

Adaptive RTP/UDP/IP Header Compression for VoIP over Bluetooth

Luca Marzegalli¹, Mirco Masa², Mario Vitiello³
¹marze@cefriel.it, ²masa@cefriel.it, ³vitiello@cefriel.it

CEFRIEL / Politecnico di Milano – Via Fucini, 2 – 20133 Milano – Italy
Ph. +39 02 23954541, Fax. +39 02 23954254

Abstract - When Voice over IP traffic is carried over a wireless link the amount of overhead due to Internet headers becomes unacceptable. Therefore, header compression techniques can be applied to reduce bandwidth waste. But standard compression algorithms have poor performances in wireless environments due to high bit error rate. In this article we intend to analyze an adaptive header compression scheme for Bluetooth, named BAHC (Bluetooth Adaptive Header Compression), which can be used on VoIP packet stream. BAHC scheme is very efficient in terms of compression ratio and is also able to work very well when many packet losses occur on Bluetooth channel. Simulation results show the performance of a G.723.1 voice codec [1] when the BAHC header compression scheme is applied to the top of an unreliable Bluetooth link with different channel conditions.

1 INTRODUCTION

Given the success of wireless and IP communications, it seems inevitable that some future voice service will support IP over wireless links. However, the RTP/UDP/IP [2][3][4] protocol overhead could be a serious problem when real-time multimedia information is transmitted in bandwidth-limited environments, such as Bluetooth [5]. As far as Voice over IP (VoIP) is concerned, headers can occupy more than half the bandwidth used by the session. For example, a VoIP packet carrying a 5.3 Kbit/s G.723 voice frame per packet consists of at least 40 bytes (60 bytes for IPv6) of IP, UDP and RTP header information, while only 20 bytes are the actual speech payload.

If we consider these numbers it is obvious that the header size must be reduced in order to increase bandwidth efficiency. The header compression schemes (Figure 1) that have been studied by Internet researchers in the last ten years can reduce the RTP/UDP/IP header size to 5 bytes when used on wired links. But they do not perform well on wireless links, which are characterized by high bit error rate and transmission delay. These schemes do not perform well because they require the Decompressor to use always the same reference packet (also called *context*) that was used by the Compressor. This reference packet, however, could be lost on the channel, because of bad channel conditions. The innovative Window-based Least Significant Bit encoding (W-LSB encoding) depicted in [11] does not impose this restriction because it uses many reference packets in the Compressor whereas the Decompressor is re-

quired to have at least one of these reference packets. In W-LSB scheme the dimension of a Variable Sliding window (VSW) defines how many reference packets have to be used in the Compressor and it defines, therefore, the compression ratio efficiency: a small value (usually one) gives high compression ratio but losses of packets are not allowed on the channel; otherwise, when VSW is large we remark high robustness but low compression ratio and low bandwidth efficiency.

Considering the features of W-LSB encoding, the header compression scheme we are going to propose in this article introduces an innovative algorithm estimating optimum VSW dimension on Bluetooth wireless channel: the Bluetooth Adaptive Header Compression (BAHC).




Figure 1: General scheme of header compression.

This paper is organized as follows: Section 2 briefly describes Bluetooth technology, Section 3 introduces reference architecture and summarizes VoIP packet structure, Section 4 describes principles of header compression, Section 5 briefly explains BAHC algorithm, Section 6 reports simulation results, which fully characterize the behavior of BAHC, and Section 7 gives concluding remarks.

2 BLUETOOTH TECHNOLOGY

Bluetooth technology, which has gained support from leading producers such as Ericsson, Nokia, IBM, Toshiba, Intel and many others, is removing the need for wires and cables in personal environments [6]. Bluetooth radio system uses Frequency Hopping – Code Division Multiple Access (FH-CDMA) technique governing channel access. The hop frequency is 1600 hops per second; the frequency spectrum is divided in 79 bands and each of them is of 1 MHz bandwidth. The frequency-hopping scheme is combined with a channel's use scheme called polling Time division Duplex (TDD) that is led by a master unit. All the other devices participating to a piconet are slave units that can transmit only when authorized by the master. Bluetooth uses

fast Automatic Repeat request (ARQ), Cyclic Redundancy Check (CRC) and Forward Error Correction (FEC) to detect and correct errors. Finally, Bluetooth provides a nominal data rate of 1 Mbit/s for each piconet [7].

Baseband and Logical Link Control and Adaptation Protocol (L2CAP) layers form Bluetooth data-link layer.

L2CAP provides upper layer protocols with connection-oriented and connectionless data services through protocol multiplexing capability, segmentation and reassembly operation, and group abstractions. It also enables higher-level protocols and applications to transmit and receive data packets up to 64 kilobytes in length.

Baseband protocols implement ARQ retransmissions of baseband packet extracted from L2CAP frames. Bluetooth defines several baseband packet types characterized by different error protection levels and theoretical transmission capacity. These packets are called DM1 (108.8 Kbit/s), DH1 (172.8 Kbit/s), DM3 (258.1 Kbit/s), DH3 (390.4 bit/s), DM5 (286.7 Kbit/s) and DH5 (433.9 Kbit/s).

Baseband retransmission mechanism is suspended when the Logical Link Control and Adaptation Protocol timeout (L2CAP timeout) expires and the IP datagram is discarded.

3 VOICE OVER IP OVER BLUETOOTH

Voice over IP has become common in wired LANs and on the Internet, while its importance in 3G wireless systems may increase in the near future. However, its application to Bluetooth technology is somehow new.

3.1 Reference architecture

The reference architecture that has been considered in this work is depicted in Figure 2.

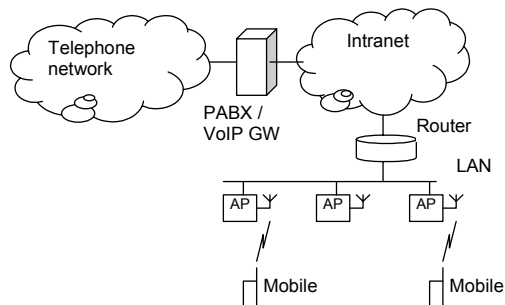


Figure 2: Reference architecture for VoIP over BT.

A Mobile handset with Bluetooth capability accesses a corporate network through a set of Bluetooth data Access Points (AP). Since it has been foreseen that 3G phones will embed an IP stack in the future, the reference architecture allows using the Bluetooth-enabled 3G mobile handset as a cordless phone, establishing

VoIP connections inside corporate Intranet. A dedicated VoIP gateway (where voice and signaling are properly translated) is then used to interface the data network with the legacy telephone network.

This work focuses on optimizing the transport of VoIP traffic on Bluetooth data link between the AP and the mobile handset.

3.2 VoIP packet structure

According to VoIP paradigm, voice samples are compressed and packetized into segments whose length does not usually exceed 20 bytes in order to avoid introducing large delays in the conversation. This payload is time-stamped using the Real Time Protocol (RTP), which introduces a 12-byte header. The resulting segment is then carried by a UDP datagram, which further adds a 8-byte header. While RTP provides the facilities for time synchronization, UDP allows several streams to be multiplexed together into a connectionless logical channel. Finally, when encapsulated into an IP datagram, the RTP/UDP packet is added a 20-byte IP header, which results in a large overhead. Things get even worse if IPv6 is used, because the IP header size increases from 20 to 40 bytes. Figure 3 shows the resulting packet structure.



Figure 3: VoIP packet.

The low efficiency in bandwidth usage is not a big problem when VoIP packets are carried over a wired LAN, but it can become a serious limitation when a wireless link is used instead.

4 HEADER COMPRESSION PRINCIPLES

When a single VoIP flow is observed, it can be noticed that only a few fields in the RTP/UDP/IP header vary between consecutive packets. This redundancy can be used to compress the headers on the wireless link.

Header compression mechanisms reduce header overhead since it is not necessary to send *static* header fields (e.g., IP addresses and UDP ports) in every packet because these *static* header fields do not change during the same session. Moreover, it is possible to efficiently handle the fields that change during the sessions (e.g., RTP timestamp, RTP sequence number and IP identification), so that the header overhead can be further reduced. In some cases, these so called *changing fields* can be foreseen from the last packet using a simple linear extrapolation. Other header fields (e.g. IP header length and UDP length) can be inferred from data-link level and it is not necessary to transmit them; these header fields are called *inferred fields*.

The Header compression scheme proposed by Casner and Jacobson [8] compresses the RTP/UDP/IP header,

but it is not designed to handle error rates and round-trip delay found over typical wireless links [9][10]. Many techniques have been proposed to adapt header compression techniques to wireless environment, such as ACE (Adaptive Header ComprESSION) [11] and ROCCO (Robust Checksum-based header COmpression) [12]. ACE introduces an innovative field encoding scheme for *changing field* (W-LSB or Window-based Least Significant Bit) that uses many reference values contained in VSW (Variable Sliding Window) in the Compressor, while the Decompressor uses only one reference. ROCCO uses a CRC to verify correct reconstruction in the Decompressor and to avoid error propagation.

The IETF RObust Header Compression Working Group (ROHC WG) [13] is studying new header compression schemes that are having good performances over wireless links of 3G cellular systems [14]. However, the schemes should also be applicable to other future link technologies with great loss and long round-trip times. Ideally, it should be possible to compress over unidirectional links. ROHC scheme uses and combines all the techniques developed by ACE and ROCCO.

5 A PROPOSAL OF HEADER COMPRESSION FOR BLUETOOTH: BAHC

Bluetooth is a radio technology for Wireless Personal Area Network. It implements segmentation and reassembly mechanisms of IP datagram at L2CAP (Logical Link Control and Adaptation Protocol) level. Every portion of L2CAP frame is encapsulated in a Bluetooth baseband packet and is transmitted on the channel. Each baseband packet is retransmitted until a correct acknowledgment is received. The retransmission mechanism is suspended when the L2CAP timeout expires and the IP datagram is discarded.

Before introducing BAHC algorithm it is important to brief the functionalities of W-LSB encoding that are most related to our work.

5.1 W-LSB encoding

Window-based Least Significant Bits encoding is used for header fields whose values are always subject to small changes among consecutive packets (*changing fields*).

With W-LSB encoding, the k least significant bits of the field value are transmitted instead of the original field value because the Most Significant Bits (MSB) remain relatively constant during VoIP session.

The value of k is calculated in the Compressor module using N reference values included in the Variable Sliding Window (VSW): for each reference value (v_{ref}), k_{ref} is determined so that the value to compress belong to the interpretation interval $f(v_{ref}, k_{ref})$:

$$f(v_{ref}, k_{ref}) = [v_{ref} - p, v_{ref} + 2^{k_{ref}-1} - 1]$$

Following, right value of k is the maximum among the N values of k_{ref} calculated above. After receiving k bits, the Decompressor derives the original value using a previously received value as reference.

It should be clear now how the dimension of VSW influences k and therefore the compression ratio too: k is typically a monotonous increasing function of the number of reference values N .

5.2 BAHC algorithm

BAHC utilizes information from Logical Link Control and Adaptation Protocol (L2CAP) to estimate channel condition and packet losses. In particular, it uses L2CAP timeout signaling: to use L2CAP timeout instead of BER measure to estimate channel conditions simplifies the algorithm and also reduces the dimension of Variable Sliding Window (VSW) update rate.

This information allows BAHC to calculate VSW dimension, which is the key parameter for compression performance. When VSW is small the scheme has high compression ratio but it cannot avoid many packet losses. Instead, when VSW is large the scheme has strong robustness but weak compression ratio and weak bandwidth efficiency.

Every time a packet enters the Compressor, Variable Sliding Window dimension is adjusted using the rules imposed by an estimation algorithm. We use a time interval called *Memory timeout* during which timeout values are registered. Indicating the actual value with $VSW(n)$ and the previous value with $VSW(n-1)$. At current time, Variable Sliding Window dimension is expressed as:

$$VSW(n) = VSW_{MIN}(n) + VSW_{BER}(n)$$

where $VSW_{MIN}(n)$ depends on the channel utilization, while $VSW_{BER}(n)$ depends on the channel state.

$VSW_{MIN}(n)$ is defined as follows:

$$VSW_{MIN}(n) = \left\lceil \beta * \frac{1.25 * \left\lceil \frac{MaxPkSize}{136} \right\rceil}{MeanArrivalTime} \right\rceil$$

where $MaxPkSize$ is the size, expressed in bit, of the biggest compressed or not compressed IP packet transmitted on Bluetooth channel; 136 bits are the maximum payload of a DM1 baseband packet; 1.25 ms is twice Bluetooth slot-time; $MeanArrivalTime$ is the mean time gap between packets in the same application flow; β is a constant bigger than one.

$VSW_{BER}(n)$ is dynamically adjusted according to the algorithm depicted in Figure 4.

The Compressor analyses the channel state by registering the number of L2CAP timeouts associated to the logical connection used to transmit VoIP compressed packets. Whenever the number of L2CAP timeouts over the memory timeout interval increases, VSW_{BER} increases accordingly to provide greater robustness.

For example, if the window dimension is 5, up to 4 VoIP frames can be lost in the window and the information can still be recovered.

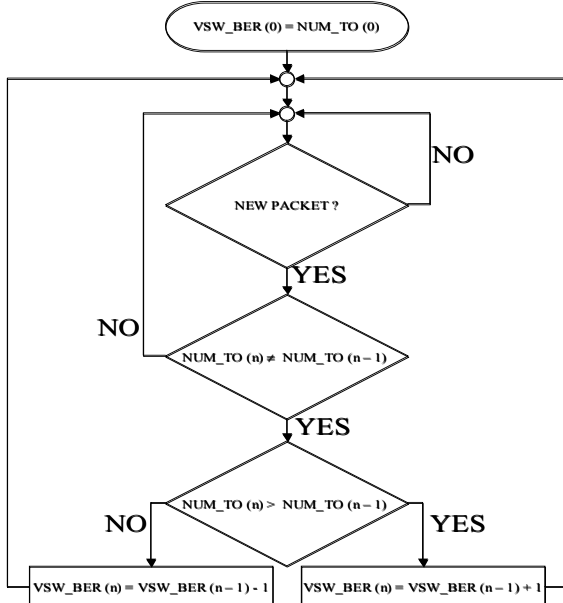


Figure 4: Algorithm adapting dynamically VSW dimension to Bluetooth channel state.

6 SIMULATIONS

The results presented in this section are based on BAHC compressed Voice over IP connection over a point-to-point Bluetooth link simulation. No other traffic sources were applied to the same Bluetooth link where VoIP traffic was carried.

OPNET Modeler was used to simulate BAHC, Bluetooth channel, network traffic, and to collect numerical results [15].

A Bluetooth L2CAP timeout of 12.5 ms was introduced to discard voice frames and both DM1 and DH1 baseband packet types were used in varying channel conditions.

Independently, uniformly distributed errors were introduced in each baseband packet at different rates and Bluetooth error control (FEC and ARQ) was accurately modeled.

The parameter β used to determine the VSW_{MIN} parameter was fixed at 3.

UDP checksum field was always transmitted.

The time during which the numbers of L2CAP timeouts were recorded, called *Memory timeout*, is 2 seconds.

Silence suppression periods used in G.723.1 codec were exponentially distributed with an average of 650 ms.

In order to give stronger reasons for our work, it is important to mention that BAHC algorithm implemented on the simulator can be seen as a simplified version of ROHC RTP profile [14]. Furthermore, BAHC uses a

new algorithm to estimate VSW dimension (see section 5.2).

The main differences between BAHC and ROHC are summarized below:

- W-LSB coding was applied to the fields: IP identification, RTP timestamp and RTP sequence number. Instead, ROHC RTP applied W-LSB only to the RTP sequence number while slightly different coding schemes were used for the rest of the fields;
- The format of the packets was simplified (see section 6.1 at page 4);
- UDP checksum field (2 bytes) was always transmitted;
- The Compressor and the Decompressor had only two states.

For these reasons, simulation results represent conservative indications of ROHC behavior, which is expected to improve the ratio and the robustness of the compression plotted in the graphs of this section. We are currently working on a ROHC profile for Bluetooth that we hope to present as soon as possible next year.

6.1 BAHC Compressor states and packet format

BAHC Compressor module has two compression states as shown in Figure 5: Initialization and Refresh state (IR) and First Order state (FO).




Figure 5: BAHC Compressor state.

IR state is used to transmit static contexts of compression session; instead, in FO state the Compressor transmits only dynamic fields encoded with W-LSB scheme.

The packet formats used in the BAHC compression scheme are depicted below.

0	7
Context ID	0000
RTP/UDP/IP header	
Application Payload	

Table 1: IR packet format.

When the Compressor is in IR state we use the packet shown in Table 1, where the full packet is transmitted so that the Decompressor may create (or update) a context identified by a Context ID.

0	7
Context ID	Profile
UDP checksum	
M bit, IP-ID, RTP SN, RTP TS	
Application Payload	

Table 2: FO packet format.

The structure of the compressed packet is shown in Table 2. The Profile parameter is used to indicate the number of bits that have to be used in the W-LSB decoding (up to 15 combinations are possible, since 0 indicates an IR packet type). UDP checksum field is inserted to allow checking the correct reconstruction of the decompressed header.

The next field has a variable length and depends on the encoding profile, i.e. the number of bits used in the W-LSB decoding for each field.

6.2 Numerical results

The following graphs show the performance assessment of BAHC header compression algorithm in terms of error robustness, compression ratio and bandwidth efficiency.

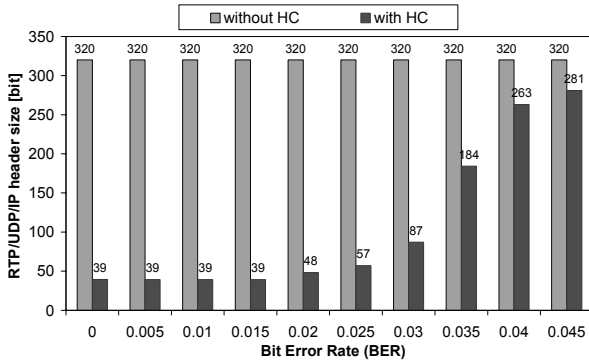


Figure 6: RTP/UDP/IP header length using BAHC.

In Figure 6 the length of the compressed header is plotted against the channel bit error rate when DM1 baseband packets are used to carry a L2CAP frame. With the best channel conditions, the RTP/UDP/IP header reduces to 5 bytes, so that the resulting L2CAP frame length is at its minimum:

- 20 bytes application payload;
- 5 bytes compressed RTP/UDP/IP header;
- 4 bytes L2CAP header.

Therefore, 29 bytes is the minimum L2CAP frame length for a VoIP packet, which means that at least two DM1 packets are needed to transmit it. Ideally, we would have liked to shrink the VoIP packet into a single baseband packet, but this solution does not seem to be

practically feasible. However, without header compression the total frame length would have been 64 bytes; therefore 2 DM1 or 1 DH1 packets would not have been enough.

The reduced length of L2CAP frame and the consequent reduction in the number of lost packets result in a higher application throughput as shown in Figure 7, where both DM1 and DH1 packet types are considered.

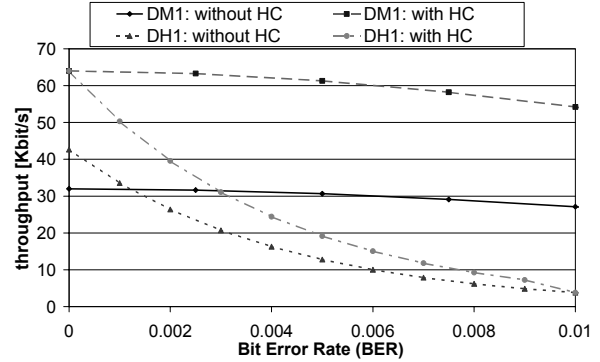


Figure 7: Maximum application throughput.

The Y-axis represents the maximum throughput expressed in bit/s available at the application layer. This allows calculating the maximum number of simultaneous voice connections in a piconet, as shown in Figure 8.

Figure 7 also allows estimating the trade-off between DH1 packets with header compression and DM1 without header compression: it can be noticed that for bit error rates above $3 \cdot 10^{-3}$ the latter choice provides higher throughput. It is also possible to conclude that using DH1 packets instead of DM1 packets is never a good choice, as evidenced by other simulation results that are not reported here.

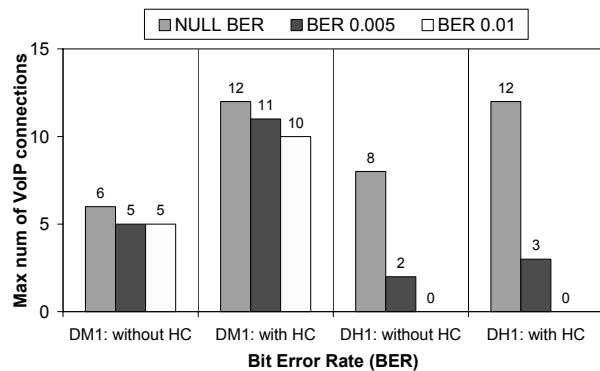


Figure 8: Maximum number of VoIP connections per piconet.

The bandwidth efficiency of a VoIP connection is a function of the bit error rate when the header compression is applied and it can be defined as:

$$\eta(BER) = \frac{\text{payload}}{\text{payload} + \text{header}}$$

