An adaptive ARQ timeout approach for audio streaming over Bluetooth

Ling-Jyh Chen*

Institute of Information Science,
Academia Sinica,
128, Sec. 2, Academia Road, Taipei 11529, Taiwan
E-mail: cclljj@iis.sinica.edu.tw
*Corresponding author

Rohit Kapoor

Qualcomm Inc.,
5775 Morehouse Drive, San Diego, CA 92121, USA
E-mail: rkapoor@qualcomm.com

Sewook Jung

Broadcom Corporation,
16340 West Bernardo Drive, San Diego, CA 92127, USA
E-mail: sewookj@broadcom.com

M.Y. Sanadidi and Mario Gerla

Department of Computer Science,
University of California at Los Angeles,
Los Angeles, CA 90095, USA
E-mail: medy@cs.ucla.edu   E-mail: gerla@cs.ucla.edu

Abstract: In this paper we study the application of audio streaming over Bluetooth and propose an enhancement to the Bluetooth link layer ARQ mechanism to compensate for channel degradation and to better support audio streaming. Specifically, our scheme adaptively sets the ARQ timeout value based on current channel conditions. We show through simulation and testbed experiments that the adaptive ARQ improves the streaming quality significantly compared to the 'vanilla' link layer of Bluetooth, especially in noisy environments. Our proposed approach is simple to implement and can actually be extended to the link layer of any wireless technologies.

M.Y. Sanadidi is an Adjunct Professor in the Computer Science Department at UCLA. He received his PhD in Computer Science from UCLA in 1982. As co-principal investigator on NSF sponsored research, he is leading research in modelling and evaluation of high performance Internet protocols. His current research interests are in path characteristics estimation and its application in congestion control and adaptive streaming in heterogeneous (wired/wireless) networks.

Mario Gerla is a Professor in the Computer Science Department at UCLA. He received his Graduate Degree in Engineering from the Politecnico di Milano in 1966, and his MS and PhD Degrees in Engineering from UCLA in 1970 and 1973, respectively. He joined the Faculty of the UCLA Computer Science Department in 1977. His current research is in the area of analysis, design and control of communications networks. His ongoing projects include the design and evaluation of QoS routing and multicast algorithms for IP domains, the design of wireless mobile, multimedia networks for mobile computing applications, and the development of measurement methods and tools for evaluating the performance of high-speed networks and applications.

1   Introduction

The last few years have seen an impressive growth in multimedia content on the internet. One striking success in this area has been MP3 audio, which has led to the development of MP3 players, such as the iPod (http://www.apple.com/ipod/). Another orthogonal area of growth has been wireless, both in wide area technologies such as 2.5G/3G and local area technologies such as 802.11b and Bluetooth. One interesting usage scenario of Bluetooth is the Personal Area Network (PAN), i.e., a network of personal devices such as laptop, PDA, camera, MP3 player etc., that a person carries with her.

A number of interesting usage scenarios arise when audio streaming (MP3) content is combined with wireless technologies, such as Bluetooth. For instance, a Bluetooth-enabled MP3 player can stream MP3 audio to a Bluetooth-enabled headset, allowing the user to be 'locally' mobile while listening to songs. In another scenario, MP3 content could also be streamed through a 2.5/3G cellphone to a Bluetooth headset. Another scenario could be an MP3 player downloading songs from a LAN access point, and playing them real-time.

These and several other wireless streaming applications are envisioned. However, the varying nature of the wireless link can make audio streaming over wireless a challenging problem. In fact, a number of sources operate in the same ISM frequency band and, thus, can interfere today with Bluetooth and 802.11 (e.g., microwave ovens, cordless phones, etc.).

Good quality audio streaming requires that audio packets reach the client at a consistent rate within regular delays, even though audio (e.g., MP3) is not typically very high-bandwidth in nature. Clearly, a bad channel can cause packets to be dropped/delayed, adversely affecting the quality of streamed audio.

A well-designed link layer can go a long way in solving some of these problems. Most wireless link layers of today employ some kind of ARQ mechanism to protect packets from a bad channel. While this can be greatly beneficial to non-real-time TCP traffic, it needs to be judiciously used when real-time/streaming traffic is being transferred.

For example, if the ARQ retransmission limit is too high, packets can get severely delayed under bad channel conditions. On the other hand, a very low value of the ARQ limit can cause a large percentage of packets to be dropped at the link layer. Either of these situations will reduce the quality of streaming traffic. It is thus important that the ARQ retransmission limit be chosen in the correct range.

In this paper, we investigate adaptive settings for the ARQ retransmission limit based on wireless channel conditions. We show that an adaptive selection of the retransmission limit by far outperforms a fixed limit. We propose a simple algorithm to calculate the retransmission limit based on the transmission history of previous packets. Finally, we implement this algorithm both in the NS-2 simulator and on a Linux Bluetooth testbed. The experiment results show that our proposed approach improves performance of audio streaming under a wide range of channel conditions. While our solution was developed for Bluetooth, it can be easily applied to other wireless ARQ link layers.

The rest of the paper is organised as follows: Section 2 describes the Bluetooth technology and the related work. Sections 3 and 4 present our approach and the details of its implementation. Sections 5 and 6 reports experiment results from simulation and the testbed. Finally, Section 7 concludes the work.

2   Background and related work

2.1   Bluetooth overview

Bluetooth (http://www.bluetooth.com) is a short-range radio technology operating in the unlicensed 2.4 GHz Industrial-Scientific-Medical (ISM) frequency band. Its original goal was to replace the numerous proprietary cables to provide a universal interface for devices to communicate with each other. But it soon became a good solution to interconnect devices to form so-called PANs (Johansson et al., 2001), primarily due to the low cost, low power and small size of Bluetooth chips.

Bluetooth employs Frequency Hopping Spread Spectrum (FHSS) to avoid interference. There are 79 hopping frequencies (23 in some countries), each having a bandwidth of 1 MHz. Frequency hopping is combined with stop and wait Automatic Repeat Request (ARQ), Cyclic Redundancy Check (CRC), and Forward Error Correction (FEC) to achieve high reliability on the wireless links.

Bluetooth units can be connected to each other to form a piconet, which consists of up to eight active units. One of the units acts as a master and the others act as slaves. All the data/control packet transmissions are coordinated by the master. Slave units can only send in the slave-to-master slot after being addressed in the preceding master-to-slave slot. Each slot lasts for 625 microseconds.

For real-time data such as voice, Synchronous Connection Oriented (SCO) links are used, while for data transmission, Asynchronous Connectionless Link (ACL) links are used. There are several ACL packet types, differing in packet length and whether they are FEC coded or not. Table 1 shows the different ACL packet types and their properties.

Table 1 Different Bluetooth ACL connection modes

Mode	FEC	Packet Size (bytes)	Packet Length (slots)	Symmetric throughput (kbps)	Asymmetric throughput (kbps)	
DM1	Yes	17	1	108.8	108.8	108.8
DM3	Yes	121	3	258.1	387.2	54.4
DM5	Yes	227	5	286.7	477.8	36.3
DH1	No	27	1	172.8	172.8	172.8
DH3	No	183	3	390.4	585.6	86.4
DH5	No	339	5	433.9	723.2	57.6

Note that, in the symmetric connection mode, both master and slave nodes will occupy the same amount of Bluetooth time slots (625 microseconds in each time slot); whereas in the asymmetric connection mode, the Bluetooth link will occupy 1/3/5 time slots (for DM1/DM3/DM5 or DH1/DH3/DH5 mode) in one direction of this link and only one time slot in the opposite direction.

However, due to operating in the unlicensed 2.4 GHz band, Bluetooth can be subject to interference from other wireless technologies, such as 802.11b, HomeRF, cordless phones and other sources such as microwave ovens. As various studies (Fainberg and Goodman, 2001; Lansford et al., 2001; Punnoose et al., 2001) have shown, these sources can severely degrade Bluetooth performance.

2.2 Audio streaming

Streaming audio (aka MP3 files) has become very popular on the internet, and improving streaming quality has been a topic of active research. While receiving streaming multimedia, users expect smooth and uninterrupted quality and real-time effect, which has very different requirements than best-effort applications, e.g., FTP and HTTP. However,

real-time data are very sensitive to network conditions, and any delays or packet losses generated by the network can degrade the stream quality. Much of the work in this area has looked at improving streaming over wired networks. For example Rejaie et al. (1999a, 1999b) and Saparilla and Ross (2000) propose that the sender should dynamically adjust its sending rate by considering the media encoding, the frame rate, and the priority of media parts. http://www.real.com allows the receiver to select the most appropriate stream when the connection is setup. Additionally, recent research has focused on utilising the available bandwidth and making the stream friendlier with other TCP flows. For example, TEAR (Rhee et al., 2000) deploys an end-to-end, TCP equation-based approach to adjust its stream sending rate. TCP-Friendly Rate Control (TFRC) (Handley et al., 2003), another equation-based approach, computes the subsequent sending rate on the sender side, instead of on the receiver side (i.e., TEAR). Video Transport Protocol (VTP) (Balk et al., 2004), a new streaming protocol inspired by TCP Westwood (Wang et al., 2002), adapts its sending rate according to the estimated 'Eligible Rate'. On the other hand, some commercial products claim to support adaptive streaming, e.g., Helix Universal Server (http://www.realnetworks.com) and Microsoft Media Server (http://www.microsoft.com/windows/windowsmedia), though lack of product disclosure and related analysis hinders independent efforts to verify the claims or to evaluate the streaming performance.

However, all the above works have an implicit assumption that packet losses and delays are mainly due to network congestion, and the delay time caused by link layer retransmissions is negligible. Such an assumption is valid for wired connections, but not for wireless links, where the bit error rate is much higher than in wired links. While operating over wireless links, the retransmissions on the link layer usually result in a significant delay for each retransmitted packet. Such a delay can limit the applicability of end-to-end approaches and degrade the streaming quality.

2.3 Link layer enhancements

A wireless communication channel typically exhibits higher error rates due to attenuation, fading, scattering, or interference from other active sources. In order to provide more reliable data transmission over wireless links, modern wireless technologies usually employ different levels of retransmission and/or redundancy techniques in their link layer level against wireless errors (Fairhurst and Wood, 2002). For example, most wireless technologies perform Automatic Retransmission reQuest (ARQ) to ensure the integrity of the link layer, and some technologies, such as 802.11a and Bluetooth, utilise redundancy techniques to enhance transmission reliability by employing FEC coding in link layer packets.

However, such techniques may not lend themselves easily to supporting audio streaming applications. For instance, as the link quality becomes worse, multiple retransmissions become more and more frequent. Setting a high retransmission ARQ limit can ensure the successful

data transmission; however, it also leads to extremely large delays in audio packets, thus degrading the audio quality. Though most audio clients use a playout buffer to alleviate such problems, this does not solve the problem when the link quality is very bad for a sustained period of time (i.e., the buffer will become overflowed, and the data will then be dropped). A low retransmission limit, on the other hand, can keep the clients catching up the speed of audio streaming (i.e., keep the real-time effect); however, it leads to a high percentage of packets being dropped at the link layer, which also results in poor audio quality.

Additionally, redundancy methods such as FEC coding are only effective in counteracting random losses, but the connection is still vulnerable to burst channel errors. While wireless channel errors are well-known bursty and dependent in occurrences, FEC coding based approaches may still need large number of retransmissions when bursty errors are present, and the streaming quality will thus become degraded.

In order to overcome the problem, Krishnamachari et al. (2003) proposed a cross-layer approach to adaptively change the maximum number of MAC layer retransmissions and FEC encoding level in the application layer using the estimated MAC layer link quality (SNR). However, the link quality estimation is based on the previous measured link qualities, which may not be appropriate and applicable to the future link quality prediction. In our testbed experiments, we have found that the link quality oscillates very often and dramatically due to the wireless network dynamics, and the prediction of future link qualities is almost impossible to be achieved barely by the measured link quality in the previous history.

Thus, in order to maintain good quality of streaming, not only should the packet loss rate be kept small, but the delays of packets should also not be allowed to increase too much. We present our link layer solution to this problem in the next section.

3   Proposed approach

In this section, we present the importance of the application-aware link layer and our proposed approach, i.e., Adaptive ARQ Retransmission Timeout (RTO).

3.1   Application-aware link layer

It is evident that a non-application-aware link layer can result in reduced performance. For instance, if one fragment of an audio packet (say a RTP packet) is dropped (due to the RTO limit being exceeded), the non-application-aware link layer will still try to send the remaining fragments of the RTP packet, even though they are of no use since the receiving link layer will never be able to reassemble the RTP packet. This will result in wasted bandwidth and eventually degrades the overall application performance.

However, some sort of "*cross-layer optimisation*" may be of great help in improving the performance of the upper layer applications. A well-tuned link layer should adapt its configurations/behaviour in accordance to the present network conditions in order to maximise the application performance. More specifically, for the audio streaming discussed in this paper, the link layer needs to be able to measure the RTT for each RTP packet and perform different levels of adaptation in the link layer. The application-aware functionality required is that the Bluetooth stack should be able to identify and deal with a RTP packet since a vanilla link layer can only deal with Bluetooth baseband packets. The details of how this functionality is implemented are explained in Section 4 and the details of link layer adaptation are described in the following subsection.

3.2   Adaptation of ARQ RTO value

Bluetooth uses a stop-and-wait ARQ at the link layer and retransmits a packet until either the acknowledgement of a successful reception is received or the RTO is exceeded, at which point the packet is dropped. However, in most current Bluetooth chipsets, the default value of the RTO is infinite. An infinite timeout value makes the link layer reliable, since packets are not dropped even when the channel conditions are very bad. However, this can lead to problems for real-time/streaming media, since packets may be severely delayed when link quality is poor.

A simple approach could be to use a fixed, finite RTO value. This will result in packets being dropped at the link layer, whenever the retransmission limit is exceeded. Since a severely delayed packet may be completely useless at the client, dropping such a packet may be a good idea anyway since subsequent packets have a higher chance of being transmitted. Thus, this approach may improve audio quality if the RTO can be selected judiciously. Therein lies the problem of using a "fixed, finite RTO" since it may not be easy to find a timeout value suitable for all the cases, such as different link qualities.

Moreover, if such a setting is not appropriate, it may again degrade the audio quality. For example, if the fixed timeout value is too small, it will increase the drop rate; if the value is too large, it will lead to the same situation as the infinite RTO setting.

To overcome such problems, we propose an adaptive ARQ RTO approach that adapts the value of the link layer RTO based on the measured properties of previous packets. Thus, if the link layer has spent too much time on the previous few packets, it should decrease the RTO setting for the next packet, since the audio client has already experienced much delay from the previous packets. In other words, it is better to risk dropping the next packet (due to the decrease in RTO) than to incur another increase in delay. On the other hand, if the link layer has sent the previous packets very efficiently with short RTTs, it should increase the RTO value since the client has already saved time on previous packets and is capable of tolerating some delay for the next one. Thus, it pays to put some extra effort into reducing packet loss (by increasing RTO).

Namely, the link layer measures the Round Trip Time (RTT) of each audio packet, say RTP (Schulzrinne et al., 1996) packet, which is the time for the whole RTP packet to

get transmitted by the link layer (this implies the use of an application-aware link layer, as we discuss later). Using the value of the RTT, a smoothed RTT, SRTT, is calculated (equation (1)), from which the RTO is calculated. The SRTT and RTO update equations are:

$$SRTT' = (1 - \gamma) \times SRTT + \gamma \times RTT \qquad (1)$$

$$RTO' = \begin{cases} \alpha \times RTO; & \text{if } RTT < SRTT \\ \beta \times RTO; & \text{if } RTT > SRTT \\ RTO; & \text{if previous packet is dropped} \end{cases} \qquad (2)$$

where we take $\gamma = 0.25$, $\alpha = 1.1$, and $\beta = 0.9$. Note that the RTO is dynamically updated using a multiplicative increase/decrease scheme following the threshold check. RTO increases when RTT decreases and vice versa. This is very different from TCP, where the RTO is increased proportionally to RTT.

In addition, we also apply upper and lower bounds to the RTO value. The lower bound RTO_{min} is taken as 2 times T_{packet}, which is the time interval between two RTP packets sent on the server side. T_{packet} can be easily derived from the packet size of the RTP packet, and we will present the calculation in the next section. We found through our experiments that if the RTO value was set close to the T_{packet}, too many packets were dropped at the link layer due to the RTO limit being exceeded. Thus, 2 times the T_{packet} proved to be a good lower bound.

The upper bound RTO is proportional to the available buffer size, as shown in the following equation.

$$RTO_{max} = T_{packet} \times Max(AvailableBuffer \times 75\%, 2) \qquad (3)$$

where *AvailableBuffer* is the system maximum input buffer size minus the used buffer size, divided by the RTP packet size. This equation takes into account the fact that if the RTO for an RTP packet is too large, it may cause new incoming RTP packets to be dropped from the input queue due to overflow. In fact, we found that for very large values of RTO, a number of packets were dropped because the buffer was full. Limiting the RTO to this upper bound prevented such packet drops.

Note that, in equation (2), we do not update the RTO if the previous packet has been dropped due to timeout. However, because contiguous packet dropping is harmful to audio quality, we reset the RTO to RTO_{max}, using equation (3), if at least two of the last five packets have been dropped.

4 Implementation

In order to evaluate our proposed approach, a Bluetooth simulation model has been implemented in NS-2 simulator (http://www.mash.cs.berkeley.edu/ns/), and a Linux-based Bluetooth testbed has been created. For simplicity, we used DH5 packets in all experiments and stream the MP3 audio using RTP protocol. The details of the implementation and experiment scenarios are presented in the followings.

4.1 Implementation of Bluetooth simulation

The Bluetooth simulation model was developed as an extension to NS-2 (Network Simulator). The simulation model implements most of the features of the Bluetooth baseband layer, such as frequency hopping, time division duplexing, multi-slot packets, fragmentation and reassembly of packets. In addition, the simulator enables scatternets and inter-piconet communication by defining gateway nodes to forward packets between piconets. Our proposed approach, as well as fixed ARQ RTO value approaches, is also included in the simulation model.

Figure 1 shows the Bluetooth network topologies used in the simulation, where topology (a) was used to simulate one audio stream in a one hop connection (in one piconet), topology (b) was used to simulate one audio stream in a two hops connection (the streaming is sent from one piconet to the other one, where both sender and receiver are master nodes), and topology (c) was used to simulate two audio streams in a three hop connection (these two streaming flow share one bottleneck link). While topologies (a) and (b) are designed to evaluate the packet loss rate and average delay, topology (c) is employed to evaluate the fairness of our proposed approach.

Figure 1 Various Bluetooth scatternet topologies

For simplicity, all simulation use DH5 packet type as the link layer packet type, and the error model is assumed to be uniformly distributed with a given bit error rate b. The link layer packet error rate P_{err} can be therefore obtained by equation (4).

$$P_{err} = 1 - (1 - b)^{339 \times 8}. \qquad (4)$$

4.2 Implementation of Bluetooth testbed

We implemented both the fixed RTO and the adaptive RTO method on our Bluetooth testbed. The testbed consists of two Linux based laptops, both equipped with a Bluetooth PCMCIA card. We installed BlueZ (http://bluez.

sourceforge.net), which is an open source Bluetooth Stack on the Linux operating system, on both laptops. We also used some other 802.11b devices to generate the interference to our Bluetooth connection during our experiments. We used DH5 packets in all our experiments. The system setup is shown in Figure 2.

Figure 2   Bluetooth testbed (see online version for colours)

The streaming protocol used in our experiments is Real-Time Streaming Protocol (RTSP) (Schulzrinne et al., 1998), which is widely used on the internet. RTSP is an application-level protocol to control the delivery of real-time multimedia data for both unicast and multicast. RTSP segments the MP3 stream into many small RTP packets at the server; the client can control the delivery (move backward, move forward, play, or stop) via the RTCP (RTP Control) protocol. Each RTP packet contains an RTP sequence number and may contain several MP3 packets.

The audio stream used in our experiments is a 128 kbps bit rate, 44.1 MHz frequency, Layer II MP3 file. The MP3 packet size of this music is 417 bytes, and each RTP packet contains 3 MP3 packets plus the header information (16 bytes). Thus, the RTP packet size is $417 \times 3 + 16 = 1267$ bytes, and the T_{packet}, the time interval between two RTP packets sent on the server side, is $1267 \times 8/128000 \approx 80$ ms. Since we use DH5 packets which have a maximum payload size of 339 bytes, each RTP packet will need at least 4 DH5 packets to be transmitted; therefore, the minimum transmission time for each RTP packets is $0.625 \times (5 + 1) \times 4 \approx 15$ ms, where 0.625 ms is the Bluetooth slot time, (each DH5 packet consumes 5 time slot in one direction and 1 in the opposite direction).

BlueZ is an open-source implementation of the Bluetooth stack on the Linux operating system. In the Bluetooth stack (shown in Figure 3), the Host Controller Interface (HCI) layer provides a command interface to communicate, access, and control the hardware layer. The L2CAP (Logical Link Control and Adaptation Layer Protocol) layer provides connection-oriented and connectionless data services to upper layer protocols with protocol multiplexing capability, segmentation and reassembly operation. The Bluetooth Network Encapsulation Protocol (BNEP) (http://www.bluetooth.com) layer lies on top of the L2CAP layer and provides support for IP.

To make the link layer application-aware, we extract header information in the BNEP layer from the received RTP packet. The BNEP layer passes the information downward the Bluetooth stack to the L2CAP layer and then the HCI layer. Each RTP packet is split into multiple fragments by the L2CAP layer. Once the fragments of the RTP packet arrive at the HCI layer, it queues all the fragments and stores the arrival time of the first fragment. The measured RTT is the time for all the fragments of the RTP packet to be successfully transmitted. In case one fragment of the RTP packet is dropped due to the RTO being exceeded, the remaining fragments of the RTP packet are removed from the HCI layer and the baseband by using the *HCI_Flush* command (which is defined in the Bluetooth specs).

Figure 3   Bluetooth stack
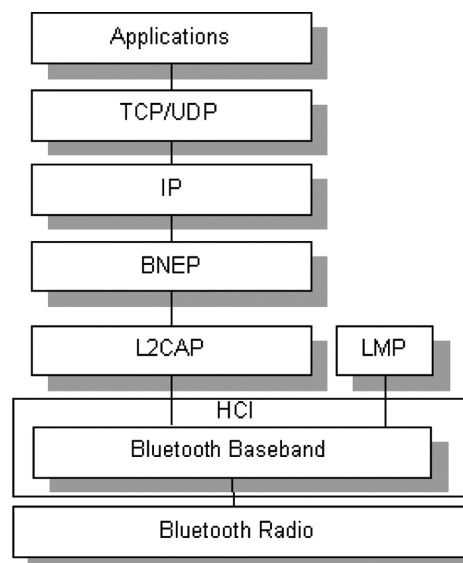
4.3   Generating interference for Bluetooth testbed

One challenge of our testbed setup was generating reliable interference. We found that it was very difficult to control the signal strength in the testbed since environmental factors make the link quality unpredictable. We used 802.11b devices to create interference for the Bluetooth connections, because both of them operate in the 2.4 GHz frequency band. Though increasing the physical distance between the two Bluetooth devices decreased the link quality, we found that it also increased the variance of the link quality and therefore makes conditions difficult to control. We were able to control the signal strength by keeping the two Bluetooth devices very close and controlling the traffic loads on interfering 802.11b connections.

To obtain the link quality from the Bluetooth chipset, we used the *Get_Link_Quality* function call. This call is defined in the Bluetooth spec (http://www.bluetooth.com) in the following manner:

Get_Link_Quality: This command returns the value of the Link Quality. It returns a number between 0 and 255, with the higher value representing a better channel. Each Bluetooth module vendor will determine how to measure the link quality.

As an example, for Bluetooth cards containing CSR (Cambridge Silicon Radio (http://www.csr.com)) chipsets, the Link Quality is calculated from the Bit Error Rate in the following manner:

If BER (Bit Error Rate) = 0, LQ (Link Quality) = 255

If BER <= 40/40000, LQ = 255 – BER × 40000

If 40/40000 < BER <= 4000/40000, LQ = 215 – ((BER/32) × 40000)

If 4000/40000 < BER <= 40000/40000, LQ = 105 – ((BER/256) × 40000).

Figure 4 shows the relation of the measured link quality value vs. bit error rate for the CSR chipset, and we will make use of such a relation to monitor and calibrate the channel and correlate to the RTO dynamics in the following experiments.

Figure 4 Link quality vs. BER for CSR chipset
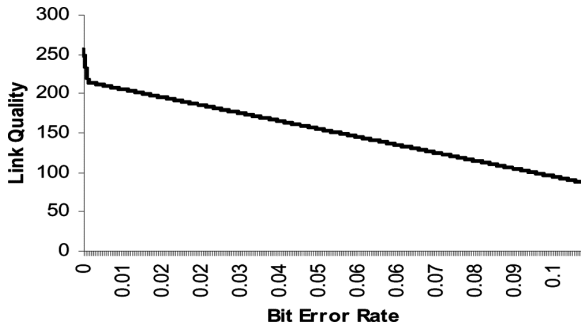

5 Simulation results

In this section, we evaluate our proposed approach using simulation method. In the following subsections A and B, we present the simulation results showing the audio streaming throughput and packet delay in a one hop Bluetooth connection (Figure 1(a)). In C, we show the results of the packet success rate of the audio streaming in both one hop and two hops Bluetooth connections (Figure 1(a) and (b)). Finally, we evaluate the fairness of our proposed approach by simulating two concurrent audio streams sharing one bottleneck link (Figure 1(c)). All audio streaming presented in this section was simulated using 128 Kbps RTP traffic, and all Bluetooth connections were using DH5 link layer packet type. In the following subsections, we use "fixed 160, 400, 1200" to stand for the fixed RTO method with timeout values 160, 400, and 1200 ms, respectively.

5.1 RTP packet throughput

Figure 5 shows the throughput results of audio streaming in one hop Bluetooth connection with different bit error rates, where the *x* axis (bit error rate) is presented using the logarithmic scale. From the simulation results, it is evident that the adaptive RTO scheme always outperforms other

fixed RTO schemes, especially when the bit error rate is higher than 10^{-5}. On the other hand, while the bit error rate is small (e.g., smaller than 10^{-7}), the throughput results of the adaptive RTO scheme and fixed RTO schemes are almost the same, since the data retransmissions are rates in this case.

**Figure 5** Average throughput on 1 hop Bluetooth connection

5.2 RTP packet delay

Figure 6 shows the average RTP packet delay in one hop topology with different bit error rates. From the results, the adaptive RTO scheme achieves the smallest packet delay in all the cases. While the bit error rate is as high as 0.0001, the adaptive scheme can reduce the delay by 50% and 20% when compared with using fixed RTO as 1200 and 400 ms.

Note that, the packet delay presented in this paper is an average of packet delays from the successfully transmitted packets, regardless of those packets which are dropped in the link layer. Therefore, though fixed RTO as 160 ms can achieve almost the same average packet delay in the simulation results, it suffers more serious packet losses due to the small RTO value. As a result, the user-perceived audio quality will become degraded.

**Figure 6** Average packet delay on 1 hop Bluetooth connection

5.3 RTP packet success rate

Figures 7 and 8 show the average RTO packet success rate on 1 hop/2 hops Bluetooth connection. The packet success rate stands for the ratio of those audio packets which are

successfully transmitted to the receiver. From the results, it clearly shows that the adaptive RTO scheme outperforms other fixed RTO scheme in all test cases. The performance improvement of using the adaptive scheme is more than 20% when operating on a 1 hop connection with 10^{-5} bit error rate, and 40% when operating on a 2 hops connection with 10^{-6} bit error rate. Additionally, while changing the Bluetooth connection from one hop to two hops, the average RTP packet success rate decreases in all schemes, especially when the bit error rate is high. However, the decrease in the adaptive RTO schemes is much less than the other fixed RTO value schemes, which means the adaptive RTO scheme is more robust with the increase of the number of connection hops.

Note that when the bit error rate is as high as 10^{-5} (on 1 hop connection) or 10^{-6} (on 2 hops connection), the large fixed ARQ RTO scheme (e.g., fixed 1200) may not get a higher packet success rate than the small fixed RTO scheme (e.g., fixed 160). This is in some way counter-intuitive since larger ARQ RTO should provide more reliable data transmission, i.e., high success rate; however, while the error rate is high, such large ARQ RTO would require more data retransmission in the link layer, and therefore results in link layer buffer overflow if the buffer size is limited.

**Figure 7**   Average RTP packet success rate on 1 hop Bluetooth connection

Figure 8   Average RTP packet success rate on 2 hops Bluetooth connection

5.4   Fairness evaluation

In the last simulation, we evaluate the fairness of the proposed adaptive ARQ RTO scheme using the topology shown in Figure (1(c)). From the results shown in Figure 9, two audio streaming flows have very similar throughput results, i.e., they are able to fairly share the bottleneck link capacity and thus achieve fairness. More specifically, the results also indicate that the modified Bluetooth link layer not only improves the performance of audio streaming, but also keeps the fairness feature, which is an inherited property of the deployed Time Division Multiple Access (TDMA) MAC layer behaviour, of Bluetooth connections.

Figure 9   Average throughput of two audio streaming flows with one shared bottleneck Bluetooth link

6   Experiment results

In this section, we present experimental results showing the improvement in real-time audio quality when using the proposed approach. The testbed setup and implementation are described in Section 5, and the experiments are repeated under different link quality conditions. In the following experiments, we use 'Adapt' to stand for our adaptive scheme and "Fixed 160, 400, 1200" to stand for the fixed RTO method with timeout values 160, 400, and 1200 ms, respectively. It should be noted here that in the default Bluetooth implementation, the ARQ timeout is infinite, i.e., the link layer never drops a packet.

Note that, because of the difficulty in controlling the link quality in the real experiments, the link quality distribution of each experiment is unique, even though the average BER might be the same. However, the average performance trend should be clearly observed.

6.1   Adaptive RTO

In the first experiment, we show the RTO adaptation behaviour of our proposed approach. In Figure 10, the solid line represents the RTT, i.e., the time between the arrival of an RTP packet at the Bluetooth baseband layer and completion of its successful transmission, and the dashed line represents the adaptive RTO value. The ARQ timeout

events are marked as circles in the figure. It can be seen that the adaptive RTO value increases as the RTT decreases and vice-versa, in accordance with equation (2).

Figure 10 RTO adaptation of the proposed approach

6.2 RTP packet success rate

Figure 11 shows the RTP packet success rate on the receiver side, i.e., the percentage of packets successfully transmitted. Different BERs were generated by varying the load on interfering 802.11 connections, as explained earlier. Our adaptive scheme outperforms the fixed ARQ timeout schemes, clearly showing the benefit of changing timeout based on channel conditions. The 'fixed 160' actually outperforms the higher ARQ timeout cases. Though this result may look counter-intuitive since a higher ARQ timeout is expected to drop fewer packets due to ARQ timeout, in reality higher ARQ timeouts also lead to a larger number of packets being dropped due to input queue overflow. This is exactly the reason why 'fixed 1200' shows the least RTP packet success rate. The vanilla BlueZ link layer, which has an infinite ARQ timeout, performs similar to the 'fixed 1200' timeout.

Figure 11 RTP packet success rate

6.3 RTP packet delay

Figure 12 shows the RTP average packet delay at the audio client. While the packet success rate determines how much data are successfully transmitted, the average packet delay represents how smooth the quality is.

From the figure, it is obvious that the adaptive approach is able to achieve smaller average delays. Even with a BER of 0.0045, the average delay is still close to the minimum

value, which means acceptable audio quality. For the fixed RTO cases, the higher the ARQ RTO timeout, the higher is the average packet delay. This is obvious since a larger ARQ timeout value leads to a larger number of retransmissions.

Figure 12 RTP packet delay
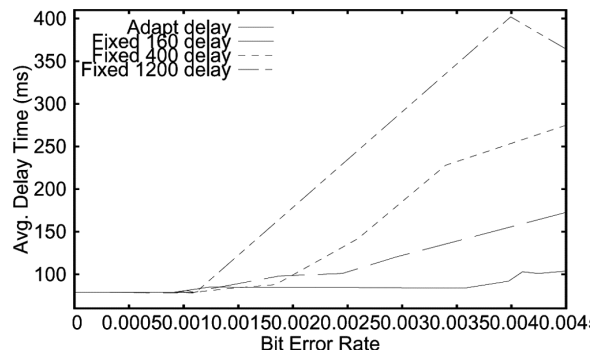
In balance, from the examination of the results in Figures 11 and 12 one concludes that the adaptive packet scheme operates adequately with BER up to 0.004 (packet loss rate less than 10% and delay less than 50 ms). The 'vanilla' scheme, with infinite RTO will fall apart for BER < 0.002. This is a remarkable improvement in performance.

7 Conclusion

In this paper, we studied the influence of link layer behaviour on streaming audio over Bluetooth. A cross-layer approach was proposed to adapt ARQ RTO value in accordance to the transmission of previous few streaming packets. We implemented a simulation model of the proposed adaptive ARQ timeout approach and created a Linux based testbed. A set of experiments was performed to evaluate the adaptive ARQ timeout scheme at the Bluetooth link layer in improving the quality of streaming audio. We compared our results with the fixed ARQ timeout methods; the results show that our method improves both average delay and the packet loss rate of RTP packets, particularly when channel conditions are bad. Moreover, the proposed approach is simple and can be easily implemented in the link layer of other wireless technologies such as 802.11b.

References

Balk, A., Gerla, M., Maggiorini, D. and Sanadidi, M. (2004) 'Adaptive video streaming: pre-encoded MPEG-4 with bandwidth scaling', *Computer Networks*, Vol. 44, No. 4, pp.415–439.

Saparilla, D. and Ross, K.W. (2000) 'Optimal streaming of layered video', *Proceedings of IEEE Infocom (Infocom00)*, Tel Aviv, Israel, Vol. 2, pp.737–746.

Fairhurst, G. and Wood, L. (2002) *RFC 3366: Advice to Link Designers on Link Automatic Repeat reQuest (ARQ)*, http://www.rfc-editor.org/rfc/rfc3366.txt, Retrieved 20 August, 2004.

Schulzrinne, H., Rao, A. and Lanphier, R. (1998) *RFC 2326: Real Time Streaming Protocol (RTSP)*, http://www.rfc-editor.org/rfc/rfc2326.txt, Retrieved 20 August, 2004.

Schulzrinne, H., Casnet, S., Frederick, R. and Jacobson, V. (1996) *RFC 1889: RTP: A Transport Protocol for Real-Time Applications*, http://www.rfc-editor.org/rfc/rfc1889.txt, Retrieved 20 August, 2004.

Rhee, I., Ozdemir, V. and Yi, Y. (2000) *TEAR: TCP Emulation at Receivers – Flow Control for Multimedia Streaming*, Technical Report, Department of Computer Science, NCSU.

Lansford, J., Stephens, A. and Nevo, R. (2001) 'Wi-Fi (802.11b) and Bluetooth: enabling coexistence', *Network, IEEE*, Vol. 15, No. 5, September–October, pp.20–27.

Fainberg, M. and Goodman, D. (2001) 'Analysis of the interference between IEEE 802.11b and Bluetooth systems', *VTC Fall*, Vol. 2, pp.967–971.

Handley, M., Floyd, S., Pahdye, J. and Widmer, J. (2003) *RFC 3448: TCP Friendly Rate Control (TFRC): Protocol Specification*, http://www.rfc-editor.org/rfc/rfc3448.txt, Retrieved August 20, 2004.

Johansson, P., Kazantzidls, M., Kapoor, R. and Gerla, M. (2001) 'Bluetooth: an enabler for personal area networking', *Network, IEEE*, Vol. 15, No. 5, September–October, pp.28–37.

Punnoose, R.J., Tseng, R.S. and Stancil, D.D. (2001) 'Experimental results for interference between Bluetooth and IEEE 802.11b DSSS systems', *Proceedings of IEEE Vehicular Society Conference (VTC01-Fall)*, Atlantic City, NJ, USA, Vol. 1, pp.67–71.

Rejaie, R., Estrin, D. and Handley, M. (1999a) 'Quality adaptation for congestion controlled video playback over the internet', *Proceedings of ACM SIGCOMM (SIGCOMM'99)*, New York, NY, USA, pp.189–200.

Rejaie, R., Handley, M. and Estrin, D. (1999b) 'RAP: an end-to-end rate based control for real time streams in the internet', *Proceedings of IEEE Infocom (Infocom'99)*, New York, NY, USA, Vol. 3, pp.1337–1345.

Wang, R., Valla, M., Sanadidi, M.Y. and Gerla, M. (2002) 'Adaptive bandwidth share estimation in TCP westwood', *Proceedings of IEEE Globecom (Globecom'02)*, Taipei, Taiwan, Vol. 3, pp.2604–2608.

Krishnamachari, S., Schaar, M., Choi, S. and Xu, X. (2003) 'Video streaming over wireless LANs: a cross-layer approach', *Proceedings of IEEE Packet Video Workshop (PV'03)*, Nantes, France.

Websites

Apple – iPod, http://www.apple.com/ipod/

Bluetooth Network Encapsulation Protocol Profile ver. 1.0, http://www.bluetooth.com

BlueZ – Official Linux Bluetooth protocol stack, http://bluez.sourceforge.net

Cambridge Silicon Radio, http://www.csr.com

Helix Universal Server, http://www.realnetworks.com

Microsoft Media Server, http://www.microsoft.com/windows/windowsmedia

Network Simulator (NS-2), http://www.mash.cs.berkeley.edu/ns/

Real Player, http://www.real.com

Specification of the Bluetooth System – Core vol.1 ver. 1.1, http://www.bluetooth.com