

A Study of Bluetooth Propagation Using Accurate Indoor Location Mapping

Anil Madhavapeddy and Alastair Tse

University of Cambridge Computer Laboratory, Cambridge, CB3 0FD, UK
avsm2@c1.cam.ac.uk, acnt2@c1.cam.ac.uk

Abstract. The ubiquitous computing community has widely researched the use of 802.11 for the purpose of location inference. Meanwhile, Bluetooth is increasingly widely deployed due to its low power consumption and cost. This paper describes a study of Bluetooth radio propagation using an accurate indoor location system to conduct fine-grained signal strength surveys. We discuss practical problems and requirements encountered setting up the infrastructure using the ultrasonic Active Bat indoor location system, and limitations of the commodity Bluetooth devices used. We conclude that Bluetooth is poorly suited to the purpose of fine-grained, low latency location inference due to specification and hardware limitations, and note that the movement speed of mobile devices is an important factor in calculating available bandwidth. We publish our data sets of signal strength samples for the community to freely use in future research.

1 Introduction

A number of recent projects have sought to use wireless protocols as accurate location estimators for mobile users. WiFi has proven to be an especially popular candidate for this purpose, as research projects such as RADAR [2] showed its viability, and commercial implementations of the technology have begun to appear in the marketplace. Place Lab [19] has even more ambitious goals by seeking to create a comprehensive location database which uses fixed commodity WiFi, GSM and Bluetooth devices as global beacons.

In this paper, we describe a systematic methodology for the evaluation of wireless signal-strength propagation. We use the accurate Active Bat indoor location system as a “location oracle” which allows us to: (*i*) perform sweeping surveys to measure signal strengths across a building-wide area; (*ii*) algorithms to correct errors from the location system present in those traces; and (*iii*) visualize and analyze the data to deduce location estimation properties of the wireless protocol. This paper also contributes guidelines which apply to other ultrasound- or radio-based location systems being used as survey tools (e.g. Cricket [18]).

In particular, we focus on the indoor location properties of the Bluetooth wireless protocol. Bluetooth has steadily gained popularity in many Ubicomp projects due to its emphasis on short-range, low-power, and ease of integration

into devices. It is most commonly used as a “cable replacement protocol” to perform connection establishment between devices without requiring physical contact between them. In the wireless location estimation space, Bluetooth is especially important due to its ubiquitous and “always-on” presence in commodity everyday devices such as mobile phones and PDAs. This is in contrast to the more power-hungry WiFi, which is generally only switched on in stationary devices (e.g. a laptop is rarely used when a person is walking around).

Although some Ubicomp projects such as Place Lab have begun to use Bluetooth as a coarse-grained “on or off” indicator in their location databases, we contribute the first systematic and detailed analysis of more accurate Bluetooth location sensing by using our survey technique. In addition to sharing our analysis, we announce the availability of our data-sets of signal-strength/location samples [14] which will be of use to future projects seeking to exploit Bluetooth location information.

Previous attempts to accurately measure location with commodity Bluetooth hardware have been difficult because: (*i*) unlike WiFi, measuring Bluetooth signal strength requires the establishment of an active Bluetooth connection; (*ii*) many common Bluetooth chipsets, especially those found in mobile phones, only support a single Bluetooth connection at a time which makes triangulation difficult; and (*iii*) Bluetooth devices use frequency hopping algorithms which make location inference more difficult (see Section 4.3).

The rest of the paper is structured as follows. Section 2 begins by introducing the basics of the Bluetooth protocol and the Active Bat indoor location system. Section 3 describes our experimental setup and methodology, and Section 4 presents our results and discussion. Finally, we cover related work in Section 5 and conclude in Section 6.

2 Background

Section 2.1 introduces some of the concepts and background for the Bluetooth wireless protocol. Section 2.2 then describes the Active Bat ultrasonic indoor location system which provided the accurate location information needed for the experiments we conducted.

2.1 Bluetooth

Bluetooth is a short-range, wireless, cable-replacement protocol operating in the license-free 2.4GHz spectrum. Unlike WiFi, which offers higher transfer rates and distance, Bluetooth is characterised by its low power requirements and low-cost transceiver chips. Over 69 million Bluetooth ICs shipped in 2003 [12] in mobile devices such as cellphones, PDAs and laptop computers, providing a ubiquitous mechanism for wireless transfer of relatively small amounts of data.

In order to ensure robustness in noisy environments, Bluetooth divides the band into 79 channels and frequency hops across them up to 1600 times per second. A connection between two devices is time-division multiplexed with $625\mu\text{s}$

time slots, coordinated by one device designated as the “master”. A hopping code is used to frequency hop between the channels for each time slot. The Bluetooth baseband defines two types of links: (i) the Asynchronous Connection-Less (ACL) link; and (ii) the Synchronous Connection-Oriented (SCO) link. The SCO link offers isochronous communication via reserved time slots and is primarily used by voice traffic. ACL links offer reliable communication via packet retransmission, and are used to build the control and data transfer protocols.

Although the Bluetooth protocol allows a master to support up to 3 simultaneous SCO links and 7 simultaneous ACL links, typical consumer hardware only supports a single SCO link and 3 ACL links. Very lightweight devices such as the current generation of Bluetooth-enabled mobile phones only support a single SCO and ACL link at a time (which is sufficient to drive a wireless headset).

2.2 Active Bat System

The Active Bat system [21,10] is a sophisticated indoor location system in which small active devices (“Bats”) periodically emit narrowband ultrasound pulses. This ultrasound is detected by multiple sensors in the ceiling which use the time-of-flight information to multi-laterate the position of the Bats—a process which is accurate to 3 cm 95% of the time.

The Bat ultrasonic pulses are scheduled via a radio channel, and can trigger up to 50 location updates per second per radio zone; there are two such radio zones in the current installation. The system adaptively schedules Bat updates and offers highly-mobile Bats increased priority by scheduling them more often. Applications can request a higher update rate from the scheduler for a set of Bats for limited periods of time.

The Active Bat system is designed to locate tags inside a room to a high degree of accuracy, but is susceptible to multipath effects and reflections arriving at the receivers. The receivers are carefully surveyed and know their positions in the building to sub-centimetre accuracy. The system processes the raw Bat updates and rejects incorrect ultrasound responses by using a multi-lateration algorithm which requires at least three receivers to agree on a triangulation.

The SPIRIT spatial indexing middleware [1] provides applications with an easy CORBA interface to get notifications when location events occur, such as a person entering or leaving a room. SPIRIT provides feeds of various granularities; for the purposes of our measurements, we use the raw location feed which eliminates readings that fail to pass the multi-lateration algorithm. In addition to location events, SPIRIT also has a record of the *world model* of the building and categorises them by type (such as desks, walls, windows, and partitions).

The Active Bat system provides higher update rates, lower latency and greater levels of accuracy than are likely to be found in any deployed commercial indoor location system in the near future. It provides an excellent research platform for conducting fine-grained measurement surveys to better understand the problem of real-world wireless radio propagation.

3 Methodology

We aim to characterise the properties of Bluetooth by gathering a large number of signal strength measurements across a typical office environment. An Active Bat is attached to a Bluetooth device enabling us to gather: (i) the link quality readings from the device; (ii) the device’s precise location from the Bat; and (iii) the time of measurement. This data collection method allowed us to cover nearly all 2D positions in every room except for those with obvious obstructions.

The experiments described analyse the signal strengths in a typical office building which consists of locations that are both obstructed from the transmitter and in line-of-sight. By varying the speed of movement during collection, we also gain insight into the link quality under dynamic conditions.

Section 3.1 defines our use of the term *signal strength*. We then detail our experimental setup in Section 3.2, and describe the post-processing performed on the measurements in Section 3.3.

3.1 Bluetooth Link Quality

The term *signal strength* is somewhat ambiguous in the Bluetooth specification. The Bluetooth Core Specification [3] mandates that vendors provide access to the link quality of established connections through the Host Controller Interface (HCI). The HCI interfaces provides access to two values: (i) the Link Quality (LQ); and (ii) the Received Signal Strength Indication (RSSI). Additionally, if a Bluetooth link is established, the remote device’s perceived RSSI can also be obtained via the “AT+CSQ” command [17] through the Serial Port Profile.

An LQ query returns an 8-bit unsigned integer that quantifies the perceived link quality at the receiver, but the exact coding of that integer is allowed to be device specific. For most Bluetooth chipsets, this number is a proportion of the Bit Error Rate (BER) recorded by the hardware. The reported RSSI value is the difference between the real RSSI in dB and the optimal receive power range, also known as Golden Receive Power Range. RSSI values reported through the HCI are also device dependent, but generally of considerably lower resolution and accuracy than the LQ.

The reported RSSI accuracy level in the Bluetooth specifications only requires a device to report the correct sign of the difference against the Golden Receive Power Range. Cambridge Silicon Radio (CSR) [4] chipsets make no guarantees about the accuracy of the magnitude. Therefore the only source of signal strength information we can rely on for Bluetooth devices is the link quality.

For our experiments we chose to use the popular USB Bluetooth adaptors from CSR as our base stations. The choice of vendor was an important one, as the resolution of link quality data available from the chip varies significantly across different manufacturers. For example, the Broadcom Bluetooth adaptors only offered updated measurements every 5 seconds, with little variation in the actual values reported. The LQ value (Q_l) ranges from $0 \leq Q_l \leq 255$ and is updated once per second by the CSR chipset. Equation 1 describes how to convert Q_l to a percentage BER β .

$$\beta = \begin{cases} 0, & Q_l = 255 \\ (255 - Q_l) \times 0.0025, & 255 < Q_l \leq 215 \\ 0.1 + (215 - Q_l) \times 0.08, & 215 < Q_l \leq 90 \\ 10.1 + (90 - Q_l) \times 0.64, & 90 < Q_l < 0 \\ 67.7, & Q_l = 0 \end{cases} \quad (1)$$

CSR chipsets report Link Quality values with lower resolution as BER increases. Figure 1 illustrates three separate resolutions for the range of reported BER values. The BER is a reported average of errors encountered over the connection over a period of approximately 10 seconds. This smoothing property restricts us to a slow survey to ensure the average has an opportunity to settle closer to a representative average at that point (discussed further in Section 4.1).

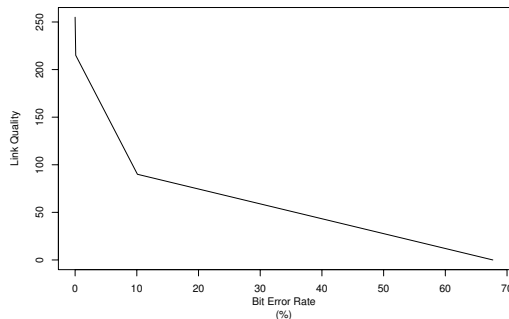


Fig. 1. Link Quality to BER relation for CSR Chipsets

Another factor in measuring link quality is the adaptive properties of Bluetooth transmission. The Bluetooth specification recommends a number of interference reduction techniques to boost reliability in crowded environments. These include power control, Channel Quality Driven Data Rate (CQDDR) and Adaptive Frequency Hopping [11]. The support for these optional enhancements are communicated when a Bluetooth link is established. Users may also explicitly query a remote device for supported features.

Bluetooth power control is an optional feature for the receiver to ask the transmitter to increase or decrease the power output according to its perceived RSSI in relation to the Golden Receiver Power Range. Power control should only be used if both parties in the connection support the feature. Class 1 devices must support this feature whereas lower powered Class 2 and Class 3 devices optionally support it. We found that both CSR and Broadcom chipsets support power control, but the Nokia Bluetooth chipsets found in their mobile phones do not advertise support for it. Using a hardware Bluetooth protocol analyser [16], we were able to verify during our measurements that no power control signalling occurred and thus was not a factor in the results presented here.

CQDDR is an adaptive scheme to select different Bluetooth packet sizes and types depending on the link quality. A receiver may signal the transmitter to use different packet types in order to compensate for bit errors on received packets. Bluetooth allows for 3 packet sizes (1-slot, 3-slot and 5-slot) corresponding to the number of time slots the packet consumes. In addition to that, each packet size can contain no forward error correction (FEC) or 2/3 FEC coding. In the presence of increased BER, a receiver may wish to sacrifice transmission speed with more robust encoding to prevent retransmission. Again, the Nokia chipset did not support this feature whereas both Broadcom and CSR did support it.

Adaptive Frequency Hopping [11] was introduced in the Bluetooth 1.2 specification in order for Bluetooth to avoid using channels which have high levels of interference. It requires that the master device detects which channels produce high BER or low RSSI, and instruct the slave device to avoid those channels. By querying the device and analysing the packets exchanged, we found that none of the Bluetooth hardware we used supported this feature. However, we anticipate that this will be an important factor to consider when performing link quality measurements for future Bluetooth devices.

3.2 Experimental Setup



Fig. 2. An Active Bat attached to a Bluetooth capable mobile phone used to correlate signal strength readings with their locations

Our experiments were conducted in a typical office-like environment where the Active Bat system is installed. For the portable client, we utilised a standard Nokia 6600 mobile phone with Bluetooth capabilities. An Active Bat was physically attached to the back side of the phone to facilitate retrieving the position of the phone in 3D space (see Figure 2). Throughout the experiment, a Bluetooth RFCOMM connection was established to the phone from a notebook PC running Linux 2.6.7 (using the Bluez Bluetooth stack) with a D-Link DBT120 Class 2 Bluetooth USB Adaptor (using a CSR Bluetooth chipset). A USB extension cable was used to place the adaptor in an open area 1 m above the floor.

The link quality of the established RFCOMM connection was queried every time a Bat sighting occurred, which is about 20-30 times per second. The RFCOMM link was used to transmit a remote RSSI query command defined in the AT specification [17]. Position information for the attached Bat was retrieved from the Active Bat system and recorded along with the most recent link quality status of the open RFCOMM connection, remote perceived RSSI and time-stamp.

The low latency and accuracy of the Active Bat system is crucial for data collection since it allows a device's current location to be correlated with the signal strength readings. The Active Bat system is traditionally used to locate tags worn by humans in the building. Context-aware applications access location information via the SPIRIT middleware, which performs aggregation and filtering of readings to cull outliers and artifacts introduced by multi-path reflections and other ultrasonic noise sources.

Using a positioning system in a surveying mode allows us freedom to perform post-processing filtering off-line rather than incur a delay for the system to perform real-time filtering. We opted to obtain the raw positioning information before any filtering had occurred other than the basic multi-lateration outlier rejections.

Link quality observed from the CSR Bluetooth adaptor is a result of a running average over previous values recorded by the chipset. We performed the experiments in a slow sweep of the area, to both ensure stability of the positioning results and also to allow time for the running average to converge closer to a stable value at that position.

Surveys were conducted during after hours in the building to minimise the impact of human interference to propagation. We took care to ensure that the surveyor did not obscure the path between the Bluetooth device and access point to ensure consistency of measurements.

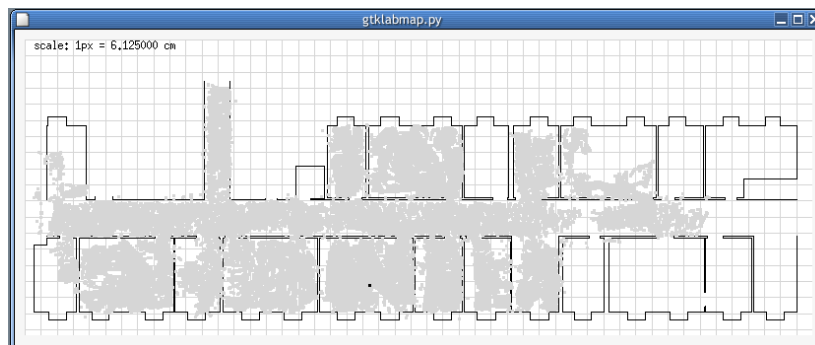


Fig. 3. A GTK application which plots the progress of a survey in real-time, ensuring good coverage of the target area by the surveyor.

We placed the Bluetooth adaptor at different places in the building. All were placed roughly 1 m from the ground at least 50 cm away from any other electronic devices. A real-time graphical map of results was produced from the collection program to ensure we covered the all the available positions in the survey area (see Figure 3).

3.3 Post Processing

Before performing analysis of the recorded measurements, the data was post-processed to remove outliers and inaccurate sightings. These errors consisted of: (i) the surveyor moving faster than the desired speed for the experiment and skewing the recorded Link Quality; and (ii) errors in the locations reported by the Bat system itself (recall that the Bat system has a reported accuracy of 3cm, 95% of the time).

Since the Active Bat system uses the time-of-flight information from ultrasonic pulses to determine the location of a Bat, it is vulnerable to indoor multipath and reflection effects from walls and other obstacles. The Bat system has a degree of built-in robustness to mitigate these effects; at least three Bat receivers must agree on a location using their multi-lateration algorithms before a sighting is considered valid. While this system works very well in the middle of rooms, it breaks down slightly when a Bat is placed very close to a large obstacle such as a wall. This is because the Bat system has limited knowledge of the world model (i.e. the presence of the wall), and thus occasionally extrapolates a reflected signal as originating from beyond the wall. Harle et al. implemented automated mechanisms for the Bat system to infer this information [9,8]; however, these algorithms work at a higher level than the raw data feed we used. The “teleportation effect” is characterised in location traces by a few sightings moving to the other side of a wall before snapping back to the correct location.

Spurious ultrasound emissions in the environment may also cause the Active Bat system to report a single sighting seemingly from a far away room. The effect can be attributed to ultrasonic noise detected by the ceiling receivers just between the time the surveying Bat is polled for an update and when the pulse from the Bat reaches the receivers.

Algorithm 1 describes the “teleportation filter” for detecting incorrect sightings that jump across walls. Given the set of sightings S of size N , we require a look-ahead window of L future sightings and the function $CheckWalls(P_1, P_2)$ to detect wall intersections (described later by Algorithm 2). The algorithm iterates over all the sightings and checks for intersections between the last known good sighting and the current sighting. If an intersection is detected (i.e., a potential teleportation), the next L sightings are also tested for interference against both the last good sighting and the current candidate. The current sighting is erroneous if it exhibits a larger number of future wall intersections than the last good sighting does (remember that a user walking near a wall will trigger the occasional teleport, but the bulk of future sightings will still be along the wall). Experimentally, using a look-ahead of $L = 50$ worked effectively to remove teleportation effects from our data.

Algorithm 1 Teleportation filtering algorithm

Require: $L \geq 0$ and $N > 2$

- 1: $LastValid \leftarrow S_1$
- 2: **for** $i = 2$ to N **do**
- 3: **if** $CheckWalls>LastValid, S_i) > 0$ **then**
- 4: $last \leftarrow \sum_{i+L}^{x=i} CheckWalls(S_x, LastValid)$
- 5: $cur \leftarrow \sum_{i+L}^{x=i} CheckWalls(S_x, S_i)$
- 6: **if** $last > cur$ **then**
- 7: $LastValid \leftarrow S_i$
- 8: $Valid(S_i)$
- 9: **else**
- 10: $Invalid(S_i)$
- 11: **end if**
- 12: **else**
- 13: $Valid(S_i)$
- 14: **end if**
- 15: **end for**

The speed of sightings was also clamped to a maximum value in order to account for the occasional incorrect distant sighting and the surveyor walking too fast. This was done by deriving the velocity v from the time-stamps and distances of subsequent sightings, and eliminating any above v_{max} (which varied depending on the target movement speed for that particular survey).

Algorithm 2 Intersection test for two line segments

- 1: $q \leftarrow (y_1 - y_3)(x_4 - x_3) - (x_1 - x_3)(y_4 - y_3)$
- 2: $d \leftarrow (x_2 - x_1)(y_4 - y_3) - (y_2 - y_1)(x_4 - x_3)$
- 3: **if** $d = 0$ **then**
- 4: $false$ {lines are parallel}
- 5: **else**
- 6: $s \leftarrow (y_1 - y_3)(x_2 - x_1) - (x_1 - x_3)(y_2 - y_1)$
- 7: **if** $d < q < 0 \parallel d < s < 0$ **then**
- 8: $false$ {not on line segment}
- 9: **else**
- 10: $x_{intersect} \leftarrow x_1 + \frac{r(x_2 - x_1)}{d}$
- 11: $y_{intersect} \leftarrow y_1 + \frac{r(y_2 - y_1)}{d}$
- 12: **end if**
- 13: **end if**

4 Survey Results

In this section, we present the results of the measurements described earlier. Section 4.1 details the tests which determined an optimal movement speed for

conducting the rest of the surveys. Section 4.2 then describes the results from our main surveys of the building with transmitters placed at different points.

4.1 Movement Speed

The purpose of the first experiment conducted was to determine the best movement pace at which to take the subsequent measurements. This was done by selecting a long corridor with no obstructions such that free-space path loss could be expected. A transmitter was placed in the middle of the corridor, and the surveyor repeatedly walked up and down at a constant pace while holding the receiver (a Nokia 6600 Bluetooth mobile phone) towards the transmitter at all times. The results were then filtered to remove outliers and portions where the surveyor walked too fast or too slowly (see Section 3.3 for more details on the filtering algorithms used).

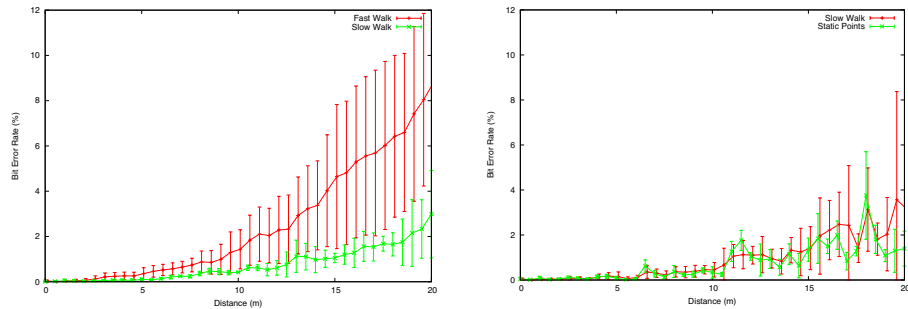


Fig. 4. On the left graph, the BER is plotted for a “fast walk” ($v_{max} = 2 \text{ m/s}$) and a “slow walk” ($v_{max} = 0.22 \text{ m/s}$), showing how unreliable signal strength readings for a fast moving user are. The right graph illustrates very little BER variation between a slow walk ($v_{max} = 0.22 \text{ m/s}$) and stationary measurements (where the receiver has been idle for over 10 seconds).

Figure 4 (*left*) shows a graph of the measured BER against the distance between the receiver and transmitter. We observe that the BER (both average and standard deviation) of an observer moving at a fast walking pace was far higher than that of a slower walk. With this in mind, all the future surveys were done at a slower pace in order to give the BER adequate time to settle.

Interestingly, an experiment to measure the BER of a completely static receiver and transmitter pair failed to achieve a better result than the slow walk. This experiment was conducted by mounting the mobile phone and Bat on a wooden stand, placing it at the desired location, and leaving it to settle for 10 seconds. After this time, BER readings were taken for another 10 seconds to measure the receiver in its “steady-state”.

Figure 4 (*right*) shows very similar measurements of the BER between a slowly moving receiver and a static receiver. A common source of problems in indoor environments is multi-path reflections, which can cause a loss of signal strength by out-of-phase signals causing destructive interference. In order to mitigate this effect, Bluetooth uses pseudo-random frequency hopping algorithms which shift frequencies between 2.4-2.4835GHz. This change of frequency (and thus, change of phase of incoming signals) causes a Bluetooth receiver to effectively “jitter” as if it were constantly in a small amount of motion in its environment. The jitter is useful to push the receiver in and out of zones of destructive interference even if the receiver is static, and accounts for the lack of differences between a slowly moving and a static receiver.

4.2 Signal Strength

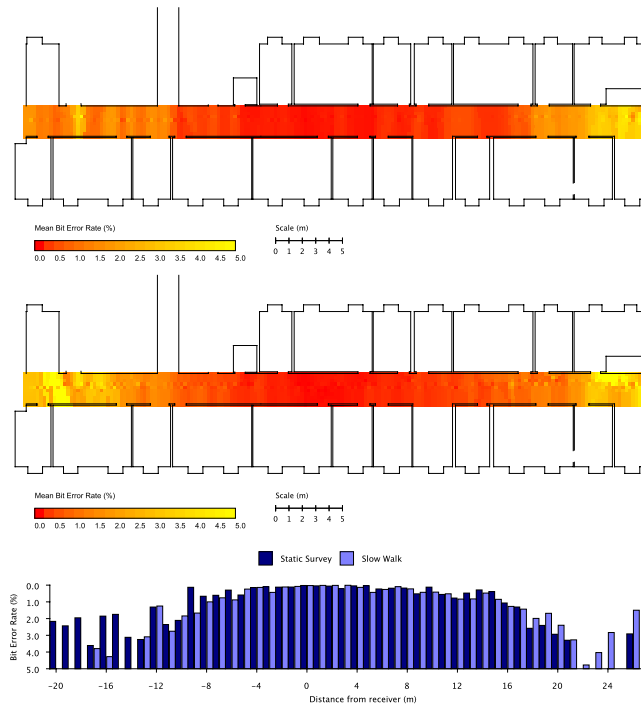


Fig. 5. Measurements of signal strength from a fixed points along the corridor (*middle*) with receiver stationary for more than 10 seconds and signal strength measured from a slow walk (*top*). A graph of BER versus distance from receiver (*bottom*) shows little difference between the two surveys until the signal degrades significantly at over 12m.

To verify our experimental setup, we first conducted an experiment with a known result. Figure 5 (*middle*) illustrates the free space path loss of a signal in a

long corridor in the building. A Bluetooth transmitter was placed in the middle of the corridor, and a mobile phone was placed at points along the corridor and left for 10 seconds. This allowed the running average BER value to settle down to a stable value. We recorded 26938 raw points, out of which 26690 (99.0%) were valid samples after post-processing. This survey confirms that our location measurements were accurate, as we expected to see the standard free-space path loss shown in the figure. The slight disruption in path loss to the right of the figure is due to large metal cupboards at that location which were difficult to move. Figure 5 (*top*) plots the free-space path loss with the surveyor moving at a slow walk surveying the same area. Figure 5 (*bottom*) confirms the observation from Section 4.1 that there is little difference between a survey with a static receiver and one moving slowly until the signal degrades significantly after separation exceeds 12m.

Figure 6 (*top*) shows the results of a survey conducted by placing a Bluetooth transmitter towards the west end of the building, and performing a slow survey with a Nokia 6600 mobile phone. This survey consisted of 79742 raw samples, of which 66815 (83.7%) were valid samples which passed the post-processing error filters. Each grid location represents the average value of all the signal strengths located inside it. The survey shows that the Bluetooth transmitter has a practical range of around 2 rooms, after which the BER exceeds 2% and the bandwidth on the connection drops below useful levels.

Figure 6 (*bottom*) plots the standard deviation of the values for each grid location. The standard deviation gives an indication of the stability of the connection, and a stable value is very useful for determining the location of the user with respect to the transmitter. The standard deviation is low when the user is close to the transmitter, as the signal strength remains consistently high. As the user moves away, the chipset reports inconsistent signal strength values (cycling between low and high values). This can be partially attributed to the frequency hopping that Bluetooth uses; as the distance from the transmitter increases, the effect of reflected waves introducing constructive or destructive interference for a particular frequency increases. Bluetooth encounters both types of interference as it hops through the 2.4GHz frequency range, accounting for the inconsistent link quality values. Adaptive Frequency Hopping has been proposed as a solution to this problem in future Bluetooth devices, but it is currently not implemented on any consumer Bluetooth chipsets such as CSR, Nokia or Broadcom, so it is irrelevant for the current generation of widely deployed hardware.

A second survey, shown in Figure 7, picked a more central area to help confirm this observation. This survey covered a wider area and consisted of 113686 samples, 91428 (80.4%) of which were valid samples after post-processing. The Nokia 6600 mobile phone could maintain a Bluetooth base-band connection without terminating it even in areas of 5% BER. However, no payload data could be transmitted over the connection during this time; the phone buffered the data until the surveyor moved closer to the fixed Bluetooth transceiver and the BER decreased.

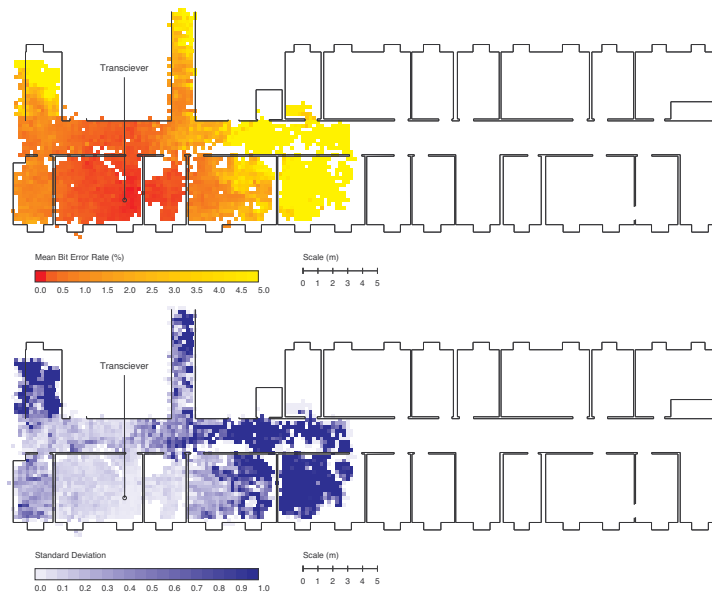


Fig. 6. Measurement of signal strengths (*top*) and the standard deviation (*bottom*) from the west end of the building.

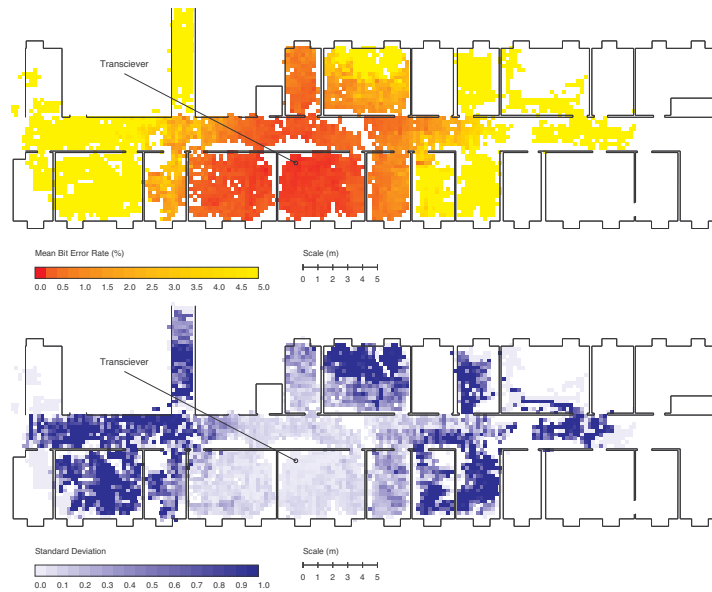


Fig. 7. Measurement of signal strengths (*top*) and the standard deviation (*bottom*) from the centre of the building.

4.3 Discussion

Recall from Figure 4 that at a regular walking speed, the average link quality behaves as expected, and BER increases as the distance between transceivers increases. However, the variance of the BER values is much greater at high speeds than when compared to slow movement or stationary readings. This suggests that users walking around a building at normal walking pace would have an adverse impact on their link quality and available bandwidth. This is of concern to location-based services deployed in public buildings such as shopping malls—in order to push high-bandwidth content reliably, the user would have to be relatively stationary rather than walking through the Bluetooth zone.

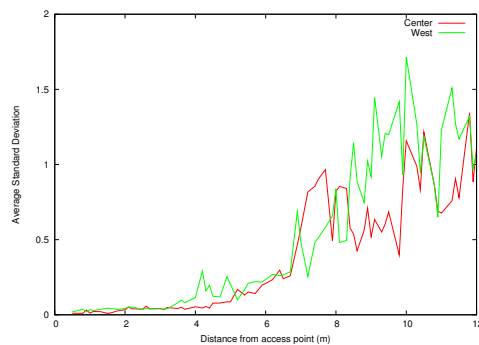


Fig. 8. Distance between the receiver and transmitter versus the average standard deviation of the BERs observed at those points. This was plotted for the data-sets shown in Figures 6 and 7 to show that they both show the trend of increased standard deviation as distance increases.

Figure 8 plots the distance between the receiver and transmitter versus the standard deviation of the BERs observed at all the points of the Bluetooth surveys shown in Figure 6 and Figure 7. Standard deviation for BER values exhibit a slow gradual increase in standard deviation between 0 to 6.5 m followed by a steep increase. At distances further than 6.5 m, BER values are unstable and thus using BER at such distances would produce unreliable distance predictions. For distances under 4 m from the transceiver, there is only a small magnitude of variance when moving at slow speeds. A low standard deviation of BER values is important for low-latency location-sensing, as it allows for a stable and accurate distance prediction of the user from the transmitter.

We observe several factors that lead us to conclude that Bluetooth is a poor choice for the use of signal strength for low-latency location sensing: *(i)* all the hardware we surveyed only exposes a running average of the BER which is updated at different intervals depending on the hardware (1s for CSR and 5s for the Broadcom dongles); *(ii)* measured BER had a high variance as the distance increased beyond about 6m; *(iii)* none of the mobile phones supported more

than a single Bluetooth connection, making triangulation difficult; and *(iv)* the reported RSSI values was of little or no use due to the lack of resolution and slow update rate. One must also take into account possible adaptive features available in future revisions of the mobile phone chipsets, which further blur location inferencing.

Although the actual propagation distance achieved is dependent on hardware (e.g. antenna type and construction), we believe our results would apply to different omnidirectional Class 2 Bluetooth devices because of the restriction in power output of the device. Limitations in consumer hardware prevented us from obtaining higher resolution data from the Bluetooth hardware itself. Compared to WiFi, none of the reported link quality values give high enough accuracy and dependability to enable location-sensing based on signal strength alone.

The Bluetooth specification could improve support for location inferencing if changes were introduced requiring devices to: *(i)* expose a fine-grained RSSI dBm value via HCI (similarly to 802.11); *(ii)* accept multiple simultaneous Bluetooth connections in order to assist triangulation; and *(iii)* update RSSI values per-packet without a significant time lag.

However, Figure 8 indicates show a positive trend of standard deviation as the distance to the transmitter varies. This implies that if a Bluetooth location system is willing to increase the latency of reported results, it could estimate the distance of a slow moving user based on the combination of standard deviation and average values of measured signal strength over a period of time. For tracking stationary objects, it may also be feasible to overcome the restriction of one Bluetooth connection at a time by polling the device from multiple co-operating base stations.

Our use of the Active Bat system as a positioning system proved valuable to understanding the useful characteristics of using indoor location systems for this purpose: *(i)* systematic stable latency between the real position and reported sightings enables easy correlation of external signal strength readings with their associated positions; *(ii)* high rate of readings allowed us to understand the dynamic behaviour of signal strength on moving devices; and *(iii)* reported location sightings must be stable before signal strength readings are taken; Bluetooth Class 2 devices have a useful range of 10 m, for which the BER varies between 0-2%, with around 50 possible link quality steps reported by the CSR chipset. If these readings were spread out linearly, it would require 0.2 m intervals to record them all distinctly. The Active Bat system offered much higher precision than is needed for a survey of a Bluetooth Class 2 device.

As reported in the literature, inaccurate world models have a significant negative impact on context-aware applications and services [8]. For our purpose, surveying and propagation measurements required more detailed world model information that would not be useful to existing context-aware applications (e.g. wall and object material composition). Other objects such as switching cabinets and server room equipment were not in the world-model, but are still variables which affect indoor radio propagation.

5 Related Work

Wireless LAN technology (WLAN) has been a popular candidate for indoor location estimation based on signal strength. Most of these methods rely on a surveying phase where checkpoints in the environment are statically surveyed to obtain calibration data. The position of a device is determined by correlating the signal strength measurements at different access points using Bayesian analysis [20] or simple triangulation [2]. Our Bluetooth survey attempt to map all physical space covered by the location system and also allows us to map the signal strength rate of change for moving targets.

The advantage of WLAN positioning is that signal strength information can be obtained at a per packet level and signal strength can be sampled from a number of access points simultaneously. Under Bluetooth, signal strength information is only obtained indirectly through the BER value (a running average over a 10 s interval). Furthermore, many mobile Bluetooth devices cannot accept more than one connection at a time, which is needed to calculate signal strength between two end points. Bluetooth devices are more attractive because of their low power usage and proliferation in current consumer devices; e.g. in previous work, we used a preliminary version of our survey technique to integrate Bluetooth mobile phones into a fast-paced location-aware game [15].

Projects such as Place Lab [19] and the Location Stack [6] have used the coarse-grained metric of Bluetooth device sightings along with other sources such as WiFi or GPS to infer user location. Our work complements this by performing more fine-grained analysis on Bluetooth indoor propagation using the Active Bat system and making these data sets freely available to other researchers interested in Bluetooth [14].

Gwon et al also propose a mechanism of location estimation for mobile and stationary users based on signal strength measurements [7]. They report their Region of Confidence (RoC) algorithm performs better using Bluetooth than WiFi, but this is a result of using a denser collection of Bluetooth nodes for the same area. Simulations of Bluetooth channel capacity [13] and its efficient coexistence with other protocols in the same frequency range (such as WiFi) have been extensively studied [5]. In this paper, we isolate Bluetooth and consider its propagation through comprehensive measurements.

6 Conclusions and Future Work

The extensive signal strength surveys using the Active Bat system have made a wealth of data available for future analysis [14]. In order to minimise the number of variables in our study, we chose to restrict the surveys in this paper to a single Bluetooth transmitter and receiver. There has been extensive simulation work into how effectively multiple Bluetooth devices can co-exist [13] with each other, or with other protocols in the 2.4GHz range (e.g. WiFi). We intend to extend our measurements to investigate the actual impact of these interactions on moving devices such as PDAs and mobile phones. Although the current generation of mobile phones do not yet support them, future versions of Bluetooth

will support Adaptive Frequency Hopping [11] and Channel Quality Driven Data Rate (CQDDR) that will have a negative impact on location inferencing based on signal strength. The survey methodology described in this paper allows researchers to easily perform an analysis of the real-world impact of these features with mobile devices of varying speeds.

In this paper, we have presented a measurement-based approach to help evaluate and visualize the accuracy of wireless signal propagation. The Active Bat ultrasonic indoor location system was used to gather a large number of samples of Bluetooth signal strengths from transceivers placed around a building. The samples exhibited a number of artifacts which are to be expected from a high-resolution, raw location feed from a multi-lateration-based location system, and we presented algorithms which filtered out bad samples to a high degree of accuracy. These samples were then analyzed for receivers moving at different speeds to simulate the movements of real users (running, walking or staying stationary). We observed a significant increase in BER variance for users moving at walking speeds or faster, implying that location-based “push services” operating over Bluetooth will have less bandwidth available to transmit multimedia content in busy public areas where consumers are not standing still.

We conclude from our measurements that Bluetooth is ill-suited for the purpose of accurate, low-latency location sensing due to: *(i)* common chipsets only exposing a running average of signal strengths and updating it infrequently; *(ii)* the high variance in signal strengths for longer distances; and *(iii)* the inability for consumer mobile phones to maintain multiple Bluetooth connections simultaneously for triangulation purposes. Finally, we have made our data-sets available on-line [14] for other researchers to benefit from the large number of samples gathered during this research.

7 Acknowledgments

Thanks to David McCall from Cambridge Silicon Radio for prompt and accurate replies to our queries about the BlueCore Bluetooth chipset. We are grateful to Andrew Rice, Kieran Mansley, Alastair Beresford, Ford Long Wong, Robert Harle, Ian Wassell and our shepherd Anthony LaMarca for their valuable discussions and contributions to this work, and to the members of the Digital Technology Group who tolerated strange people repeatedly sweeping through their offices armed with mobile phones and Bats. This work was partly funded by Intel Research Cambridge.

References

1. M. Addlesee, R. Curwen, S. Hodges, J. Newman, P. Steggle, A. Ward, and A. Hopper. Implementing a sentient computing system. *IEEE Computer*, 34(8):50–56, August 2001.
2. P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *IEEE Infocom*, March 2000.

3. Bluetooth Special Interest Group. Bluetooth core specification 1.1. <https://www.bluetooth.org/spec/>.
4. Cambridge Silicon Radio. BlueCore Bluetooth chipset. Online at <http://www.csr.com/products/bc-family.htm>.
5. N. Golmie, R. E. V. Dyck, A. Soltanian, A. Tonnerre, and O. Rebala. Interference evaluation of Bluetooth and IEEE 802.11b systems. *Wireless Networks*, 9(3):201–211, 2003.
6. D. Graumann, J. Hightower, W. Lara, and G. Borriello. Real-world implementation of the Location Stack: The Universal Location Framework. In *Fifth IEEE Workshop on Mobile Computing Systems and Applications*, October 2003.
7. Y. Gwon, R. Jain, and T. Kawahara. Robust indoor location estimation of stationary and mobile users. In *IEEE Infocom*, March 2004.
8. R. K. Harle. *Maintaining World Models in Context-Aware Environments*. PhD thesis, University of Cambridge, 2004.
9. R. K. Harle, A. Ward, and A. Hopper. Single Reflection Spatial Voting: A novel method for discovering reflective surfaces using indoor positioning systems. In *First International Conference on Mobile Systems, Applications and Services*, pages 1–14, May 2003.
10. A. Harter, A. Hopper, P. Steggles, A. Ward, and P. Webster. The anatomy of a context-aware application. In *5th Annual ACM/IEEE International Conference on Mobile Computing and Networking*, 1999.
11. C. Hodgdon. Adaptive frequency hopping for reduced interference between Bluetooth and wireless LAN, May 2003. Ericsson Technology Licensing.
12. In-Stat/MDR. Bluetooth-enabled products poised to swim in the mainstream, April 2004. Order Code R97-1590.
13. A. Kumar and R. Gupta. Capacity evaluation of frequency hopping based ad-hoc systems. In *Proceedings of the 2001 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 133–142. ACM Press, 2001.
14. A. Madhavapeddy and A. Tse. Bluetooth propagation datasets. Available online at <http://www.cl.cam.ac.uk/Research/DTG/research/bluetoothsurvey/>.
15. K. Mansley, D. Scott, A. Tse, and A. Madhavapeddy. Feedback, latency, accuracy: exploring tradeoffs in location-aware gaming. In *SIGCOMM 2004 Workshops: Proceedings of ACM SIGCOMM 2004 workshops on NetGames '04*, pages 93–97, New York, NY, USA, 2004. ACM Press.
16. Mobiwave. Bluetooth protocol analyser BPA-D10. <http://www.mobiwave.com/bpa.htm>.
17. Nokia Mobile Phones. AT command set for Nokia GSM products. Online at <http://www.forum.nokia.com/>.
18. N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The Cricket Location-Support System. In *6th ACM MOBICOM*, Boston, MA, August 2000.
19. B. Schilit, A. LaMarca, G. Borriello, D. M. William Griswold, E. Lazowska, A. Balachandran, and J. H. V. Iversen. Challenge: Ubiquitous location-aware computing and the Place Lab initiative. In *First ACM International Workshop of Wireless Mobile Applications and Services on WLAN*, September 2003.
20. P. Tao, A. Rudys, A. M. Ladd, and D. S. Wallach. Wireless LAN location-sensing for security applications. In *Proceedings of the Second ACM Workshop on Wireless Security (WiSe)*, San Diego, CA, Sept. 2003.
21. A. Ward. *Sensor-driven Computing*. PhD thesis, University of Cambridge, 1998.