in Figure 7, namely peer manager, query processor, and data collector. The peer manager implements the peer discovery/connection protocol and the peer selection protocol. The query processor resolves user queries by searching its index database as well as its neighbors'. For a given file, the data collector uses a file swarming protocol to opportunistically collect the file piece by piece. Finally, we provide a simple mathematical analysis to show the feasibility of BlueTorrent in a dynamic environment with high churn rate such as in subways or streets.

4.1 Peer Manager

The peer discovery/connection protocol is presented in Section 3. Given this, the peer manager provides peer selection to the upper layer (i.e., query processor and data collector). It manages its physical neighbors among which it tries to choose the best peer. This is important because making a connection takes almost the same amount of time as discovery. In BlueTorrent, a node keeps track of connections for efficient peer selection, by logging connection frequency and duration. For a given node, the log can be kept for a certain threshold period of time, or can be exponentially aged. The main purpose is to rank the physical neighbors by using the connection frequency/duration. A node is less helpful if the number of connections is high, and the average connection duration is long. Thus, BlueTorrent give more chances to those freshly encountered nodes.

Although the above procedure gives us the peer with more data to download, the contact duration is determined by the physical distance among peers. It is preferable to select a peer that has the longest residual contact duration, which can be calculated by the direction/distance among peers. However, it is nontrivial to support the scheme without GPS or localization. Instead, BlueTorrent can be configured to use the measured RSSI values of the inquiry responses to estimate the distance among peers. RSSI values can be periodically collected as a result of peer discovery. Peers are then labeled with their distance estimates. Thus, the peer manager

uses the distance estimate as the first criterion for peer selection. A group of nodes within a certain threshold distance are then considered, and the best peer is selected using the frequency/duration metric. Note that one can implement the distributed localization using RSSI values given that the density of BlueTorrent users is *high*, resulting more accurate peer selection.

4.2 Query Processor

The prerequisite of content distribution is to know where the content is. Content can be searchable through indices which include a unique ID (e.g., 32 bit hash value of the content), title, producer, media type, etc. A static BT AP is a publisher of content, and it pushes index to mobile nodes. Users who are interested in specific information can proactively query other peers. A user can express his request for contents as a simple query string. For example, those who are interested in downloading a movie preview clip of “Pirates of the Caribbean” will prepare a query string such as “Pirates & Caribbean” with media type “avi/mpg.” We assume that BlueTorrent is equipped with a lightweight database for index management and search. Since users have a limited, often local scope of files of interest, the size of the indices is small. Thus, we can assume that the storage overhead of index dissemination is minimal compared to the actual content size. Whenever a node discovers another node, it first sends the query. Upon receiving a query, the target node will look up its database to find an exact match of keys. The node could employ sophisticated similarity measures such as [14]. The meta-data match will be automatically reported to the query originator, and the user may decide whether to initiate downloading using the ID of the content.

4.3 Data Collector

BlueTorrent shares contents by using file swarming, mainly due to the limited bandwidth and the short contact duration. The Bluetooth channel tends to be error-prone in the urban streets due to multipath, WiFi interference etc. In addition, the short communication range and mobility of users result in short link/contact duration. For example, two peers moving opposite direction with 1m/s have 10 seconds of link duration. Assuming that peer discovery and connection take 4 seconds, the maximum data size that they can symmetrically transfer is 286KB with DM5 mode. Therefore, it is infeasible for mobile Bluetooth nodes to share a relatively large file without using file swarming.

In BlueTorrent, the data source (i.e., static BT-AP) divides a file into K pieces or blocks. BlueTorrent peers can download pieces from static BT-APs and other peers. Each node has a bitmap of the available pieces for efficient piece reconciliation. Whenever a connection is available, peers first exchange their bitmaps to find out

missing pieces through simple bit operations. The size of a typical bitmap is as small as tens of bytes even if there are more than 100 blocks. To be more precise, given K pieces, the size of a bitmap table is $\lceil K/8 \rceil$ bytes. For example, the size of a bitmap table for 100 blocks takes 13 bytes. Given 286KB of average data transfer per contact with DM5 mode, the overhead is negligible. It is interesting to note that when only one party has data to transfer, it uses asymmetric mode whereas a symmetric mode is used when both have data to transfer.

The size of a piece should be carefully selected based on the characteristics of Bluetooth bandwidth and mobility patterns. If the block is too big, peers cannot download a single block during their contact duration. A block can be divided into very small sub-blocks, say, with size 1KB. However, this costs additional bandwidth and computation overhead for sub-block level reconciliation. Through a simple mathematical model we show that in the dynamic environment with limited bandwidth, sharing a large file, i.e., a large number of blocks, leads to ineffective file swarming [15].

4.4 Dissemination Analysis

The behavior of our protocol can be analyzed using a simple mathematical model. Due to space limits, we present several key results. For the sake of analysis, we assume that in the system there exists only a single file. When a connection is available, peers resolve their queries, exchange a *single* piece (instead of multiple blocks), and finally disconnect. We have N nodes interested in downloading the file, and K is the number of blocks. Let λ and μ denote the rate of rendezvous (i.e., meeting) among those peers and the rate of departing the system (or content sharing area) respectively. Our goal is to find how fast the index/data spreads with such a dynamic environment as well as their survivability conditions. This can be analyzed via a simple epidemic model used in [26]. We present the analysis of index dissemination followed by content sharing. In both cases, let I denote the number of infected peers, i.e., those who have a specific index or a piece.

A single infected node meets other susceptible nodes (i.e., nodes without an index or a piece) with rate $\lambda(N - I)$. Since I infected nodes are independently infecting others/leaving the system, the total rate of infection and departure is $\lambda(N - I)I$ and μI respectively. Let us first calculate the number of peers having a specific index. Since the rate of change solely depends on the total meeting rate, we have $\dot{I} = \lambda(N - I)I - \mu I$. This differential equation is separable and can be solved with the initial condition I_0 (i.e., the number of BT-APs at the beginning) to give the solution:

$$I(t) = \frac{N - \mu/\lambda}{1 + (N - I_0 - \mu/\lambda)e^{-(\lambda N - \mu)t}} \quad (1)$$

The exponent of the denominator shows that as long as $\lambda N > \mu$, a random index dissemination can reach all the peers. The initial number of nodes I_0 does not have a

significant impact on the overall index dissemination compared to other parameters such as λ and μ . In fact, it is more sensitive to the popularity of content (i.e., N).

The number of peers having a specific piece can be computed using a similar model. Unlike index dissemination, an infected node cannot infect other encountered nodes with probability 1 since we are swarming K blocks of a file and each contact allows only a single piece exchange. At the very beginning, a node without any piece will download a random piece with probability $1/K$. On the other hand, as the node collects blocks, it only tries to download remaining blocks. Since $I(t)$ nodes have a piece at time t and K pieces are spreading independently, a node has a random piece with probability $I(t)/N$. The expected number of remaining blocks is given as $K(1 - I(t)/N)$. At time t a random node encountered by an infected node will download a specific piece with probability $\frac{1}{K(1-I(t)/N)}$. Therefore, the rate of change can be given as $\dot{I} = \lambda(N - I)I\frac{1}{K(1-I(t)/N)} - \mu I = (\lambda\frac{N}{K} - \mu)I$. The equation gives the following solution with the initial condition I_0 .

$$I(t) = I_0 e^{(\lambda\frac{N}{K} - \mu)t} \quad (2)$$

From the equation, we conclude that the key condition of a piece dissemination is $\lambda\frac{N}{K} > \mu$. If the rate of departure is greater than $\lambda\frac{N}{K}$, a piece cannot survive in the network. If $N \gg K$, we have a similar results as before, $\lambda N > \mu$. Otherwise, for fixed λ and μ , the number of blocks K plays a key role in the survivability of a piece. In other words, in the dynamic environment with limited bandwidth, sharing a larger number of blocks (i.e., potentially large files) leads to ineffective file swarming.

5 Simulation

In this section, we evaluate the performance of BlueTorrent using UCBT NS-2 extension, a Bluetooth simulator that is publicly available and open source.⁵ UCBT implements the majority of the protocols in the Bluetooth including baseband, LMP, L2CAP, and BNEP.

5.1 Simulation Setup

Mobility – We assume Bluetooth device users are in a corridor that has a fixed boundary. Users are moving with specific waypoint as from East to West or from West to East. Waypoints are randomly chosen in the initial stage. Maximum speed

⁵ UCBT: <https://www.ececs.uc.edu/cdmc/ucbt>
NS-2: <http://www.isi.edu/nsnam/ns/>

Moving Area (X_r, Y_r)	$[25, 50, 100] \times [3, 5] m^2$
Number of Nodes ($ N $)	25, 50, 75, 100
Number of AP Nodes (N_{AP})	2
Moving speed of nodes (S)	$[0.0, 0.4, 0.8, 1.2, 1.6] m/s$
Packet type (P)	DH5
Block Size/Number (B_S, B_N)	[(6KB, 200), (12KB, 100), (24KB, 50)]
Transfer time (T_t)	10 sec

Table 1
Simulation Parameters

(0.0, 0.4, 0.8, and 1.6 m/s) is predefined to limit node’s speed. 1.6 m/s is selected because this speed is 5.76 km/h and it is about 1.5 times faster than regular walking speed. To add a random factor, direction is changed periodically with an offset in the range $[-10, 10]$ degrees with respect to the original direction. When a node reaches North or South bound of a simulation area, it is mirrored back in the simulation area. When a node reaches East or West bound, we regard that the node moves out of corridor. So, we eliminate that node and regenerate a new node at the other bound (for example, if a node went out of the East bound, a new node starts at the West bound). The regenerated node has same speed and direction of the eliminated node.

When a node is eliminated and a new node regenerated to substitute the eliminated node, we choose *reset* or *no-reset* mode. The reset mode deletes all stored data in the eliminated node and the regenerated node starts with empty data. So, it simulates the situation where one node is going out of an area and a new node is coming into the area. The no-reset mode, on the other hand, transfers all stored data from the eliminated node to the regenerated node. So, it is as if the node going out of one side re-enters from the other side.

Test Scenarios – There are two types of nodes: static APs and mobile nodes. APs have all data set and transfer data to mobile nodes. They are randomly located in the area. Mobile nodes start without data, and they pull data from APs or other mobile nodes. In the *AP mode*, only the AP does inquiry and connects to mobile nodes and transfers data. In the *P2P mode*, every node (AP or mobile node) can connect to the other nodes and exchange data. Initially, only APs can transfer data. Received data blocks from APs are then re-distributed in a P2P fashion. In the simulations, APs distribute a single file. Since the transfer of meta-data is negligible compared to data transfer, we assume that each node already has meta-data of a file.

Metrics – We measure the *download finish time* for *no-reset* mode. In the case of *reset* mode, nodes usually are not able to download a file completely, instead, we measure the progress using *download percentage* of all the nodes that have passed the simulated area during the simulation. By dividing the summation of the downloaded blocks by the nodes, we can calculate the expected number of downloaded blocks. As time tends to infinity, by the strong law of large number, this is equal to the ensemble average of the number of blocks that a random node

can download while passing by the simulated region.

Inquiry and Transfer Stages – In the Inquiry stage, nodes perform inquiry and inquiry scan periodically (i.e., P2P mode). When a node discovers others, it chooses one of the nodes as a master. Peer selection could be a potential problem. To deal with this issue, each node keeps a connection log, which contains a list of peers with discovered time and last connection time. If there are nodes without any connection, the node randomly chooses one. Otherwise, it chooses the earliest connected node, which potentially has met others and carries more fresh blocks than others. This selection is automatically logged. Transfer stage begins after creating a connection to the selected node. Nodes first exchanges block bitmaps, which contain flag bits of data block, to identify useful blocks. If there are useful blocks, data blocks are transferred. Otherwise, the connection is terminated. The master node then selects another node and restarts the overall procedure. If there no more node to connect to, it goes back to the Inquiry stage.

Parameter Summary – We choose as moving area a corridor model that has a relatively long length compared to width. Unless otherwise mentioned, an area of $100 \times 5 \text{ m}^2$ is used by default. The number of nodes varies to change density while the number of AP is fixed to two. Nodes move at the speed up to 1.6 m/s . DH5 packet type is used because this has the highest data rate among all ACL packet types in Bluetooth v1.2. We use 1.2MB size file to transmit and this file is divided as 200, 100, or 50 blocks. Transfer time is set as 10 seconds to tradeoff between reducing inquiry overhead and reducing out-of-range possibility during communication. When transfer time is long, inquiry overhead is reduced but nodes can be exit the communication range more frequently during transmission. When transfer time is short, inquiry overhead is increased. Details of parameters are shown in Table 1.

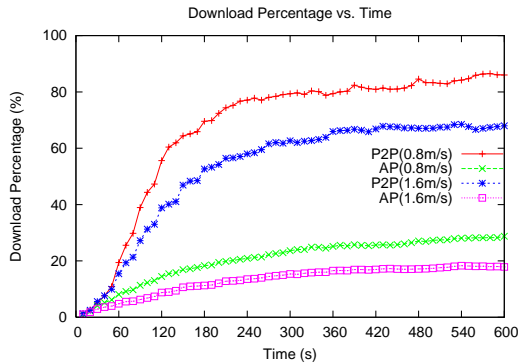


Fig. 8. Download Percentage (reset)

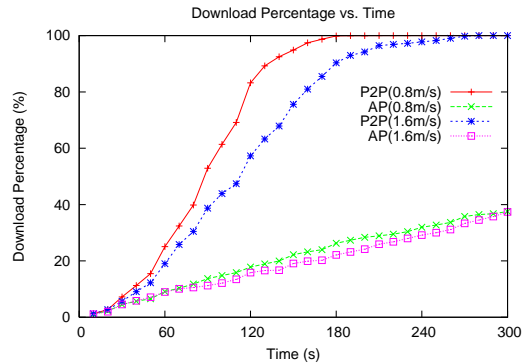


Fig. 9. Download Percentage (no-reset)

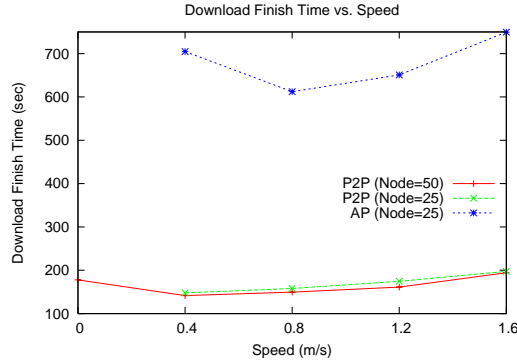


Fig. 10. Finish Time vs. Speed

5.2 Simulation Results

In this section, we mainly compare the performance of P2P and AP modes. The overall download progress is first presented to show the benefits of the P2P mode. We then show the impacts of various parameters, namely the average speed of mobile nodes, the number of blocks, and the corridor length.

Overall download progress – Figure 8 and Figure 9 show average download percentages of both *reset* and *no-reset* modes for every 10 seconds. Total 50 nodes are moving in an area of $100 \times 5 \text{ m}^2$, and one hundred blocks are used. The figure shows that the P2P mode has three times better performance than the AP mode. Multiple peer-to-peer connections increase transmission chances and throughput. In the case of *no-reset*, although at the beginning nodes collect blocks almost linearly over time, after passing around 50%, we see that it takes progressively longer to collect new blocks. This problem is known as a coupon collection problem: the more the coupons you collect, the higher is the probability of collecting an overlapping coupon. This problem can be mitigated by using coding techniques, namely source or network coding [8][20]. Evaluating these techniques is part of our future work. In [18], readers can find some preliminary results of using network coding in mobile ad hoc networks.

Impact of speed – Figure 10 shows average download finish time of nodes. $100 \times 5 \text{ m}^2$ moving area and 100 Blocks are used. The figure shows that as speed increases the download finish time increases. When nodes are static (i.e., speed = 0.0m/s), the density of nodes is critical; if it is below a certain threshold, some nodes may not be able to download any blocks. In the case of the P2P mode, the average finish time is not sensitive to density of nodes. Interestingly, the AP mode has an optimal speed that minimizes the download finish time. Two APs are shared among all the nodes. If nodes move slowly, it is possible that at some point AP is idle (owing to low density, $N=25$). The idle period is mainly dependent on the speed; as speed increases, the idle period decreases. However, if a node moves too fast, the effectiveness of a connection, which is calculated by dividing the data transfer time

by the total connection duration (i.e., discovery+data transfer), decreases. Thus, fast mobility adversely affects the performance. The figures shows that when nodes move at the speed of 0.8m/s on average, the AP mode has the best performance.

Figure 11 shows average numbers of helpful connections and unhelpful connections of nodes. 25 nodes are moving in an area of $100 \times 5 m^2$ and share a 100 block file. The number of unhelpful connections for P2P mode significantly decreases, as speed increases. Fast mobility blends nodes well, thus reducing the chances of having unhelpful connections. The AP mode also experiences a slight decrement of the number of unhelpful connections, since the total number of connections is much smaller than that of P2P case. On the other hand, the number of helpful connection increases up to a certain threshold (P2P case 0.8m/s, AP case 0.4m/s). Thus, we conclude that the gain of fast mobility comes from the relative increment of helpful connections.

Figure 12 shows average lifetime of nodes. $100 \times 5 m^2$ Moving area, 50 nodes, and 100 Blocks are used for this test. As speed increases, lifetime decreases because node reset is more frequent. As corridor length increases lifetime increases because node reset is less frequent.

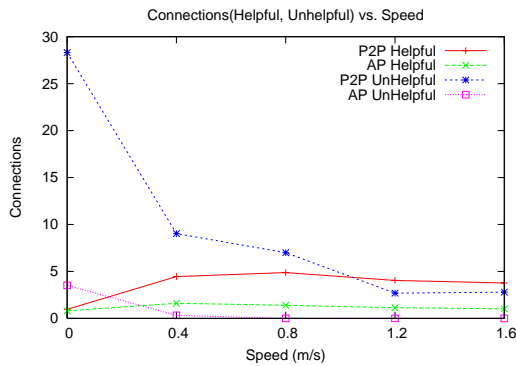


Fig. 11. Connection status

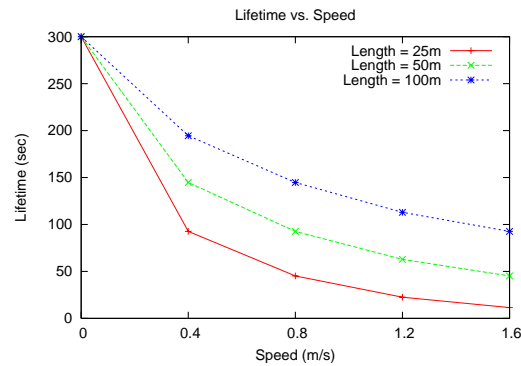


Fig. 12. Lifetime

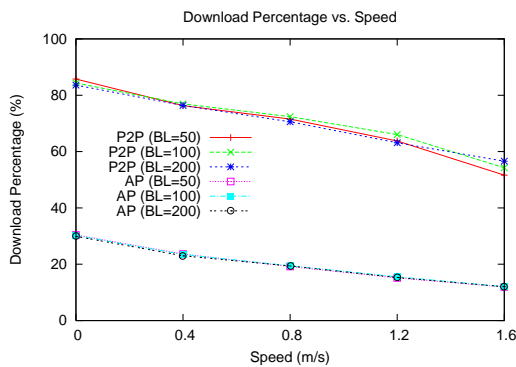


Fig. 13. Number of blocks

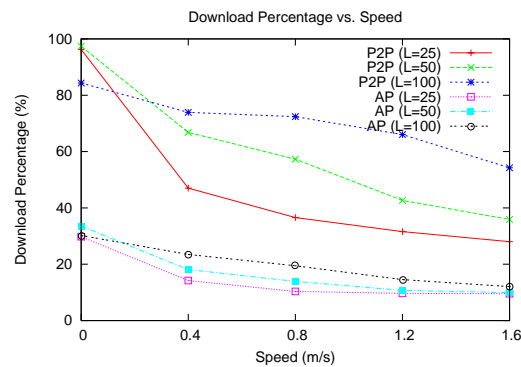


Fig. 14. Corridor length

Impact of the number of blocks – Figure 13 shows average download percentage of nodes at the 200 second mark as a function of speeds with different number

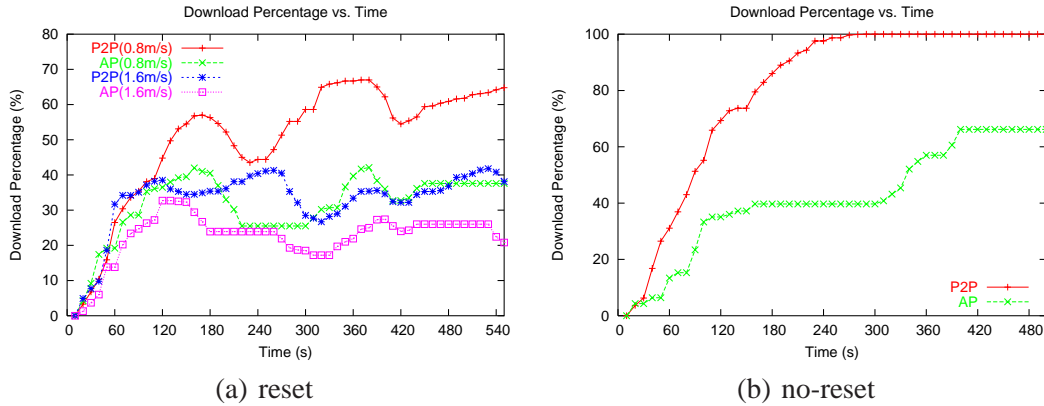


Fig. 15. Download Percentage (Testbed)

of blocks. In general the download percentage is not sensitive to the number of blocks. Although the overhead of L2CAP level acks increases as the number of blocks increases, its impact is not significant. Instead, mobility has a greater impact. With fast mobility a node will experience more frequent disconnections, causing cancelation of the currently downloaded block. If the size of a block is large, this will incur non-negligible performance degradation. Thus, when the speed is 1.6m/s, the largest block size (BL=200) performs the best.

Impact of the corridor length – Figure 14 shows download percentage of nodes at the 200 second mark with different corridor length. When nodes are static (0.0 m/s), there is no mobility; thus, node reset does not happen. So, length of corridor only affects density of nodes and short corridor length shows better performance. When nodes are mobile (0.4-1.2 m/s), short corridor length makes more frequent node reset than longer ones therefore decreases download percentage. Note that density increment in the P2P mode has much greater impact on the performance. As explained before, the performance of the AP mode is not significantly improved by the density after a certain threshold.

6 Experiments

In this section, we show our preliminary testbed implementation results. In the experiment, we use BlueZ Bluetooth protocol stack for Linux ⁶ and Bluetake BT009Si (Silicon Wave, Bluetooth v1.2). BlueZ consists of HCI Core, HCI USB device driver, L2CAP protocol module, and testing utilities. We used 3 desktops and 5 laptops which have RedHat 9.0 with similar configurations of Pentium IV with 512MB RAM. We place two nodes around the AP, and the rest of the nodes are located about 10m away from the AP.

⁶ <http://www.bluez.org>

To emulate mobility we use the following method. The AP is up for a certain percentage of a given period, and for the rest of the period it is down. This emulates nodes moving out of the AP’s communication range. During this period (i.e., when AP is down), only P2P mode can transfer data. Similar to the simulation section, we use the lifetime of the nodes that are moving at maximum speeds of 0.8m/s and 1.6m/s in a 100m length corridor. Nodes are reset at the end of their lifetime to emulate the maximum moving speeds.

Figure 15 shows average download percentage of nodes every 10 seconds. For Figure 15(a), nodes are reset according to the lifetime model. But, for Figure 15(b), they are not reset. P2P mode shows about twice better performance than AP mode because multiple simultaneous peer-to-peer connections increase transmission chances and throughput. In Figure 15(a), the download percentages of both mode increase fast at the beginning, because all nodes have no data. After a while, AP case download percentage drops because AP is down (i.e., nodes are out of AP’s communication range). However, P2P case download percentage continues to increase because of simultaneous transmissions among nodes. In the figure, the AP down period is shown by non-changing download percentage of AP mode. When reset happens, new nodes are generated and thus the average download percentage drops for both AP and P2P cases. In AP mode, this drop shows steeper slope than in P2P case if AP is down during this period. Reset does not happen in Figure 15(b) and thus decrement of download percentage does not happen. However, we still have AP down period in which AP mode shows no change in download percentage during this period. P2P case continuously increases and the average download percentage reaches 100% far earlier than AP mode. AP mode does not reach 100% average download percentage during the test time.

7 Discussion and Future Work

Power Consumption: Inquiry and inquiry scan are the most power consuming state of BlueTorrent and thus it is important to minimize unsuccessful device discovery attempts. The selection of optimum settings shown in Section 3 allows to reduce such failed attempts to discover the highest number of available devices in communication range. Note that power consumption is also different among states of the BT device as shown in [19]. Let us analyze the power consumption of both the client/server and P2P models. Let t denote the total amount of time spent for peer discovery for both cases and t' denote the total amount of time for data transfer. In the client/server model, the master performs the inquiry for t seconds and thus, it will spend $t \times C_{inq}$. The slave node only scans during t seconds, and its power consumption is $\lfloor \frac{t}{1.28s} \rfloor \times 11.25 \times 10^{-3}s \times C_{scan} + [t - \lfloor \frac{t}{1.28s} \rfloor] \times 11.25 \times 10^{-3}s \times C_{Idle}$. After a connection is established, both cases will spend additional $t' \times C_{data}$ for data transfer. These equations clearly show that a master (i.e., inquiry) consumes much more energy than a slave (i.e., scan).

Power consumption of the P2P mode can be analyzed as follows. In the P2P mode, each node repeats inquiry and inquiry scan. In Section 3.2, we show that the length of each trial can be represented as $T_{w_inq} + T_{diff}/2$. In other words, for a given period t , a node spends $\alpha_{inq} = \frac{T_{w_inq}}{T_{w_inq} + T_{diff}/2}$ fraction of time t for inquiry and $\alpha_{scan} = 1 - \alpha_{inq}$ for inquiry scan. Thus, the total power consumption is given as follows:

$$\alpha_{inq} \times t \times C_{inq} + \left\lfloor \frac{t(1 - \alpha_{inq})}{1.28} \right\rfloor \times C_{scan} \times 11.25 \times 10^{-3} s + (1 - \alpha_{inq}) \left(t - \left\lfloor \frac{t}{1.28} \right\rfloor \times 11.25 \times 10^{-3} s \right) \times C_{idle} + t' \times C_{data} \quad (3)$$

Given the above equations, we now present the feasibility of our application on mobile devices. We use the following measured values for calculation as shown in [19]: $C_{inq} = 38mAs$, $C_{idle} = 20mAs$, $C_{scan} = 48mAs$, and $C_{data} = 35mAs$. Let α_{inq} be $1/2$. Assume we used our application one hour each of data transfer and inquiry/scan in P2P mode. We spend $126,000mAs$ for data transfer and $104,858.71mAs$ for inquiry/scan. The total power consumption is $230,856.71mAs$ that is approximately $64.13mAh$. The amount of power used by the application is quite reasonable and it does not drain the battery since the most of recent cellular phone such LG chocolate phone have $800mAh$ standard battery size and Motorola MOTOKRZR K1m have $880mAh$ battery size. The power consumption of LG chocolate phone during two hours usage is approximately $457.14mAh$ and compare to our result, our application's power consumption is very small. Therefore, this result shows that our application is feasible on mobile devices.

Since inquiry consumes more energy than inquiry scan, we propose the following adaptive peer discovery scheme. Initially, all nodes are in their scan modes. Nodes passing by an AP opportunistically are connected. This makes the connected clients to activate the P2P mode. Activated peers can connect/activate other peers. The inquiry and scan intervals are dynamically adjusted based on the contact frequency. Note that evaluating this simple energy conserving scheme is part of our future work.

Another factor for power consumption is Bluetooth packet types for data transmission. According to [9], the more the slots for a given packet, the higher is the transmission power. The energy consumption per bit, however, is higher in the case of using multiple slot packets. Thus, it is preferable to use multiple slot packets, if possible.

User Cooperation: BlueTorrent uses peer-to-peer file sharing that requires users' cooperation. In [24], 25% of Gnutella clients do not share any files and 40-60% of Napster peers share only 5-20% of the shared files, which indicates that there is a large amount of free-riding.

To prevent this selfish peers, several incentive methods were proposed. Cash-based

systems are too complex while lighter-weight credit mechanisms have not provided strong incentives for cooperation. In [13], to solve free-rider problem in peer-to-peer file sharing construct a formal game theoretic model and analyze equilibria of user strategies under several novel payment mechanisms. In [2], exchange-based incentive mechanisms were proposed for cooperation in peer-to-peer file sharing networks. Peers give higher service priority to requests from peers that can provide a simultaneous and symmetric service in return. They showed exchange-based incentive mechanisms can provide string incentives for sharing and reduce free-riders. In [17], Signature-Seeking Drive (SSD), a secure incentive framework was proposed for commercial ad dissemination in Vehicular ad hoc networks. SSD does not rely on tamper-proof hardware or game theoretic approaches, but using Public Key Infrastructure to provide secure incentives for cooperative nodes.

For BlueTorrent, general P2P or Bluetooth P2P specific incentive methods can be applied to reduce free-riders and increase cooperation. We will implement incentive mechanism in the future and improve BlueTorrent more efficiently.

8 Related work

The P2P mode introduced in our paper is also known as symmetric discovery mode since each device alternates their roles as master and slave. To this end, *Periodic_Inquiry_Mode* HCI command was introduced as of Bluetooth specification v1.0. Inquiry window size ($T_{w.inq}$) is fixed and the inquiry period (i.e., the interval in between two consecutive inquiries) is chosen uniform randomly over $[T_{inq.min}, T_{inq.max}]$. To find average discovery latency, Salonidis et al. [23] simplified the model, making it analytically tractable: both inquiry and inquiry scan interval are treated random, and inquiry scan interval ($T_{inq.scan}$) is not considered. But the model hardly reflects a real Bluetooth device, and it cannot be used to optimize the periodic inquiry mode. In fact, it is non-trivial to analytically model the system. In this paper, we accurately simulate the mode and find the optimal parameter configuration. Moreover, we show that the inquiry scan interval, which was neglected in [23], has a great impact on the average latency. Instead of optimizing the periodic inquiry mode, Siegemund et al. [25] proposed a cooperative peer discover protocol. A few nodes per piconet perform cooperative discovery; i.e., each node inquires others and then, proactively exchanges the results. Our finding can be applied to this scheme to further expedite the discovery.

Bohman et al. [4] proposed a variant of the periodic inquiry mode: for each round, the duration of inquiry and inquiry scan states is randomly selected over $[min, max]$. Similarly, in [10], the duration of inquiry and inquiry scan states is chosen over $[C_{inq}, C_{inq} + 2V_{inq}]$ and $[C_{scan}, C_{scan} + 2V_{scan}]$ respectively, where $C_{inq/scan}$ represents the fixed part and $V_{inq/scan}$ the variable part. We assume that these variants were proposed because the authors were unaware of the periodic inquiry mode in

the specification.

Aalto et al. introduces a B-MAD system for delivering permission-based “location-aware” mobile advertisements to mobile phones using Bluetooth positioning and Wireless Application Protocol (WAP) Push [1]. They installed nine Bluetooth sensors (Nokia 3650 phones running the Bluetooth Sensor software) in the display windows of eight retail stores. The store produced eleven advertisements containing some special offers and discounts. They measured positioning time (time to discover a new device), positioning accuracy, scalability, and latency.

BlueCasting is a widely accepted proximity marketing system.⁷ BlueCasting servers can be deployed at poster sites, retails, etc. such that it can identify Bluetooth users and deliver tailored messages. The advertisement log is managed in a central database in order to prevent receiving redundant advertisements. BlueBlitz Magic Beamer supports similar functionalities.⁸ In particular, it provides two way messaging/advertisements and the Internet access as a hotspot. Although these devices support class 1, i.e., 100m communication range, typical Bluetooth devices only support class 2, i.e., 10 m, and thus, mobility of users has a huge impact on the performance. As mentioned before, with an error-prone channel due to multipath, WiFi interference etc., downloading bandwidth is limited; downloading several mega byte files without stopping near the APs is not feasible. BlueTorrent remedies this problem by letting mobile users cooperatively carry and forward data to minimize downloading delay.

A Bluetooth Content Distribution (BCD) station supports content distribution in a “static” environment such as a bus [16]. A BCD station is also equipped with WiFi, and with help of opportunistic connections, it can be synchronized. Once synchronized, data will be distributed to the Bluetooth end users while they are on board. Like other previous products, BCD does neither consider P2P-based content distribution, nor mobility of users.

9 Conclusion

In this paper we proposed BlueTorrent, a cooperative content sharing for mobile Bluetooth users. We analyzed and found the optimum setting for the periodic inquiry mode in order to minimize inquiry/connection time. We then proposed and analyzed index dissemination and file swarming protocols in dynamic, sparse networks. Simulation results showed that cooperative P2P file sharing achieves a greater performance improvement in download time compared to the traditional AP only mode. Finally, via testbed experiments we showed the feasibility of BlueTorrent.

⁷ BlueCasting: <http://www.bluecasting.com>

⁸ BlueBlitz: <http://www.blueblitz.com>

References

- [1] L. Aalto, N. Gothlin, J. Korhonen, , and T. Ojala. Bluetooth and WAP Push Based Location-Aware Mobile Advertising System. In *MobiSys'04*, Boston, NY, Jun. 2004.
- [2] K. Anagnostakis and M. Greenwald. Exchange-based incentive mechanisms for peer-to-peer file sharing. In *24th International Conference on Distributed Computing Systems (ICDCS'04)*, Tokyo, Japan, March 2004.
- [3] BBC NEWS: Music trial taps into Bluetooth. <http://news.bbc.co.uk/1/hi/technology/4392534.stm>.
- [4] D. Bohman, M. Frank, P. Martini, and C. Scholz. Performance of Symmetric Neighbor Discovery in Bluetooth Ad Hoc Networks. In *German Workshop on Mobile Ad-hoc Networking (WMAN'04)*, Ulm, Germany, Dec. 2004.
- [5] Bluetooth SIG. Bluetooth Specification v2.0, 2004.
- [6] J.-C. Cano, J.-M. Cano, E. Gonzalez, C. Calafate, and P. Manzoni. Power Characterization of a Bluetooth-based Wireless Node for Ubiquitous Computing. In *ICWMC'06*, Bucharest, Romania, July 2006.
- [7] CBS Goes Bluetooth To Promote Fall TV Line-Up. <http://www.telecomweb.com/tnd/18847.html>.
- [8] P. Chou, Y. Wu, and K. Jain. Practical Network Coding. In *51st Allerton Conf. Communication, Control and Computing*, Allerton, Oct. 2003.
- [9] DARFON Bluetooth Module Data Sheet Aug 2005 v1.1. <http://www.darfon.com>.
- [10] C. Drulă. Fast and Energy Efficient Neighbour Discovery for Opportunistic Networking with Bluetooth. *Master Thesis, Dept. of Computer Science, University of Toronto*, 2005.
- [11] The Economist: Bluetooth's quiet success. http://www.economist.com/science/tq/displayStory.cfm?story_id=7001843.
- [12] J. Ghosh, S. Yoon, H. Q. Ngo, and C. Qiao. Sociological Orbit for Efficient Routing in Intermittently Connected Mobile Ad Hoc Networks. In *Technical Report TR-2005-19*, University at Buffalo, Apr. 2005.
- [13] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge. Incentives for sharing in peer-to-peer networks. *Lecture Notes in Computer Science*, 2232:75+, 2001.
- [14] T. Hofmann. Probabilistic Latent Semantic Analysis. In *UAI'99*, 1999.
- [15] S. Jung, U. Lee, A. Chang, D. Cho, and M. Gerla. BlueTorrent: Cooperative Content Sharing for Bluetooth Users. Technical report, UCLA, Dec. 2006.
- [16] J. LeBrun and C.-N. Chuah. Feasibility Study of Bluetooth-Based Content Distribution Stations on Public Transit Systems. In *ACM MobiShare'06*, Los Angeles, CA, Sept. 2006.

- [17] S.-B. Lee, G. Pan, J.-S. Park, M. Gerla, and S. Lu. Secure incentives for commercial ad dissemination in vehicular networks. In *The Eighth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc '07)*, Montreal, Canada, September 2007.
- [18] U. Lee, J.-S. Park, J. Yeh, G. Pau, and M. Gerla. CodeTorrent: Content Distribution using Network Coding in VANETs. In *MobiShare'06*, Los Angeles, CA, Sept. 2006.
- [19] J. B. L. Negri and M. Dyer. The Power Consumption of Bluetooth Scatternets. In *IEEE Consumer Communications and Networking Conference*, pages 519–523, Jan. 2006.
- [20] P. Maymounkov and D. Mazieres. Rateless Codes and Big Downloads. In *IPTPS'03*, Berkeley, CA, Feb. 2003.
- [21] B. S. Peterson, R. O. Baldwin, and J. P. Kharoufeh. Bluetooth Inquiry Time Characterization and Selection. *IEEE Transactions on Mobile Computing*, 5(9):1173–1187, Sep. 2006.
- [22] S. M. Ross. *Introduction to Probability Models*. Academic Press, 2002.
- [23] T. Salonidis, P. Bhagwat, L. Tassiulas, and R. O. LaMaire. Distributed Topology Construction of Bluetooth Personal Area Networks. In *INFOCOM*, Anchorage, AK, Apr. 2001.
- [24] S. Saroiu, P. K. Gummadi, and S. D. Gribble. A measurement study of peer-to-peer file sharing systems. Technical Report UW-CSE-01-06-02, Dept. of Computer Science and Engineering, Univ. of Washington, July 2001.
- [25] F. Siegemund and M. Rohs. Rendezvous Layer Protocols for Bluetooth-Enabled Smart Devices. *Personal and Ubiquitous Computing Journal*, 7(2):91–101, Jul. 2003.
- [26] T. Small and Z. J. Haas. The Shared Wireless Infostation Model - A New Ad Hoc Networking Paradigm (or Where There is a Whale, There is a Way). In *ACM MOBIHOC'03*, Annapolis, Maryland, USA, June 2003.
- [27] W. Stallings. *Wireless Communications and Networks*. Prentice Hall, 2005.

Sewook Jung received B.S. in Physics and M.S. in Computer Engineering from Seoul National University in 1996 and in 1998, respectively, and received Ph.D. degree in Computer Science from University of California, Los Angeles in 2007. He worked at Samsung Electronics from 1998 to 2003. He will start working at Broadcom Corporation from 2007. His research interests are Wireless Personal Area Networks (Bluetooth, ZigBee, Wibree, and UWB), HealthNet, Ubiquitous Computing, and Network Protocols.

Uichin Lee received the B.S. in computer engineering from Chonbuk National University in 2001, and the M.S. degree in computer science from Korea Advanced Institute of Science and Technology (KAIST) in 2003. He is currently a Ph.D. candidate with the Department of Computer Science, University of California, Los Angeles. His research interests include "mobile" wireless sensor networks, e.g., vehicular and underwater sensor networks, and user behavior studies, e.g., P2P and web search. He is currently working on vehicular sensor networks with his advisor Prof. Gerla in Network Research Lab (NRL).

Alexander Chang received the B.S. degree in Computer Science and Applied Mathematics from University of Washington in 2003, and the M.S. degree in Computer Science from University of California, Los Angeles, in 2006. He is currently pursuing the Ph.D. degree in Computer Science at University of California, Los Angeles. His research interests are Wireless Sensor/Personal Networks and Wireless Network Security.

Dae-ki Cho received BS degree in Computer Science from University of California, Los Angeles in 2006. He is currently pursuing master's degree in the Department of Computer Science at University of California, Los Angeles. His research interests are Wireless Sensor/Personal Networks, Underwater Networks, and Data Minings. He is currently working on Bluetooth P2P with his advisor Prof. Gerla in Network Research Lab (NRL).

Mario Gerla received a graduate degree in engineering from the Politecnico di Milano in 1966, and the M.S. and Ph.D. degrees in engineering from UCLA in 1970 and 1973. He became IEEE Fellow in 2002. After working for Network Analysis Corporation, New York, from 1973 to 1976, he joined the Faculty of the Computer Science Department at UCLA where he is now Professor. His research interests cover distributed computer communication systems and wireless networks. He has designed and implemented various network protocols (channel access, clustering, routing and transport) under DARPA and NSF grants. Currently he is leading the ONR MINUTEMAN project at UCLA, with focus on robust, scalable network architectures for unmanned intelligent agents in defense and homeland security scenarios. He is also conducting research on scalable TCP transport for the Next Generation Internet (see www.cs.ucla.edu/NRL for recent publications).



Fig. 1. Sewook Jung



Fig. 2. Uichin Lee



Fig. 3. Alexander Chang



Fig. 4. Dae-ki Cho



Fig. 5. Mario Gerla

List of Changes

Reviewer #1: This is a very well-written paper, covering well most of the relevant aspects of using Bluetooth to support "torrent"-like sharing mechanisms.

I'm not an expert in communication and discovery mechanisms, so I will decline to comment on those issues. But by being an expert on HCI issues related to mobile computing, it seems that this paper fails to address a fundamental issue that has huge issues in terms of adoptance: power consumption. Unlike in BitTorrent, which often runs on desktop machines directly connected to power, a "torrent"-like system in a mobile device has the potential to draw all the energy of one's battery... for the benefit of other people! So unless the power consumption of such a system is reasonable to the point of not impacting a user's everyday usage, it is very unlikely that people will adopt this technology. And since battery capacity is not increasing in a fast rate, that would mean that this technology would be dead before even starting.

So I would like the authors to include a sub-section discussing power issues related to having a mobile device (such as a cell-phone) doing bluetooth sharing, and how the battery span of such a device would be impacted. Some basic estimation is needed here to make an argument, considering for instance, how much data could be shared before draining 20% of a typical cell phone in your simulations, and how long that would take. This would give the reader a better understanding of the real possibilities of this technology in the real world.

Section 7 is added and discuss the power consumption of BlueTorrent compared to that of cellular phone calling. BlueTorrent can be used more than 12 hours in regular cellular phone and power consumption is negligible.

Reviewer #2: The authors look at the problem of implementing a mobile P2P network over multiple Bluetooth devices. In particular, the paper analyzes the Bluetooth discovery process, and finds the optimum inquiry and connection time settings that in simulation, and in preliminary experiments, work for this application.

The paper is very well written, and includes a nice overview of Bluetooth with particular attention to the discovery process. I believe the paper is self-contained, in that someone who is not familiar with Bluetooth should be able to work through the paper.

The biggest weakness of the paper is that their approach has only been implemented in a very contrived setting (3 desktop computers and 5 laptops, with AP's being brought up and down periodically to simulate mobile devices coming in and out of

range). It's hard to say whether the paper's findings would hold up under a more realistic deployment.

Of course, it is very difficult to deploy such a system in a more realistic way, since software to support the system would have to be installed on a reasonable percentage of the mobile devices before this particular P2P network would be useful. It also remains to be seen whether (1) people will consent to having their phones used in this way, (2) people desire to pull down media files and other files in this fashion and (3) that providers would cooperate (I think that is probably required). But these concerns are primarily outside the scope of the present work; one should consider this paper as an examination of whether Bluetooth is up to the task, given that system software developers can tweak some of the timings here and there. I think the conclusion of the paper is that it is up to the task, but that there are some lingering doubts about power consumption.

Section 7 describes some incentive methods already proposed for P2P file sharing network and Car-to-Car network in the other papers. These already proposed or Bluetooth specific incentive methods are applied to BlueTorrent in the future, and increase cooperation of users.

Ultimately, I think the paper is a useful analysis of the Bluetooth discovery process and is illuminating in how it brings out some of the relevant parameters that are available to system software developers. An analysis of how power consumption is affected by the various settings would complete the picture, but perhaps that is too much for one paper.

Section 7 is added and discuss the power consumption of BlueTorrent compared to that of cellular phone calling. BlueTorrent can be used more than 12 hours in regular cellular phone and power consumption is negligible.