Accelerating Bluetooth Inquiry for Personal Area Networks

Gergely V. Záruba[1], and Imrich Chlamtac[2]
[1]Center for Research in Wireless Mobility and Networking (CReWMaN)
Department of Computer Science and Engineering, The University of Texas at Arlington
Arlington, TX, U.S.A. zaruba@uta.edu
[2]Center for Advanced Telecommunications Systems and Services (CATSS)
Department of Computer Science, The University of Texas at Dallas
Richardson, TX, U.S.A. chlamtac@utdallas.edu

Abstract - **The recent industry standard Bluetooth promises low cost replacement of communication cabling with moderate symbol-rate, short-range wireless links. The same specification also addresses the establishment of point-to-multipoint piconets and the interconnection of several of these piconets into scatternets, enabling Bluetooth to be used as a technology for realizing personal area networks. Establishing Bluetooth piconets requires nodes to discover each other by completing an inquiry phase. This paper investigates the inquiry phase; shows the shortcomings of the current inquiry procedure in multi node – PAN scenarios, and outlines and analyses a backwards Bluetooth compliant modification to accelerate Bluetooth inquiry. Extensive simulations comparing the original and the proposed Bluetooth inquiry schemes show improvements of more than an order of magnitude in device discovery times.**

I. Introduction

Wireless communications gained significant acceptance and importance in the last decade of the previous century. The wireless revolution was ignited by 1st and 2nd generation cellular standards, but at the end of the last decade, dropping prices of novel wireless technologies helped to create new wireless networking paradigms such as wireless local area networking (WLAN) and short range wireless personal area networking (WPAN). It is generally predicted that this and the next decade will bring an even higher wireless penetration mainly due to WLANs and WPANs.

One of the emerging short-range wireless networking technologies is the recent industry standard Bluetooth (BT) [1]. Bluetooth evolved from the need to replace wires in short-range communication, e.g., cables between mobile handsets and their headsets, or serial cables between computers and peripherals, with short-range wireless links. Derivable from the maximum transmission power and receiver sensitivity of the specification [1], a class-3 (most pervasive) BT has a transmission range of approximately 10 meters in a free propagation environment at a nominal ISM band frequency of 2.4GHz. BT employs frequency hopping over 79 carrier frequencies[1], spaced 1MHz away with a spectral efficiency of 1bps/Hz. The main communication structure of Bluetooth – called *piconet* – is a point-to-multipoint star topology, with a *master* node in the center and *slave* nodes at the perimeter of the star. Piconets can be interconnected to *scatternets* by nodes taking on master or slave roles in more than one piconet concurrently. The above-described characteristics make BT a viable candidate for establishing inexpensive personal area networks.

Before BT devices can exchange information among each other, they have to go through a three-phase link establishment procedure. In the first phase, the devices have to scan and search through a subset of the hopping frequencies to determine whether there are other devices in their transmission range. This *first phase* serves as the *neighbor discovery* process, and is referred to as *inquiry* in BT terminology. In the subsequent step, devices that are already aware of each other's proximity, initiate a handshake process by which they exchange crucial information for piconet formation; this phase is also referred to as *paging*. The third phase deals with the setting up of a virtual channel for further control information exchange and negotiating communication parameters relating to link management issues. It is implied, by the above three-phase process, that only because devices may be in each other's transmission-range they do not necessarily have the means to communicate with each other (unlike with Wi-Fi).

This paper concentrates on the first – inquiry – phase pointing out the limitations of the mandatory inquiry process while proposing and describing a technique to accelerate device discovery. The rest of this section describes the mandatory inquiry process of BT in more detail, outlines previous work, and introduces our approach for accelerating the inquiry process.

A. The Bluetooth Inquiry Process

The purpose of the BT inquiry process is for devices in each other's transmission range to become aware of this proximity, i.e., a successful inquiry "handshake" between two nodes results in the initiating node acquiring knowledge on the responding node's identity and clock offset. Such

[1] This paper is restricted to the 79-hop system, yet the proposed approach can easily be applied to the 23-hop system as well.

 0-7803-7974-8/03/$17.00

handshakes require some of the nodes to be in an inquiry (initiating) state while other nodes to be in an inquiry scan (responding) state.

In the inquiry state, devices transmit very short (68µs) ID packets every 312.5µs. The short duration and unique bit pattern of these ID packets does not only enable a receiving node to efficiently correlate its receiver to the ID packet but makes the division of regular 625µs slots into two 312.5µs "half-slots" possible. Consequently, in even numbered slots the inquiring node will send out two ID packets at different frequencies, while in an odd numbered slot, the same node will tune its receiver to the corresponding response carriers. The number of frequency carriers used for inquiring is reduced from 79 to 32 inquiry and 32 inquiry response carriers for increased discovery efficiency. These 32 carriers are divided into two 16-carrier trains, i.e., A- and the B-train. According to the mandatory inquiry scheme a single train has to be repeated for at least 256 times before changing to the other train.

In the inquiry scan state a node is listening for at least 2048 slots (1.28s) on one of the 32 different inquiry-scan frequencies waiting to overhear an ID packet from an inquiring node. If an ID packet is overheard, the node generates a uniform random number b from [0,1023] and suspends the inquiry process for a duration of b slots to reduce the chances of colliding inquiry responses from different nodes listening on the same frequency. Once this timeout has expired, the device reenters the inquiry scan state and responds to the first ID packet it overhears. The response packet is a so-called FHS (frequency hopping sequence) packet containing the ID and clock values of the responding node. The inquiring node can now make a note of the clock and ID value of the responding node and can either continue the inquiry process or initiate the paging process (or assume its original state).

It has been shown [1] that proper inquiry and inquiry scan state holding times ensure that two devices discover each other in less than 10 seconds, yet that does not imply any performance metrics for cases where there are more devices competing. In a PAN scenario, where all nodes are in each other's transmission range, we can identify five problems slowing down the inquiry process. The first problem corresponds to the fact that each time an ID packet is overheard the nodes spend a random time drawn from a *static interval* in a back-logged state. The second problem corresponds to the infrequent change of inquiry trains, which reduces the probability that an inquiring node is transmitting on a frequency a scanning node is listening to. The less obvious problems include: i) the possible case where by the time a node returns to the scan mode after backing-off, the corresponding inquiring node has changed its state and is not transmitting ID packets anymore; ii) when a node returns from the back-off to the inquiry scan state it will reply to the first ID packet it overhears, which may be transmitted by a different device than earlier, and iii) since the ID packet is a unique bit pattern, it is impossible for devices to detect who the originator of the transmitted packet was, and thus it is possible that a scanning node will reply to the same inquiring node again and again.

B. Previous Work

Previous research work on Bluetooth has mainly focused on distributed scatternet formation in PAN and ad hoc environments (e.g. [2,4] respectively), on scheduling policies of master nodes in pico- and scatternets, and on performance and interference measurement in the presence of other piconets or interfering electro-magnetic forces (e.g., Wi-Fi networks, or microwaves). Some scatternet formation approaches assume [4], that device discovery has already been taken place, thus all nodes are aware of the identities of all neighboring devices, further motivating our work.

The authors in [3] derived strategies based on different state holding probability distributions for the case when there is one inquiring and one scanning node present assuming that inquiry trains change after each train. At the time of this writing, the authors are not aware of any other work published addressing BT device discovery.

C. Our Approach

In the following sections we describe a novel approach to reduce device discovery times based on measuring the number of idle slots during inquiry scans and estimating the number of contending nodes by this measurement. The number of contending nodes is important metric for making the random backoff selection adaptive, thus reducing the number of "wasted" slots/time during inquiry.

The rest of the paper is organized as follows: in the next section we present our accelerated service discovery technique using a top-down approach; in the succeeding section we will show using simulations, what effect the proposed measurement's resolution has on the estimation of backoff and what performance improvement can be expected from our scheme. Finally we draw our conclusions and identify future work. An early, "research in progress" version of this paper has appeared in [5] identifying the problem and outlining a solution but not showing performance evaluation results on the proposed acceleration technique.

II. ACCELERATING INQUIRY

Looking at the problems with the BT inquiry procedure, we imply that its inefficiency comes mainly from the static maximum backoff period B that is set to 1023. In order to accelerate the inquiry process, B should reflect the number of nodes actually contending for the attention of the inquiring node. In this section we are going to show how this number can be made adaptive and how it should be calculated to achieve lower device discovery times. We will consider typical WPAN scenarios where all nodes are in each other's proximity; we will assume that no other interfering sources are present.

A. Calculating the Best Backoff Value - B

Let us assume that we already possess the knowledge on the number of nodes that are in inquiry scan mode listening on the same frequency, and denote this population by n_{sf}. All these nodes are going to overhear an inquiring node transmitting on the corresponding carrier frequency, thus all of them are going to generate a uniform random number from [0,B]. In order to minimize the collisions while reducing the number of wasted replying opportunities, it can be shown [5] that B should be set to n_{sf}.

B. Number of Nodes Scanning on the Same Frequency

In the acceleration scheme we will change the train sequence after every train, thus inquiring not only on 16 but on all $I_s=32$ designated inquiry frequencies in a row. Nodes in the inquiry scan mode can listen on any of these I_s frequencies, the exact frequency determined only by their native clock. Thus the probability $P_{is}(k)$, that there are exactly k nodes listening on the same frequency, assuming the knowledge on the number of nodes n_s in the inquiry scan mode, can be derived from an $(ns, 1/I_s)$ parameter binomial distribution:

$$P_{is}(k) = \binom{n_s}{k}\left(\frac{1}{I_s}\right)^k\left(1-\frac{1}{I_s}\right)^{n_s-k}, \text{ where } I_s = 32$$

The most likely value of an (n,p)-parameter binomial distribution is $\lfloor (n+1)p \rfloor$. Thus, a good approximation for the number of listening nodes and the maximum backoff B is:

$$B = n_{sf} = \left\lfloor \frac{n_s+1}{I_s} \right\rfloor$$

C. Determining the Number of Nodes in Inquiry Scan State

Let us denote the number of nodes in the inquiry state by n_i, and assume that all devices work according to the same inquiry strategy and assume that by knowing this strategy a probability distribution function $P_{IQ}(s)$ can be derived that corresponds to the probability that $s=n_s/n_i$. Let us also assume that $P_{IQ}(s)$ has a finite expected value $E(P_{IQ}(s))=Q$. If nodes in the inquiry scan mode can determine or estimate the number of nodes in the inquiry mode - n_i, then they can estimate the number of nodes in the inquiry state mode they are contending with by: $n_s=Q*n_i$, thus:

$$B = \left\lfloor \frac{Q*n_i+1}{I_s} \right\rfloor$$

D. Determining the Number of Nodes in the Inquiry State

Let us make the assumption that nodes change from inquiry to inquiry scan state and vice versa less frequently than they change between trains, i.e., a change in number of nodes in the inquiry state while two consecutive inquiry trains take place is negligible (i.e., the number of nodes performing a state change is small in a $t_{TR}=625\mu s*32=20ms$ period). Appropriate mean values for inquiry and inquiry scan durations can validate this assumption. If nodes that have just started the inquiry scan operation spend the first 32 slots by listening at the channel only, and assuming that there are no collisions among inquiry transmissions of inquiring nodes, then they could count the number of ID packets received, which would determine n_i. Yet, this latest assumption is likely not to hold, thus a good estimation on n_i allowing collision should be sought. To be able to reduce the complexity of this problem, let us assume that all nodes are synchronized, i.e., there is no drift between clocks and clock ticks are synchronized. Later on we will show how this assumption may be relaxed. If a node is listening for t_{TR} duration on a given frequency, then bit-zero of its clock will change $s=64$ times. Devices in the inquiry mode will start their transmission exactly on these "half-slot" boundaries. The listening node will be able to tell in how many of these s slots there was radio energy present on the channel, thus be able to determine how many of these s slots were idle slots (denoted by s_i). The question is: can the number of inquiring nodes n_i be estimated by knowing s and measuring s_i? For the synchronized case a relatively simple estimation can be given: let us reverse the problem and assume knowledge on n_i and s thus attempting to determine s_i. The probability ($P_e(k)$) that a slot was used for exactly k transmissions can be determined by a (n_i, $1/s$) parameter binomial distribution:

$$P_e(k) = \binom{n_i}{k}\left(\frac{1}{s}\right)^k\left(1-\frac{1}{s}\right)^{n_i-k}$$

Thus the probability that a randomly selected slot is empty is $P_e(0)=(1-1/s)^{ni}$. The probability $P_s(i)$ that exactly i slots are idle can be modeled again by a binomial distribution with parameters: $(s,P_e(0))$:

$$P_s(i) = \binom{s}{i}\left(1-\frac{1}{s}\right)^{i*n_i}\left(1-\left(1-\frac{1}{s}\right)^{n_i}\right)^{s-i}$$

The expected value of this function is the product of its parameters: $E(P_s(i))=s*(1-1/s)^{ni}$. Thus it is expected that if n_i nodes have to contend for s slots, there will be $s*(1-1/s)^{ni}$ slots empty, providing with a good estimation for s_i. Now let us solve the above equation for n_i, thus calculating a good estimate on the number of inquiring nodes employing channel measurements with synchronized clocks:

$$n_i = \ln(s_i/s)/\ln\left(1-\frac{1}{s}\right)$$

The synchronization assumption can be relaxed by enhancing nodes to measure the time t_E in which energy is present on the channel. Depending on the sophistication of the node's hardware, the resolution r_e of the measurement of t_E can be from a couple of microseconds up to the duration of an ID packet $t_{ID}=68\mu s$. We argue, that a measurement resolution of $1\mu s<r_e<10\mu s$ can be easily achieved without generating a major cost increase of Bluetooth chips. In this case nodes measure energy on the channel during s slots with r_e resolution, i.e., in t_{TR}/r_e micro-slots. If the radio energy on the channel during a micro-slot does not go above a given threshold, then s_i will be incremented. The selection of the resolution time r_e has a significant impact on the accuracy of n_i's estimation as we are going to see in the next section. The estimated value of n_i with r_e resolution of measurements and s_m idle micro-slots can be given as:

$$n_i = \frac{\ln(s_m * r_e / t_{TR})}{\ln(1 - r_e / t_{TR})} * \frac{r_e}{t_{ID}}$$

Now the optimal backoff B can be determined by:

$$B = \left\lfloor \frac{Q * \frac{\ln(s_m * r_e / t_{TR})}{\ln(1 - r_e / t_{TR})} * \frac{r_e}{t_{ID}} + 1}{I_s} \right\rfloor$$

Leaving Q as the only variable to be evaluated.

E. Proportion of Nodes in the Inquiry Scan State

Obviously, the inquiry policy will have a significant impact on the proportion of nodes in the inquiry state versus nodes in the inquiry scan state. Thus the value of Q will depend on this policy. Here we will calculate Q for a probabilistic inquiry scheme relying on more or less uniform distribution of state holding times. In [3] it has been shown that deterministic state holding times will result in device discovery times with an infinite expected value.

Let us determine the value of Q for a simple state holding time distribution. It is necessary for the previous assumptions to hold, that each node has a lower bound on the number of slots it spends in each of the inquiry states. Obviously, the more time they spend in the individual states at a time, the more likely that they discover each other. On the other hand, since the device discovery procedure is asymmetric (i.e., node i discovering node j, will not provide node j with knowledge on node i's identity), nodes have to change states frequently. Additionally, the only way to avoid continuously overlapping transmissions of different inquiring nodes at the same frequency (of which the likelihood is growing with the population) is to make sure they eventually and independently switch states. The determination of the best state holding times and distributions is beyond the scope of this paper.

In our simple probabilistic model each node is required to spend k_0 time slots in inquiry state and then an additional k_1 slots that is randomly chosen from the interval [*0, 2K*]. Furthermore every node has to assume the inquiry scan state for l_0 time slots and then an additional l_1 slots that is randomly chosen from the interval [*0, 2L*]. Another simplified way of describing this model is that after spending k_0, and l_0 time slots in the appropriate states, each node is changing to the other state with a probability of K and L into the inquiry scan/inquiry states respectively. Thus in average a node spends k_0+K-1 slots in the inquiry and l_0+L-1 slots in the inquiry scan state. Consequently the estimated value for the ratio of the state holding times can be given as: $Q=(l_0+L-1)/(k_0+K-1)$. Therefore the optimal backoff value can be estimated by:

$$B = \left\lfloor \frac{\frac{l_0 + L - 1}{k_0 + K - 1} * \frac{\ln(s_m * r_e / t_{TR})}{\ln(1 - r_e / t_{TR})} * \frac{r_e}{t_{ID}} + 1}{I_s} \right\rfloor \quad [1]$$

Note, that if all nodes work according to the described strategy, then all variables are known in Equation 1 after the first time a device spends t_{TR} time in the inquiry scan state.

III. Simulation Results

Simulation efforts to the evaluation of the proposed scheme are two fold: i) determination of the accuracy of the prediction of the inquiry state population of nodes with different slot measurement resolutions; ii) evaluation of the performance gain in the inquiry acceleration scheme compared to the original BT inquiry scheme. The first simulation effort was carried out using Matlab scripts, while for the second we have implemented a C++ based discrete event simulation, modeling the inquiry behavior of BT devices. For both simulation approaches we have evaluated the results' statistics on the fly, thus executing an adequate number of simulation instances to claim that the confidence intervals of our simulations are 95-5, i.e., we are 95% sure, that our results have less than 5% error.

A. Accuracy of Prediction

Intuitively, the higher the resolution of the measurement is (i.e., the smaller t_r is), the better the estimation of the population of inquiring nodes (n_i) will be. To study the effect of the resolution on n_i and to evaluate the estimation given by Equation 1, we have simulated N number of nodes accessing a channel randomly during t_{TR}=*20ms* duration. Figures 1 and 2 draw the real and estimated number of nodes in function of the empty micro-slots: s_m with the dotted lines representing simulation, while solid lines representing calculated results (for more results and discussion consult [5]).

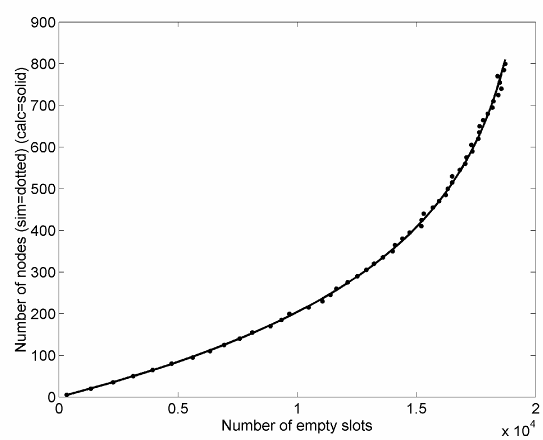

Figure 1. Simulated and calculated population - 1µs resolution

Figure 1 depicts the case where t_r was set to 1µs, a value significantly less than the duration of an ID packet t_{ID}, the figure clearly shows that Equation 1 provides a good estimation in this case. Figure 2 shows the results for t_r=10µs concluding that the estimated value does not significantly differ from the real population even for a measurement resolution an order of magnitude higher than before. Figure 1 and 2 imply that the measurement method will limit the population to below two thousand nodes.

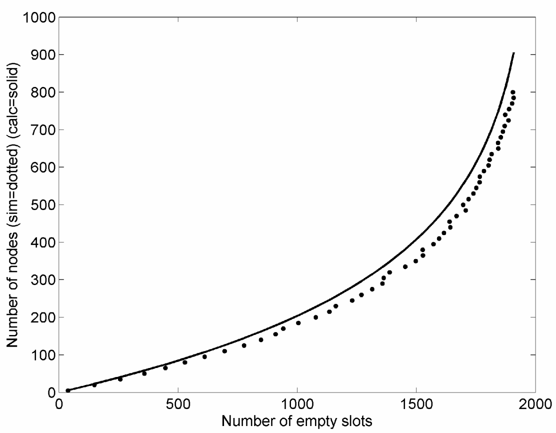

Figure 2. Simulated and calculated population - 10µs resolution

B. *Improvement to the Original BT Inquiry Scheme*

Figures 3 and 4 show our results of device discovery times using the original BT scheme (dotted line), our accelerated scheme with 1µs resolution (solid line) and a third curve representing a scenario where a train sequence change is enforced after each train with the original BT inquiry. Figure 3 depicts the average time between two consecutive non-redundant discoveries in function of the population, while Figure 4 provides insights in the total time required for all peers to discover each other.

It can be observed that more than an order of magnitude of improvement in the device discovery times can be achieved using our scheme compared to the original BT scheme, and that the improvement is significantly due to the adaptive back-off not only to the immediate train changes.

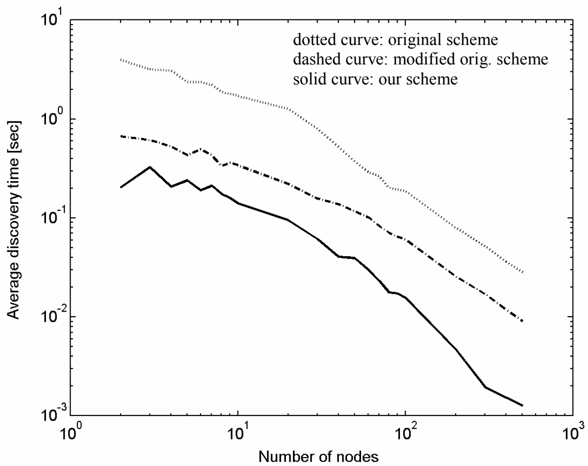

Figure 3. Average inter-discovery times of devices.

IV. Conclusions and Future Work

In this paper we presented an approach to reduce the device discovery times in Bluetooth PANs, where all nodes are assumed to be in each other's proximity. Our technique is based on the observation that in the original inquiry specification, the backoff value does not reflect the number of contending nodes on the channel, but is set to a static high value. We have provided with a top-down analytical approach to make the backoff value adaptive requiring only a small modification to the nodes: enabling them to measure the RF channel. We have demonstrated with simulations that a good approximation on the number of contending nodes can be achieved with appropriate measurement resolutions. Furthermore, we have shown by simulations that the accelerated inquiry approach outperforms the original BT inquiry scheme significantly in means of discovery times,.

Future work includes i) considering channel changes of listening nodes, where they have to keep track of devices they already have replied to (another situation where measurements can provide with reduced device discovery times); ii) evaluation of the impact of symmetric discovery, where nodes completing an inquiry handshake in turn initiate a connection set-up, has on device discovery times.

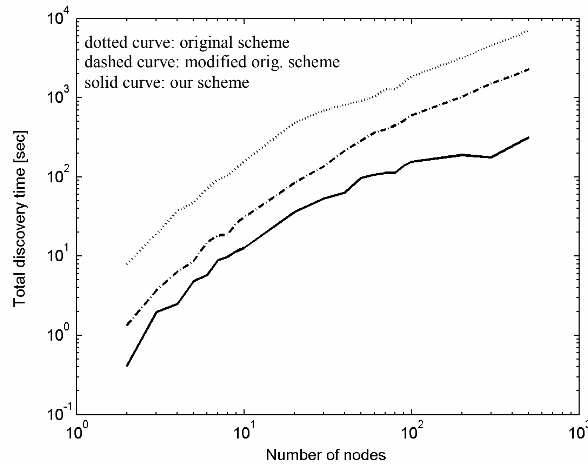

Figure 4. Total discovery time for all peers.

References

[1] Bluetooth Special Interest Group, "*Specification of the Bluetooth System, Core Version 1.1*", Specification Volume 1, http://www.bluetooth.com, February 2001.

[2] T. Salonidis, P. Bhagwat, L. Tassiulas, and R. LaMaire, "*Distributed Topology Construction of Bluetooth Personal Area Networks*," Proc. IEEE INFOCOM 2001, pp.1577-1586, Anchorage, AK, April 2001.

[3] T. Salonidis, P. Bhagwat, L. Tassiulas, and R. LaMaire, "*Proximity Awareness and Ad Hoc Network Establishment in Bluetooth,*" Institute of Systems Research, Univ. of Maryland Tech. Report, 2001.

[4] G.V. Záruba, S. Basagni, and I. Chlamtac, "*Bluetrees-Scatternet Formation to Enable Bluetooth-based Ad Hoc Networks*," Proceedings of the IEEE International Conference on Communications, ICC 2001, vol. 1, pp.273-277, Helsinki, Finland, June 2001.

[5] G.V. Záruba, "Neighbor Discovery in Bluetooth Based Personal Area Networks," IASTED PDTPA 2002 Workshop, June 2002.