Click here to view PowerPoint presentation; Press Esc to exit

Ensuring Quality of Service for IP Communication over Radio Links in Tactical Networks

Markus Drasdo
Landshuterstrasse 26
85716 Unterschleissheim
GERMANY

Markus.Drasdo@eads.com

1 INTRODUCTION

In the civil world of computer networks, the Internet Protocol (IP) turned out to be the dominating network protocol. The pace of hardware development, which supports this protocol, accelerated due to the sheer size and strong competition within the civil market.

Recently Commercial Off The Shelf (COTS) products have also penetrated the military communications market. The reason for this shift from customer designed to COTS based solutions can be found in the development costs, which are associated with these complex products, the decreased defence budgets, and also because of the incompatibility of the various solutions, which causes tremendous interoperability problems for joint missions of different nations.

If COTS products are to be used for military applications, a series of problems arise. Besides the ruggeadisation of the equipment to withstand the hard environmental conditions, the usually physical layer available in the military environment is one of the toughest challenges:

Contrary to the civilian network world, where fibre or at least copper connections are available, most military networks rely upon a wireless infrastructure, such as LOS.

All wireless links, e.g. LOS, satellite and mobile communication, are susceptible to the introduction of bit errors into the transmission stream. Although this problem is as old as the existence of wireless transmission schemes, its consequences are worsened in the case of IP transmission.

IP was not originally designed for wireless communication and hence bit errors will cause

- Degradation in application quality or
- Decrease in application data rate.

In severe cases a network can become blocked or the connection will fail. In particular affected wireless IP networks include those which have to rely on links with the currently available quality, must have a near 100% availability or should work under extreme environmental conditions.

2 OPERATIONAL SCENARIOS LOS COMMUNICATION

2.1 Fixed LOS communication

Within this document, a *fixed LOS communication* link refers to a microwave radio link, which is *geographically fixed*. That is, the link is operated between pair antennas which are mounted on fixed masts.

Drasdo, M. (2006) Ensuring Quality of Service for IP Communication over Radio Links in Tactical Networks. In *Military Communications* (pp. 16-1 – 16-14). Meeting Proceedings RTO-MP-IST-054, Paper 16. Neuilly-sur-Seine, France: RTO. Available from: http://www.rto.nato.int/abstracts.asp.

A fixed link has the *advantage*, that the propagation issues can be *investigated in detail* (under the assumption that enough time is available). This can be achieved with the help of propagation calculation tools, on-site investigations, or experimental transmission trails.

According to these link investigations, a link budget calculation can be made, which outputs certain link attenuation. Traditionally, certain link availability estimation can be made based on environmental models and system parameters. If the link availability is too low counteractions can be taken, such as increasing the antenna size or transmit power.

Although a fixed communication link is static in terms of geography, *it is by far not static in its propagation behaviour*. Several factors influence the propagation, such as the rain attenuation (for frequencies > 10GHz), k-factor variations (especially for long distances and extreme climate conditions), multipath propagation, etc. Therefore, the receive signal level is time varying.

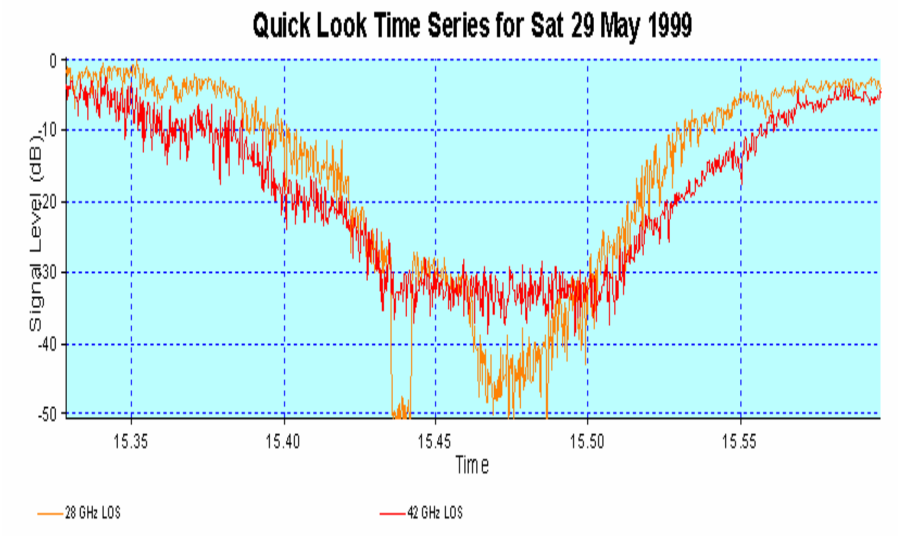

Figure 1: Receive level variation (Source: Rutherford Appleton Laboratory, RCA, UK)

Above it was mentioned that a link calculation tool estimates the link availability. However, a real radio link does not only have two states (available or unavailable). Instead*, with decreasing link receive level, it experiences a fluent degradation of the link quality* (that is an increasing residual bit error rate: BER) up to a certain level, where the radio losses synchronisation. Hence, *any receive level variations result in a varying BER*. As a consequence, the block based IP transmission is even more susceptible to link errors.

For a fixed link design, this increased error susceptibility can be taken into account and compensated by an additional link gain (for example from an increased antenna size) as long as the link permits such an improvement (for example the mast allows for larger antennas). If not, the effective range is reduced and a repeater might be required.

Even though an additional required link gain could be compensated with traditional methods, the link gain, provided by the radio components (radio system gain, antenna gain) is static. *This means that radio links, which have to run with a near 100% availability, require an extreme high system reserve (fade margin).* This dramatically reduces the range of the radio.

Therefore, fixed LOS links, which operate at the limits of their link range and/or have to have a high reliability of network connectivity call for a solution, which compensates the non-static propagation effects.

2.2 Deployable LOS communication

Within this document, a microwave radio link which is operated from a deployable mast is called a *deployable LOS communication link*. Such systems are especially common in tactical communication networks.

Due to the nature of such systems and contrary to fixed LOS links, no exhaustive link calculation can be made. The system has to be operational, wherever the mission requires the antennas to be. Also, the size of the antenna and hence the antenna gain is fixed. No additional link gain can therefore be obtained through the usage of larger antennas.

Therefore, even the path attenuation without the effects of rain attenuation, k-factor variations, etc. is not static. Obviously it is a function of the distance between the antennas and the shadowing of the propagation space. The following figure depicts the free space path attenuation for various frequencies. It can be seen, how the path loss increases with increased link distance and that a doubling of the distance results in 6 *dB* additional free space loss.

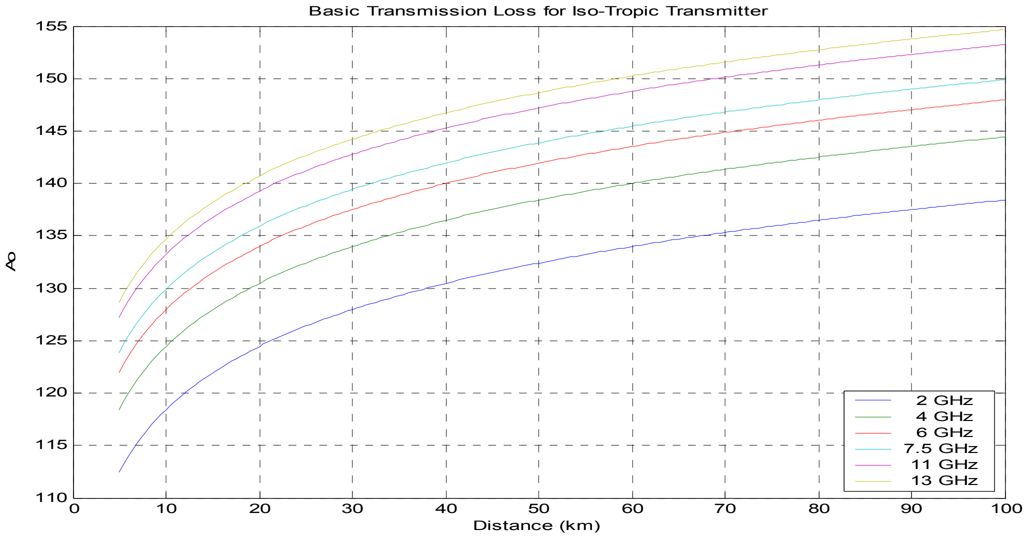

Figure 2: Free Space Propagation Attenuation for Various Transmission Frequencies.

The requirement to be operational as much as possible, independent of the length (and hence the quality) of the link calls for a solution which transforms the error prone wireless link into a channel with acceptable block error rate (BLER).

3 WIRELESS IP PROBLEM

3.1 Error vulnerability

3.1.1 IP

Contrary to standard fixed line networks, which use fibre or copper connections, wireless links, such as LOS links, have the inherent problem of introducing errors into the transmitted bit stream. Although this problem is not new and has existed since the creation of wireless transmission schemes, its consequences are exacerbated in the case of IP transmission. This effect is illustrated in the figure 3.

A single residual transmission bit error in a complete IP packet (including UDP or TCP packets) results in a complete IP packet loss.

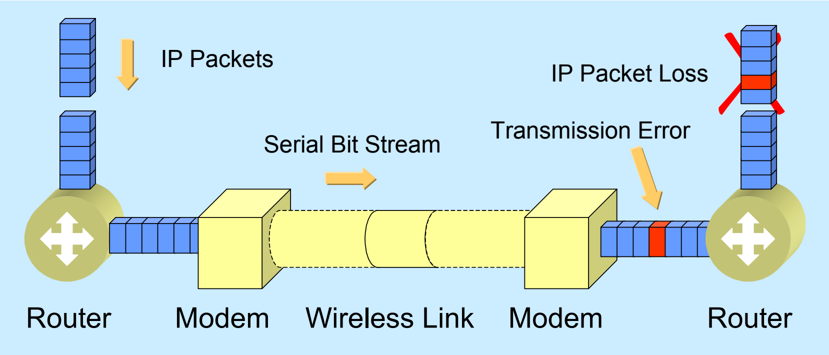

Figure 3: Schematic Block Diagram of a Digital Wireless Transmission Scheme, Illustrating the Wireless IP Problem.

Consequently, the IP packet loss rate is magnitudes higher than the residual channel bit error rate (BER).

In general IP packet losses are caused by a failed IP header error check. At each routing device, the IP header check of each received IP packet is performed. In case of passed test, the time to live counter is decremented, which makes a new header check calculation necessary. Thereafter the IP packet with the new header is routed and released to the network. If, on the other hand, the header checksum verification fails, the router discards the complete IP packet.

UDP/TCP
As discussed, a single transmission error in an IP packet results in an IP packet loss. Depending on the transport protocol this has different effects for the application.

3.1.2 UDP

The User Datagram Protocol (UDP) is defined in the transport layer (layer four) of the ISO-OSI Reference Model. This layer differs from the layers below it in that it provides true end-to-end communication. This means that a protocol within this layer directly communicates, using port addresses, with the corresponding protocol at the receiving host. In comparison protocols in layers one to three use a chain system to indirectly communicate to their corresponding peer, i.e. they communicate to their immediate neighbour, e.g. a router, which in turn forwards the message on.

UDP provides an unreliable, connectionless service, which means that packets are not guaranteed to be transmitted and delivered error free, i.e. there is no feedback system for packet acknowledgement, and that a defined connection route does not have to be established prior to communication, therefore packets can take independent paths across a network. These functions allow real time applications to send data whereby prompt delivery is more important than accurate delivery, for example speech or video transmission.

UDP uses a checksum to verify that the transmitted packet, i.e. header and payload, has been correctly transmitted. This is different from the IP checksum (as described), because the IP checksum only checks the header, although the algorithm principles used in both cases are the same.

UDP packets losses are caused by failed UDP packet error checks. After an IP packet has been accepted by a receiving host, the UDP packet (IP payload) is passed to the UDP entity, which performs the error check.

For the UDP protocol an IP packet loss results in a UDP packet loss. The following probability of a UDP/IP packet loss was obtained for the binary symmetric channel (BSC):

$$P(UDP/IP\_Packet\_Error) \cong L_{IP\_Packet} \cdot BER \quad \text{for} \quad BER < \frac{1}{L_{IP\_Packet}}$$

where L_{IP_Packet} is the length of the UDP/IP packet. Hence, the probability of an IP packet loss is roughly L_{IP_Packet} times higher than the channel BER. Figure 4 (a) and (b) depict the resulting UDP/IP packet losses for the BSC and the Additive White Gaussian Noise (AWGN) channel for binary phase shift keying (BPSK) with various packet lengths.

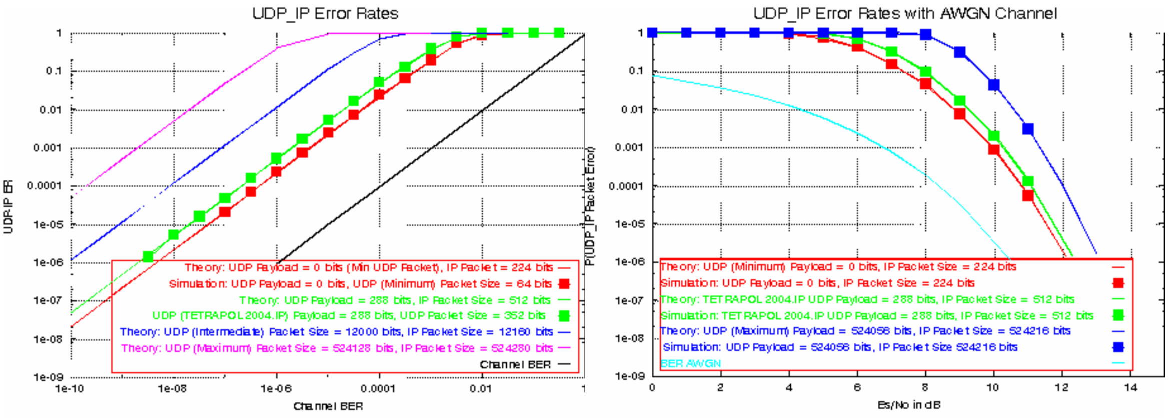

Figure 4: UDP/IP Packet Loss Probability for the BSC (a) and AWGN Channel (b).

3.1.3 TCP

The Transmission Control Protocol (TCP) is defined in the transport layer (layer four) of the ISO-OSI Reference Model and provides a reliable byte stream over an unreliable network, which may have different parameters, e.g. bandwidths or delays, in different parts of the network. This means that each data unit (information packet) is guaranteed to be transmitted and received error free, through the use of a feedback system, which sends an acknowledgement back to the sender for every packet that is received.

The service is connection-orientated, which means that a connection must first be established between the sending and receiving hosts before information can be transmitted. This is done by each host creating an end point, called a socket, which comprises of the IP address of the host and a local port number. Each connection requires two sockets, i.e. one for the sender and the other for the receiver, and transmitted data is then directly addressed to a socket. Connection set-up is not currently investigated, because it is a complicated process, which falls outside of the scope of this document.

These services are required by applications where accurate error free delivery is of high importance, for example data file transfer.

TCP is able to send multiple packets over a connection, therefore a fast sender must be prevented from swamping a slow receiver and hence TCP is responsible for handling flow control. This is done using a sliding window scheme. The window size represents the receiving host’s buffer size and indicates the number of packets that the receiver is able to accept.

If the TCP/IP protocol suite is used over a wireless infrastructure, channel errors result in IP or TCP packet losses and after the expiration of the associated TCP packet timer in retransmissions. This retransmissions lead to an inefficient utilisation of the channel data rate, which is provided by the wireless communication system. As a consequence, the average end-to-end TCP/IP data rate steadily decreases with the likelihood of channel errors:

This degradation of the effective user data rate is a major problem for wireless TCP/IP communication. Two major contributing effects can be noticed:

Firstly the repetition itself causes a degradation of the user data rate. For illustration consider the case that on average two TCP/IP packets have to be transmitted to deliver a TCP payload to the sink. Clearly, the effective data rate drops by 50% compared to a channel without errors.

Secondly, to make full usage of the channel data rate, the sender has to transmit continuously. The TCP transmitter, however, fall into a stop-and-wait mode when its window size is exceeded. If the packet to be acknowledged or its acknowledgement has been lost, the transmitter remains idle until the RTT timer expires.

Both effects contribute to an inefficient utilisation of the channel data rate and result in a decreased TCP data rate R_{TCP} and increased TCP delay.

The channel data rate is a measure of the number of bits, which are accurately received over a given time. Figure 5 shows the effects of BER on throughput for different channel delays. The throughput is illustrated using the percentage of possible channel data rate.

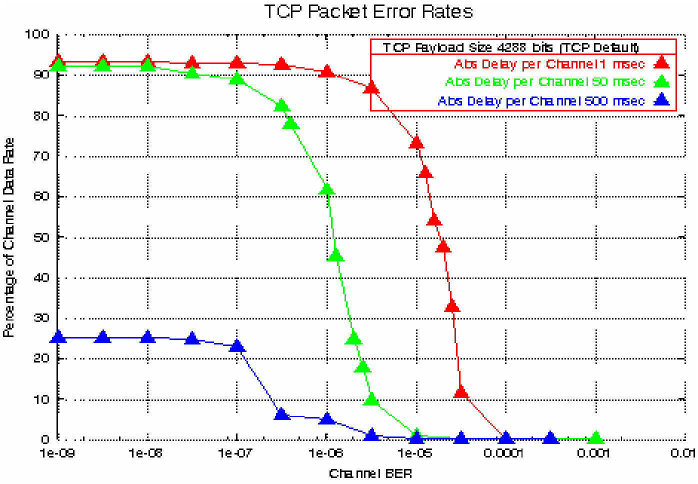

Figure 5: Relative degradation in data rate for the standard TCP/IP packet length and three different network delays for the BSC

In general it can be seen that the throughput rapidly decreases as the BER increases. This is because the RTO increases exponentially if several packets are retransmitted for the first time and therefore it takes longer for packets to be accurately received and for the system to recover. Eventually the BER causes the system to collapse, due to numerous retransmissions, and hence the system is blocked.

Under a normal network delay of 50msec, the channel is blocked when the BER .> 10^{-5}

An increase in channel delay causes a large decrease to the maximum obtainable throughput. The RTT is increased, but the large decrease in throughput is due to the TCP not being able to pass an optimal number of packets to the channel, because of TCP memory limitations and the fact that the delay means that the channel is unable to process further data for extended periods of time.

Due to the automatic retransmission request (ARQ) behaviour of TCP, no packet loss is observed. Instead, the effective data rate is decreased by the IP packet losses. This effective data rate (throughput) degradation is heavily dependent on the IP packet length as well as on the network delay. Due to the complexity of the retransmission protocol, no analytical evaluation is possible. Instead simulation was used to evaluate the throughput performance.

3.2 Interface delay

3.2.1 Serialisation delay

Serialisation delay is caused by the network to serial link conversion. When data packets arrive from the network at the interface they are queued until they can be sent along the serial link. The serial link can only transmit one packet at any given time and the time to transmit (*t*) depends on the packet length and link data rate (R_{data}), as expressed by:

$$t = \frac{PacketLength}{R_{data}}$$

Large packets take time to be transmitted and hence can hold up the transmission queue, e.g. a 1500 byte TCP packet takes 12msec to be transmitted along a 1Mb/sec serial link during which time all other packets have to wait. The following table shows further transmission time examples.

Example for TCP transmission time

Data Rate (Mbit/s)	Transmission time (536 byte packet)	Transmission time (1500 byte packet)
1	4 msec	12 msec
2	2 msec	6 msec
8	0.5 msec	1.5 msec
34	0.1 msec	0.4 msec
155	0.02 msec	0.07 msec

Transmission time can be improved by decreasing the transmitted packet size, but this increases some problems as described in the previous sections.

3.2.2 Bottleneck Delay

Bottleneck delay is caused by the data rate miss-match between the network (typically 100 Mb/s) and the serial link (typically 2 Mb/s). The reduction in data rate causes queuing delays at the interface, which in turn causes degradation in the quality of service (QoS).

Classification and scheduling schemes can be used to reduce delay.

Diffserv is an example of a classification scheme, which uses several classes to categorise and hence prioritise packets depending on the individual QoS requirements. The packets can be stored in queues according to prioritisation and are forwarded using a scheduling scheme.

Round-robin is an example of a scheduling scheme, which sequentially allows each queue to forward data in that queue, then moves to the next queue. Each queue can be individually weighted (known as weighted round-robin) so that only a limited number of packets (byte count) can be forwarded from that queue in any one cycle. Thus packets with a high priority classification will be placed in a queue, which has a high byte count so that it can reach the destination with the shortest possible delay.

3.3 Efficiency problem

This problem is caused by protocol overhead and mainly affects stream data (e.g. voice). Each packet of data carries a header and payload data. The header contains information to get the payload data to its destination, but it is not actually part of the data required at the destination. The transmission efficiency is the ratio of the payload data size to the total packet size (header and payload data). At each layer of network architecture, a new header is appended to the packet. A standard header has a fixed size and then often has the possibility to append options at the end depending on the protocol.

As an example, text data transmission (e.g. file transfer) will use the TCP/IP protocols. The default TCP payload size is 536 bytes and the standard headers (TCP&IP) together are 40 bytes, thus the transmission efficiency is 93%. In general the transmission efficiency will be good, providing that the payload size is much larger than the header size, but each network router has the possibility to change the transfer unit (packet) size to suit the network capability. If the transfer unit is reduced, the transmission efficiency will degrade, as shown in the table below.

Transmission Efficiency for different TCP Payload sizes

TCP Payload Byte Size	Transmission Efficiency
536	93%
400	91%
250	86%
100	71%
50	55%
20	33%

In comparison, stream data transmission (e.g. voice) uses the UDP/IP protocols. For VoIP applications UDP normally carries RTP packets and in general the transmission efficiency is quite poor due to the small payload to packet size ratio. Figure 6 shows an example for an IP stream (RTP) data packet, that has a payload size of 24 bytes and the headers (RTP/UDP/IP) together are 40 bytes, hence channel utilisation is only 38% efficient.

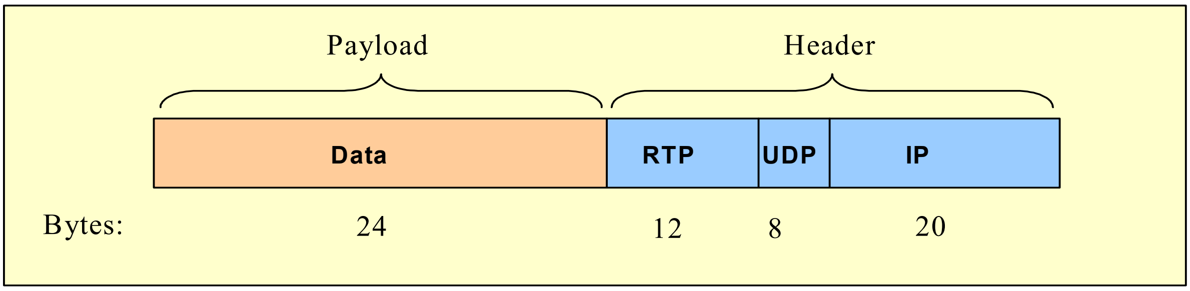

Figure 6: Typical IP packet for VoIP

In order to maintain good transmission efficiency the payload to packet size ratio must be kept high. This is achievable by reducing the header size or using large transfer units. Large transfer units will affect delay problems and various networks have different handling capabilities.

Header compression schemes are a way in which the transmitted header size can be reduced. The basic scheme relies on the fact that many header fields remain constant during a transmission session and therefore these fields, in theory, only need to be transmitted once during the session.

4 SOLUTION INVESTIGATION

4.1 Adaptive error control

4.1.1 FEC

To improve the channel link quality various error control techniques were investigated. For UDP traffic, adaptive FEC based on punctured convolutional codes were used. The amount of redundancy is controlled by the puncturing matrix. Without puncturing the maximum redundancy is introduced. Using this technique, the effective UDP/IP packet loss can be kept below a customisable threshold. The figure below depicts the resulting UDP/IP packet loss and the resulting effective data rate for a given convolutional code and a certain set of puncturing matrices.

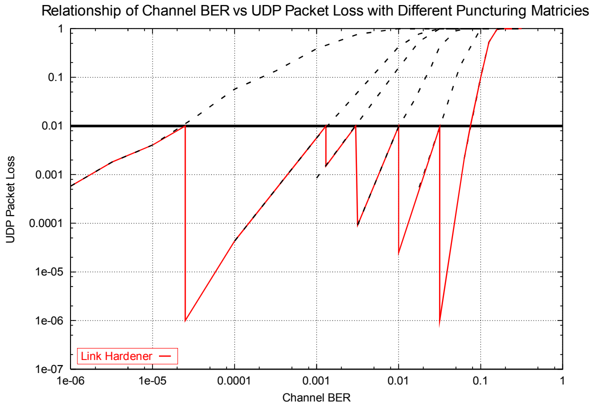

Figure 7: Relationship of Channel BER vs. UDP Packet loss with different puncturing matrices

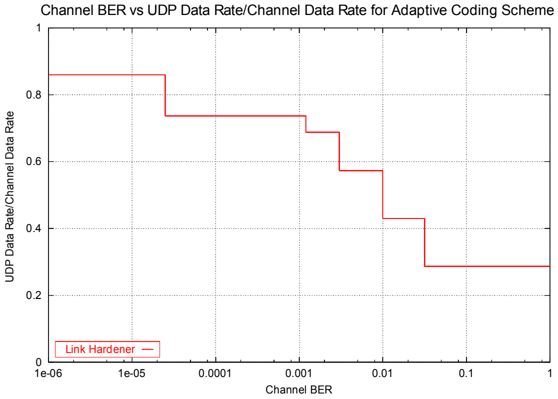

Figure 8: Channel BER vs. UDP Data Rate / Channel Data Rate for Adaptive Coding Scheme

Adaptive FEC (AFEC) can also be used to control the number of TCP/IP packet losses.

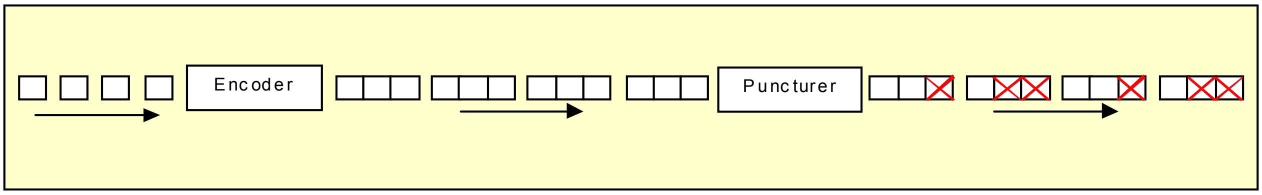

Figure 9: Adaptive FEC scheme, Puncturing a Rate 1/3 Mother Code to a Rate 2/3 Code

The following figure depicts the resulting throughput for various puncture matrices.

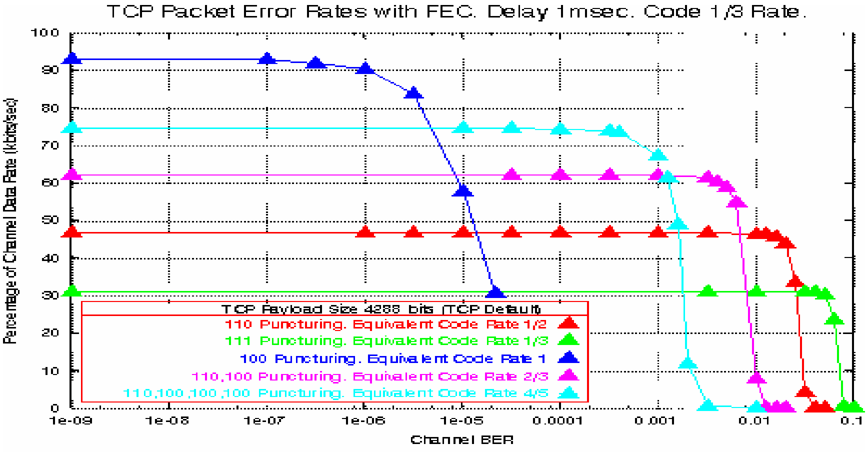

Figure 10: Data Rate vs. BER for 1msec delay *with adaptive FEC (Puncturing)*

4.1.2 HARQ

The problem with AFEC and TCP/IP traffic is that the ideal switching points between the individual puncture matrices is delay dependent. Hybrid ARQ (HARQ) on the other hand does not suffer from this drawback. Therefore, HARQ is the recommended method for reliable transmitting TCP traffic.

The results from a study show that for TCP a BER of 10^{-7} to 10^{-8} can be tolerated for the BSC. In traditional link design, a reference BER of 10^{-3} to 10^{-6} is usually used for the link calculation. The difference in *dB* from the traditional point of calculation to the point where the radio has a TCP critical residual error rate depends on the used modulation format and the applied FEC.

The Figure below depicts the BER curve of BPSK for the AWGN channel without FEC. With a operation point of 10^{-4} and a critical TCP value of 10^{-7} this difference is about 3 *dB*. Hence, to compensate the IP effect an additional 3 *dB* fade margin is required. This is equivalent to doubling the transmit power or reducing the range to 70% of its original value! If a operation point of 10^{-3} and a critical TCP value of 10^{-8} is used, the range expands even further. On the other hand, if the radio has a built in FEC, this range is slightly reduced.

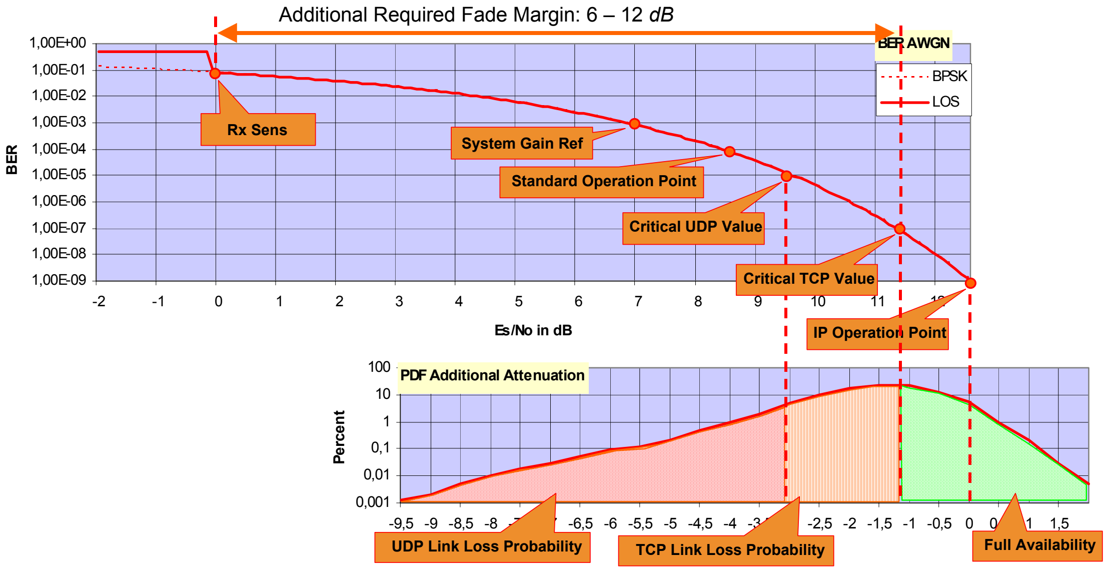

Figure 11: BER curve of BPSK for the AWGN channel without FEC

Instead of increasing the antenna size or the transmit power, AFEC and HARQ algorithm could be used to compensate the additional required link margin. For 3 *dB* only a marginal reduction of the effective data rate is required. Yet it seems that an increase in antenna size is the more cost effective solution, wherever this is possible.

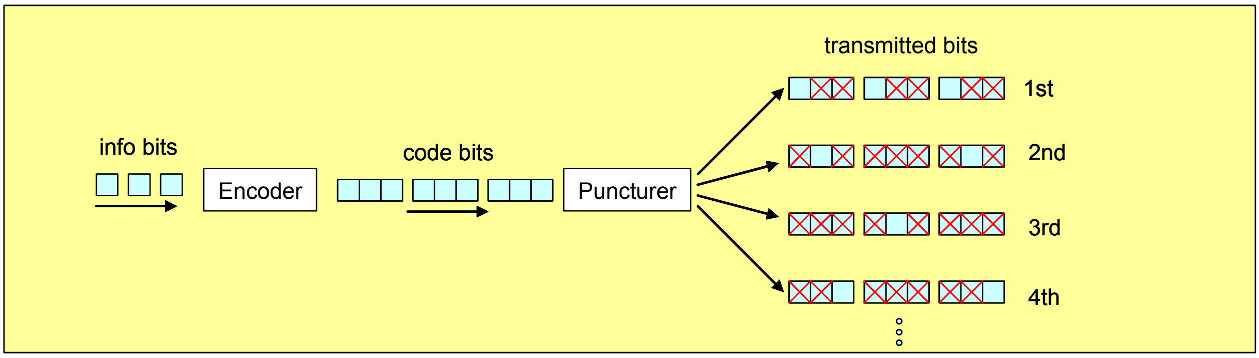

Figure 12: Hybrid ARQ scheme

4.2 Delay Management

Delay management controls the packet congestion at the network-serial interface. Packets are classified according to their priority and these priority assignments determine which queue to store a packet and hence the orders in which packets are transmitted. To optimise performance, several types of queues can be created and scheduled (forwarded) in accordance with traffic load and priority.

4.2.1 Classification

The classification scheme uses the packet priority to determine, which queue the packet should be stored in. Diffserv is a classification scheme, which uses several classes to categorise and hence prioritise packets depending on the individual QoS requirements.

4.2.2 Scheduling and Queuing

The Weighted Round-Robin (WRR) Scheduling uses several customised queues. The scheduler sequentially cycles though the queues (in a round-robin fashion), allowing each queue to transmit data. Each queue is customised with an allocated amount of bandwidth that it can use during a cycle. If a queue is empty, the scheduler passes onto the next queue in the sequence, which has data ready to send.

The allocated bandwidth is determined by the number of bytes (or packets) that can be forwarded by a queue. The number of bytes (byte count) per queue is user specified. The scheduler allows a queue to forward packets until the byte count is exceeded. Once this value is exceeded, the packet that is currently being transmitted is sent and then the scheduler moves to the next queue.

The weighted round-robin scheme prevents multiple large packets blocking the transmission link and guarantees bandwidth at congested points. In order to prevent unintended bandwidth allocation or long cycle delays, it is important to optimise the byte count based on each protocols packet size.

Four prioritised FIFO queues, i.e. high, medium, low and best effort priorities, are considered. The low priority queue will be served by a weighted round-robin scheduler, which will manage a secondary combination of four customised FIFO queues. The bandwidth allocation of the secondary combination is user configurable. See the following figure for a diagram of the proposed scheme.

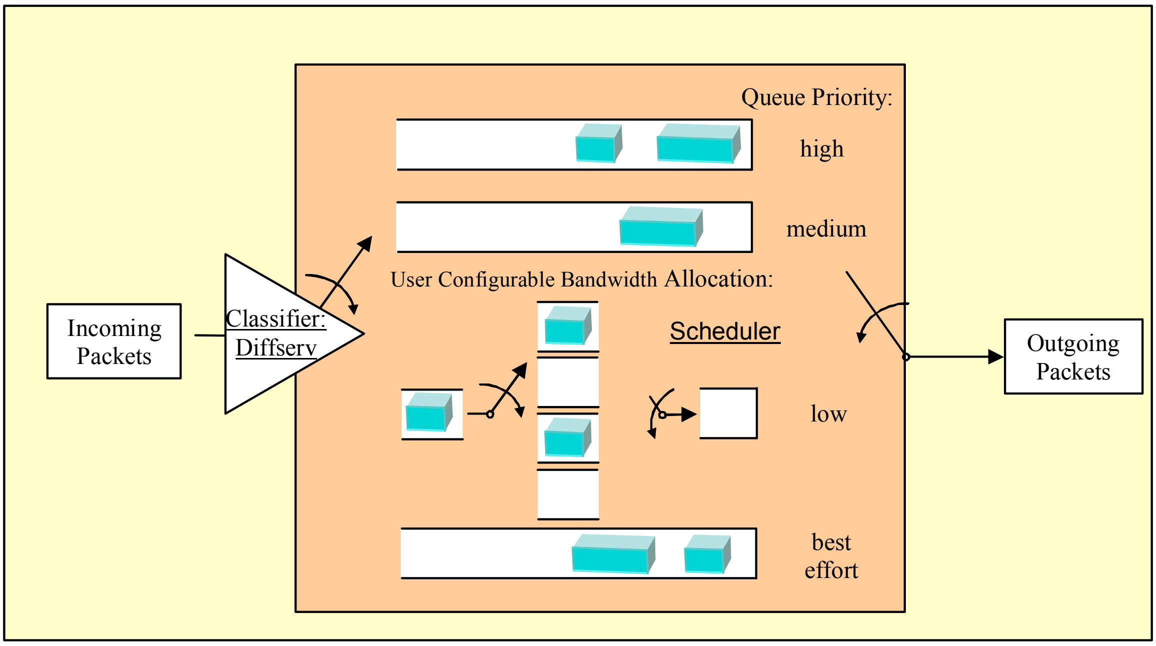

Figure 13: Scheduling and Queuing

4.3 Header Compression Algorithm

Header compression is used to reduce the total packet size and hence alleviate the problems associated with poor transmission efficiency, serialisation delay and error vulnerability. For example, an RTP/UDP/IP header can be reduced from 40 bytes to 2-4 bytes in this way.

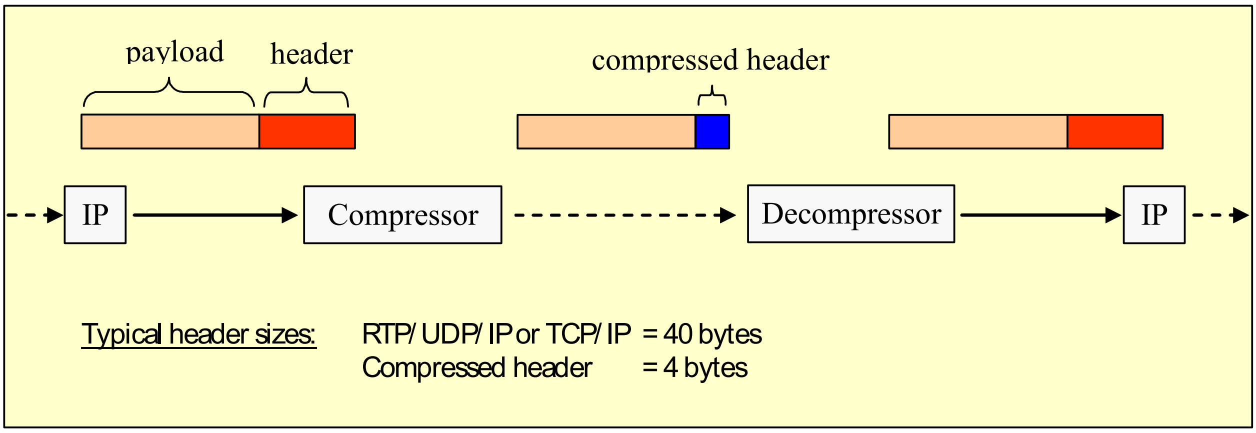

Figure 14: Header compression

The following section describes the general concepts and basic algorithm used for RTP/UDP/IP header compression. It is anticipated that this frame work can also be applied to TCP/IP transmission

Concept

The basic header compression scheme relies on the fact that many header fields remain constant during a transmission session and therefore these fields, only need to be transmitted once during initial session communication and then they can be removed from the following compressed headers. In this document, these fields are known as *Static Fields*.

Further reduction is also possible from the fact that, although several header fields change in every packet, the difference from packet to packet is often constant and therefore the second-order difference is zero. By storing the uncompressed header and the first-order differences, for a particular session, at the compressor and decompressor, all that needs to be communicated is that the second-order differences are zero. The decompressor can reconstruct the original header by adding the first-order difference to the stored uncompressed header. In this document, these types of header fields are known as *Predictable Fields*. An *Unpredictable Field* is a header field, which can change within a session context by an unpredictable amount.

As an example, Figure 14 shows a standard format IPv4 packet header. The minimum length is 20 bytes (as depicted), but there is also the possibility to have option fields appended at the end. The total length, packet identification (ID) and header checksum fields will normally change; all other header fields are *Static Fields* for a particular session context. The total length field is an *Unpredictable Field*, but it is made redundant, because the packet length is provided by the link layer protocol. The header checksum field is also an *Unpredictable Field*, but it is unnecessary, because firstly, it is protecting a header that is not actually transmitted and secondly, the compression scheme relies on good error detection from the link layer. Packet ID usually increments by one for each new packet, therefore it is a *Predictable Field*. Only changes in the packet ID need to be transmitted.

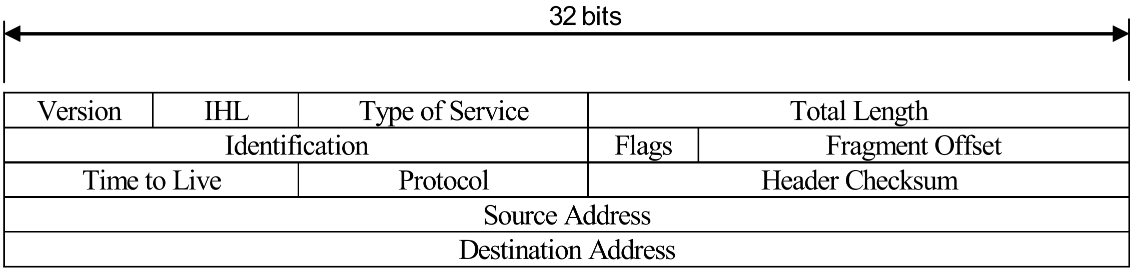

Figure 15: IPv4 packet header

5 SUMMARY

To counteract the problems associated with wireless IP transmission several solutions or a combination of several solutions were investigated.

Data classification and scheduling scheme at the data link layer to reduces the effect of delay and hence a good quality of service (QoS).

Packet header compression to maximise channel data rate utilisation.

Adaptive error correction methods to improve IP transmission quality and counteract the error vulnerability caused by wireless communication.

Link quality reporting to enable management of complex networks and observation of service level agreements.

Some algorithm are already implemented in network devices (radios, router etc.) like FEC, classification and scheduling scheme or compression algorithm and of course, many of the problems can be solved by a appropriate network or link design (if possible).

But if no exhaustive link calculation can be made or the network or parts of the network must be operated under extreme environmental conditions the combination of the investigated solutions will be helpful to ensure the required network availability with a adequate Quality of Service.