

Maximizing Transfer Opportunities in Bluetooth DTNs

Marc Liberatore Brian Neil Levine
Department of Computer Science
University of Massachusetts, Amherst
Amherst, MA USA
{liberato,brian}@cs.umass.edu

Chadi Barakat
Planète Research Group
INRIA Sophia-Antipolis
Sophia-Antipolis, France
chadi.barakat@sophia.inria.fr

ABSTRACT

Devices in disruption tolerant networks (DTNs) must be able to communicate robustly in the face of short and infrequent connection opportunities. Unfortunately, one of the most inexpensive, energy-efficient and widely deployed peer-to-peer capable radios, Bluetooth, is not well-suited for use in a DTN. Bluetooth's half-duplex process of neighbor discovery can take tens of seconds to complete between two mutually undiscovered radios. This delay can be larger than the time that mobile nodes can be expected to remain in range, resulting in a missed opportunity and lower overall performance in a DTN. This paper proposes a simple, cost effective, and high performance modification to mobile nodes to dramatically reduce this delay: the addition of a second Bluetooth radio. We showed through analysis and simulation that this dual radio technique improves both connection frequency and duration. Moreover, despite powering two radios simultaneously, nodes using dual radios are more energy efficient, spending less energy on average per second of data transferred.

Categories and Subject Descriptors

C.2.1 [Computer-Communications Networks]: Network Architecture and Design—*Wireless communication*; C.4 [Performance of Systems]: Reliability, availability, and serviceability

General Terms

Design, Performance

Keywords

Bluetooth, disruption tolerant networks, neighbor discovery

This work was done in part while the authors visited Intel Research Cambridge. This work was also supported in part by the U.S. National Science Foundation under grants CNS-0519881, CNS-0133055, and ANI-0325868, and by DARPA under contract W15P7T-05-C-P213.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CoNEXT'06, December 4–7, 2006, Lisboa, Portugal.
Copyright 2006 ACM 1-59593-456-1/06/0012 ...\$5.00.

1. INTRODUCTION

“Never underestimate the bandwidth of a station wagon full of tape hurtling down the highway.”
– A. Tanenbaum [25]

Never overestimate the bandwidth available between two station wagons as they pass each other hurtling down the highway. Such in-motion data *transfer opportunities* play a fundamental role in *disruption tolerant networks* (DTNs), which are networks that attempt to route data despite intermittent and infrequent link-layer connectivity. Mobility in DTNs is likely the largest source of disconnections, so radios in DTNs must be capable of fast *discovery* of new nodes and efficient transfers to those radios once discovered.

Bluetooth is the most ubiquitous peer-to-peer capable wireless radio carried by humans. It is increasingly integrated into such mobile devices as phones, PDAs, and laptop computers [13], yet the installed base of devices is poorly suited to the conditions found in a DTN. Others have noted the difficulty of supporting mobility applications using Bluetooth [3] — it was originally designed for replacing cables between stationary devices. In particular, the *inquiry* protocol for discovering other radios is time consuming. Successful inquiry can take as long as 10 seconds, and in our experiments required about 3.5 seconds on average. Additionally, nodes cannot themselves be discovered while they are inquiring, which can delay discovery further. Given that moving nodes will eventually pass out of range of each other, this is a fundamental problem for using Bluetooth successfully in DTNs.

As we have shown in previous work [6], other radios such as 802.11b have more desirable characteristics when power is not an issue, such as in automotive DTNs. In this paper, we examine scenarios where communication devices are carried by individuals, and are thus power-constrained. As Bluetooth is an inexpensive and widely-deployed radio technology, it is the focus of our study.

In this paper, we show through analysis and simulation that devices with a *single* Bluetooth radio present a major impediment to node-to-node transfers and therefore end-to-end throughput because of the long, half-duplex discovery process at each node. We propose the use of a second, off-the-shelf Bluetooth radio to provide a full-duplex inquiry channel for neighbor discovery. The goal of our *dual* radio approach is to enhance the performance of Bluetooth in DTNs without modifications to the standard, enabling more and better uses of the millions of devices already deployed.

In our simulations comparing the single and dual radio approaches, we found that dual radios reduce the time until

nodes discover one another by 12–25%, which results in an increase in the number of transfer opportunities successfully discovered to 170–440% of the single radios (depending on the speed of the nodes). For this reason, dual radio nodes have mean time spent transferring data to neighbors that is 220–240% of the time achieved by single radios. Moreover, despite powering two radios simultaneously, nodes using dual radios are more energy efficient, spending 7–27% less energy on average per second of data transferred. We found that single radio devices have a limited effectiveness in the network that cannot be solved with unlimited energy resources. Unlike the dual radio approach, single radios do not always increase the length and number of transfer opportunities nodes receive in a DTN when the energy spent on discovery is increased.

In addition to the performance gains that dual radios provide, this solution is simple, cost-effective, and incrementally deployable. Many portable devices contain expansion slots that can support a second radio.

We begin by presenting a model of the Bluetooth inquiry process between two nodes. This model predicts the performance gains of our dual radio solution over the single-radio approach independent of radio ranges. We validate the analytical model with an empirical simulation, and then extend the simulation to use more realistic assumptions about the inquiry process that are difficult to model analytically. It is also important to quantify the performance gains of our dual radio solution in the context of multi-node DTN. Therefore, we extend the simulation again to model multiple moving nodes, which allows us to determine the effects of delays between transfer opportunities.

In the next section we give an overview related work. In Section 3, we present our model and evaluation of the inquiry process between two nodes. In Section 4, we extend the model and performance evaluation to a multiple moving nodes. We conclude and explore directions for future work in Section 5.

2. BACKGROUND

In this section we survey existing work on DTNs. Much of this work is focused on routing rather than improving individual contact opportunities. We also provide an overview of Bluetooth radio technology and work related to radio performance.

2.1 DTNs

There is a large body of work examining a message ferrying approach to solving connectivity problems in partitioned Mobile Ad-hoc NETWORKS (MANETs) and to improving network performance. One of the first presentations of message ferrying is given by Davis et al. [11]. Ammar et al. have also written several papers on message ferrying [28, 29]. Similar algorithms have been proposed by Sarafijanovic-Djukic and Grossglauser [16, 22], Burns et al. [7], and Burgess et al. [6], exploiting structure in mobility to improve overall routing performance. Burns et al. [7] also propose the use of robotic devices for increasing the performance and capacity of DTNs. Zhao et al. [30] suggest the placement of battery-powered radios with storage to improve DTN performance. Fall et al. [18] formalize the problem of routing in a DTN, and give an argument as to why such routing is distinct from and more difficult than routing in traditional fixed networks.

Ott and Kutscher investigated the performance of 802.11b

| | 802.11b | Bluetooth |
|-----------------------------------|------------|-----------|
| Chipset power consumption: | | |
| Idle | ⊥ | 0.65 mW |
| Inquiring | 180–280 mW | 1.2 mW |
| Tx/Rx | ⊥ | 21 mW |
| Discovery-related times: | | |
| Inter-inquiry delay | 100 ms | variable |
| Inquiry duration | 3 ms | 10.24 s |
| Avg successful inquiry | 3 ms | 3.5s |
| Performance: | | |
| Bandwidth | 11 Mbps | 721 kbps |
| Range | 250 m | 10 m |

Table 1: A comparison of class 2 Bluetooth and 802.11b radios.

and TCP/IP on the Autobahn [19]. Their investigation finds that properly configured base stations and mobile radios could achieve nontrivial data transfers during short, intermittent periods of connectivity.

2.2 Radio Technologies

Bluetooth [5] and 802.11b [1] are two widely available radio technologies, and thus worth investigating for use in DTNs. 802.11b offers performance comparable to traditional wired Ethernet. Bluetooth consumes less power and has lower range. It was originally designed as a peripheral cable replacement technology, but is now widely deployed in mobile devices. Discovery of newly in-range radios in Bluetooth takes much longer than in many comparable systems. Table 1 shows the values of key characteristics for a specific class 2 Bluetooth [4] and 802.11b [26] radio.

In this paper, we are concerned with aspects of neighbor discovery, which is known as *inquiry* in Bluetooth. The Bluetooth inquiry process works roughly as follows. A device interested in discovering its neighbors transmits an inquiry. Other devices reply, and are thus discovered. The key problem present in Bluetooth 1.0 and 1.1 is that a device performing inquiry cannot be discovered by other inquiring devices – it is a half-duplex process. The Bluetooth 1.1 specification states that an inquiry process for neighbor discovery should last about ten seconds, and we have fixed our inquiry period to this length in our study. This inquiry delay is due to a repeated sequence of scans through the entire allotted frequency range. It would be much faster to have a fixed set of frequencies for device discovery, but Bluetooth must perform this long scan to comply with ITU and FCC regulations [15]. The regulations guarantee that no single technology dominates the license-free 2.4Ghz spectrum by preventing devices from monopolizing any specific frequency for substantial amounts of time. Newer Bluetooth specifications reduce the inquiry time to half of that in Bluetooth 1.0 or less. However, after failing to discover neighbors, newer devices fall back to the legacy mode. As described later in this paper, our contribution makes the deployment of DTNs more practical in the face of legacy and legacy-mode devices, as it provides an energy-efficient, legacy-compatible system with the small added cost of a second, inexpensive Bluetooth radio.

There are several studies of Bluetooth performance relevant to this paper. Salonidis, et al. [21] analyze the Bluetooth discovery protocol and look for parameters that con-

trol its length. Their technique allows for a minimization of the expected discovery time, but relies upon changes that are not possible in software for many Bluetooth implementations. Basagni et al. [3] evaluate the Bluetooth specification for various topologies of Bluetooth scatternets. As in our work, they find that the inefficiency of device discovery is a major problem with Bluetooth. Peterson et al. [20] devise a complex model of Bluetooth inquiry and validate it against empirical measurements. In the process, they identify a modification to the inquiry process made by the hardware vendor that retains specification-level compatibility while improving inquiry time, though still requiring seconds on average. Duflot et al. [14] also provide an exact model of Bluetooth inquiry. Chakraborty, et al. [10] found that Bluetooth device discovery depends on the density of the devices and the random back-off time.

Bluetooth is actively studied as a component of DTNs. Hui, et al. [17] and Chaintreau, et al. [9] utilized Bluetooth connection data gathered as part of the HAGGLE project [12]. Among their contributions is a characterization of the duration of contact opportunities and the delay between such opportunities in a DTN-like environment.

Our work is also related to previous work on using multiple tiers of radios. For example, Shih et al. [23] and later work by Bahl et al. [2] have found that using multiple radios per device can reduce power consumption, and that larger mesh network capacity can benefit as well. Similarly, Sorber et al. [24] have proposed multi-tier platforms for power management in mobile devices. Woodings, et al. [27] and Busboom, et al. [8] proposed the use of secondary systems, such as integrated infrared transceivers or RFID tags to speed Bluetooth discovery.

3. MAXIMIZING INDIVIDUAL OPPORTUNITIES

In this section, we present a model of the Bluetooth inquiry process with the goal of maximizing the utility of individual transfer opportunities. Transfer opportunities that occur intermittently but that last minutes or hours are trivial to support at the link layer. However, the combination of mobility and low node density can result in short-duration transfer opportunities that can be missed. Making the most of these opportunities is therefore of paramount importance.

We present a mathematical model of available off-the-shelf radio technology, and use this model to develop a tunable strategy for maximizing the frequency and length of individual transfer opportunities. We validate the model using a simulator we developed independently. We make the simplifying assumption that all radio links are symmetric.

Bluetooth radios are extremely widely deployed, inexpensive (less than \$5 per unit), and peer-to-peer capable. Thus we developed our model to examine Bluetooth’s neighbor discovery performance. However, our model is general and can be adapted to other radios. In particular, we apply our model to traditional Bluetooth nodes, and to nodes carrying two independent Bluetooth radios. We refer to the latter as a *dual radio* configuration. The main purpose of this configuration is to provide a full-duplex channel for discovery.

In this section, we are concerned only with the perspective of a single node and its discovery of other nodes. In Section 4, we present a more global view and examine how discovery and node density affect a network composed of an

| | |
|--|--|
| Feature of radio technology used: | |
| D | Inquiry duration |
| Configurable by peers: | |
| I | A random variable representing idle time between inquiries |
| $1/\lambda$ | Mean of the idle time distribution |
| Assumed or used by our analysis: | |
| $G(x)$ | Probability of inquiry of length x succeeding |
| Derived by our analysis: | |
| T | Mean time from start of analysis until detection succeeds |

Table 2: Table of variables.

arbitrary number of nodes.

3.1 Radio Model

We define a strategy as an *idle time distribution*, I , and a *inquiry length distribution*, D . I is the inter-inquiry delay, and during idle mode, the radio is listening for other’s inquiries. In this section, we examine through a mathematical model the performance of two device strategies:

1. Devices with a single Bluetooth radio with strategy $\{(I; D = 10s)\}$, where I is an exponentially-distributed random variable with mean $1/\lambda$;
2. Devices with two Bluetooth radios with a strategy $\{(I; D = 10s), (I = \infty, D = 0s)\}$, with the first radio’s I as above. In this case, one radio inquires periodically, and the other is always on and always listening.

These strategies correspond to actual implementations of Bluetooth, with an inquiry duration of 10 seconds and a configurable inter-inquiry delay. Note that we choose an exponentially-distributed inter-inquiry delay to avoid the problem of synchronization; other models for delay and their analysis correspond to the one presented here. Our goal is to find for each strategy the idle time distribution I that minimizes the time until a neighbor discovery occurs, so that we can lengthen the resulting transfer opportunity and increase the number of bytes transferred. Individual Bluetooth radios are half duplex, and so cannot be found while inquiring. Thus having a single radio system inquire continuously is ineffective. Avoiding this problem through a full-duplex channel for discovery is what originally motivated our dual radio proposal.

To develop and compare strategies, we give a model of radio performance specialized to the case of discovering other radios as they come into range. In this section, we present our model, and the resulting analysis for the single and dual radio case. Under our assumptions, we are able to derive a closed-form expression for the effects of the inter-inquiry time for both single and dual radios. This enables us to choose a time that maximizes transfer opportunities.

3.1.1 Single Radio Analysis

We model each radio as a state machine. In our model, a radio is always in one of two states: *idle* or *inquiring*. When a radio is idle, it is listening for the transmission of another radio’s inquiry. When a radio is inquiring, it is transmitting, and unable to receive other radio inquiry. We define a step function $G(x)$ as the probability that an inquiry without interference of length x succeeds. We elide the details of




Figure 1: Predicted time until connection according to our model.

our model here; they can be found in Appendix A. One important assumption of the model is that successful inquiries require D time (i.e. $G(x)$ steps at $x = D$) — in practice, successful inquiries typically take time less than D . The result of this model is a closed form expression for T , which is the expected time from when two nodes are in range until a successful inquiry completes. (Recall $1/\lambda$ is the mean of the idle time distribution.)

$$T = \frac{D \lambda D(12 + 18\lambda D + 5(\lambda D)^2 + (\lambda D)^3)}{12(1 + \lambda D)^2} + e^{\lambda D}(1 + \lambda D)^2/(2\lambda)$$

From this expression, one can easily derive that the optimal mean inter-inquiry delay I is equal to $2.66D$. In Section 3.2, we validate this model and examine the effect of our assumptions with an empirical simulation.

3.1.2 Dual Radio Analysis

By installing two radios in a single device, we create a full-duplex channel for discovery. In the absence of power constraints, the optimal idle time for the first radio is zero: if a radio is always inquiring, then the detection will occur as soon as possible. This strategy is not the most energy efficient strategy when power is a limited resource, a topic we return to in Section 4. To provide a meaningful comparison with the single radio system, we determine the expected time until success for the dual radio case. This result is general for all inter-inquiry delays, and can be used to optimize system design for specific sets of power requirements and battery capacities.

As above, we leave the details of analysis to Appendix B, and present here our result for T , again the expected time from when two nodes are in range until a successful inquiry completes:

$$T = \frac{3 + 6\lambda D - 6\lambda^2 D^2 + 2\lambda^3 D^3}{6\lambda(\lambda D + 1)^2} + D$$

This is the average time until two dual radio devices detect each other, and is much better than the single-radio result with the same inter-inquiry delay, as shown in Figure 3.




Figure 2: Validation of the mathematical model.

3.1.3 Results

Graphs of the results of our analysis of neighbor discovery for single and dual Bluetooth radios are shown Figure 1. Our choices of 3.5 and 10 seconds for D are not arbitrary: in field measurements, we found D to be about 3.5 seconds in the absence of external interference, and the Bluetooth specification requires a full inquiry to last 10.24 seconds. For reasonable values for Bluetooth, dual radios are dramatically better than single radios, and as the expected value of I (i.e., $1/\lambda$) increases, the performance degrades.

Our model is useful because it shows bounds on the performance of the single-radio case independent of radio range (i.e., class 1, 2, or 3 Bluetooth). Our analysis below shows that supporting DTNs with a single Bluetooth radio is a challenge; given our assumptions, peers must be in range for about 12.5 seconds for an expected successful connection, a number that is independent of radio range.

Selecting a radio range does allow us to take away a concrete example from the model. Given a radio range, we can use our model to determine *speed limits* for expecting successful discovery. For two peers passing each other head-on with a single class 1 Bluetooth radios (100m range) each, they cannot be traveling faster than 8m/s relative to one another (about 18mi/h each, the speed of a bicyclist). When each has a single class 2 Bluetooth radios (10m range), which are common to phones and PDAs, the peers cannot travel faster than about 1.7mi/h each (a slow walking speed) as they walk head-on. These speed limits are only for a successful discovery and leave no time for data transmission. Interestingly, even when peers do not have constrained power resources, these speed limits do not change.

3.2 Simulation

We also constructed an event-based simulator to determine the time until a successful inquiry. This simulator serves two purposes. First, it allows us to validate the results of our mathematical model. Additionally, it allows us to relax some of our constraints to more closely conform with reality, and also to compare results between 802.11b and Bluetooth radios. Specifically, for Bluetooth, we set the trigger on the step function determining the probability of detection at 3.5, corresponding with the value we observed in the absence of interference.




Figure 3: Expected time until discovery.

The behavior of the simulator is as follows. Two nodes equipped with idealized radios perform inquiries periodically. The nodes move into range of one another at some point during this cycle of inquiries. We neglected external interference and details of the Bluetooth stack. Once nodes are in range during the simulation, they stay in range, and we calculate for each strategy how long the nodes must be in range for a successful neighbor discovery.

We first ran the simulation with parameters identical to those implied by the assumptions of our mathematical model. The results are shown in Figure 2. The agreement between simulation and mathematical model is almost exact, giving us some degree of confidence that our models are correct. Also shown is the result of a more realistic $G(x)$ function: the probability that inquiry without interference of length x succeeds. Our model requires that $G(x)$ step when $x = D$, but in our simulation we can set the value for the step function to conform with the mean time for successful inquiry that we measured with real radios.

The graph in Figure 3 compares the results of our simulated Bluetooth radios in single and dual configuration, and of an 802.11b-like radio. “BT single radio” is for a $\{(I, D = 3.5)\}$ strategy. “BT dual radio” is the result when two Bluetooth radios are installed in each peer using a $\{(I; D = 3.5), (I = \infty, D = 3.5)\}$. As before, two Bluetooth nodes discover one another when one node receives 3.5 seconds of inquiry with transmitting its own inquiry. “802.11b” is for a radio with strategy $\{I : \text{normal} (\mu = 0.1, \sigma = \sqrt{\mu}), D = 0.003\}$. The 802.11b discovery function $G(X)$ steps at its D , 0.003. As previously stated, we assume Bluetooth is configured to use an exponentially-distributed inter-inquiry delay to avoid synchronization; we assume a normal inter-inquiry delay distribution to reflect the fixed beaconing of 802.11b radios. For all strategies, the mean of I , ($1/\lambda$ for Bluetooth, μ for 802.11b), is a free variable on the x-axis. The optimal value for mean inter-inquiry delay for single-radio Bluetooth is approximately 12.5 seconds, and that dual radios can achieve this performance with a mean delay of 25 seconds.

4. MULTIPLE NODE EVALUATION

In this section, we examine the relative performance of the radio technologies and techniques we have discussed thus

far in scenarios involving up to 100 nodes. Our results show that, as compared to the single radio case for Bluetooth, our proposed dual radio solution provides more opportunities for transfer that are longer in duration on average. Moreover, dual radio solutions are a most energy efficient solution.

To evaluate these performance metrics, we developed our own event based simulator of mobile nodes carrying Bluetooth devices. This section details our simulation model and presents our results.

4.1 Simulation Model

Our simulation is event driven and uses a mobility model based on random, predetermined paths which we consider a form of the freeway mobility model. In our model, the appearance of nodes is dynamic: Nodes enter and leave the simulation at consistent points, and follow paths between these points. This models pedestrian and vehicle movement along walkways and streets.

Nodes move in the *walking scenarios* with a speed drawn uniformly from the interval (1.0, 2.0) m/s and in the *biking scenarios* with a speed drawn uniformly from (2.0, 9.0) m/s. We did not simulated mixtures of biking and walking speeds in the same scene. We believe that self-propelled vehicles would not typically use Bluetooth as they have the power available for more energy-intensive radios, and so we did not simulate automotive vehicle speeds for this initial study.

We generated eight random *scenes*, four each for biking and walking consisting of 25, 50, 75, and 100 nodes in scene at all times (i.e., whenever a node left, a new node joined). Each scene had a new, random set of paths across a flat, 1000m by 1000m square area. Each scene lasted 15,000 simulated seconds (i.e., 4 hours). This means, for example, that over 5,000 nodes were simulated in the 100-node biking scenario and 1,375 in the 100-node walking scenario. The latter used fewer nodes because it takes longer for nodes to leave the scene.

In the results we present, the wireless range of all nodes was 10m. Results for 100m (i.e., representative of class 1 Bluetooth) are not significantly different and are thus not shown.

We made some simplifying assumptions that should not affect the relative performance of the dual and single radio strategies. First, we assumed that nodes can not be discovered while transferring data, although in practice this is not true. This assumption results in a consistently greater delay for nodes to discover one another when already in contact with another node, but allowed us to avoid making connection policy decisions and greatly simplifies the simulation. We capped all data transfers at 15 seconds, with the assumption this was sufficient time to transfer all necessary data. We assumed nodes require exactly 3.5 seconds for a successful inquiry of another node, while there is more variance to discovery in practice. We had transfers begin as soon as discovery occurred. All unsuccessful inquiries lasted 10 seconds, which is the time specified by the Bluetooth specification.

We did examine various inter-inquiry delays, but unless otherwise noted the results are for inter-inquiry delays with a mean of 12.5 seconds, which is the predicted best-performing case for the single-radio strategy.

Given our movement model and radio range, we measured the mean amount of time that nodes were in wireless range of one another in our model. Nodes in the biking scenarios were

| Nodes | Mean time until discovery (seconds) | | | |
|-------|-------------------------------------|------|---------|------|
| | biking | | walking | |
| | Single | Dual | Single | Dual |
| 100 | 5.2 | 4.6 | 6.9 | 5.3 |
| 75 | 5.4 | 4.6 | 6.8 | 5.2 |
| 50 | 5.3 | 4.5 | 6.8 | 5.1 |
| 25 | 5.5 | 4.6 | 6.0 | 5.2 |

Table 3: Mean time until discovery from when nodes are first in range (or from last transmission).

in range for an average 3.9 seconds for the 100-, 75-, 50-, and 25-node simulations, and the mean ranged from 13.0 and 15.2 seconds for walking scenarios. Although the mean time in range should not differ much between the different number of simulated nodes (which ran for equal simulated time), the difference arises because there are fewer data points for walking scenarios (especially for simulations with only 25 nodes).

4.1.1 Performance Metrics

We examined four performance metrics. By each metric, dual radios were more successful and efficient than single radios.

- *Time until discovery:* The time from when a node is in range and not connected to another node until it is discovered.
- *Number of successful neighbor discoveries:* The number of times a node discovers another node.
- *Transfer duration:* The time from when a node discovers a node until when it moves out of range. As noted above, we capped this value at fifteen seconds.
- *Energy cost per second of transfer:* The total amount of energy used by a node divided by the total transfer duration of that node. We cannot state the number of millijoules/byte because we measure transfer duration, not transfer size.¹

4.1.2 Performance Results

The mean time until successful discovery for all scenarios is shown in Table 3. The mean time until discovery is larger than the mean time in wireless range because unsuccessful discoveries (i.e., when the time in range is shorter) are not included. Dual radios reduce discovery time by 12–25%, depending on the scenario, compared to single radios. The standard deviation of the means are much larger than the one second difference between the means — however, this does not indicate the means are not significantly different. Nodes can take many different paths in the simulation geography, which directly affects the mean. Because we use the same scene file for tests of dual and single radio schemes, we are able to use a paired t-test to determine that difference of the means are significant; the p-values for all reported differences are less than 0.0001.

Not only is the mean time until discovery lower for the dual radio case, but the number of transfer opportunities

¹Millijoules per second is a unit of power, usually denoted milliwatts. We keep the former units to remind the reader of the focus on transfer duration.

| Nodes | Total num. of opportunities discovered | | | |
|-------|--|--------|---------|-------|
| | biking | | walking | |
| | single | dual | single | dual |
| 100 | 5,877 | 10,486 | 660 | 2,687 |
| 75 | 3,943 | 6,797 | 396 | 1,739 |
| 50 | 4,137 | 8,483 | 450 | 1,995 |
| 25 | 2,836 | 5,697 | 310 | 1,285 |

Table 4: Total opportunities discovered.

dual radios are able to discover is significantly larger. Table 4 lists the number of discovered transfer opportunities for each method across node densities. At biking speeds, the number of discoveries using dual radios increases to 170–210% of the single radios. At walking speeds, dual radios increase the number of discoveries to 410–440% of single radios.

Our results show that the mean time spent transmitting per node during the simulation with dual radios is significantly higher than single radios. Dual radios increase the mean time spent transmitting to 240–250% of single radios in the biking case (Figure 4 (top)) and by 220–230% in the walking case (Figure 5 (top)). This performance increase is due mainly to two factors. First, because neighbors are discovered sooner on average, the time they spend transmitting is longer. Second, because they have two radios, they can transmit up to twice the data. Bluetooth is a frequency hopping protocol, and so we would not expect interference among the two pairs of radios. Informal experiments between two pairs of Bluetooth devices showed no significant drop in throughput. To be conservative, we assumed only 1.75-times the throughput of a single radio in our results. Again, paired t-tests show means are significantly different with p-values less than 0.0001.

The final performance question to ask is if the energy spent on discovery and transfer is less per second of data transfer in the dual or single radio case. To calculate energy costs, we used the power specifications of the low-power BlueCore3-ROM single chip Bluetooth radio [4]. The chip requires a 1.8V power supply. The idle state draws 0.36 milliamperes (mA), inquiry draws 0.66mA, and data transfer draws 11.66mA. Let r be the time a node was powered, i be the time a radio was idle, q be the time a node was inquiring, and t be the time a node was transferring data. For the single radio case, $r = i + q + t$, and the energy usage (power consumption) of the single radio case is:

$$E_s = (0.36\text{mA})(1.8\text{V})i + (0.66\text{mA})(1.8\text{V})q + (11.66\text{mA})(1.8\text{V})t$$

For the dual radios, one radio is idle except when transmitting. The second radio is the same as the single radio. Therefore, the energy used by the dual radios is:

$$E_d = (0.36\text{mA})(1.8\text{V})(r - t) + (0.36\text{mA})(1.8\text{V})i + (0.66\text{mA})(1.8\text{V})q + (11.66\text{mA})(1.8\text{V})(2t)$$

Figures 4 (bottom) and 5 (bottom) show energy costs (in milliwatts per second) of transferred data for biking and walking speeds, respectively. For dual radios, we conservatively assume that operating two radios incurs twice the energy cost but provides only 1.75 times the bandwidth over a single radio. The figures show that dual radios use less en-

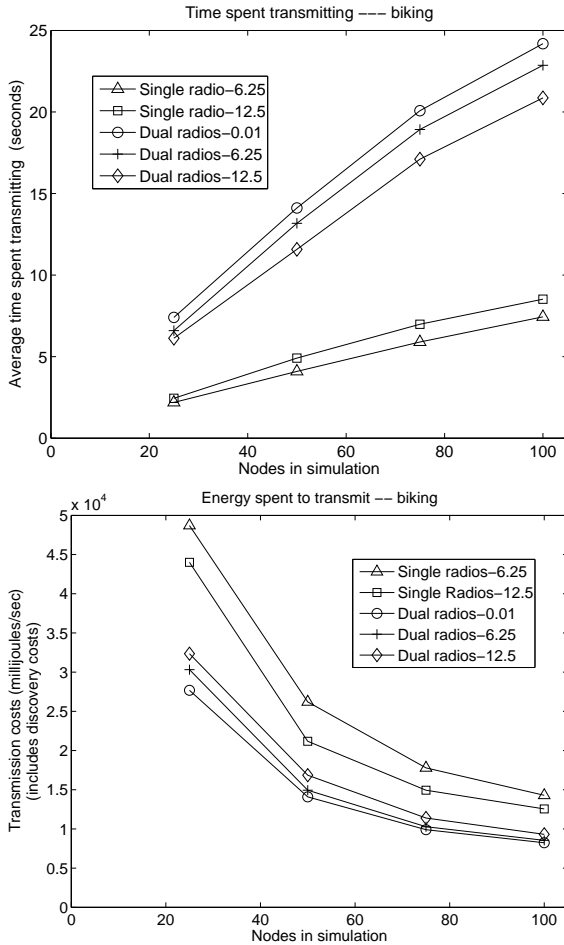


Figure 4: Biking speed: (top) Mean time spent transmitting per node. (bottom) Mean millijoules spent per second of transmission time, including costs for discovery and transmission. The legend states the mean inter-inquiry time in seconds for each line.

ergy even though they achieve a higher average time transmitting data. In the biking scenario, dual radios reduce energy used per second of transfer time to 19–27% of single radios, and to 7–15% in the walking scenario.

4.1.3 Energy Efficiency

There are diminishing returns for the dual radio scheme as the expected time between transfer opportunities increases. Figures 4 (bottom) and 5 (bottom) show that as the population increases in density, a node will more frequently come into wireless range of another peer, and powering two radios gains in efficiency. However, for sparser populations, the dual radio approach is still more efficient than that single radio approach in our experiments.

It is interesting to note that the difference in time spent transferring between the dual and single radio cases keeps widening, while the difference in efficiency narrows. This can be explained by examining the inter-opportunity time in each scenario, shown in Figure 6. As the mean time between transfer opportunities decreases, the chance of a suc-

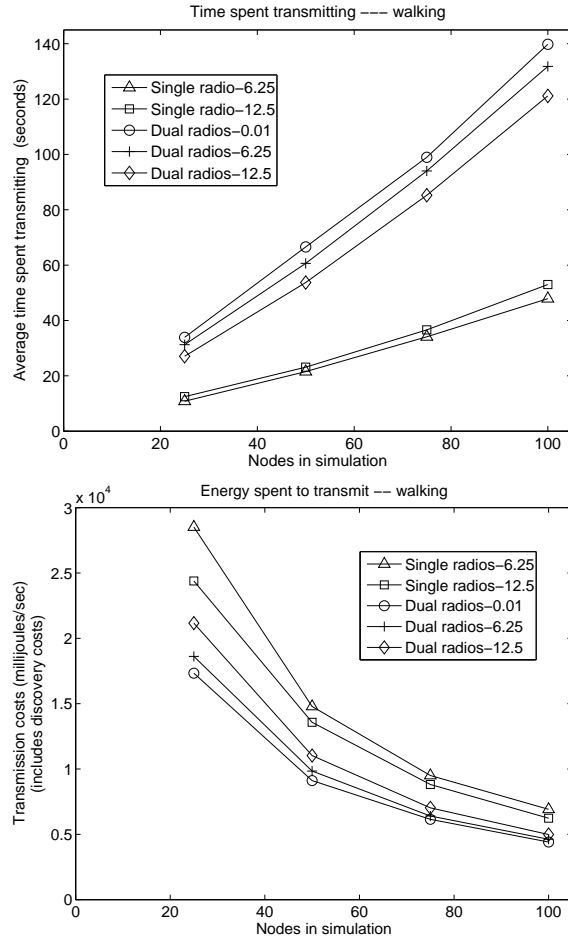


Figure 5: Walking speed: (top) Mean time spent transmitting per node. (bottom) Mean millijoules spent per second of transmission time, including costs for discovery and transmission.

cessful discovery increases since there is likely to be a node present. Eventually, even the single radio case is efficient — but the dual radio case is able to take greater advantage of the increased number of transfer opportunities.

5. CONCLUSION AND OPEN PROBLEMS

In this paper, we have shown that supporting a DTN with off-the-shelf Bluetooth radios is a challenge. As an incrementally deployable and simple solution, we have proposed a simple modification to Bluetooth-enabled devices, the addition of a second radio. We found that this dual radio approach and the full-duplex discovery that it enables allows for longer and more frequent transfer opportunities in a simulated DTN. Remarkably, this dual radio strategy is also more energy efficient than the single radio strategy in the scenarios we evaluated. Our main contribution is to show how DTN designers can leverage the millions of consumer Bluetooth devices currently deployed with minimal and inexpensive hardware.

Several areas for further work follow naturally from what we have presented. A real-world implementation of the dual radios would test the efficacy of the design in practice. Examination of mixed dual and single radio scenarios would al-




Figure 6: Mean time between transfer opportunities for biking and walking speeds.

low us to quantify the benefits, both individual and network-wide, of incremental deployment. A study that investigates link layer protocols and performance would give better estimates of energy usage per byte rather than per time, and knowing the routing characteristics of DTNs could result in further refinements to the dual radio design.

6. ACKNOWLEDGEMENTS

We thank Christophe Diot for his encouragement and for discussion of the ideas that led to this paper.

7. REFERENCES

- [1] 802.11 Specification. Available at: <http://grouper.ieee.org/groups/802/11/>.
- [2] P. Bahl, A. Adya, J. Padhye, and A. Wolman. Reconsidering wireless systems with multiple radios. *ACM SIGCOMM Computer Communications Review*, 34(5), October 2004.
- [3] S. Basagni, R. Bruno, G. Mambrini, and C. Petrioli. Comparative performance evaluation of scatternet formation protocols for networks of Bluetooth devices. *ACM/Kluwer Wireless Networks*, 10(2):197–213, March 2004.
- [4] BlueCore-3 ROM. Documentation available at: <http://www.csr.com/products/bc3range.htm>.
- [5] Bluetooth 1.2 Specification. Available at: <http://www.bluetooth.org/>.
- [6] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. In *Proc. IEEE INFOCOM*, April 2006.
- [7] B. Burns, O. Brock, and B. N. Levine. MV routing and capacity building in disruption tolerant networks. In *Proc. IEEE INFOCOM*, pages 398–408, March 2005.
- [8] A. Busboom, I. Herwono, M. Schuba, and G. Zavagli. Unambiguous Device Identification and Fast Connection Setup in Bluetooth. In *Proceedings of European Wireless Conference (EW)*, Feb 2002.
- [9] 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 *Proc. IEEE INFOCOM*, Apr 2006.
- [10] D. Chakraborty, G. Chakraborty, S. Naik, and N. Shiratori. Discovery and delay analysis of bluetooth devices. In *International Conference on Mobile Data Management (MDM)*, pages 114–114, 2006 2006.
- [11] J. Davis, A. Fagg, and B. N. Levine. Wearable Computers and Packet Transport Mechanisms in Highly Partitioned Ad hoc Networks. In *Proc. IEEE Intl. Symp on Wearable Computers (ISWC)*, pages 141–148, October 2001.
- [12] C. Diot (coordinator). The Huggle Project. <http://www.huggleproject.org/>.
- [13] R. Dodes. When a stranger calls, from afar or nearby. *New York Times*, March 2005.
- [14] M. DufLOT, M. Kwiatkowska, G. Norman, and D. Parker. A formal analysis of Bluetooth device discovery. In *International Symposium on Leveraging Applications of Formal Methods (ISOLA)*, 2004.
- [15] FCC Rules and Regulations, Part 15. Available at: <http://www.fcc.gov/oet/info/rules/>.
- [16] M. Grossglauser and D. Tse. Mobility increases the capacity of ad hoc wireless networks. In *Proc. IEEE INFOCOM*, April 2001.
- [17] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot. Pocket switched networks and human mobility in conference environments. In *Proc. ACM Workshop on Delay-Tolerant Networking*, pages 244–251, 2005.
- [18] S. Jain, K. Fall, and R. Patra. Routing in a Delay Tolerant Network. In *ACM SIGCOMM*, August 2004.
- [19] J. Ott and D. Kutscher. Drive-thru Internet: IEEE 802.11b for “Automobile” Users. In *Proc. IEEE INFOCOM*, March 2004.
- [20] B. S. Peterson, R. O. Baldwin, and J. P. Kharoufeh. A specification-compatible Bluetooth inquiry simplification. In *Hawaii International Conference on System Sciences*, 2004.
- [21] T. Salonidis, P. Bhagwat, and L. Tassiulas. Proximity awareness and fast connection establishment in bluetooth. In *Proc. of the First Annual ACM Workshop on Mobile and Ad Hoc Networking and Computing*, August 2000.
- [22] N. Sarafjanovic-Djukic and M. Grossglauser. Last encounter routing under random waypoint mobility. In *IFIP-TC6 Networking Conference*, May 2004.
- [23] E. Shih, P. Bahl, and M. J. Sinclair. Wake on wireless: an event driven energy saving strategy for battery operated devices. In *Proc. ACM MobiCom*, pages 160–171, 2002.
- [24] J. Sorber, N. Banerjee, M. D. Corner, and S. Rollins. Turducken: Hierarchical power management for mobile devices. In *MobiSys*, Seattle, WA, 2005.
- [25] A. Tanenbaum. *Computer Networks*. Prentice Hall, 3rd edition, 1996.
- [26] UniFi-1 Consumer. Documentation available at: <http://www.csr.com/products/bc3range.htm>.
- [27] R. Woodings, D. Joos, T. Clifton, and C. D. Knutson. Rapid heterogeneous connection establishment: Accelerating Bluetooth inquiry using IrDA. In *Proc. IEEE Wireless Communications and Networking (WCNC)*, 2002.
- [28] W. Zhao and M. Ammar. Message ferrying: Proactive routing in highly-partitioned wireless ad hoc networks. In *IEEE Workshop on Future Trends in Distributed*

Computing Systems, May 2003.

- [29] W. Zhao, M. Ammar, and E. Zegura. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *ACM MobiHoc*, May 2004.
- [30] W. Zhao, Y. Chen, M. H. Ammar, M. Corner, B. N. Levine, and E. Zegura. Capacity Enhancement using Throwboxes in DTNs. In *Proc. IEEE Intl Conf on Mobile Ad hoc and Sensor Systems (MASS)*, Oct 2006.

APPENDIX

A. SINGLE RADIO ANALYSIS

We assume that when two devices are in range, the probability that the two devices detect each other is a function of the length of the non-overlapping inquiry. We define a function $G(x)$ as follows: if one device is transmitting a inquiry of length $x < D$ and if the other device is listening during this time, then the function $G(x)$ is the probability of detection during x . $G(D)$ is equal to 1. We also define $g(x) = \frac{dG(x)}{dx}$. $g(x) dx$ can be seen as the probability of detecting the neighbor in a short interval dx around x .

To simplify our analysis, we make two main assumptions. First, we assume that the time between two inquiries by the same device is exponentially distributed according to I with mean $1/\lambda$. Second, we assume that inquiries have fixed duration D .

Consider the following general scenario. Suppose that at time $t = 0$, one device starts inquiring with duration x . We wish to compute two functions: $B(x)$, the average time until the end of this period; and $P(x)$, the probability that this busy period results in a detection. By setting $x = D$ we have our original problem.

The length of the busy period depends on the actions of the radios. If only one device is inquiring, the length is equal to D and in this case we are sure that two devices detect each other (since $G(D) = 1$). The analysis is much more complex if two or more inquiries overlap in a busy period. We analyze this case below. Unfortunately, obtaining a closed form expression is quite difficult when more than two inquiries overlap in a busy period. A closed expression can be obtained (and the expression is rather complex) when we assume that at most two inquiries can overlap in a busy period.

1. In the first case, the other device does not inquire during time x . This happens with probability $e^{-\lambda x}$, and we have for this case

$$\begin{aligned} B(x) &= \int_0^x ug(u) \frac{du}{d+} x(1 - G(x)) \\ P(x) &= G(x) \end{aligned}$$

2. In the second case, the other device inquires during x , which happens with probability $1 - e^{-\lambda x}$. Let y be the time at which the second device inquires. Again we have two cases: with probability $G(y)$, a detection happens by time y ; or, with probability $1 - G(y)$, no detection happens at all. If there is a detection by time y , we can write:

$$\begin{aligned} B(x) &= (1/G(y)) \int_0^y ug(u) du \\ P(x) &= 1 \end{aligned}$$

If there is no detection in time y , we can write

$$\begin{aligned} B(x) &= x + B(D - (x - y)) \\ P(x) &= P(D - (x - y)) \end{aligned}$$

Summarizing each in one equation we can write:

$$\begin{aligned} B(x) &= e^{-\lambda x} \left(\int_0^x ug(u) du + (1 - G(x))x \right) \\ &+ \int_0^x \lambda e^{-\lambda y} \left(\int_0^y ug(u) du \right. \\ &\left. + (1 - G(y))(x + B(D - (x - y))) \right) dy \end{aligned} \quad (1)$$

and

$$\begin{aligned} P(x) &= e^{-\lambda x} G(x) + \int_0^x \lambda e^{-\lambda y} (G(y) \\ &+ (1 - G(y))P(D - (x - y))) dy \end{aligned} \quad (2)$$

With the functions $B(x)$ and $P(x)$ thus computed, we can set $x = D$ to solve our original problem. Let $t = 0$ be the time when the two devices are in radio range. The expected time, T , until either one detects the other depends on what occurs at time zero.

At time zero we have one of the following cases.

1. *No device inquiring*: Let $T = T_0$ in this case. The probability of this case is $(1/(\lambda D + 1))^2$. We can write:

$$\begin{aligned} T_0 &= 1/(2\lambda) + B(D) + (1 - P(D))T_0 \\ &= \frac{1}{P(D)} \left(\frac{1}{2\lambda} + B(D) \right) \end{aligned} \quad (3)$$

2. *One device is inquiring*: The probability of this case is $2\lambda D/(\lambda D + 1)^2$

$$T = \frac{1}{D} \int_0^D (B(x) + (1 - P(x))T_0) dx \quad (4)$$

3. *Both devices inquiring*: The probability of this case is $(\lambda D/(\lambda D + 1))^2$. Therefore we have

$$\begin{aligned} T &= \frac{1}{D^2} \int_0^D \int_0^D (\min(x, y) + B(\max(x, y)) \\ &\quad - \min(x, y)) + (1 - P(\max(x, y)) \\ &\quad - \min(x, y))T_0) \frac{dx}{d} dy \\ &= \frac{1}{D^2} \int_0^D \int_0^y (x + B(y - x)) \\ &\quad + (1 - P(y - x))T_0) dx dy + \\ &\quad \frac{1}{D^2} \int_0^D \int_y^D (y + B(x - y)) \\ &\quad + (1 - P(x - y))T_0) dx dy \end{aligned} \quad (5)$$

It is difficult to solve the above model due to the implicit integral. Here, we try to illustrate the results for the very simple example of a step function:

$$G_{\text{step}}(x) = \begin{cases} 0 & \text{if } x < D \\ 1 & \text{if } x = D \end{cases}$$

In other words, to correctly detect a device, we need to listen from it for a time longer greater than or equal to D . The step function is chosen for both the simplicity of the analysis and the lower bound on performance that it yields. For a given λ and D , any other function $G(x)$ will yield a T less than the one generated by this step function. In other words, $G_{\text{step}}(x)$ is in some sense a worst case: if we choose λ such that T is less than some desired value with this step function, we can be sure this will be the case for any other function $G(x)$.

Eqs. (1) and (2) give for $0 \leq x \leq D$,

$$B(x) = x + \int_0^x \lambda e^{-\lambda y} B(D - x + y) dy \quad (6)$$

and

$$P(x) = \delta_x(D) e^{-\lambda D} \quad (7)$$

For equation (6), if we differentiate with respect to x , we get for $x \in [0, D]$,

$$\frac{dB(x)}{dx} = 1 + \lambda x + \lambda(B(D - x) - B(x)) \quad (8)$$

We define the function $B(x)$ as zero for negative values of x and for values of x larger than D ; we are interested in the values of $B(x)$ between 0 and D . We apply Laplace Transform to the above differential equation (8), and let $B^*(s)$ be the Laplace transform of $B(x)$, which gives:

$$\begin{aligned} sB^*(s) + B(D)e^{-sD} &= \\ \frac{1}{s}(1 - (1 + \lambda D)e^{-sD}) &+ \\ + \frac{\lambda}{s^2}(1 - e^{-sD}) &+ \\ + \lambda(e^{-sD}B^*(-s) - B^*(s)) & \end{aligned}$$

By setting $s = 0$ and using L'Hopital's rule, one can easily prove that $B(D) = D + \lambda D^2/2$. Hence, using Eqs. (3) and (7), we get:

$$T_0 = e^{\lambda D}(1 + \lambda D)^2/(2\lambda)$$

However, to solve for T , we need the entire function $B(x)$ between 0 and D , that is, we need to solve Eqs.(4) and (5). The above implicit transformed equation can be shown to have as a solution:

$$\begin{aligned} B^*(s) &= -\frac{B(D)e^{-sD}}{s} + \\ &\frac{(1 + \lambda B(D))(e^{-sD} + 1) + \lambda D e^{-sD}}{s^2} \\ &+ \frac{\lambda(1 + \lambda D)(e^{-sD} - 1)}{s^3} \end{aligned}$$

We invert the Laplace transform, take the values of x between 0 and D , and we substitute $B(D)$ by its value provided above. This gives:

$$B(x) = (1 + \lambda D + \lambda^2 D^2/2)x - \lambda(1 + \lambda D)x^2/2$$

This equation satisfies the above differential equation.

We come now to the computation of T . We consider the above three cases: no device inquiring, one device inquiring, and both devices inquiring. For case (1), T is equal to T_0 . For case (2), we have

$$\begin{aligned} T &= \frac{1}{D} \int_0^D B(x) dx + T_0 \\ &= D\left(\frac{1}{2} + \frac{\lambda D}{3} + \frac{\lambda^2 D^2}{12}\right) + T_0 \end{aligned}$$

As for case (3), all inquiries are lost. A simple computation shows that for this case:

$$T = D\left(\frac{5}{6} + \frac{\lambda D}{4} + \frac{\lambda^2 D^2}{12}\right) + T_0$$

We sum over all the three cases to obtain a closed-form expression for T :

$$\begin{aligned} T &= \frac{D}{12} \frac{\lambda D(12 + 18\lambda D + 5(\lambda D)^2 + (\lambda D)^3)}{(1 + \lambda D)^2} \\ &+ e^{\lambda D}(1 + \lambda D)^2/(2\lambda) \end{aligned}$$

Note that the value of T is a function of the product λD and is proportional to D . Thus, to minimize D , one needs to find the minimal value of this function:

$$\frac{x(12 + 18x + 5x^2 + x^3)}{12(1 + x)^2} + e^x(1 + x)^2/(2x)$$

which occurs when $x = 0.376$. So for a given D , the optimal value of λ to use is given by $\lambda D = 0.376$. Thus $1/\lambda$ (which is the optimal average time between two inquiries) is equal to $2.66D$.

Also note that this analysis is weaker when $D \approx 1/\lambda$, and improves when $D \ll 1/\lambda$.

B. DUAL RADIO ANALYSIS

We find the expected time from when the radios move into range of one another until the first successful inquiry from either radio and add D , giving us the time until the two radios detect one another. Let Y_1 be the time at which the first device inquiries and Y_2 be the time at which the second device inquiries. We are interested in computing $T = E[\min(Y_1, Y_2)] + d$.

We compute the distribution of Y_1 , as Y_2 has the same distribution. We have two cases. Either the device is inquiring at time 0 or it is not inquiring. The first case happens with probability $\lambda D/(\lambda D + 1)$, the second with probability $1/(\lambda D + 1)$. Hence, for $y > D$,

$$\begin{aligned} \Pr\{Y_1 > y\} &= \frac{1}{\lambda D + 1} e^{-\lambda y} + \frac{\lambda D}{\lambda D + 1} \int_0^D e^{-\lambda(y-x)} dx \\ &= \frac{e^{-\lambda(y-D)}}{\lambda D + 1} \end{aligned}$$

For $y < D$, we have

$$\begin{aligned} \Pr \{Y_1 > y\} &= \frac{1}{\lambda D + 1} e^{-\lambda y} \\ &+ \frac{\lambda}{\lambda D + 1} \int_0^y e^{-\lambda(y-x)} dx + \frac{\lambda D}{\lambda D + 1} \frac{D-y}{D} \\ &= \frac{1 - \lambda(D-y)}{\lambda D + 1} \end{aligned}$$

Then,

$$\begin{aligned} \Pr \{\min(Y_1, Y_2) > y\} &= \Pr \{Y_1 > y\} \Pr \{Y_2 > y\} \\ &= \Pr \{Y_1 > y\}^2 \end{aligned}$$

and $E[X] = \int_0^\infty \Pr \{X > x\} dx$. Hence,

$$\begin{aligned} E[\min(Y_1, Y_2)] &= \int_0^D \frac{(1 - \lambda(D-y))^2}{(\lambda D + 1)^2} dy \\ &+ \int_D^\infty \frac{e^{-2\lambda(y-D)}}{(\lambda D + 1)^2} dy \\ &= \frac{3 + 6\lambda D - 6\lambda^2 D^2 + 2\lambda^3 D^3}{6\lambda(\lambda D + 1)^2} \end{aligned}$$

and $T = E[\min(Y_1, Y_2)] + D$. This is the average time until the two devices detect each other, and is much better than the single-radio result for the same inter-inquiry delay. As for the single radio case, this analysis is more exact when $D \ll 1/\lambda$.