Noname manuscript No.
(will be inserted by the editor)

TCP/IP link layer error mitigation for MIMO wireless links

Ian Vince McLoughlin · Harsha Sirisena

Received: 06 June 2008 / Accepted: tbc

Abstract This paper investigates the effect of link-layer error mitigation support in MIMO wireless linking systems, and compares connection approaches between SISO and MIMO at different BER operating points substantiated by analysis of captured channels. In particular, this paper concentrates on a packet-based TDD approach, with a link-layer error mitigation scheme based on selective-repeat ARQ of segmented IP packets. Analytical expressions are derived for transfer efficiency over such a system, and simulation results presented to verify performance in terms of application delay experienced by users under various error conditions. This is repeated for SISO and for three alternative MIMO connection arrangements. Result show the degree of improvement available through the incorporation of link-layer error mitigation based upon the selective repetition of erroneous sub-IP packets, and in particular that presenting decomposed MIMO bit-pipes exhibiting diverse error conditions to the link layer, may be advantageous.

Keywords MIMO · networking · TCP · RLP · wireless linking

Ian Vince McLoughlin
School of Computer Engineering
Nanyang Technological University
Block N4, Nanyang Avenue
Singapore 639798
E-mail: mcloughlin@ntu.edu.sg

Harsha Sirisena
Dept. of Electrical and Computer Engineering
University of Canterbury
Christchurch, New Zealand 8020
E-mail: sirisehr@elec.canterbury.ac.nz

1 Introduction

Over recent years, MIMO technology has become an extremely popular research area due to its inherent promise of providing higher spectral efficiencies, and thus greater goodput over constrained radio frequency bandwidth to the end user. Various MIMO-based products are already available, and MIMO-based systems are increasingly likely to be deployed. It is also probable that many of these will provide wireless Internet connectivity.

This paper does not need to introduce TCP/IP as the preeminent Internet data conveyance standard. The convergence of wireless technology with Internet standards means that it is becoming rare to see new wireless systems that do not employ TCP/IP in some way. It is however well known that the TCP/IP protocol suffers from several disadvantages when applied to wireless links. This issue will be explored in Section 2.1, along with a discussion of methods to overcome these effects.

One such method, segmentation and reassembly with selective repeat (SAR-SR), has recently been applied to MIMO systems [21]. This, as with many other alternative approaches, could equally be applied to both SISO and MIMO systems, and has been shown to be useful for both.

The techniques of space-time block coding or MIMO coding, tend to turn multi-antenna wireless systems into units that, from a network-layer perspective, appear as a single data pipe with BER characteristics better than an equivalent SISO system. The actual BER experienced by that pipe is determined by the strength of coding and diversity (if any) to offset the effects of fading and interference in the transmission channel.

In a simple case, SISO BER varies largely in response to varying SNR at the receiver. A MISO system with two transmitters could use both transmitters at half power to achieve similar results, on average, to the SISO alternative. A better method would be to apply transmit antenna selection to always choose the antenna that results in the better BER at the receiver [19], thus employing spatial diversity to improve goodput [14]. An Alamouti or time-reversal space-time block coded (TR-STBC) system goes further by using those two transmitters simultaneously to achieve improved goodput [13].

The important point is that each of these examples, from a network layer perspective, provide a single data pipe for the conveyance of TCP/IP data, hiding the reality of multiple underlying physical or logical channels. Previous work tends to maintain the single pipe connectivity at the link layer and above [21].

This paper by contrast, considers the situation where the underlying MIMO data pipes are made available at the data link layer. Data transmitted by each pipe may then be conveyed as separate logical or physical bitstreams, either through decomposed MIMO channels, or through physically separate links. In such a scenario, TCP/IP data streams conveyed through different pipes will be subject to different degrees of time varying BER. This observation has also recently been considered for ARQ strategies [4] in 2×2 MIMO systems, and in the scheduling of MIMO ad-hoc networks [5], both with good results. Both effectively extend the antenna-selection strategy [19] into the link layer, utilising channel-state information (CSI) fed back from the receiver to transmitter to direct future transmission strategy, since both are directed by receiver signal-tonoise.

The intention of this paper is different. In this research, we do not rely upon CSI feedback, or propose a particular ARQ or scheduling scheme. Instead, we demonstrate, using both simulated MIMO systems and results from an experimental MIMO testbed, that treatment of decomposed MIMO channels with separate radio link protocol (RLP)-like error mitigators can have advantages over a combined mitigation approach under certain circumstances. In addition, rather than focus on any one particular MIMO system implementation, we extend the approach by noting that the characteristics of such systems are sufficiently described by perpacket BER time evolution for individual channels, to determine the benefit of using a particular connection strategy.

There exist three possible connection classes in MIMO systems that allow the link-layer to access decomposed channels. We explore and compare each of these, through simulations using real captured channel characteristics (which are themselves briefly analysed in Section 3),

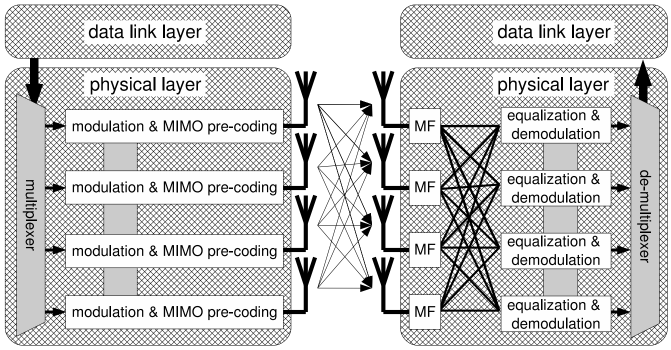

Fig. 1 Block diagram of 4×4 link data handling structure.

in addition to the more standard MATLAB simulation models.

The real channels were derived from a MIMO product which provides internet connectivity at distances up to 40 km (near-line-of-sight) at up to 288 kb/s within a 25 kHz UHF channel [21]. In fact the technique has also been validated for a 4×4 MIMO system with wideband microwave channel, DQPSK modulation, and decision feedback equalization [10]. Obviously such systems are inherently similar in the discussed aspect (namely multiple underlying physical or logical wireless connectivity) to widely deployed MIMO-based systems such as 802.11n. A block diagram of the underlying structure of a generic MIMO system is given in Fig. 1, showing the data link layer along with multiplexers feeding data through MIMO channels at the physical layer.

This paper begins by presenting the physical testbed in Section 2, then discusses channel characteristics in Section 3. From this, system simulations are developed in Section 4 and results shown in Section 5 before Section 6 concludes the paper.

2 Wireless MIMO Linking System

MIMO coding, channel estimation, equalisation and so on occur at the physical layer, with some form of data interface between this and the data link layer. The nature of this interface will form the basis of discussions in Section 2.2.

Consider a system using the link layer mitigation system (LLEM) derived in [21], namely packet segmentation and reassembly (SAR), with a selective per packet repeat ARQ mechanism. The LLEM in question was optimised for an average raw-channel BER in the region of 10^{-3}, and designed to deliver a corrected data rate for TCP packets of 10^{-7}. Based on these constraints, a sliding window protocol with selective request and resend operating on small packets, was adopted rather than a FEC-based scheme which might be better suited to alleviating higher experienced BER.

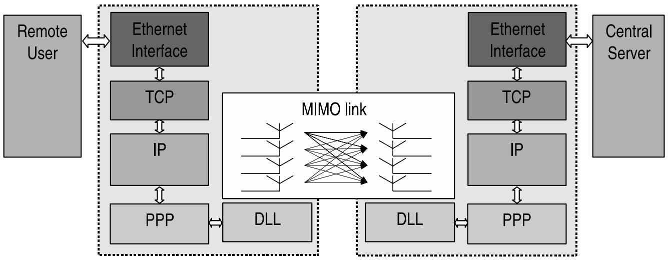

Fig. 2 Block diagram of the baseline system from a network services perspective.

Note that a similar segmentation technique is used in the radio link protocol [9], although in the system used here, the control scheme is a selective-repeat sliding window protocol based on error detection of segmented frames through a CRC byte, and of PPP encapsulated IP packets through the encapsulation CRC.

The block diagram of the system from a physical layer perspective is shown in Fig. 1, whereas a network services perspective of the system is provided in Fig. 2.

Both OPNET and MATLAB simulations were constructed of this basic system, with the aim of exploring how goodput and other factors change, based upon LLEM structure and interconnection arrangement. Results from this study will be presented in Section 5 after the system structure and channel characteristics have been discussed and defined.

2.1 Provision of wireless IP services

Wireless environments are generally characterized by higher error rates than for wired Ethernet transport. Because of this, the transport protocol usually used for Internet traffic on traditional Ethernet networks, designed and optimised for operation with low error rates, will be suboptimal. In a wireless environment, TCP tends to assume packet loss as being caused due to congestion rather than through bit-errors, and responds according to this incorrect assumption [18]. TCP invokes congestion control for wireless channel losses, thereby degrading overall system performance.

Many TCP modification schemes have been proposed for wireless systems, including split connection approaches [2][8], and the more widely used technique of link layer modification, including the Berkely Snoop protocol [3]. Most assume that the physical layer characteristics are transparent to higher layers, although the snoop protocol advertises a cross-layer approach which employs a snoop agent at a wireless base station to decouple the potentially more numerous Ethernet hops from the single wireless hop, to 'hide' the characteristics of the wireless segment from the wider network, enable fast retransmits, and allow for explicit congestion notification when congestion does really occur.

In the scenario under consideration, a PPP session is established between two endpoints prior to the data transfer phase and maintained until the data transfer is complete. The data link layer is responsible for providing a reliable transfer service in the sense of low BER traded off against variable latency according to the requirements of the higher layer protocols. But at the same time it must hide errors on the wireless link from the higher layer protocols in a similar way to Snoop. In order to achieve this, the LLEM system provides facilities for local segmentation/reassembly of the datagrams, error detection and recovery [21]. To provide reliability to the higher layer protocols it also offers in-order delivery of datagrams and uses active queue management to minimise latency and buffer memory requirements. It is designed to be transparent to both TCP and IP.

At the TCP layer, a senders slow start algorithm uses TCP segments to control the flow of data. This may require the segment size to be large to quicken the data flow. But this kind of flow is disadvantageous when used over a wireless link, since the larger the packet, the higher the probability of erroneous reception.

Segmentation and reassembly (SAR) implemented at the link layer allows segmentation of packets arriving from the higher layer into small sized link layer radio frames which are then transmitted over air more efficiently. This efficient transmission of the radio frames is managed by the link layer error recovery mechanism which uses a Selective-Repeat ARQ algorithm [1]. In such an algorithm, only the packets which are received in error are to be retransmitted by the sender. It is possible to implement this as either an ACK-based or a NACK-based system, but in this case the NACK-based system was selected due to being more straightforward to implement. The efficient SR-ARQ, working in the data link layer, retransmits erroneous frames to prevent frequent invocation of TCP retransmit strategies.

Although SR-ARQ proves to be efficient, it does not provide for in-order delivery of all frames. Out-of-order frame reassembly would result in a corrupted IP packet, potentially triggering the TCP fast retransmission and recovery algorithm – definitely an unfortunate consequence in terms of performance. To overcome this, radio frame reordering is performed in the receive link layer such that only in-order frames are delivered to the higher layer. This consequently requires frames to be indexed, and thus imposes a small packet length overhead.

Fig. 3 shows part of a PPP packet subdivided into two wireless transmission frames of fixed size 32 bytes,

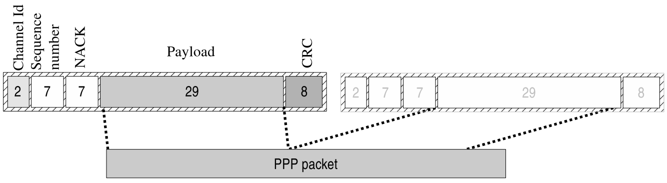

Fig. 3 Over air radio frame transmission format

comprising 29 bytes of payload and 3 bytes of header/tail. Thus a 580 byte IP packet encapsulated by PPP will be segmented into 21 radio frames. A 7-bit sequence number accompanies each frame, and a NACK field allows frames to be units of negative acknowledgment, as well as transmission or retransmission (where necessary).

Frames queued for transmission are buffered in the link layer, then copied and stored into a retransmission buffer. The sending side maintains a transmission window which is incremented when a frame is transmitted and shortened by one when the receiver receives the frame correctly. At the receiver, the received frames are checked for error. Uncorrupted frames are passed on to a reassembly buffer. However, if any of the frames is received with an error, the receiver replies with a NACK, specifying the sequence number to resend.

When the sender receives a NACK, it extracts the corresponding frame from the retransmission buffer and loads it into the transmission queue, with higher priority than the current transmitting frames. Counters track retransmission attempts and limit the number of retries for each corrupted frame, after which frames are dropped.

Like TCP, the link layer at both the sender and receiver maintains a sliding window, identified by set of three variables, continuously updated to track outgoing and incoming frames and to identify the current window limit. On the receiver side, the receiver window checks arriving frames for sequence number. If the sequence number falls within the window and it is not already stored, then it is accepted into the re-sequencing buffer.

Training frames are exchanged by both ends of the link to maintain synchronization, and to calibrate the wireless channel equalizers. This training exchange is repeated periodically, triggered when a link is lost, when performing error detection and recovery, and after a large string of errors indicating ageing channel estimates.

In the system described, IP packets are thus conveyed over the air interface by being split into smaller radio frames which are managed by a sliding window protocol. Control frames are redundantly Hamming coded with parity to give two copies of each control word per frame, and a CRC within each of the Hamming coded blocks. Only control frame halves with zero unrecoverable errors are accepted.

These measures bring final control frame error rate down toward 10^{-12} for a raw BER of 10^{-4}. Control frames are sent on change, and always repeated multiple times. A single missing control frame will result in a single data radio frame being corrupted, and necessitate resend of at worst two IP packets if the transition between these happens to occur within the lost block. Since the error rate for control packets is so low, and the effect of erroneous reception so localised, the simulations and test conditions described in this paper assume in each case that control packets are communicated error-free, thus simplifying the model construction.

The system is essentially symmetrical, modelled as two computers connected to LANs, exchanging IP datagrams over a wireless link bridging the two LANs.

2.2 Physical layer connectivity

Although previous work applied the LLEM system to a MIMO wireless link, the LLEM system itself operated on a single data pipe [21]: the physical layer presented a single data interface to the data link layer, and underlying MIMO processing within the physical layer similarly reduced the multiple antenna transmission arrangement into a single composite data channel (represented by the multiplexers in Fig. 1) albeit with improved BER.

However there is an alternative arrangement possible where more than one data pipe is made available by the MIMO hardware (the term ‘MIMO channel’ will be used in this context as a shorthand to describe either separate physical communications links, or separate logical communications links such as eigenchannels) [22]. In a single data pipe arrangement, individual bits from the data stream would be multiplexed, coded or water-filed over underlying logical or physical MIMO channels. The effect of differential error rates exhibited by different MIMO channels would be spread over all transmitted bits.

A multiple data pipe arrangement, by contrast, allows different bit streams to experience the different underlying error rates. Multiplexing can then be employed at a radio-frame or IP packet level, with alternating frames or packets being conveyed in their entirety over the same MIMO channel. Thus at any one time, several radio frames may be conveyed bitwise in parallel. Recently, this idea has also been explored by ElBatt [5], directed by receiver SNR.

We argue that typical MIMO multiplexing, coding, water-filling and so on take no account of TCP packet structure or radio frame structure, and hypothesise that taking the radio frame structure into account when multiplexing packets across multiple channels can improve performance in some circumstances.

Two connectivity alternatives exist, based on whether the MIMO system can provide a single data pipe, or can provide separate data pipes to the link layer. The first allows the data-link layer to operate as if a SISO system were in use, but one in which the capacity and error rates that the system provide by the physical layer have benefited from MIMO coding improvements. The second allows the data-link layer to multiplex datagrams between each of the data pipes presented to it at the physical layer for transmission, and must of course then reassemble these during reception. Furthermore, if the underlying MIMO techniques being used provide separate channels with differing error or rate properties, then this must be compensated for in the data-link layer rather than (or maybe as well as) in the physical layer.

Three mitigation arrangements are considered on the basis of these two connectivity alternatives, namely the single-, dual- and paired-LLEM arrangements.

The single-LLEM structure, shown in Fig.4a applies LLEM by simply multiplexing symbols on to the available transmit antennae in a round-robin fashion. A single sliding window transmit buffer, and reassembly buffer are provided in this system. Radio frame resends are handled by the next available transmission slot. This basic structure was paired up to create the dual-LLEM system of Fig. 4b in which independent radio link protocol blocks are employed for each channel, and the transmission multiplexing takes place prior to the link-layer processing. This means that bytes from each IP packet are passed to the independent processing units to form, and transmit, radio frames. Otherwise, processing is performed in much the same way as in the single-LLEM structure, with NACKs are handled on a per-channel basis.

The final arrangement, the paired-LLEM system, employs a single transmission processing block for all antennae, but maintains the per-antenna transmission buffer of the dual-LLEM system. For a two-channel example, radio frames are transmitted over channel 1 or channel 2 in a round-robin fashion based on slot availability. Retransmissions due to NACKs are always transmitted on the next available slot irrespective of channel. Thus a transmission error over channel 1 may result in a retransmission on channel 2. A block diagram is shown in Fig. 5, where the output multiplexing for both transmit channels is indicated at the output of the transmit processing blocks.

All are MIMO systems, designed to out-perform their SISO counterparts in various ways. Identical low-latency

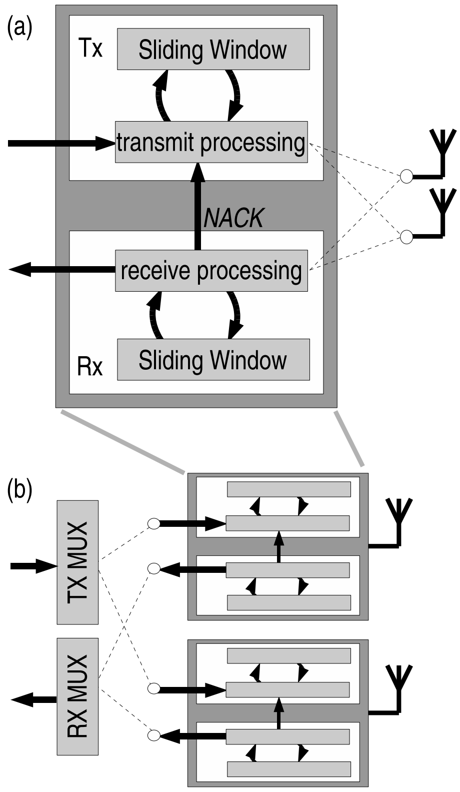

Fig. 4 Single (a) and dual (b) LLEM structures for two-channel processing

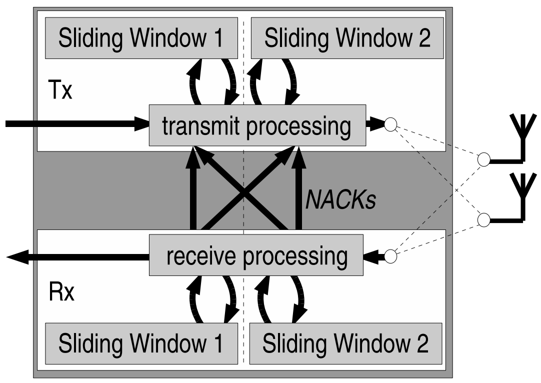

Fig. 5 Paired LLEM structure for two-channel processing

SR-ARQ algorithms with SAR were implemented just below the PPP layer at either end of the wireless link for each arrangement. The difference lies in the underlying connectivity, and the number of re/transmission buffers required. All structures were simulated in handling TCP/IP traffic over various wireless channels.

3 Channel characteristics

Multipath signals from each of N transmit antennae to each of M receive antennae will have different characteristics which evolve over time - usually slowly in a fixed system, and more rapidly where mobility is involved. Studies performed for antenna selection reveal that channel characteristics are often unequal [14] - in a two channel system, one channel typically has significantly lower BER than the other. In fact this is one of the motivating factors for considering antenna selection.

In a MIMO system, at some particular instant it is likely that all channels are exhibiting different levels of BER, and can be ranked in an order of goodness, and that the ranking changes from instant to instant. The speed at which the ranking changes is related to the stationarity of the channels.

3.1 Simulated channels

In a simulation system, the path from transmitter to receiver may be modelled by a combination of additive white Gaussian noise (AWGN) at given signal-to-noise ratio (SNR) and some flat or frequency-selective fading. In the absence of inter-symbol interference (ISI), inter-carrier interference (ICI), phase offset, equalization error or other distortion, AWGN alone contributes to received errors, such that the probability of error in a received bit for a BPSK modulated signal is shown in (refeq:Pe) for a ratio $\gamma = 2 \times Eb/No$ of energy per bit to spectral noise density [11]:

$$Pe = \frac{1}{2} \times \; erfc(\sqrt{\gamma}) \qquad (1)$$

AWGN is added to transmitted symbols in a simulation which are then detected, resulting in occasional erroneous reception. A more realistic scenario for real systems is fading channel modelled as Rayleigh probability density function in (2) for fading envelope variance σ^2:

$$p(r) = \begin{cases} (r/\sigma^2)e^{-r^2/2\sigma^2} & r \geq 0 \\ 0 & r < 0 \end{cases} \qquad (2)$$

In practical terms such a channel is generated using a normally distributed selection of random real and imaginary components, scaled to ensure that the channel causes neither energy loss nor gain, and spread over a channel delay length (effectively the ringing time of the channel response). The probability of error for BPSK modulation is then given by (3) over a slow fading Rayleigh channel [11]:

$$Pe = \frac{1}{2} \times \; \left(1 - \sqrt{\frac{\gamma}{2+\gamma}}\right) \qquad (3)$$

In a typical simulation, randomly generated channels are created and maintained for multiple symbol periods, before a new random channel is created (slow fading). In some cases there may be many thousands of symbols transmitted over a particular channel . In other cases, the channel may be updated once per symbol, or even faster (fast fading). Either way, a great many repetitions of symbols will need to be transmitted over such a system before a statistically relevant measure of BER can be obtained for a given SNR, especially at high values of SNR.

3.2 Practical channels

Even though channels may have a Rayleigh probability when averaged long enough, the short-term excursions of different channels differ markedly from the mean. In fact this effect is evident in simulations, and is the reason why lengthy simulations are required to obtain smooth BER vs. SNR plots.

In part, it is precisely this short-term variability in experienced BER, causing the unwanted spikiness of BER vs. SNR plots, that antenna selection algorithms hope to take advantage of. This effect is not confined to simulations: it occurs in real channels. In both cases, a long-term averaged channel BER plot hides significant short-term variability.

Furthermore, although it is recognised that channels change either fast or slow with respect to transmitted data [7], any real wireless system would generally experience periods of rapid change interspersed with periods of slow change [13]. Equalization and channel estimation mechanisms would have been designed to cope with a certain rapidity of channel change (a worst case), but would be over-specified for typical usage. In general, the length of a transmission frame prefixed with training symbols would be set to ensure that channel estimation occurs just frequently enough to match specified rates of channel variation. Any variation more rapid than this would be treated as excess noise, and thus contribute to increased BER.

In a static multi-channel system, channels will undergo largely independent slow variation in channel characteristics due to environmental changes, temperature effects on electronic components and so on, with occasional rapid excursions often due to physical phenomena: the most usual example experienced by the author in an office environment results from a person walking past an antenna array. Sometimes this will cause strongly correlated dropouts on all channels, whilst at other times some channels will be affected more severely than others. Whatever the reasons, such variations exist, and are examined in the following subsection.

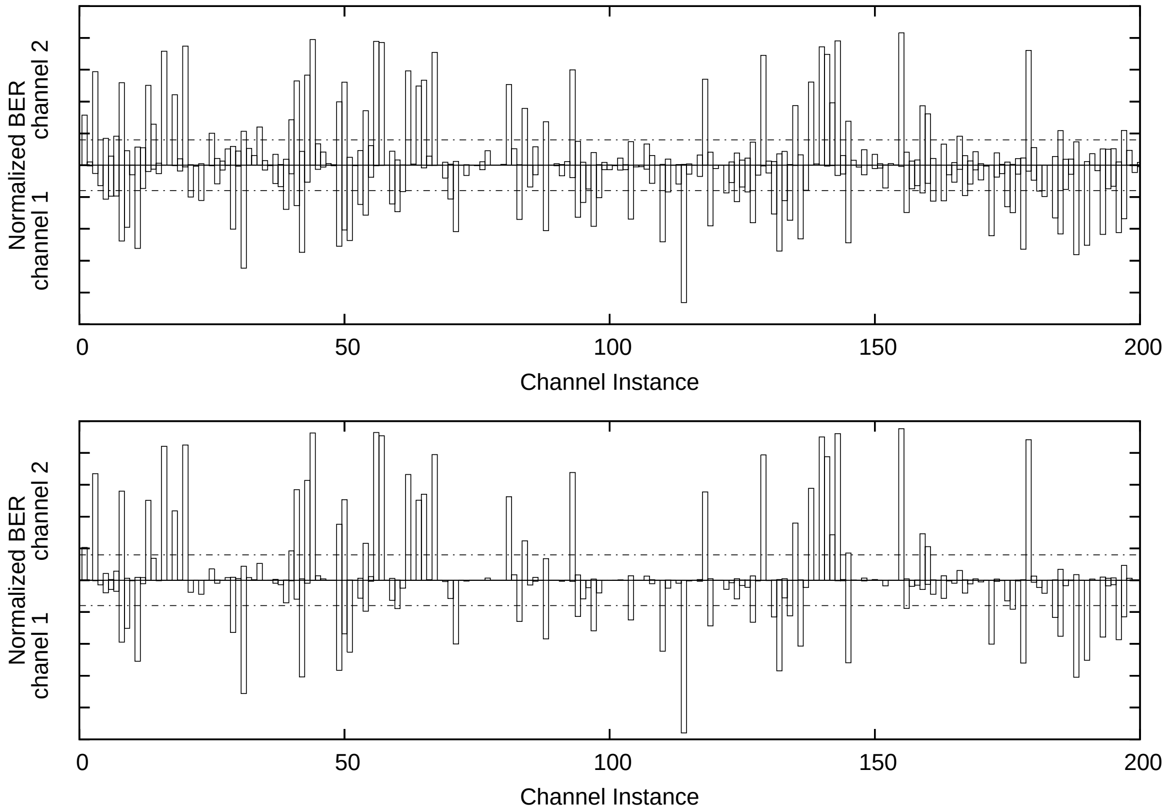

Fig. 6 Normalized per channel BER for real captured packets at 0 dB SNR (top graph) and 5dB SNR (bottom graph) with first and second channel plotted back-to-back in each case, the centre axis denoting zero BER, and each y-axis extreme denoting maximum BER. The mean BER has been overlaid as a dotted horizontal line for each channel.

3.3 Captured channels

MIMO channel sounding tests were conducted to explore the statistics of individual channels for three systems: an indoor microwave link employing the Alamouti-like time-reversal space-time block coding (TR-STBC) method over a 2 MHz wide microwave channel [15], a 4x4 MIMO multi-variate decision feedback equalized digital link [10] operating at 2.45 GHz with 17 MHz RF bandwidth over short urban distances, and a 30 km near-line-of-sight narrowband UHF MIMO link [14].

Scatterplots from the 4x4 MIMO system in [12] clearly indicate that irrespective of average performance, every channel had at least a certain duration of good performance. Furthermore, the 2x1 TR-STBC system channels statistics were examined to determine the BER excursions of the individual channels. The results are shown graphically in Fig. 6 where the BER for symbols transmitted through each of the two channels at both 0 dB and 5 dB SNR are plotted. The mean BER per channel at each level of SNR is shown as dotted horizontal lines. What is important to note is firstly that the distribution of BER at each particular channel time instance is highly variable about the mean, and that secondly, apart from some instances where both channels exhibit high BER simultaneously (presumably due to correlated interference), most of the time there is one notably good channel and one notably poor channel.

To explore further, the histogram in Fig. 7 classifies BER results at both 0 dB and 5 dB into 50 bins of difference between the channels. While it can be seen that often the normalized difference between the two channels is within 2% (the high central peak), a very significant proportion of the results lie outside the 2% range.

3.4 The impact of these characteristics

Clearly, both these captured channels exhibit significant short-term excursions even where their averaged performance is identical. We will demonstrate later that these dynamic characteristics lead to advantages in employing alternative connection arrangements for MIMO LLEM systems.

It is also instructive to note that, for the purpose of TCP throughput and goodput calculation, it is necessarily only to consider the degree and time evolution (statistics) of BER over a channel, rather than the underlying causes of those BER differences. Thus, since

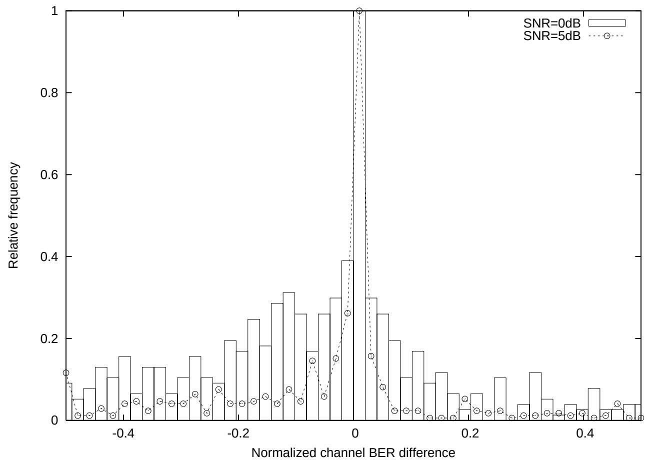

Fig. 7 A histogram showing the per-channel difference in normalized BER performance for 600 complex channel matrices.

the present paper is concerned with the effects on TCP of differing BER over MIMO channels, the discussion of channel characteristics henceforward will be described in terms of instantaneous BER rather than in terms of SNR.

This implies no loss of generality, since almost all published MIMO methods describe in detail their relationship between SNR and BER. In fact, this will allow any MIMO or space-time coding technique (from SNR vs. BER curves/equations), to easily be compared to the LLEM system in this paper.

4 System Evaluation

We will first describe a single-LLEM baseline system, before extending to the MIMO case. The algorithm of Section 2 was implemented just below the PPP layer at either end of a wireless link using ARQ operating on higher layer packets [6]. Error detection and recovery was performed on a per frame basis.

4.1 Analytical Derivation

In the LLEM system, SAR is responsible for packet segmentation while ARQ takes care of local error recovery. PPP packets are segmented into N fixed size radio frames of length P^{seg}. Each frame is prepended with its own header size of H^{seg} by the link layer to complete the frame $P^{seg} + H^{seg}$ to be transmitted over air. The Packet Error Rate (PER) is dependent upon frame size and BER [16]:

$$PER = 1 - (1 - BER)^{(P^{seg}+H^{seg})} \tag{4}$$

For an optimal ARQ protocol such as SR-ARQ, efficiency depends on packet size as well as error rate:

$$\mu = \left\{ \frac{P^{seg}}{P^{seg} + H^{seg}} \right\} \left\{ \frac{1}{(1 - PER)^{-1}} \right\} \tag{5}$$

Applying a semi-persistent version of ARQ, we allow a maximum r retransmissions per radio frame. This implies that error free delivery cannot be guaranteed [18], but in practical terms, is related to TCP timings. Assuming a random occurrence of radio frame errors, the probability that a radio frame is received after r retransmission attempts is $(1 - PER^r)$. This is a worst case scenario - in reality, radio links suffer from fading effects and dropouts, a non-uniform error distribution - generally a better scenario for a packet-based ARQ system.

The probability of N radio frames corresponding to a higher layer PPP packet correctly received after link layer recovery is $(1 - PER^r)^N$. Therefore the PPP packet discard rate can be computed as $1-(1-PER^r)^N$.

Since a packet in error is retransmitted up to r times, the average number of retransmissions required per packet will be $(1 - PER^r)/(1 - PER)$ and the average data throughput for the N packets would then be:

$$N \left\{ \frac{1 - PER^r}{1 - PER} \right\} (P^{seg} + H^{seg}) \tag{6}$$

4.2 Investigative method

The original LLEM simulation system was implemented using OPNET version 11.0 to evaluate performance over lossy wireless channels [20], the simulation model being similar to that shown in Fig. 2, and was essentially symmetrical between central network and remote user endpoints. Further simulations were performed with MATLAB/Octave.

The physical link was responsible for implementing SAR and SR-ARQ. At the central network endpoint, servers connect to the link through 100 Mbps Ethernet with negligible delay and no errors assumed. The wireless channel had a user bandwidth of up to 288kbps with negligible physical propagation delay and channel characteristics described in terms of average BER.

Test traffic was generated with a random distribution over several thousand runs for each parameter setting within the investigation: and the random seed was reset each time a parameter was changed, such that repeated tests were based on identical baseline conditions with the pseudo-random simulation data progressing in sequence over the many simulated hours and days of testing [17].

Table 1 Simulation Parameters

Attribute	Value	
TCP Flavour	Reno	
Data Rates	288, 256, 224, 192, 160 kbps	
BER Range	10^{-7}, 10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}	
Radio Frame	32 Bytes	
LL Sliding Window	61 Bytes	
Max. No. of Retx.	5 radio frames, 15 TCP	
Application	**HTTP**	**POP**
Object size	Lognormal	Lognormal
mean	10.5 KiB	22.7 KiB
std. dev.	24.5 KiB	200.3 KiB
max.	2.0 MiB	none
Objects/page	Pareto	n/a
mean	5.64	n/a
max.	53	n/a
Interarrival time	exponential	exponential
mean	30 s	360 s

The model was used to investigate the performance of the system for given channels, and thus derive LLEM system parameters including optimal radio frame size, sliding window buffer size, maximum allowed retransmission attempts allowed per radio frame, plus TCP parameters such as maximum segment size and buffer provision [21]. For a single-pipe MIMO system (which appears to the link layer like enhanced-performance SISO), the main simulation parameters are listed in Table 1.

5 Results

The effect of BER on user data rates for various channel BERs was investigated. In particular, the criteria of latency and throughput were examined for different traffic types, such as HyperText Transfer Protocol (HTTP) and Post Office Protocol (POP).

5.1 single-LLEM

Simulated data transfers were made from a server at a central office location to the remote user over a wireless channel for cases both with and without link-layer mitigation in the presence of wireless channel BERs of 10^{-5} to 10^{-3} and channel capacities of 160 to 288 kbits/s. The mean end-to-end data transfer times are plotted in Fig. 8. The figure shows POP and HTTP traffic for both raw TCP performance (no LLEM), and where a single-pipe mitigation scheme is enabled (with LLEM).

For both transport methods, the transfer time is given in seconds, and relates the time taken for the entire data set of items in the evaluation to be transferred. When errors occur, delays result from radio frame failures and retransmissions and, in the worst case, retries reaching the maximum allowed and thus being dropped. Naturally higher channel capacities will result in faster transfers, and it can be seen that in most cases the transfer times was higher without link-layer mitigation – a clear indication of the effectiveness of the LLEM scheme.

It is evident that page load time for the HTTP tests at the higher BER is excessive: approaching 100 seconds, but is improved significanty through the use of single-pipe LLEM.

These tests underscore previous work which has described the benefits of link-layer error mitigation in general [9] and the particular MIMO LLEM scheme [21] employed. In this paper, we now investigate a new interconnection arrangement which is available for this LLEM scheme (and indeed other similar schemes), namely to make availabile physical sub-channels to the link layer.

5.2 single, dual and paired-LLEM

When physical sub-channels are made available to the link layer, the connection arrangments of Section 2.2 become possible. In order to explore the effect of these arrangments, simulations were constructed in which 500 MiB of data was transmitted through each structure at maximum rate, and conveyed over two channels experiencing equal mean BERs. Important metrics for this fixed-data simulation were the goodput and overall number of bytes transmitted. Up to five radio frame retransmissions were allowed, and IP packets were always discarded if any of the constituent radio frames were received corrupted. The system did not implement early-termination of lost IP packets, in other words the remainder of an IP packet will continue to be transmitted even if the first radio frame has been lost (a more efficient implementation would be to detect IP delimiters, and immediately drop the transmission of an IP packet that has already been corrupted).

Goodput reduction is thus primarily through the mechanism of IP packet loss, which in turn relates to radio frame retransmission exceeding the maximum allowed attempts. The second metric, the overall number of bytes transmitted over air, is an indication of how 'hard' the system must work to deliver its payload, and determines overall system efficiency in the presence of BER.

Goodput versus mean BER is plotted in Fig. 9 which shows that at low BER, each structure achieves a near-perfect goodput, meaning that all IP packets were trans-

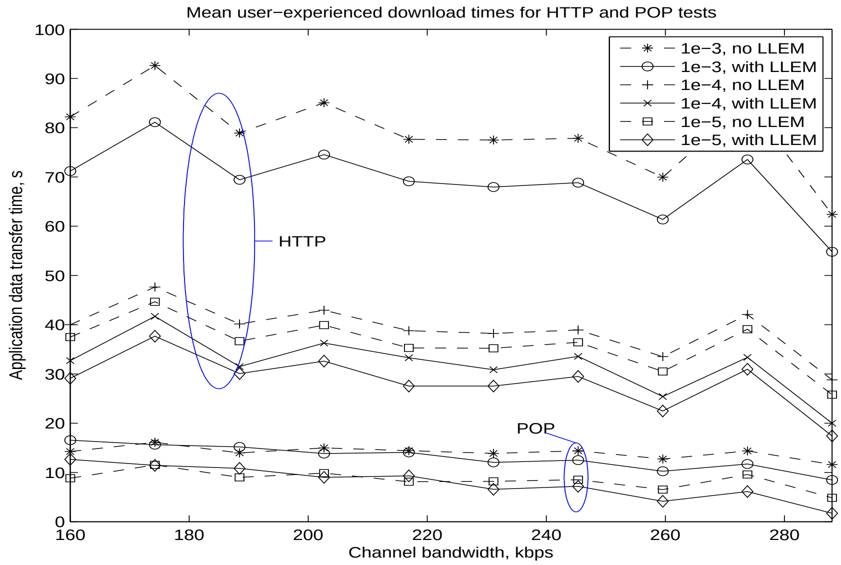

Fig. 8 User experienced HTTP page load and POP download times for various raw channel data rates and three channel BER conditions (for HTTP, only two shown for POP to preserve clarity) both with and without LLEM support.

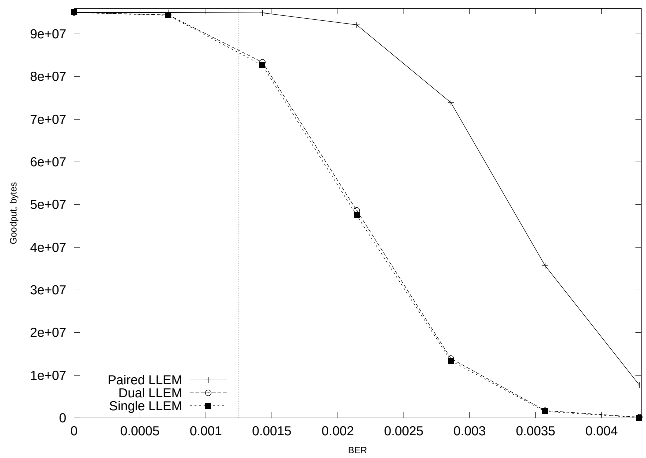

Fig. 9 Goodput in bytes for the single, dual and paired LLEM structures plotted against bit error rate. A vertical line denotes the fixed BER of 0.00125 that has been chosen for several of the subsequent plots.

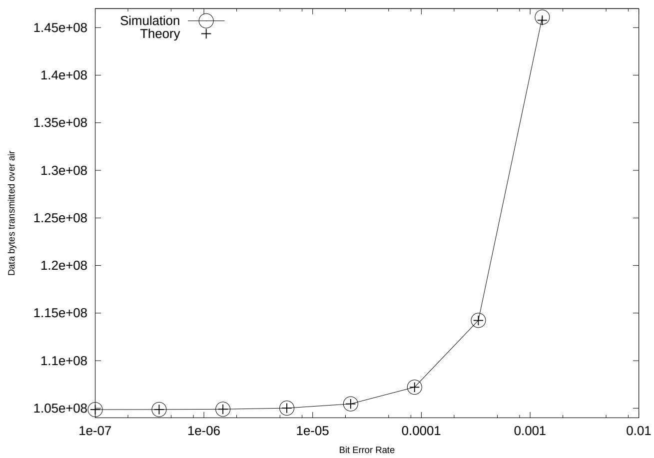

Fig. 10 Total over-air transmission in bytes plotted against bit error rate, showing both the simulated and theoretical results.

mitted successfully. However as BER increases, the number of lost IP packets increases, and hence goodput reduces. It can be seen that the performance of the dual-LLEM improves slightly over the single-LLEM system, but the paired-LLEM method performs significantly better. The primary reason for the improvement is the spreading of resent radio frames across both channels, and the consequent rapid servicing of NACKs in the first available transmission slot.

When considering the total number of bytes transmitted over air, Fig. 10 shows how the system suffers from increasing BER conditions: the number of resends, and thus the total number of bytes transmitted in attempting to transfer the entire set of test data, increases smoothly with increasing BER. This figure also plots the expected total number of bytes transmitted based upon the average number of retransmissions given in Eqn. (6). Note that the throughput results for single, dual and paired LLEM arrangements are identical (errors are randomly and independently distributed): the real difference between the three will be seen later in the goodput results.

In real MIMO system deployments, error rates vary between channels (Section 3). So exporing the MIMO-LLEM goodput performance at fixed BER (of 0.00125, chosen to clearly differentiate the performance of each

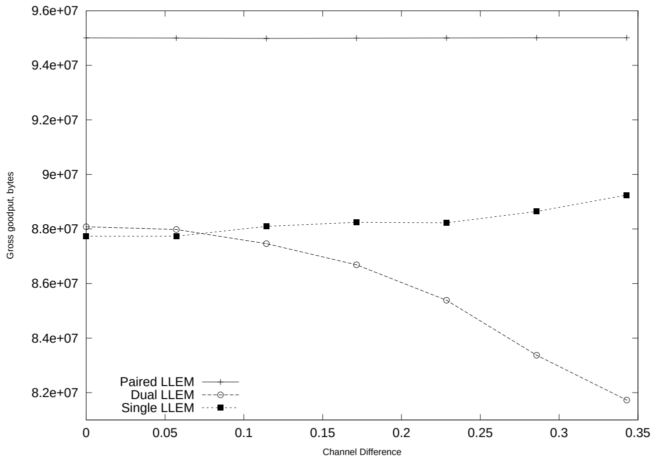

Fig. 11 Goodput in bytes for the single, dual and paired LLEM structures plotted against channel difference at BER=0.00125.

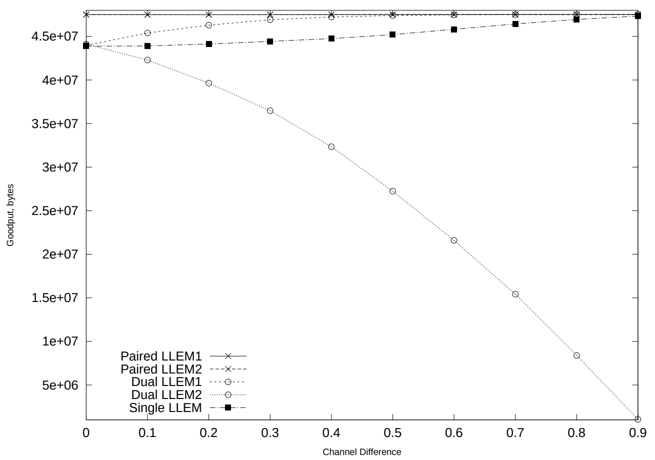

Fig. 12 Per-channel goodput in bytes for the single, dual and paired LLEM structures plotted against channel difference at BER=0.00125.

structure as seen in the plot of Fig. 9) with respect to instantaneous per-channel BER difference, yields Fig. 11). This shows the effect of exploiting the unequal attribution of instantaneous BER across channels, which otherwise may have equal long-term average BER, on the three LLEM connection arrangements.

In fact, Fig. 11 is very interesting: whilst the paired-LLEM structure remains the highest performing arrangement at 0.00125 BER, the slender advantage of the dual-LLEM structure over the single-LLEM structure is easily lost as the channel BER differential increases above about 8%. In fact, such a situation is likely to occur in practice (see Fig. 7 in Section 3). It can thus be predicted that no advantage will be gained through employing dual-LLEM over single-LLEM in a practical system.

To explore further, the differential performance is plotted in Fig. 12 on a per channel basis. This compares each structure and plots the different goodput of individual channels as the channel error difference increases. The dual-LLEM structure exhibits very different levels of per-channel performance as channel difference increases: typically one channel starts to experience high error rates while the other channel performs well. However, the paired-LLEM structure, which cross-resends erroneous NACKed radio frames, has a per-channel performance which is relatively constant. In effect, the better channel is compensating for the poorer channel as and when necessary (such as when a poorer channel radio frame is NACKed), but otherwise is able to experience excellent goodput. Overall, performance is improved through this mechanism.

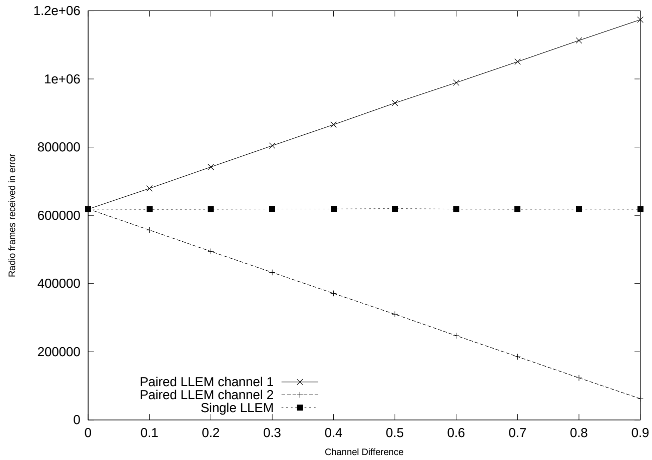

Fig. 13 Radio frames received in error per-LLEM for the single and paired LLEM structures plotted against channel difference at BER=0.00125.

Interestingly, the same trend is notable in the figure for the single-LLEM structure at the very high channel difference extreme where one channel totally fails to operate. Both structures are able to continue operating by transferring retransmits to the *good* channel.

This is illustrated further in Fig. 13 showing the number of erroneous radio link frames for each of the LLEM units in the paired-LLEM structure and for the single-LLEM structure. Clearly, the error performance is, as expected, linear with the differential BER exhibited per channel, but the single-LLEM structure experiences the average BER. The performance advantage of dual-LLEM between the two extremes of channel difference is tied to the ability of the structure to best exploit one good channel when available.

6 Conclusion

This paper has investigated how real time TCP/IP data can be delivered over a MIMO link with the help of local error recovery mechanisms implemented in the wireless interface data link layer. By making the error recovery process semi-persistent, operating on segmented radio frames, it is possible to quickly recover from errors without invoking TCP layer error control.

The MIMO case was investigated by firstly analyzing channel statistics likely to be found in a practical implementation, and then applying these to simulations using three different structures of basic link-layer error mitigation connectivity. These differed primarily in the location of the per-channel multiplexing operation (in other words the visibility or invisibility of independent bit-pipes to the link-layer error mitigation unit), and in the handling of radio frame retransmits triggered by receipt of a NACK.

Investigations into the effect of bit error rate, and differential performance of the data transmission channels in the MIMO system, indicate the effects of employing the error mitigation mechanisms in different ways. In all cases, improvements were noted through the use of link-layer support, and in particular, through the judicious and speedy handling of frame retransmissions. In a real-world practical implementation of a MIMO system where it is possible to apply error mitigation onto separate independent channels, it has been shown that where those channels exhibit differences in bit-rate, a cross-linked paired-LLEM mitigation connection arrangement can perform better, at some levels of SNR, than either a single multiplexed LLEM or dual independent LLEM units.

Acknowledgment

Thanks are due to D. McConnell and Dr K. Mehrotra of MiMOMax Ltd for their assistance in both data collection, and modeling of the system with their award-winning narrowband digital radio link. Thanks are also due to M. Thakur for his prior work on the single LLEM scheme.

References

1. Y. Bai, A. T. Ogielski, and Gand. Wu. Interactions of TCP and radio link ARQ protocol. In *Proc. Of VTC'99*, pages 1710–1714, September 1999.
2. A. Bakre and B. R. Badrinath. I-TCP: Indirect TCP for mobile hosts. In *Proc. 15th Int'l. Conf. Distributed Commputing Systems (IDCS)*, May 1995.
3. H. Balakrishnan, V. N. Padmanabhan, S. Seshan, and R. H. Katz. A comparision of mechanisms of improving TCP performance over wireless links. In *IEEE\ACM Transactions on Networking*, number 6, December 1997.
4. E. de Carvalho and P. Popovski. Strategies for ARQ in 2x2 MIMO systems. *Communications Letters, IEEE*, 12(6):441–443, june 2008.
5. Tamer Elbatt. Towards scheduling MIMO links in interference-limited wireless ad hoc networks. In *IEEE Mil. Comms. Conf.*, pages 1–7, February 2008.
6. Gorry Fairhurst and Lioyd Wood. Advice to link designers on link automatic repeat request (ARQ). In *draft-ietf-pilc-link-arq-issues-04.txt*, March 2002.
7. A. Grant. Rayleigh fading multi-antenna channels. In *EURASIP Journal on Applied Signal Processing*, volume 3, pages 316–329, 2002.
8. K.Brown and S. Singh. M-TCP: TCP for mobile cellular networks. In *Computer Communication Review*, October 1997.
9. Farooq Khan, Sarath Kumar, Kamesh Medapalli, and Sanjiv Nanda. TCP performance over cdma2000 RLP. In *Proc. Of VTC'2000*, pages 41–45, May 2000.
10. S H Kuo, J Dowle, and I V McLoughlin. A reconfigurable platform for MIMO research - realtime implementation of a 4x4 adaptive multi-variate DFE. In *Virginia Tech Symposium on Wireless & Personal Communications*, Virginia, USA, June 2004.
11. Jean-Paul Linnartz (Ed). Wireless communication. *http://wireless.per.nl/reference/contents.htm*, October 2007.
12. J LoMing and I V McLoughlin. Performance investigation and implementation of a real-time adaptive MIMO-DFE system. In *The Tenth IEEE International Conference on Communications Systems*, Singapore, November 2006.
13. Marjan Baghaie A., S H Kuo, and I V McLoughlin. FPGA implementation of space-time block coding system. In *4th IEEE Conf. Mobile and Wireless Communications*, Shanghai, China, May 2004.
14. Marjan Baghaie A., I V McLoughlin, P Martin, K Mehrotra, and D Taylor. Transmit antenna selection for UHF MIMO linking. In *Spring IEEE Vehicular Technology Conference*, Melbourne, Australia, May 2006.
15. I V McLoughlin and K Mehrotra. Time reversal space time block coding with channel estimation errors. In *Int. Conf. on Integrated Circuits and Systems*, Singapore, December 2003.
16. Eytan Modiano. An adaptive algorithm for optimizing the packet size used in wireless ARQ protocols. In *Wireless Networks 5 (1999)*, pages 279–286, 1999.
17. K. Pawlikowski, H-D J Jeong, and J.-S. R. Lee. On credibility of simulation studies of telecommunication networks. In *IEEE Communications Magazine*, volume 40, pages 132–139, 2002.
18. Kostas Pentikousis. TCP in wired-cum-wireless environments. In *IEEE Comm. Surveys*, Fourth Quarter 2000.
19. S. Prakash and I. V. McLoughlin. Predictive transmit antenna selection with maximal ratio combining. In *Global Telecommunications Conference, GLOBECOM 2009*, pages 1–6, December 2009.
20. Milind Thakur. *Wireless internet protocol for advanced space-time MIMO radio system: MEng Thesis.* The University of Canterbury, Christchurch, New Zealand, 2005.
21. Milind Thakur, Harsha Sirisena, and I V McLoughlin. Link layer error mitigation in rural UHF-MIMO linking systems. In *2005 IEEE Region 10 conference*, Melbourne, Australia, November 2005.
22. R. Vaughan and J. B. Andersen, editors. *Channels, Propagation and Antennas For Mobile Communications.* IEE, London, 2003.