A Performance Study of

Multiframe Interleaving for

CDMA2000 1X Cellular System

by

Jia Fu Cen

Submitted to the Department of Electrical Engineering and Computer Science

in Partial Fulfillment of the Requirements for the Degree of

Master of Engineering in Electrical Engineering and Computer Science at the

Massachusetts Institute of Technology

May 20, 2004

Author_____
Department of Electrical Engineering and Computer Science
May 20, 2004

Certified by_____
Walid Hamdy, Ph.D.
VI-A Company Thesis Supervisor

Certified by_____
Professor Dina Katabi
M.I.T. Thesis Supervisor

Accepted by_____
Arthur C. Smith
Chairman, Department Committee on Graduate Theses

A Performance Study of

Multiframe Interleaving for

CDMA2000 1X Cellular System

by

Jia Fu Cen

Submitted to the Department of Electrical Engineering and Computer Science

May 20, 2004

In Partial Fulfillment of the Requirements for the Degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

Multiframe Interleaving (MFI) is the support of 40 and 80ms supplemental channel frames for CDMA2000 1X. Longer frames improve interleaver performance by separating adjacent symbols further in time and thus require less E_b/N_t, in burst noise conditions, to maintain a target FER at a given data rate. The overall benefit is to improve throughput capacity for today's high bandwidth cellular services. This thesis analyzes the average power usage and data throughput with longer frames for various channel models and radio configurations. Through experimental and simulative testing in a laboratory setting, MFI is shown to use less power per unit of throughput at the link layer in slow fading, single path channel environments. However, gains over 20ms frames diminish for faster fading and multipath, and the limitation of the forward link scheduler in accurately predicting channel conditions handicaps any gains at the link level.

VI-A Company Thesis Supervisor: Walid Hamdy, Ph.D.
Title: Director of Engineering, QCT Systems Engineering

Thesis Supervisor: Professor Dina Katabi
Title: Assistant Professor, MIT Computer Science and Artificial Intelligence Lab

Acknowledgements

I am grateful to my co-workers at QUALCOMM who were always available to lend support and technical advice. I thank Levent Aydin very much for his guidance and instruction, and Walid Hamdy for his encouragement and mentoring. Much appreciation also goes to Bao Nguyen for his indispensable help with the implementation and setup, Neset Yalcinkaya and Mohan Kanthaswamy for their kind assistance in the laboratory, Anna Lau for her dedicated support on the CSMApps side, and Etienne Chaponniere for his regardful tutelage. I also thank Shim Patel, Jasper Jiang, Ling Hang, and Lei Shen for their insights and advice, and Jordi de los Piños, Pradeep Gowda, Pradeep Hasabettu, Arif Kahn, Sunil Kandukuri, Yoh-Hahn Noh, and Saritha Yaramada for their help on various things that made my job easier.

I would also like to extend my deepest appreciation to Dina Katabi for taking a risk in working on something slightly outside her field of interest. I also thank her for her support and advice throughout the project.

And big thanks to the interns, past and present, of Qualcomm who made living and working in San Diego a fun and memorable experience.

Finally I would like to thank my family for affording me the opportunity to follow my dreams. I am indebted to my mom and dad for their years of sacrifice and support. They may not understand one sentence in this thesis, but every word is possible only because of their love and encouragement. I also thank my sister for her companionship and loyalty. Lastly, I thank my grandmother who always believes I can do great things.

Contents

List of Figures

List of Tables

Acronyms

Abbreviation	*Meaning (page where term is defined or used)*
BER	Bit Error Rate (p. 41)
BS	Base Station (p. 19)
BSC	Base Station Controller (p. 19)
BTS	Base Transceiver Station (p. 18)
CDMA	Code Division Multiple Access (p. 17)
DMSS	Dual Mode Subscriber Software (p. 35)
DTX	Discontinued Transmission (p. 21)
ESCAM	Extended Supplemental Channel Assignment Message (p. 22)
FCH	Fundamental Channel (p. 21)
FL	Forward Link (p. 20)
IWF	Interworking Function (p. 19)
MFI	Multiframe Interleaving (pp. 29-32)
MS	Mobile Station (p. 18)
MSC	Mobile Switching Center (p. 19)
MSM	Mobile Station Modem (p. 35)
PDU	Packet Data Unit (p. 49)
PSTN	Public Switched Telephone Network (p. 19)
RC	Radio Configuration (p. 21)
RDA	Rate Determination Algorithm (p. 36)
RL	Reverse Link (p. 24)
RLP	Radio Link Protocol (p. 37)
SCH	Supplemental Channel (p. 21)
SDU	Service Data Unit (p. 30)
STDSO	Simple Test Data Service Option (p. 39)
SURF	Subscriber Unit Reference (p. 37)

1 Introduction

1.1 Increasing Demands for Bandwidth

The introduction of the third generation (3G) of wireless systems provided a framework for employing high-speed data services and other new features. QUALCOMM's[1] 3G technology, CDMA2000[2] 1X, supports both voice and data services over a standard CDMA channel and provides twice the capacity of its predecessor, with peak rates of up to 307 kbps. Nevertheless, the increasing bandwidth demands and growing popularity of today's data-based phone services can strain 1X systems. Without resorting to the more expensive and less mature 1xEV set of technologies, we look for ways to optimize or improve our current systems.

1.2 Multiframe Interleaving for Improved Performance

One technique that shows promise for improving capacity and throughput is called Multiframe Interleaving (MFI). MFI is an extension to the physical layer frame structure in CDMA2000, allowing for longer frame lengths. While the idea is nothing novel—it has been supported in the IS-2000 standard since 1999—it has yet to be studied extensively. However, simulations done in previous studies have shown that MFI could help in situations where forward link capacity is power limited at the base station by lowering power requirements for slow moving mobiles or mobiles experiencing slow fading environments. The drawback to this feature is added latency and slower power control.

The big picture is that power saved due to MFI may be allocated to new users for increased capacity or current users for increase data rates. The impact on users will be an increase in end-to-end delays.

---

[1] QUALCOMM® is a registered trademark of QUALCOMM Incorporated.
[2] CDMA2000® is a registered trademark of the Telecommunications Industry Association (TIA-USA) in the United States.

1.3 Purpose and Scope

The purpose of this research is to assess the performance benefits of a CDMA 1X system with support for longer frame lengths. After implementation and testing, the research will focus on forward link capacity and performance. Capacity and performance will be judged by average sector throughput[3], average user throughput, and average power requirements. All testing and analysis will be done in a laboratory setting using channel emulators and other test equipment.

1.4 Experimental Approach

The experimental approach of this study is to start testing from the lowest layers and gradually increase the complexity of the tested system. In our suite of tests, we look at demodulator performance, link level performance, application layer performance without scheduling, and application layer performance with scheduling. All results are compared to a baseline which may be theoretical values, simulation data, or results from relevant projects.

1.5 Summary of Results

Multiframe Interleaving is shown to have the greatest benefit at slow speeds and for single path fading environments. The improvements are observed at the link level and in throughput efficiency. However, improvements diminish as speed increases and as diversity is provided by other means, namely multipath. Also, even when using slow speed and single path fading for scheduler tests, the overall gain in sector throughput is minimal. This small boost is due to the difficulty in accurately predicting future channel conditions and thus a difficulty in optimally scheduling users. More study remains to be done to fully define MFI's benefit in improving capacity and throughput.

[3] Sector throughput is the sum of all user throughputs served by a given sector.

1.6 Outline

This paper is organized into five main sections: this introduction, a background chapter, an experimental environment chapter, performance analysis, and references.

The coverage of each section is as follows:

1. Introduction (Chapter 1)—The introduction poses the need for increased capacity and throughput in today's CDMA cellular system and describes Multiframe Interleaving's role in possibly improving performance. The section also describes the scope of the work and summarizes the results.
2. Background (Chapter 2)—The background section briefly explains CDMA concepts and describes CDMA system components. It also discusses the intuition behind MFI, describes the working system, and presents outside studies done on related topics.
3. Experimental Environment (Chapter 3)—This section documents the changes made to the code base in order to support MFI and describes the lab setup used to conduct the experimental analysis.
4. Performance Analysis (Chapters 4 & 5)—Chapter 4 describes the series of performance tests taken to assess the benefits of the new feature and analyzes the collected data. Chapter 5 is a summary of the test results.
5. References (Chapters 6-12)—This section includes the bibliography and appendices, which contain supplementary material and collected data.

2 Background

2.1 CDMA

2.1.1 Fundamentals

Code Division Multiple Access (CDMA) is a digital technique for sharing frequency spectrum. All users occupy the same time and frequency channels and are distinguished by unique codes. A common analogy is the international cocktail party. Every couple at the party is speaking in a different language. While the room may be filled with conversations, a person can clearly understand their partner because they speak the same language and every thing else sounds like noise. CDMA's codes are like the languages.

The first CDMA cellular standard, IS-95, was published in 1993. The end-to-end wireless system based on IS-95 and related specifications is called cdmaOne[1]. In 1999, the International Telecommunication Union adopted an industry standard for third-generation (3G) wireless systems, including three operating modes, CDMA2000, WCDMA, and TD-SCDMA. CDMA2000 is based on the IS-2000 standard, the 3G successor to IS-95.

CDMA2000 supports two spreading rates[2], 1 and 3. Spreading rate 1, also known as 1X, is the mode used for this thesis. CDMA is deployed in two frequency band classes, 800 MHz (Cellular or Band Class 0) and 1900 MHz (PCS or Band Class 1). Each channel in a band class is 1.23 MHz wide for spreading rate 1, or 3.68 MHz wide for spreading rate 3.

2.1.2 Cellular System Architecture

There are four main components to a cellular system: the network topology, mobile units and base stations, base station controllers, and the wired subsystem. The topology consists of cells, or coverage areas, commonly laid out in a hexagonal grid, as

---

[1] cdmaOne™ is a trademark of the CDMA Development Group.
[2] Spreading rate is the range of frequency over which either the forward or reverse link operate. Spreading rate is equivalent to the chip rate.

shown in Figure 1. The base transceiver station (BTS) contains the antennas and RF modems. The point in a cell where the BTS is located is called the cell site. The BTS may contain directional antennas such that the coverage area of one antenna is only a portion of the cell, called a sector. Mobile units or mobile stations (MS) consist of cellular telephones and other portable wireless devices.

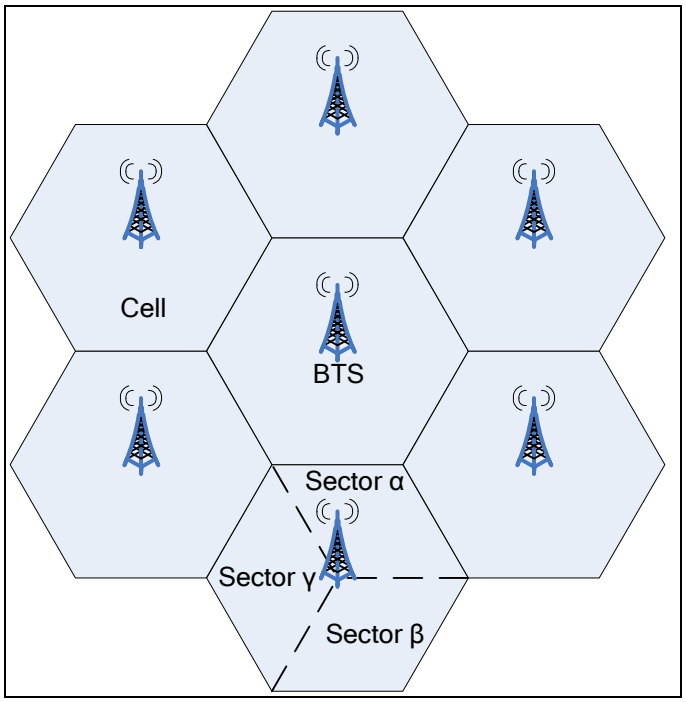

Figure 1: Diagram of cellular topology.

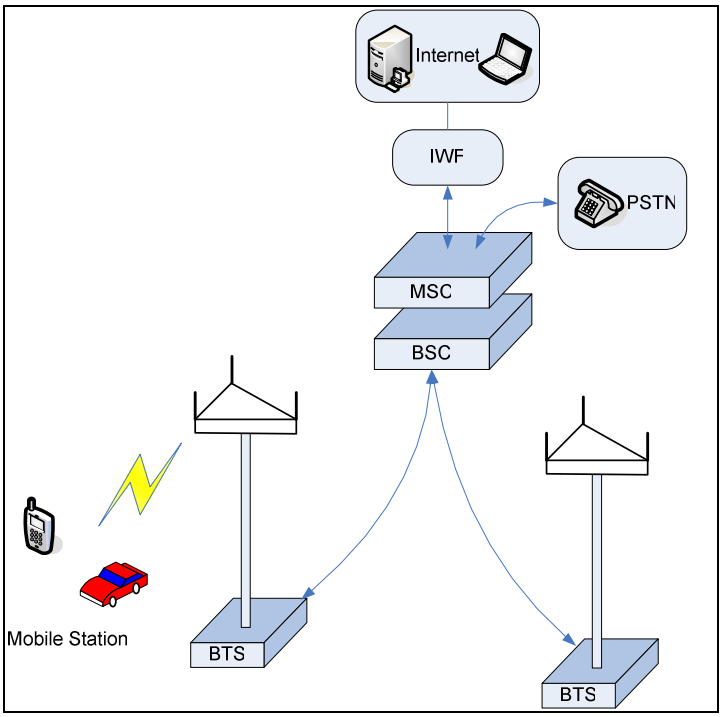

Figure 2: Diagram of the network components of the cellular system.

The base station controller (BSC) controls one or more BTSs and coordinates radio frequency (RF) resources. The BSC and BTS, collectively, are commonly referred to as just the base station (BS). The wired subsystem includes the public switched telephone network (PSTN), and the Internet. The mobile switching center (MSC) is the interface between dozens of cells and the wired subsystem. The interworking function (IWF) joins the MSC to the internet. Figure 2 shows how all the different components of a cellular system interact.

2.1.3 Mobile Radio Environment

2.1.3.1 Propagation Effects

As a wireless signal travels from the transmitter to the receiver, it experiences approximately three separable effects: multipath fading, shadowing, and path loss. Multipath fading encompasses envelope fading (nonselective), Doppler spread (time-selective), and time-delay spread (frequency-selective).

2.1.3.2 Envelope Fading

Multipath fading will cause the signal envelope to fade fast and deep. In urban environments, the fluctuated signal envelope exhibits a Rayleigh distribution, such that fades of 30dB or more below the root-mean-square of the signal envelope occurs 0.1% of the time. Moreover, the fading rate is proportional to the vehicle speed and center carrier frequency.

For emulating multipath channels, a Rayleigh or Rician model is commonly used. The Rayleigh model applies to non-line of sight conditions where the receiver gets multiple weak signals. In the Rician model, the receiver gets a strong specular path and multiple weaker ones.

2.1.3.3 "Signal to Noise" Terms

The total power transmitted from the base station is normalized and denoted as I_{or}. The fraction of the power allocated to a specific mobile is E_c/I_{or}, where E_c is the energy per chip. Geometry is the ratio of I_{or} to I_{oc}, where I_{oc} is the interference power from other cells plus thermal noise. I_{oc} is modeled as zero-mean additive white Gaussian noise (AWGN).

Since the mobile cannot tell the difference between I_{or} and I_{oc}, the interference it sees is $I_o = I_{or} + I_{oc}$. Pilot channel E_c/I_o is used to measure signal strength, which is commonly reflected in the number of "bars" displayed on the user interface. However, traffic channel signals are usually measured by E_b/N_o, classically defined as the energy per bit divided by the noise power spectral density at the receiver. Communication engineers evaluate demodulator performance using E_b/N_o values, although in practice, it is hard to measure. Instead, E_b/N_t is used, where N_t is the post de-spreading interference power.

When the channel geometry and multipath fading scenario is known, one can convert between E_c/I_{or} and E_b/N_t. This is useful when only one quantity is accessible.

2.1.4 Physical Layer

The forward link (FL), the communication link from the base station to the mobile, contains up to 128 code channels, which are distinguished by Walsh codes.

Overhead channels (pilot, paging, and sync) have fixed Walsh assignments, and the remaining Walsh codes are available for traffic channels.

Forward traffic channels are used to transmit user data and signaling information. A forward traffic channel may consist of one fundamental code channel and one or more supplemental channels. The fundamental channel (FCH) is the primary channel for all traffic communications and exclusively carries signaling information. The supplemental channel (SCH) is also used to deliver user data but at much higher data rates. The exact data rates depend on the radio configuration (RC) which specifies physical layer parameters such as rate set, channel coding, modulation, etc. RC3 and RC4 are used for this study. Table 1 summarizes these configurations.

RC	FCH Data Rates (kbps)	SCH Data Rates (kbps)	FEC Code Rate	FEC Encoding	Modulation
3	1.5 (Eighth rate) 2.7 (Quarter rate) 4.8 (Half rate) 9.6 (Full rate)	9.6 (1x[3]) 19.2 (2x) 38.4 (4x) 76.8 (8x) 153.6 (16x)	½	Convolutional or Turbo	QPSK
4	same as above	same as above plus 307.2 (32x)	¼	Convolutional or Turbo	QPSK

Table 1: Table summarizing radio configurations for CDMA2000.

The fundamental channel can vary its rate depending on how much data it has to send but data must always be sent. When the FCH has no user data to send, it transmits eighth rate frames. On the other hand, the supplemental channel has a fixed data rate. The data must either be transmitted at the prescribed rate or not transmitted at all, which is known as "discontinued transmission" (DTX). In the case of DTX, the receiver will only detect noise but it cannot tell this apart from a burst of noise that erased a frame. The process for distinguishing the two cases is called the Rate Determination Algorithm (RDA).

[3] Not to be confused with 1X, as in spreading rate 1.

2.1.4.1 Encoding

2.1.4.1.1 Channel Coding

Forward error correcting (FEC) coding introduces redundancy into the transmitted data stream such that certain bit errors can be recovered. Two types of FEC encoders are used: convolutional and turbo. Turbo encoders, while they require less E_b/N_o to achieve the same bit error performance as convolutional encoders, are also more computationally complex and only have an advantage for large block sizes (greater than 360 bits/frame).

2.1.4.1.2 Interleaver

The interleaver shuffles encoded symbols around so that burst errors are spread amongst symbols of different data bits. In CDMA systems, a block interleaver is used: data is written into a matrix row by row, and read out column by column.

2.1.4.2 Burst Assignments

By default, users have only one fundamental channel during a call. When the base station determines that it has too much data to send to a user through just a FCH, it will assign that user a supplemental channel[4]. The control message that sets up this assignment is the Extended Supplemental Channel Assignment Message (ESCAM). The ESCAM specifies the start time, duration, data rate, and other SCH parameters. Durations may be as short as one frame or as much as 256 frames or even "infinite", meaning the assignment does not expire until another ESCAM is sent to deallocate it. The length of the assignment is always measured in units of 20ms frames regardless of the actual frame duration. Supplemental channel assignments are also referred to as burst assignments, since such an assignment provides a "burst" of data transmission for the mobile.

---

[4] This decision is operator specific. For this study, SCH assignments are either static (set according to the specific test) or dynamic (done according to the forward link scheduler).

2.1.4.3 Walsh Codes

Walsh codes provide a means to uniquely identify each user on the forward link. Walsh codes have a unique mathematical property—they are orthogonal. In other words, Walsh codes are unique enough that data spread with a Walsh code can only be recovered by a receiver applying the same Walsh code. All other signals are discarded as background noise.

The Walsh space can be represented as a balanced binary tree where each node is orthogonal to all the nodes on the same tree level. A Walsh tree is shown in Figure 3. The Walsh code for a parent node can be thought of as the span of its two child nodes. Thus, the use of a particular Walsh code prevents the use of any if its "ancestors." Furthermore, the higher the data rate needed, the higher the Walsh code needs to be in the tree.

In RC3, there are 64 Walsh codes resulting in 64 1x channels, or 32 2x channels, or 16 4x channels, etc., all the way to 4 16x channels[5]. (RC4, release A has 128 1x channels with a maximum of 4 32x channels.) Although even higher rate Walsh codes theoretically exist on the tree, they are never available due to allocation of overhead channels (i.e. pilot, sync, and paging channels).

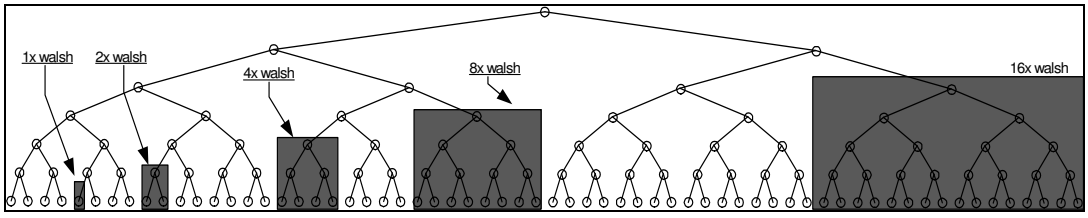

Figure 3: Diagram showing example of Walsh tree with 64 Walsh codes.

2.1.4.4 Power Control

Power control on the forward link is a process to maintain a certain error rate while transmitting the least amount of power. There are two important values that influence power control: E_b/N_t, and the Frame Erasure[6] Rate (FER).

---

[5] In an actual system, the number of traffic channels available will be less due to allocation of overhead channels.

[6] An erasure is a frame that cannot be error recovered.

The actual power transmitted by the base station is controlled by the feedback from the power control bits sent by the mobile on the reverse link (RL), forming what is called the "inner loop." These bits instruct the base station to increase or decrease transmit power such that the received E_b/N_t at the mobile is as close as possible to the target E_b/N_t or setpoint.

The value of the setpoint is adjusted in what is called the "outer loop" of power control. If a frame is a non-erasure then the setpoint is lowered by a small fixed amount[7], pushing the setpoint to the limit of optimality. However, once an erasure is received, the setpoint is raised by a large fixed amount[8]. This outer loop takes place on the mobile station.

Power control on the forward link operates in many modes. The most common one used in this study is mode 1, in which half of the power control bits provide feedback for the FCH and half for the SCH.

2.1.5 Forward Link Scheduler

2.1.5.1 Overview

The number of users and the data rates transmitted to each user on the forward link is limited by the number of Walsh channels and the maximum transmission power at the BTS. Thus, a scheduler is used to allocate resources, i.e. SCH assignments, to data users every scheduling period. Certain parameters, such as the burst duration, are static, while other parameters, such as data rates and number of assignments per period, are dynamic.

To allow for computation and setup latency, the scheduler is called approximately six frames before the start of the assignment period. The scheduler uses a scheme called modified round robin. Modified round robin differs from traditional round robin in that multiple users may be assigned simultaneously. While the strict equal-time property of round robin no longer exists, system resources are better utilized at the cost of additional scheduler complexity. Scheduling is done based on two factors: the

[7] This small fixed amount is a function of the target FER.
[8] This large fixed amount is also a function of the target FER.

required rate and the supported rate. The following sections discuss these two decisions.

Also, the scheduler is channel insensitive, i.e. it ignores channel conditions. This choice to disregard channel conditions stems from the large delays in setting up SCHs and from the lack of a fast channel state feedback system from the mobile.

Figure 4 shows the base station transmit power allocation as a function of time. A certain percentage is allocated to overhead and fundamental channels. The remainder is shared amongst the users.

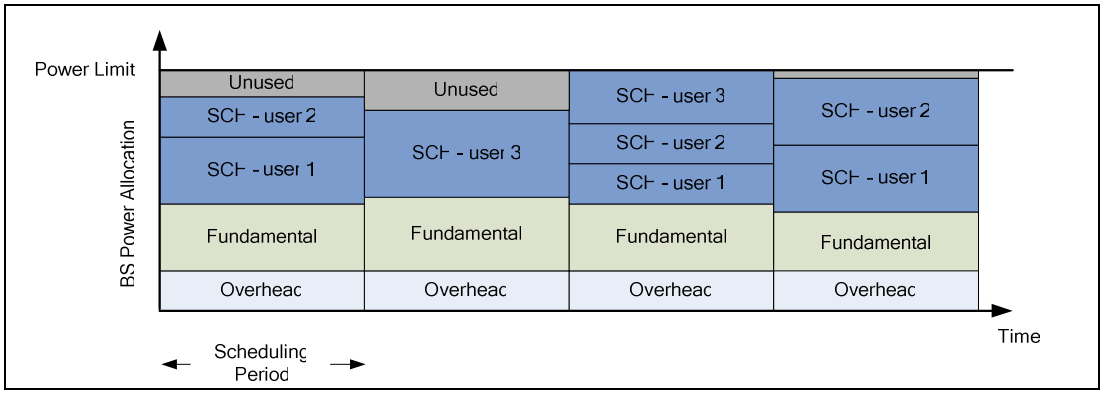

Figure 4: Diagram of sample BS power allocation vs. time.

2.1.5.2 Required rate decision

The required rate for a particular user is decided based only on the state of the data buffer at the base station. The type of client application is not known. While it makes sense to assign a high rate to a user with a large amount of data in the buffer, the answer is not so clear when a user has little data to send. Assigning a low rate for a user with little data in the buffer will degrade the perceived user experience and in applications running over TCP, a low rate might mislead TCP into thinking the available bandwidth is poor, which may not be the case. On the other hand, assigning a high rate might enhance the perceived user experience and improve TCP efficiency but capacity is wasted if the user runs out of data to send before the end of the assignment duration.

For our studies, we ignore this issue by assuming a full buffer, which is reasonable given that the tests involve a large file transfer through FTP. Thus, the required data rate assignment policy is to assign the highest rate possible.

2.1.5.3 Supported rate decision

After establishing the required rate, the scheduler must estimate if the sector will be able to support the rate. Whether a sector can support the rate depends on power availability and Walsh space availability.

2.1.5.3.1 Power Availability

The amount of power required to transmit at a certain rate can only be estimated since the actual assignment takes place some time in the future. The estimate is calculated from a moving average of power requirements for the fundamental channel. Scaling factors are applied to obtain estimated power requirements for a given data rate. Deriving these scaling factors is discussed further in Section 4.4.

With an estimate of the required power, the scheduler can then check for available power. If power is available, the scheduler reserves that block of the power pool. Otherwise, the scheduler tries a lower rate until it finds a suitable one, which may ultimately be none (zero-rate assignment).

2.1.5.3.2 Walsh Space Availability

The final step in scheduling a user for a particular rate is to confirm that there is space in the Walsh tree. As mentioned in Section 2.1.4.3, higher rates require codes higher up in the tree. Since allocating a Walsh is equivalent to allocating the subtree rooted at that Walsh, all codes in the subtree must be free and cannot be allocated while the root Walsh is still reserved. The scheduler checks for space in the Walsh tree, trying lower rates until there is an accommodation.

2.1.5.4 First user for next round

Each user is scheduled using the above process until one of the following conditions is true:

- Every active user requiring a non-zero SCH rate has been assigned a rate.
- All the available power has been secured.
- All the available Walsh codes have been secured.

Since the first user scheduled has the highest probability of getting the required rate, a policy must be implemented for selecting the first user to ensure fairness. One

simple policy is to give priority to the user(s) assigned the lowest rate SCH in the last period. A better policy is to prioritize users by their deprivation, defined as:

$$deprivation = required\_rate - allocated\_rate,$$

where *required_rate* and *allocated_rate* are from the last assignment period. The scheduler uses the former policy since sector throughput is not affected with either policy.

2.2 Time Diversity

Multiframe Interleaving may be broadly categorized as a diversity technique. Diversity techniques are based on the idea that transmission errors occur when the channel attenuation is large. However, if the receiver is given several copies of the same signal transmitted over independent fading channels, then the probability that all the replicas are lost decreases dramatically. While many diversity techniques exist, the most common ones fall under three categories: space, frequency, and time diversity. Space diversity relies on sending the signal through different spatial channels or paths. Frequency diversity places the signal components on different carrier frequencies. Finally, time diversity, a method whereby the replica signals are transmitted in different time slots, is the basis behind MFI.

Most channel encoders (e.g. convolutional encoder or block encoder) work under the assumption that the transmission medium is memoryless. However, if the channel has memory, errors can no longer be characterized as randomly distributed and independent from each other. But is it well known that interleaving symbols, a type of time diversity, improves error performance in time-varying channels for systems with encoding [9].

Interleaving the encoded data causes bursts of channel errors, after de-interleaving, to be spread out in time. The errors thus appear random and independent, which works to the decoder's advantage. Interleavers usually take a block structure or convolutional structure, ostensibly named after the type of encoding under which it works best.

The performance of an interleaver depends largely on how far in time it can separate consecutive codeword symbols. It is reasonable to choose this interval to be

greater than the coherence time, T_C, of the fading channel to ensure nearly independent channel states [10]. The coherence time is a measure of the period of time after which the correlation function of two samples of the channel response[9] falls below some threshold. Coherence time is normally discussed with the maximum Doppler frequency, or Doppler spread, f_D, and the two are usually related by $T_c \approx \frac{1}{2f_D}$. The channel is said to be slow fading if the symbol time duration is smaller than the channel's coherence time; otherwise, it is said to be fast fading [4].

To get another perspective of coherence time, let us consider the level crossing rate of the signal envelope function, $R(t)$, which can be used to estimate the average fade duration. Let R be modeled as a Rayleigh distributed random variable with mean square $2\sigma^2$. The level crossing rate, $N(R_T)$, generally defined as the expected rate at which the signal envelope crosses a specified level R_T in the positive direction, is given as [4]:

$$N(R_T) = \sqrt{2\pi} f_D \rho \exp(-\rho^2) \quad , \quad \rho = \frac{R_T}{\sqrt{2}\sigma} . \tag{1}$$

The value ρ can be seen as the ratio of the threshold level to the average signal level. The average value of the fade durations τ_{R_T} below the threshold value R_T can be expressed as [4]:

$$E[\tau_{R_T}] \equiv \tau(R_T) = \frac{1}{N(R_T)} \Pr[R \le R_T] . \tag{2}$$

Substituting (1) and the Rayleigh distribution [5] into (2) yields:

$$\tau(R_T) = \frac{\exp(\rho^2) - 1}{\sqrt{2\pi} f_D \rho} . \tag{3}$$

[9] Taken at the same frequency but different time instances.

For positive ρ much less than unity, we can approximate $\exp(\rho^2)$ as $1+\rho^2$. Therefore, the final form of the average Rayleigh fade duration is

$$\tau(R_T) \approx \frac{1}{f_D}\frac{\rho}{\sqrt{2\pi}} \tag{4}$$

If we consider the case where the threshold level, R_T, is equal to the average signal level, then the average fade duration is approximately the coherence time. In other words, the coherence time is approximately the expected time that the signal envelope stays below its RMS value. Hence, if one signal component experiences a fade and if our separation time is greater than the coherence time, then it is expected that our diversity component(s) will be out of the fade. In general, our diversity components experience nearly independent fading states.

2.3 Multiframe Interleaving

In this thesis, we study the CDMA2000 version of Multiframe Interleaving (MFI), which refers to the support in the IS-2000 standard for longer frames—40 and 80ms complementing the established 20ms. With longer frames, the interleaver is larger, and thus able to separate symbols even farther in time, which leads to better system error performance. However, this feature introduces longer delays as frames take two or four times as long to be transmitted

Multiframe Interleaving is only supported on supplemental channels and the studied system only implements it for forward link supplemental channels (F-SCH).

To ease the writing and reading of this paper, all comparisons of MFI features to 20ms frames will only refer to 40ms. The comparison to 80ms frames should be obvious from the context or else it will be explicitly stated.

First, a few notes about frames and rates:

1. A supplemental channel has a certain maximum bandwidth, which we will call the base (data) rate. These base data rates are 9.6kbps (1x), 19.2kbps (2x), 38.4kbps (4x), 76.8kbps (8x), 153.6kbps (16x), and 307.2kbps (32x)[10].
2. A supplemental channel is also associated with a certain frame duration[11]. Standard sizes are 20, 40, and 80ms.
3. The effective bandwidth of an SCH, R_{eff}, is a function of the base rate, R_b, and the frame length, L, and is given by $R_{eff} = \frac{R_b}{n}, \quad n = \frac{20}{L}$.
4. From hereon in, SCH configurations will be denoted by an ordered pair $(R_b/n, L)$, where R_b is understood to be in units of kilobits per second and L in units of milliseconds.

The following table from the IS-2000 standard summarizes the bit and symbol rates for different SCH rates with MFI.

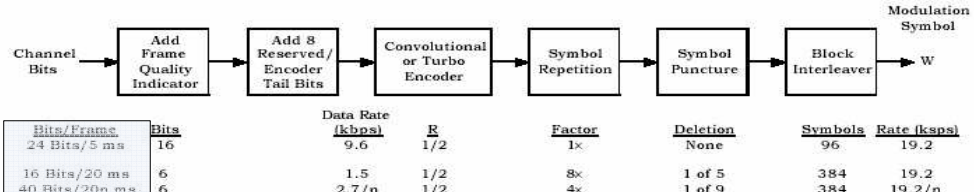

Bits/Frame	Bits	Data Rate (kbps)	R	Factor	Deletion	Symbols	Rate (ksps)
24 Bits/5 ms	16	9.6	1/2	1×	None	96	19.2
16 Bits/20 ms	6	1.5	1/2	8×	1 of 5	384	19.2
40 Bits/20n ms	6	2.7/n	1/2	4×	1 of 9	384	19.2/n
80 Bits/20n ms	8	4.8/n	1/2	2×	None	384	19.2/n
172 Bits/20n ms	12	9.6/n	1/2	1×	None	384	19.2/n
360 Bits/20n ms	16	19.2/n	1/2	1×	None	768	38.4/n
744 Bits/20n ms	16	38.4/n	1/2	1×	None	1,536	76.8/n
1,512 Bits/20n ms	16	76.8/n	1/2	1×	None	3,072	153.6/n
3,048 Bits/20n ms	16	153.6/n	1/2	1×	None	6,144	307.2/n
6,120 Bits/20n ms	16	307.2/n	1/2	1×	None	12,288	614.4/n

Notes:
1. n is the length of the frame in multiples of 20 ms. For 40 channel bits per frame, n = 1 or 2. For more than 40 channel bits per frame, n = 1, 2, or 4.
2. The 5 ms frame is only used for the Forward Fundamental Channel, and the Forward Fundamental Channel only uses from 16 to 172 channel bits per frame with n = 1.
3. Turbo coding may be used for the Forward Supplemental Channels with 360 or more channel bits per frame; otherwise, K = 9 convolutional coding is used.
4. With convolutional coding, the Reserved/Encoder Tail bits provide an encoder tail. With turbo coding, the first two of these bit are reserved bits that are encoded and the last six bits are replaced by an internally generated tail.

Figure 5: Specification of forward fundamental channel and forward supplemental channel structure for RC4 from IS-2000 physical layer standard [7].

The MFI scheme under study does not involve interleaving multiple frames, as the name might suggest. It is best to think of the system as running on a slower clock—data is fed through the system at a lower frequency. A unit of data (called an SDU) is passed to the physical layer every 20*n* ms, as specified in the first column (highlighted)

[10] A base data rate of 307.2 kbps is only supported in IS-2000 Release A and in radio configuration 4.
[11] In this paper, frame duration is synonymous with frame length and frame size.

in Figure 5. The SDU plus the Frame Quality Indicator—implemented with a cyclic redundancy check (CRC)—and encoder tail bits complete the frame. Thus, without going further, we can immediately determine the data rate since we know the number of outgoing bits per unit of time. Also, note that the Walsh length depends on the symbol rate such that the final chip rate is 1.2288Mcps.

For example, given a maximum SDU size of 6120 bits, the maximum data rate is $307.2/n$ kbps[12]. This is the reason for a reduction in peak data rates for longer frame lengths. However, note that for the same effective data rate, configurations with longer frames use a larger SDU, and thus, a larger interleaver/deinterleaver. This is the reason for improved deinterleaver performance. For example, (19.2/1, 20) and (38.4/2, 40) and (76.8/4, 80) all have an effective rate of 19.2kbps but the number of bits that constitute one frame is more for longer frame lengths.

Even though the IS-2000 standard for MFI supports SCH configurations with effective data rates below 9.6kbps (e.g. (19.2/4, 80)), this study does not explore those configurations. The primary reason is that the mobile chipset does not support such rates. A second reason is that SCH assignments are a valued resource and configurations with sub 9.6kbps effective rates only boost initial bandwidth capabilities, i.e. FCH bandwidths, by 12% to 50%. Such a meager supplement is usually not worth the resources and overhead.

As with most modifications to a system, there are advantages and disadvantages. Multiframe Interleaving offers improvement in link performance for the case of single-path fading models. For deployments without transmit diversity, MFI should benefit the capacity of sectors with slow moving mobiles. However, MFI is a cheaper means of diversity gains than transmit diversity schemes such as multiple antennas since it involves only software upgrades in both mobile and base stations. Also, compared to other frame interleaving schemes, which are presented in the next section, MFI does not increase the interleaver memory requirements.

---

[12] $\left(\frac{6120bits}{SDU}+\frac{24CRC\_tail\_bits}{SDU}\right)*\frac{1SDU}{1Frame}*\frac{1Frame}{20n\_ms}*\frac{1000ms}{1s}=307.2/n\,kbps$

Outside of the added complexity, the disadvantages of MFI are threefold. MFI has a negative impact on end-to-end latency and is not advisable for real-time applications such as video streaming. Some applications can tolerate the added delay without a negative impact on performance while others cannot. Thus, the use of MFI should ideally depend on the application. The other disadvantage is that the outer loop of forward power control reacts slower to changing channel conditions since it runs every frame boundary. This drawback can lead to higher BS power consumption in fast fading channels. Finally, as mentioned before, the maximum data rates are lower for longer frames, which may be disadvantageous when auspicious channel conditions allow for the highest data rate.

2.4 Prior Work

While we could not find any outside research that relates to our particular setup, there are two systems worth mentioning. One uses a symbol splitting technique [14, 15] and the other involves bit interleaving of either different multiple frames or a single frame repeated [9, 10, 12, 13].

2.4.1 Symbol Splitting

The symbol splitting technique is considered as an extension to the IS-95 standard. In standard IS-95 systems, chips belonging to one symbol are transmitted consecutively. The author of [14, 15] proposes combining chips into groups which are then interleaved and transmitted. He has shown significant performance gains, especially for frequency non-selective channels.

Symbol splitting divides a symbol with duration T_S into r subsymbols, each with a duration T_S / r. These subsymbols are interleaved with the subsymbols of surrounding data symbols. The time interval between subsymbols of one data symbol, T_D, is chosen according to the expected coherence time of the channel, T_C. If $T_D > T_C$, then chips from consecutive subsymbols (i.e. before interleaving) experience nearly independent fading.

The performance gains were studied through simulations. Please see [15] for the full simulation parameters. In summary, for a frequency-nonselective single-path Rayleigh fading channel with Doppler spread f_D = 200Hz, an improvement of

approximately 2dB in E_b/N_0 is seen at $P_b = 10^{-3}$, where P_b is the probability of bit error. At f_D = 50Hz, improvement is approximately 3dB at $P_b = 10^{-3}$. Finally, for four-path Rayleigh, f_D = 200Hz, the performance gain is minimal, due to the diversity gain provided by the multipath.

Important to note is that this symbol splitting system does not impact net data rate or signal bandwidth. Moreover, the frame size does not change, so latency is not affected.

2.4.2 Multiframe Interleaving (Jun, et al)

The other system, proposed by a TIA subcommittee working at Nortel, is closely related to our MFI system. However, due to the draft quality of the papers, it is not totally understood how their system is implemented. In their system, there are two types of interleaving. In the first case, the interleaving is performed by bit reverse interleaving over bits from *n* frames. In the second case, bit reverse interleaving is applied to a single frame but frames are repeated and interleaved; this is called frame repetition interleaving.

In the bit-level interleaving scheme, the interleaver is *n*-times larger, where *n* is the number of frames interleaved. The Walsh spreading happens as in the normal system. Thus, our data rates do not change with *n*. Rather, the latency increases with *n* as the interleaver takes a longer time to fill up.

In the frame repetition interleaver, a single frame is encoded and interleaved. However, "frames may be transmitted multiple times, where the frames are transmitted in an interleaved fashion." We interpret this as follows: if the symbols for one frame are ABCD and symbol interleaving makes it ACBD, then the symbols that are transmitted are AACCBBDD (assuming an interleaving factor of two.)

The authors only give performance gains for the bit-level multiframe interleaver but state that both methods give good system performance. For a data rate of 19.2kbps, 1% FER, 3kmph, and interleaving over four frames, the performance gain in E_c/I_{or} is 4dB in single-path fading. There is little improvement in 2-path fading models and improvements are best for mobile speeds of 1kmph to 30kmph.

3 Experimental Environment

3.1 Implementation

The Multiframe Interleaving feature was implemented on QUALCOMM's MSM6100[1] chipset through modifications to its software code base, known as the DMSS (Dual Mode Subscriber Software). All coding is done in the C programming language.

3.1.1 Decoder/Deinterleaver Setup

The crux of the modifications to the DMSS is in the timing of the deinterleaver. While data streams out from the demodulator, it is the deinterleaver that ultimately divides the data into frames. The other important change involves data frame processing at the MUX level.

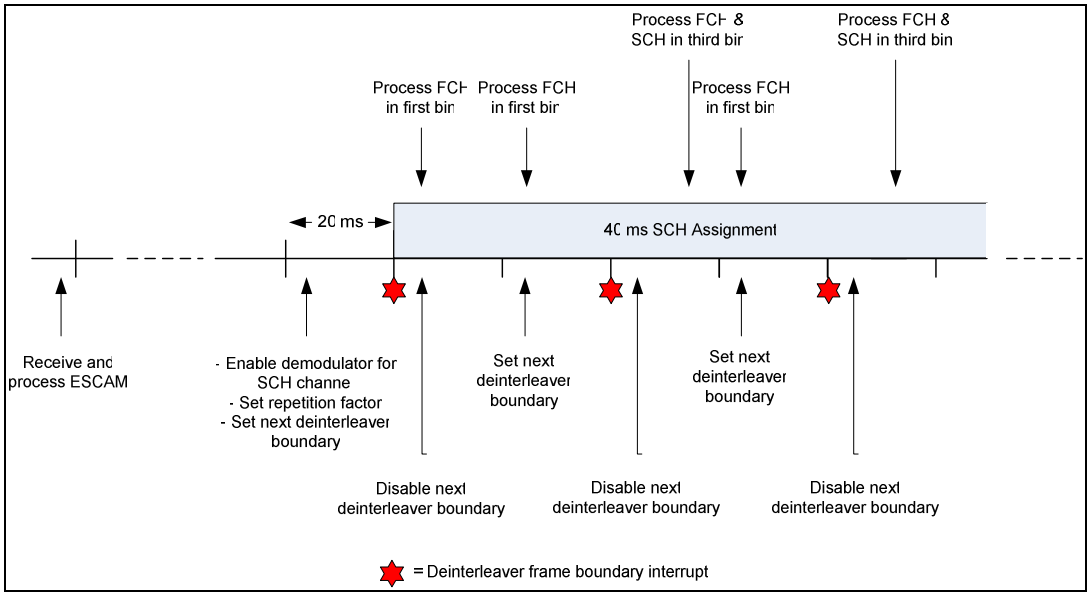

Figure 6: Diagram of deinterleaver and data processing timing for 40ms frames.

Figure 6 depicts the timing of the deinterleaver and the frame processing tasks for a frame duration of 40ms. First, the ESCAM specifying the SCH assignment is

[1] MSM6100™ is a trademark of QUALCOMM Incorporated.

received six or seven (20ms) frames before the start of the assignment. One frame before the start, the mobile will enable the demodulator and initialize the channel according to the parameters specified in the ESCAM. The deinterleaver works off of interrupts, which are generated by a hardware task. When the interrupt is set, it will go off every 20ms. However, to support longer frames, the interrupt is toggled on and off so that it goes off on the desired frame boundary (every 40ms).

Tasks in the software occur during one of four "bins," much like how a basketball game or American football game is divided into quarters; bins are only applicable to 20ms frames. When a traffic channel consists of only a FCH, the FCH data is processed during the first bin. However, if there is a SCH, processing of the FCH and SCH happen together in the third bin[2]. This optimization works fine with 20ms frames, but with longer frames, FCH and SCH frame boundaries do not always correlate. The solution is to toggle in which bin an FCH frame is processed, as indicated in the figure. Moreover, the first 40ms of the assignment is a special case where no SCH frame is processed since it is still being received.

3.1.2 RDA Issues

RDA or Rate Determination Algorithm is responsible for distinguishing the difference between DTX frames and erasures. The current RDA involves empirical data collected for 20ms frames. Therefore, when using longer frames, the RDA runs with a higher probability of misdetection, possibly causing the true FER to not converge. Such behavior is unacceptable especially when MFI is meant to improve link level performance. This issue needs to be solved before MFI can be used. However, solving this problem is beyond the scope of this study. Here we use a workaround by forcing the transmit buffer to be always full (e.g. transfer a large file) so that the base station is unlikely to DTX, and then disabling DTX detection on the mobile side, i.e. have the RDA classify all frames with failed CRCs as erasures.

---

[2] High rate SCH frames require more time to be demodulated which is why SCH frames are processed in the third bin.

3.1.3 RLP Issues

One side effect of MFI is that longer frames can cause the Radio Link Protocol (RLP) layer to function improperly, thus precluding any data testing with that frame size. However, demodulator performance testing, which does not involve RLP, is still possible. This RLP issue is discussed in detail in Appendix A.

3.2 Lab Setup

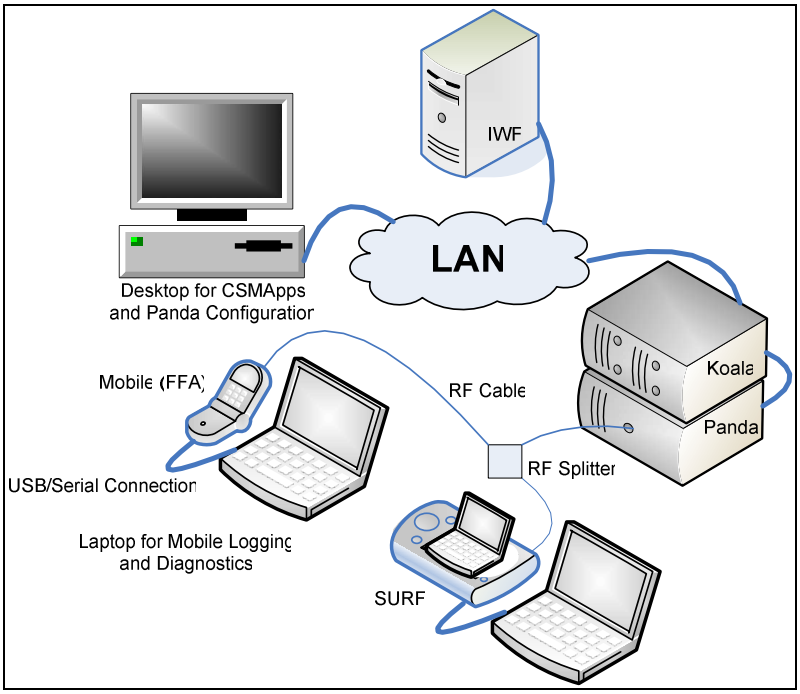

Figure 7: Diagram of lab test setup.

All analysis is done in the lab using various equipment to emulate the whole CDMA system: mobile station, over-the-air (OTA) environment, and base station. Figure 7 shows the lab equipment and setup. For the mobile station, a SURF or FFA is used. A FFA resembles a typical cellular phone but has more debugging features. A SURF (Subscriber Unit Reference) is a customizable development tool used for verification and system software development. The mobile is usually attached to a laptop for logging and, in the case of packet data testing, for acting as the PPP client. Channel emulation is done through a device called Panda. Panda is capable of modeling Rayleigh and Rician fading, geometry, transmit diversity, and more. The base station

hardware is housed in the Koala box. A desktop computer controls all settings for the Panda and Koala.

4 Performance Analysis

We try to establish performance of the system with MFI starting from the lowest layers and gradually increasing the complexity of the tested system until we reach something that resembles a mobile network. All results are compared to a baseline which may be theoretical values, simulation data, or results from relevant projects. First, we look at demodulator performance, a view of the basic physical layer. Then we examine link level performance by introducing power control. Next, we utilize upper layers, namely the RLP layer and transport/network layers, in the throughput performance analysis. Finally, we use multiple subscriber units on a single sector, with each user sharing resources through a scheduler.

Each suite of tests is described below in its own section, which includes an overview, a full description of the experiment, any relevant background information, a summary of collected data, and an examination of the data.

4.1 Demodulator Performance

In this analysis, we examine basic performance of the physical layer by looking at the probability of error versus E_c/I_{or} which is the classic waterfall curve known to communication engineers. The results show that for comparable data rates, performance is worse for longer frames. This behavior is explained by the larger number of bits in a longer frame and, thus, the lower probability of zero bit errors per frame. Therefore, targeting higher FERs for longer frame lengths could improve performance results for our subsequent tests.

4.1.1 Test Setup

In these tests, a Simple Test Data Service Option (STDSO) call is established between a SURF and the base station. The E_c/I_{or} is fixed at a certain value and the FER is measured for the SCH channel after 50,000 20ms frames (approximately 17 minutes). Given a set of parameters, each E_c/I_{or} value yields one data point; a few appropriately chosen E_c/I_{or} values will produce a nice waterfall curve.

The parameters for this analysis are summarized in the following table:

Parameter	Value(s)
Radio Configuration	RC4
Band Class	PCS (1900 MHz)
Power Control	Off
Coding	Convolutional, Turbo
Fading	AWGN, 1p
Geometry[3]	-1 dB
Data Rates (effective)	1x, 2x, 4x, 8x, 16x, 32x [4]
Frame Length	20ms, 40ms, 80ms
Number of Users	1
Test Duration	50,000 x 20ms

Table 2: Table of parameters for demodulator performance tests.

4.1.2 Data

The following graphs show a summary of the results for the above experiments. Graphs plotted against data collected from simulations [16, 17] are shown in Appendix B.

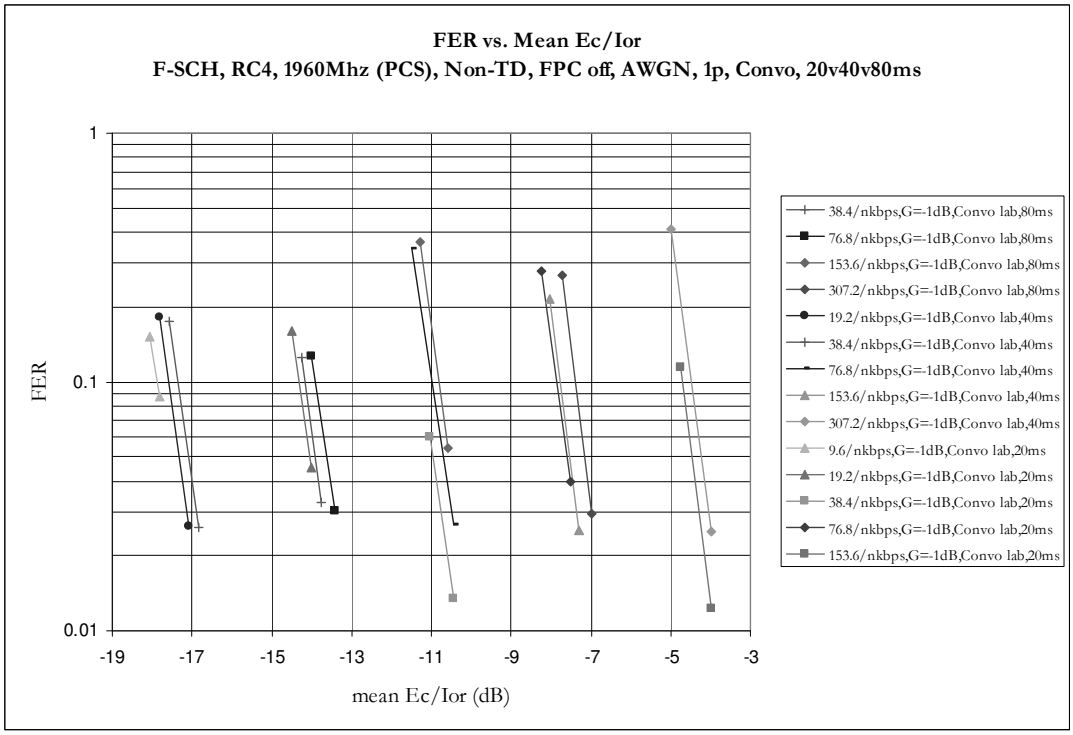

[3] Geometry = I_{or}/I_{oc}. See section 2.1.3.3 for more details.
[4] 1x not supported for 20ms turbo, 32x not supported for 40ms and 80ms, 16x not supported for 80ms.

Figure 8: Graph of demodulator performance. All points are lab results and 20, 40, and 80ms data are shown. Convolutional encoding is used.

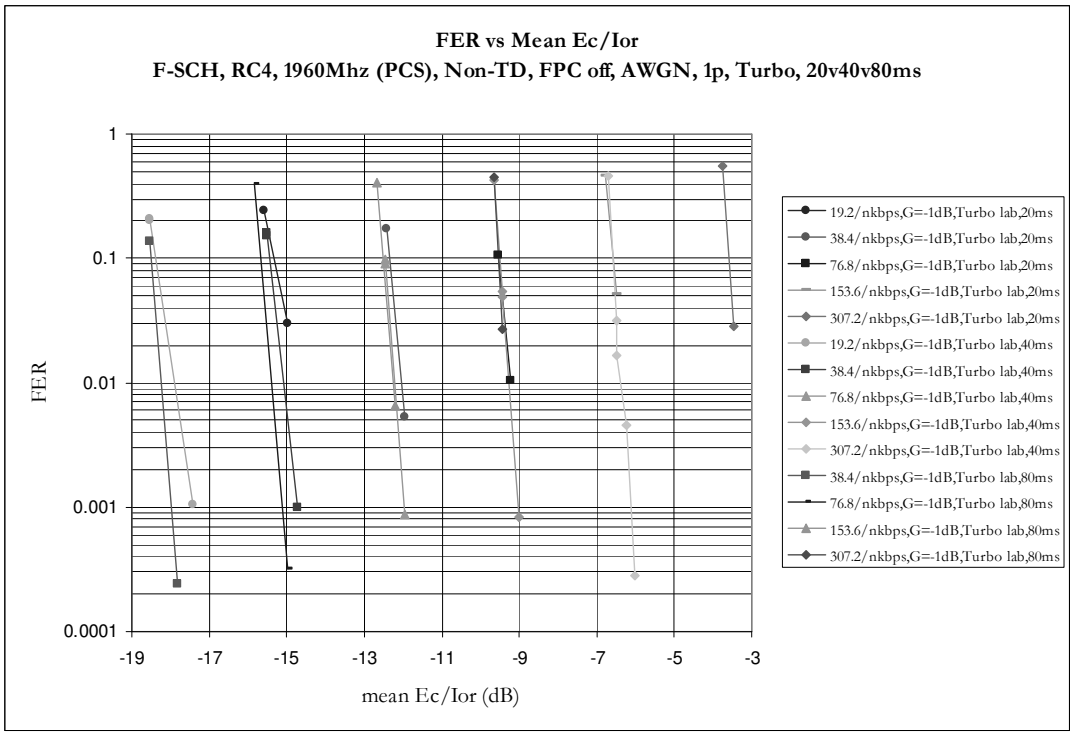

Figure 9: Graph of demodulator performance for turbo encoding.

4.1.3 Analysis

The results actually show worse demodulator performance for longer frames when comparing the same effective data rates—for the most part, results for longer frames are worse by no more than a half dB compared to 20ms frames. A reasonable argument for this behavior goes as follows: For the same effective data rate, the number of bits per frame increases for longer frame lengths. Assuming the likelihood of a frame containing no bit errors decreases with more bits per frame, longer frames will require more power to achieve the same FER.

More formally, let there be N·n bits per frame (so the effective data rate is 50·N bps) and assume that each bit in a frame has a probability of being received in error equal to BER, independent of other bits. Then the probability of an erasure is FER = 1-Pr(no errors) = $1-(1-BER)^{N \cdot n}$. For small BER, then simply

$$FER \approx N \cdot n \cdot BER \tag{5}$$

Although we have made a strong assumption about the independence of bit errors, the data seems to support the above theory. Figure 10 shows that FER and BER are linearly related for small BER and that for the configurations where N·n are equal, their FER vs. BER curves are the same.

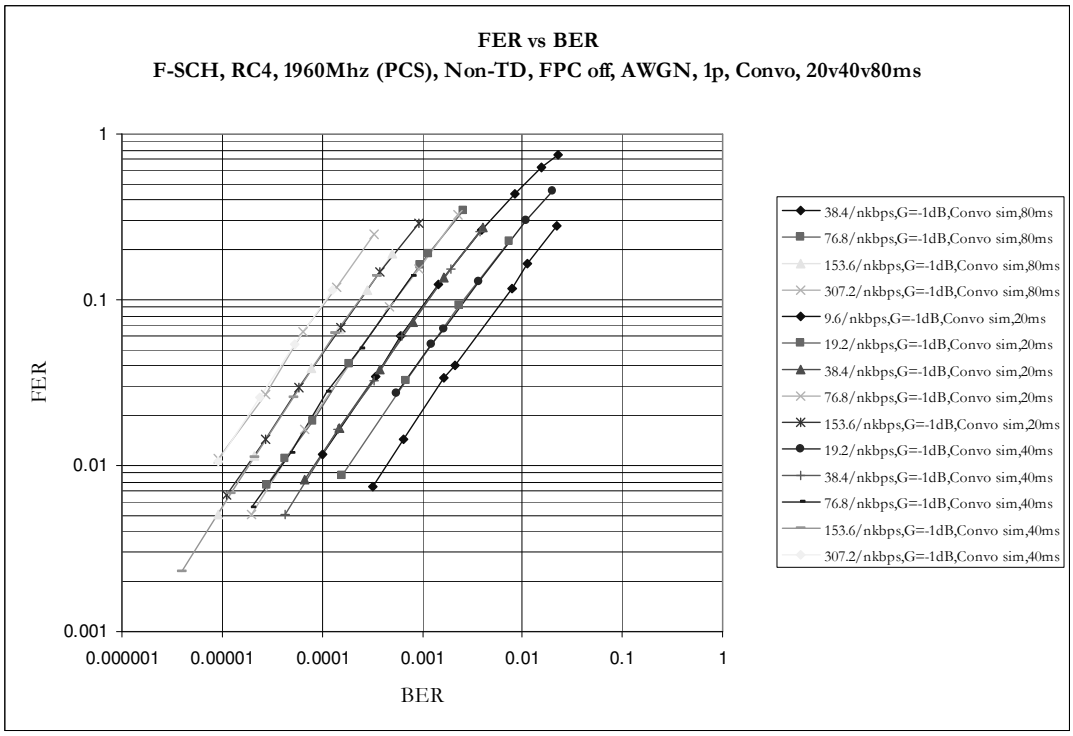

Figure 10: Graph of FER vs. BER data from simulations for convolutional encoding.

4.1.4 BER vs. FER

This relationship between BER and FER says that targeting higher FERs for longer frame lengths could produce the same bit error rates. Since at a fundamental level, we care about probability of bit error, this result suggests using higher target FER values to obtain better performance with longer frames.

4.2 Link Level Performance

In this suite of performance tests, we are still evaluating physical layer performance but power control is enabled and more channel conditions are considered. The results were within reason of simulation data, which shows the greatest gains for 1 path, slow fading conditions at low geometry. However, the gains are very small with multipath and at higher speeds.

4.2.1 Test Setup

A STDSO call is established between the SURF and the base station. A certain channel model is chosen and the geometry of the call is fixed. Target FER is set at 5% and the average E_c/I_{or} is measured for the SCH channel after 50,000 20ms frames (approximately 17 minutes). Each geometry value produces a point and several geometry values are chosen to produce a curve.

The parameters of these tests are summarized in the following table:

Parameter	**Value(s)**
Radio Configuration	RC4
Band Class	PCS (1900 MHz)
Power Control	On, Target FER = 5%, Mode 1
Coding	Convolutional, Turbo
Fading	(AWGN, 1p), (Rayleigh, 1p, 3kmph), (Rayleigh, 1p, 30kmph), (Rayleigh, 2p, 3kmph), (Rayleigh, 2p, 30kmph)
Geometry	-5 dB to 13 dB
Data Rates (effective)	1x, 2x, 4x, 8x, 16x, 32x
Frame Length	20ms, 40ms, 80ms
Number of Users	1
Test Duration	50,000 x 20ms

Table 3: Table of parameters for link level performance tests.

4.2.2 Data

The following charts show the collected data for a 1path, AWGN channel plotted against simulation data from a previous project [16, 17].

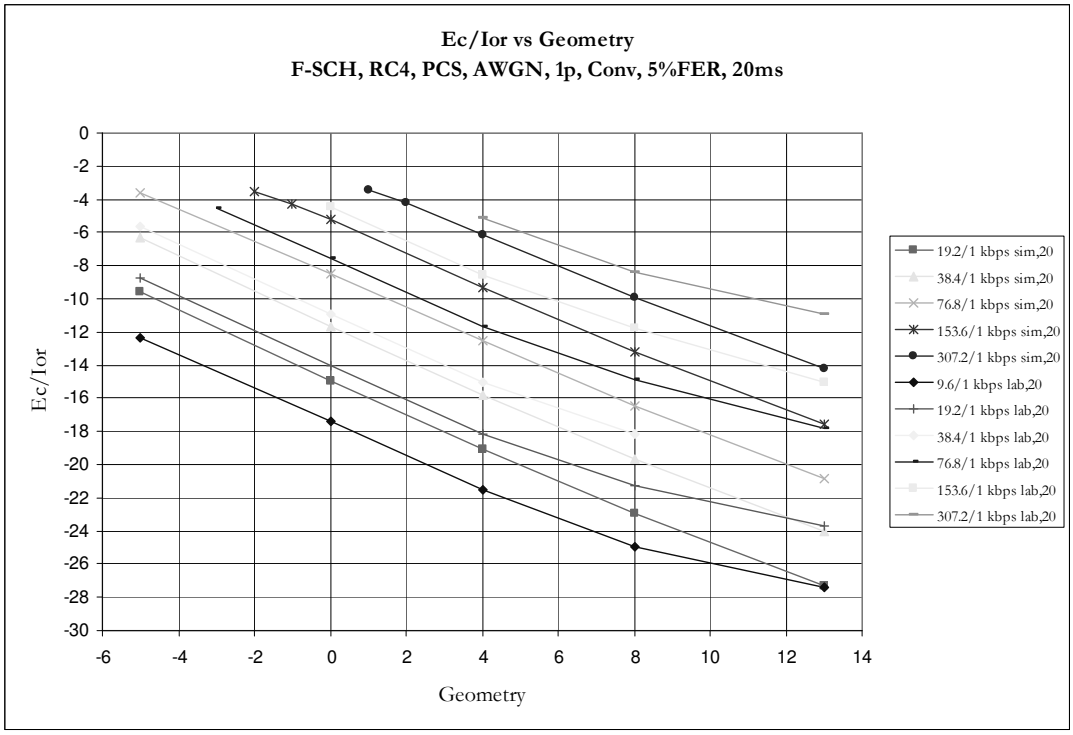

Figure 11: Graph of link level curves for 20ms in AWGN, 1path for convolutional encoding.

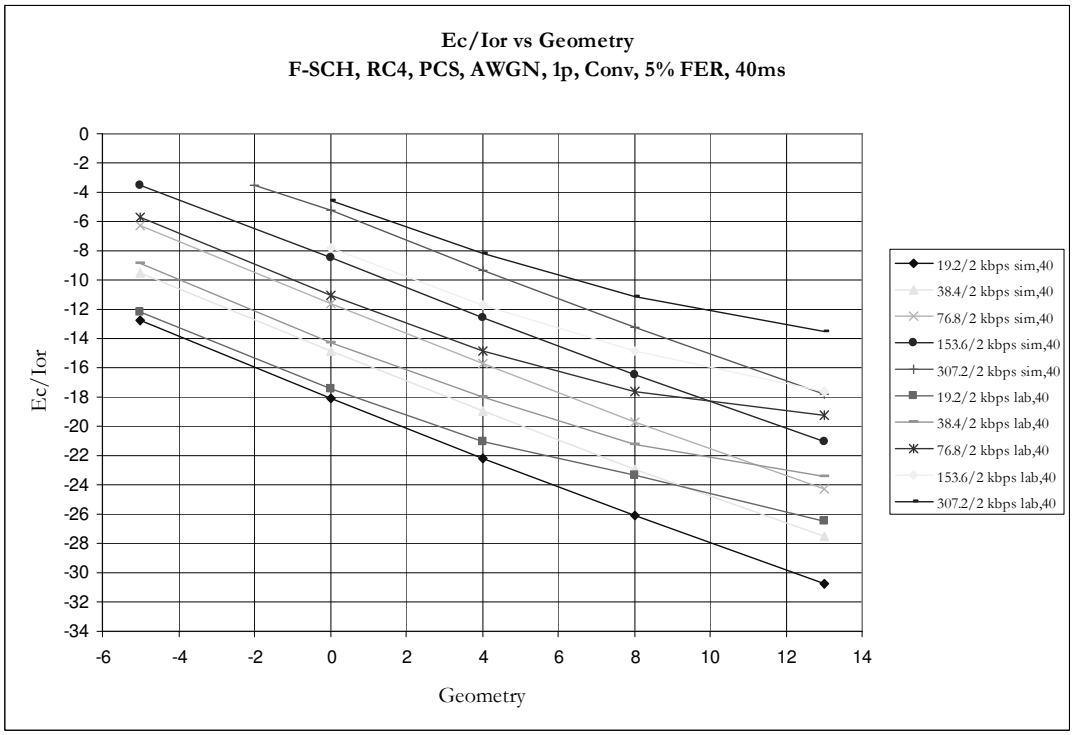

Figure 12: Graph of link level curves for 40ms in AWGN, 1path for convolutional encoding.

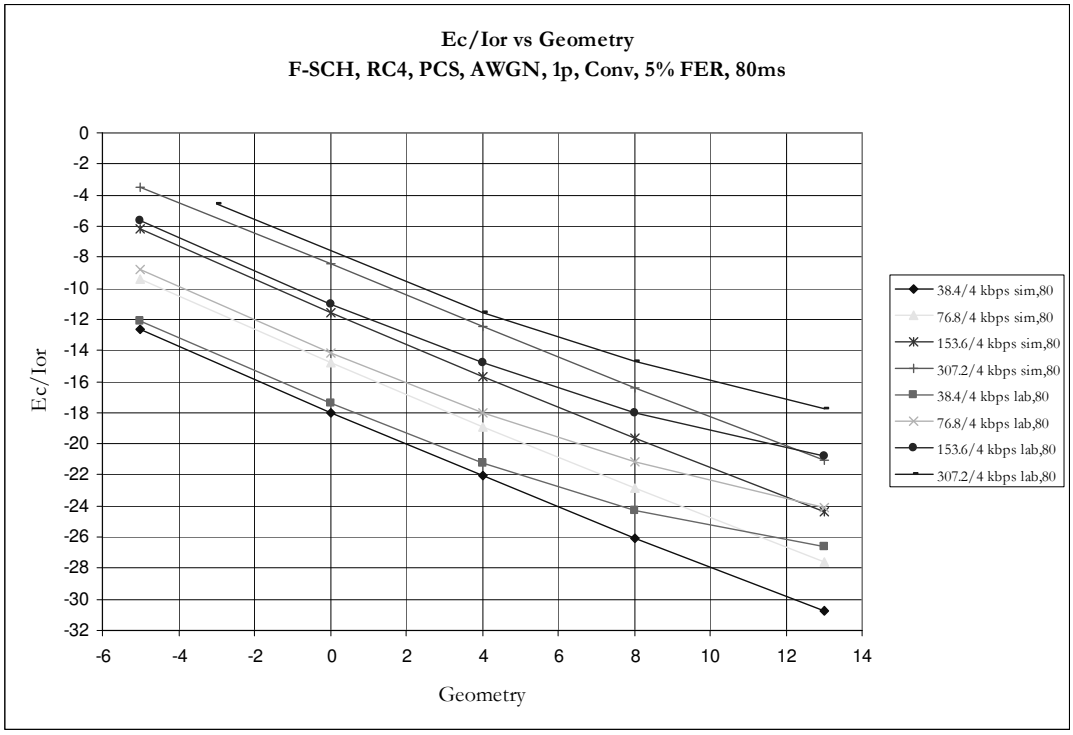

Figure 13: Graph of link level curves for 80ms in AWGN, 1path for convolutional encoding.

Results for turbo encoding in AWGN and all fading channel models can be found in Appendix C.

4.2.3 Analysis

When comparing lab values with simulation data, lab values show worse E_c/I_{or} performance given the same geometry, but by no more than one dB. This deficit is acceptable given the idealizations of simulations.

One very noticeable anomaly is the divergence of the lab data from simulation data at high geometries and in AWGN conditions. This behavior, while undesirable, is well-documented and is caused by false finger assignments to side lobes incorrectly detected due to the high geometry. Despite the large degradation in performance, this problem only seems to manifest itself in AWGN channels which is not something normally experienced in the field.

While plotting the data against simulation values is useful for sanity-checking purposes, comparing performance across frame lengths is more useful. Figures 14 and

15 show link level curves for all frame lengths grouped by effective data rates. The other fading cases are available in Appendix C.

The greatest gains are seen in 1path, 3kmph fading at low geometries—more than 2dB difference between 20ms and 80ms. However, that difference is negligible at 2path, 30kmph fading.

In general, longer frames give better E_c/I_{or} performance. Additionally, three patterns emerge from this view of the data:

1. Gains decrease at high geometries. As the channel gets better, MFI has less of an impact.
2. Gains decrease with higher mobile speeds. This behavior supports the prevalence of MFI in slow fading or slow speed conditions.
3. Gains decrease with more multipath. Since multipath provides diversity to the receiver, the additional diversity from MFI is less significant. This behavior was also seen in [14, 15]'s symbol splitting proposal.

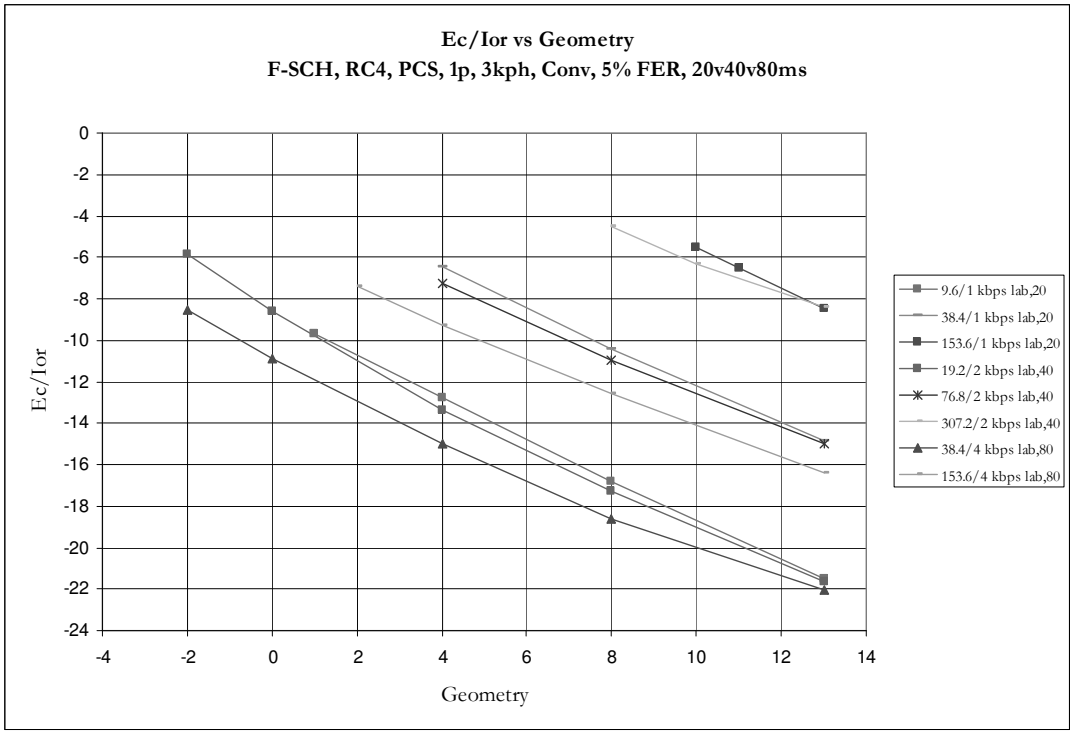

Figure 14: Graph of link level curves for all frame lengths in 1path, 3kmph with convolutional encoding.

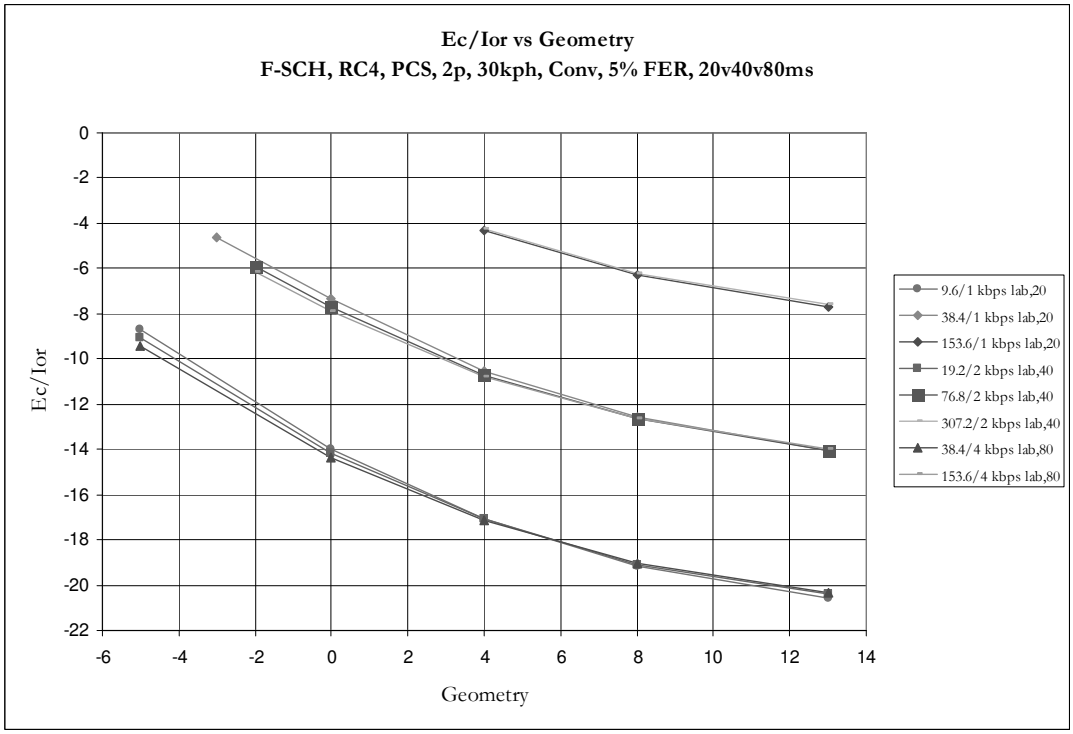

Figure 15: Graph of link level curves for all frame lengths in 2path, 30kmph with convolutional encoding.

4.3 Throughput Performance

In the throughput performance tests, we perform a file transfer over FTP, which involves upper layers in the protocol stack, including RLP and TCP/IP. At this level, we can start to appreciate the benefits of Multiframe Interleaving in terms of decreasing power requirements for a given level of throughput. The results show that at slow speed and single path fading, longer frames afford lower power requirements and/or higher throughputs; the overall effect is a gain in throughput efficiency. Again, there is less gain at higher speeds.

4.3.1 Test Setup

A single SURF is connected as a USB modem to a laptop. Through Windows Dialup Networking, a PPP connection is established through the SURF with the lab's IWF server. Then an FTP session is started and a large file—enough for a two to three minute transfer time—is downloaded. At least three trials are completed and the throughputs and E_c/I_{or} are averaged for the download. Table 4 summarizes the parameters in each test.

Notice that with the inclusion of the RLP layer in our tests, we can no longer test 80ms frames because of RLP's mishandling of delayed frames (see Appedix A.) Also, as a sanity check, tests were performed in a clean channel (no noise) with fixed E_c/I_{or} (no power control). These conditions should produce results very close to the theoretical maximum, which is discussed in the following section.

Parameter	Value(s)
Radio Configuration	RC3, RC4
Band Class	PCS (1900 MHz)
Power Control	On, Target SCH FER = 5%, FCH FER = 1%, Mode 1
Coding	Convolutional, Turbo
Fading	Model A, Model B, Model C, Model D, Model E[5]
Geometry	(Large enough so that all data rates can be run; values range from 9 to 13 dB)
Data Rates (effective)	1x, 2x, 4x, 8x, 16x, 32x
Frame Length	20ms, 40ms
Number of Users	1
Test Duration	2-3 minutes
MUX PDU	3, 5
Burst Duration	Infinite (option 0xF)
TCP/IP Header	40 bytes
TCP MSS	1460 bytes
TCP Window Size	64 bytes
PPP Header	5 bytes
Compression	None

Table 4: Table of parameters for throughput performance tests.

4.3.2 Maximum Application Throughput

The maximum application throughput will depend on the values of the parameters above. Appendix D gives a table of maximum theoretical throughputs for 0% FER and various operation settings. In general, turbo encoding will give higher throughputs because of less LTU CRCs within an SDU; MUX PDU 5 will give higher throughputs because of less overhead; and longer frames with the same effective physical layer throughput will have higher application throughputs because of amortization of headers and CRCs (recall that larger frames at the same effective data rate have larger SDUs).

[5] Channel models A through E are defined in the EV-DV evaluation methodology [8].

4.3.3 Data

The following data shows the fractional difference of the recorded throughputs[6] below baseline for our sanity check. The shorter the bar, the closer the throughput was to the theoretical maximum. The best and average value over four trials is shown.

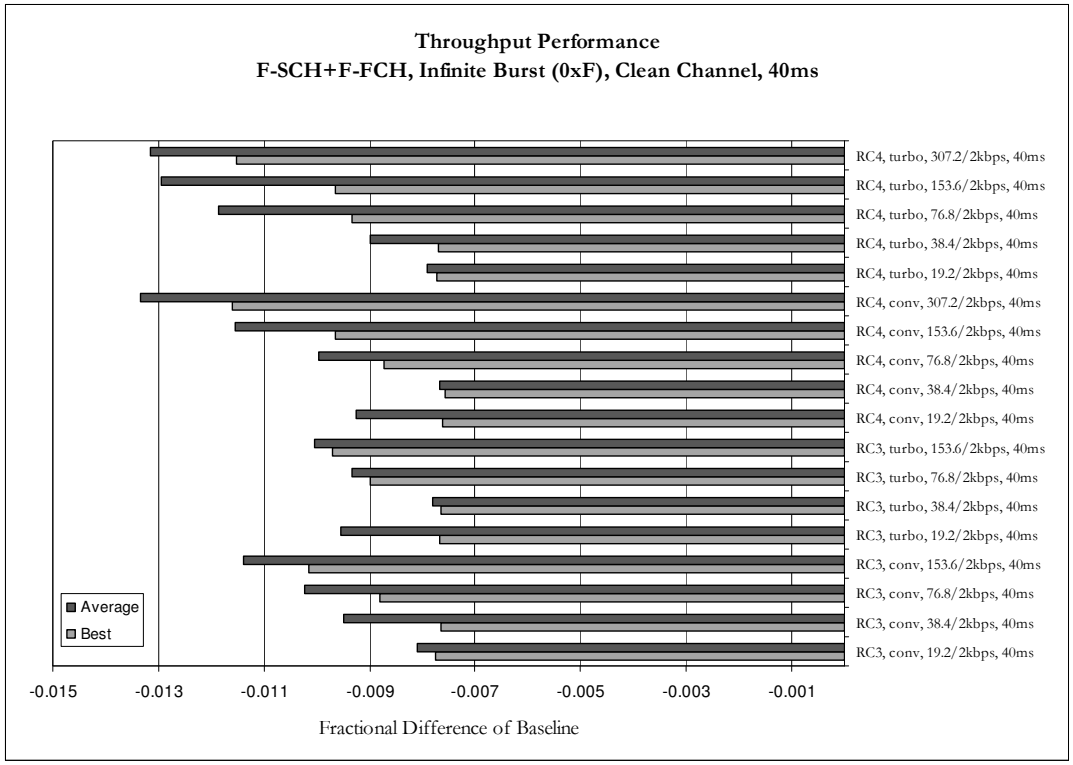

Figure 16: Graph of throughput performance for infinite burst assignments, measured relative to theoretical maximum.

The next set of data shows the throughputs versus E_c/I_{or}. Each graph contains points for a given channel model, radio configuration, and MUX PDU option. Each point on a graph represents the average throughput and average E_c/I_{or} for trials using the same encoding, frame duration, and data rate. The vertical "error" bar represents the maximum theoretical throughput. The difference between the bar and the actual point is due to the 5% target FER. More results for other channel models are available in Appendix D.

---

[6] The measured throughput is calculated from the total number of data bytes transferred divided by the total transfer time, both as reported by the FTP client.

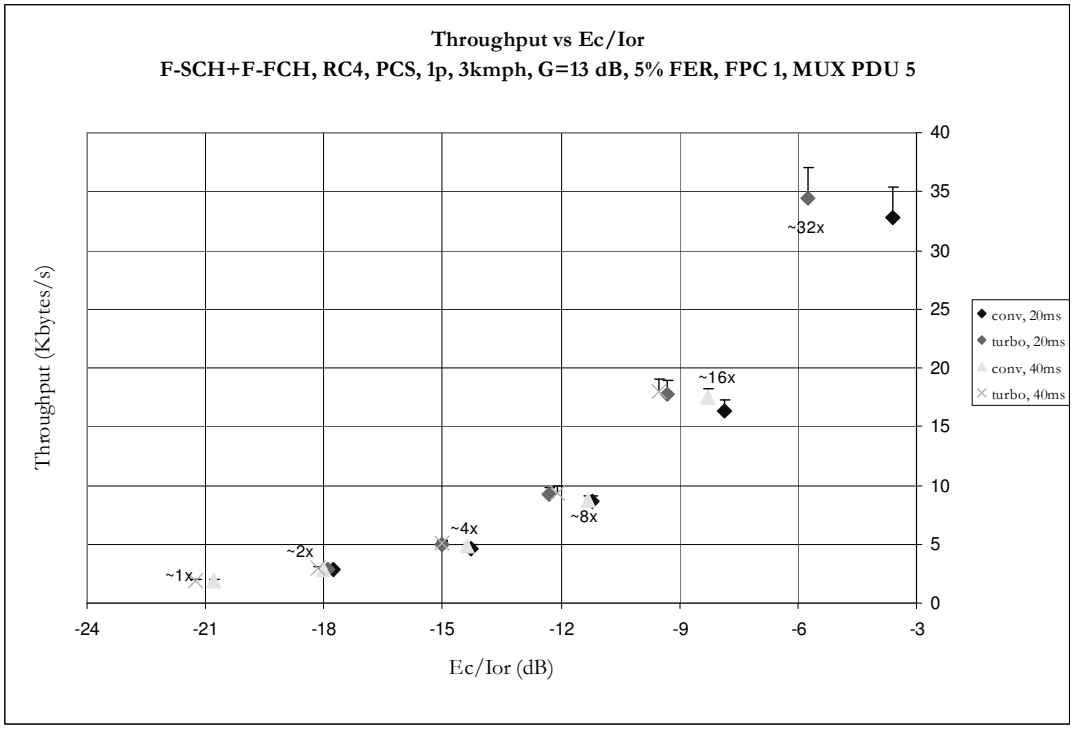

Figure 17: Graph of throughput vs E_c/I_{or} for a 1path, 3kmph, fading channel with MUX PDU 5.

We also looked at MFI performance in single path scenarios as a function of mobile speed. Each point on the following graph shows the percent difference in throughput efficiency between the 40ms case and the 20ms case. Throughput efficiency is a metric that combines throughput and E_c/I_{or} performance and is defined as average throughput divided by mean E_c/I_{or} (in linear units). Thus, this metric is directly proportional to throughput and inversely proportional to E_c/I_{or}. Moreover, the graph shows results for both 5% and 10% FER.

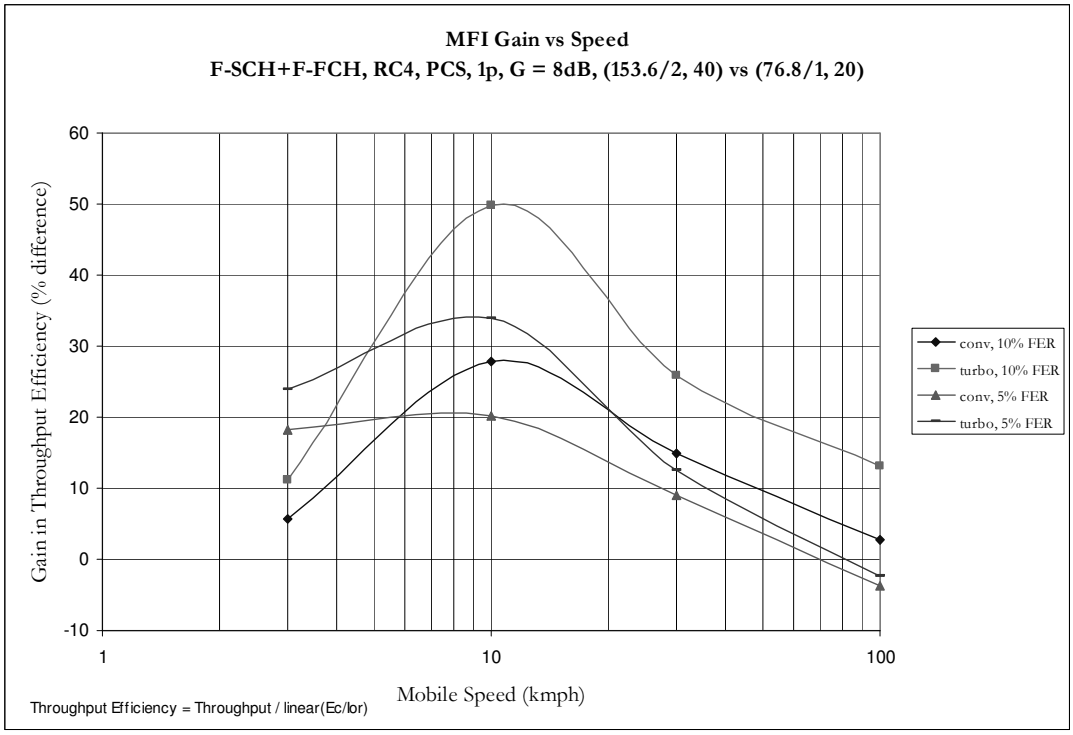

Figure 18: Graph of throughput efficiency gain vs. speed for 1path, 8dB geometry.

The use of throughput efficiency comes from the need to consider throughput and E_c/I_{or} performance at the same time with both given equal weight. Also, throughput efficiency allows us to compare related points on the throughput versus E_c/I_{or} graphs above. For example, we can compare one point with high throughput and high E_c/I_{or} to one that has lower throughput and lower E_c/I_{or}.

4.3.4 Analysis

The sanity check was designed to isolate the protocol stacks by eliminating the channel (no noise and maximum transmission power). The results were very good, all within 1.5% of baseline. There was one problem with the base station setup script which thankfully this sanity check revealed.

There are a few features to point out in the throughput versus E_c/I_{or} graphs:

- Points with the same effective data are clustered together. This behavior has been observed in our other tests as transmitting at the same data rate requires about the same E_b/N_t.

- Throughputs are approximately 5% away from their maximum values. This difference is directly caused by power control targeting a 5% frame error rate. However, in the case of "~32x" the difference may be more than 5% if the setpoint ever hits the rail of -3dB.
- In general, turbo encoding requires less E_c/I_{or} and, for MUX PDU 5, produces higher throughputs.
- With the exception of Model E[7], all other models showed little (if not inconsistent) difference between 20ms and 40ms results.

In light of the last observation, we went back to the "bread-and-butter" channel model for MFI: 1path, 10kmph Rayleigh. Although this channel model is not much different from Model A, the results showed significant improvements of 40ms over 20ms.

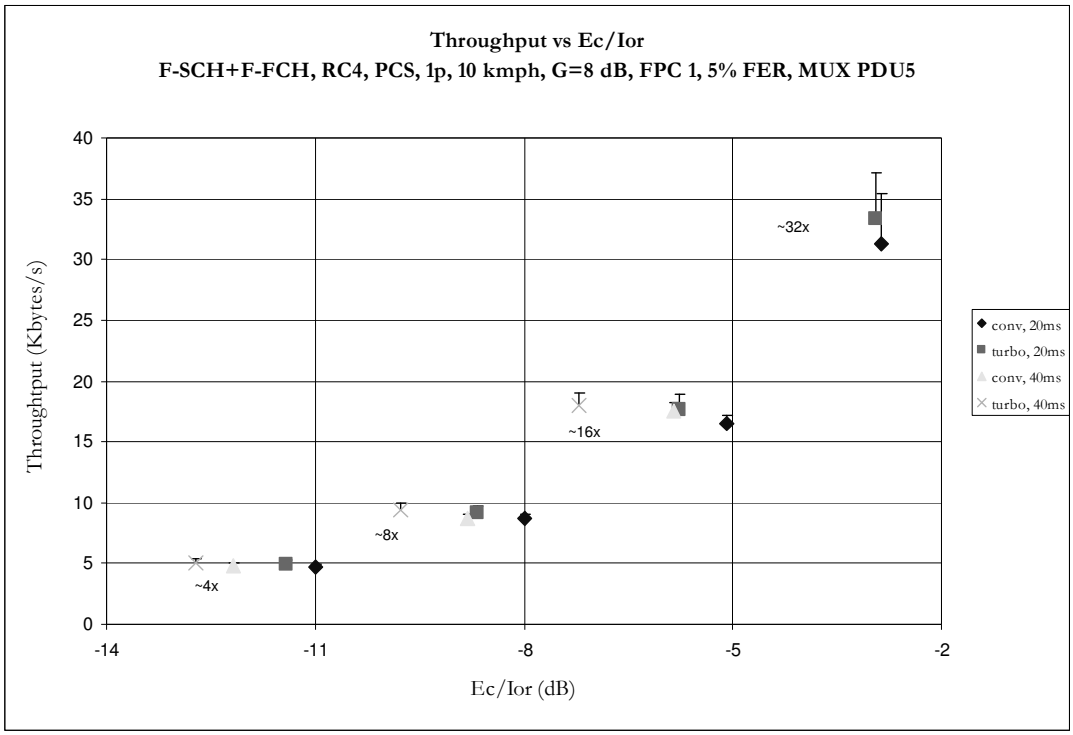

Figure 19: Graph of throughput vs. E_c/I_{or} in a 1path, 10kmph fading channel.

[7] Channel models A through E are defined in the EV-DV evaluation methodology [8].

In all cases, 40ms required less E_c/I_{or} than the comparable 20ms. We are not sure why such a small change in channel model would produce such a large increase in gains; it could be due to unverified features in the Panda emulator.

These latest results show promise for our scheduler: if 40ms frames require less E_c/I_{or}, hopefully the savings can be applied to other users for higher assigned data rates.

The analysis of throughput efficiency at varying speeds (Figure 18) confirms our initial intuition that MFI has the most benefit at slow speeds, with the greatest gain around 8-20kmph. At very slow speeds, the coherence time is about 40ms[8]—more time than even a 40ms frame structure can separate its symbols and for higher speeds, coherence time becomes less of an issue with faster fading. Also, note that at 10% FER, the gain of 40ms throughput efficiency over 20ms is larger than that at 5% FER. It had been discussed earlier that running longer frames at higher FER would be beneficial since that equated to targeting (almost) the same BER. In any case, studies have shown that targeting about 10% FER is optimal in terms of throughput efficiency.

4.4 Scheduler Performance

Performance testing with the scheduler brings us another step closer to assessing the benefits of MFI in a realistic mobile environment. In these series of tests, additional users are introduced into the sector and the shared resources of the base station are appropriated by the scheduler described in Section 2.1.5. Results show that total sector throughput increases for longer frames, although not by much. SCH assignments indicate the application of power savings toward other users or for higher data rates. In any case, the difficulty of predicting channel conditions limits the effectiveness of the scheduler in thoroughly capitalizing on the link level gains of MFI.

[8] The coherence time is calculated for a carrier frequency of 1900MHz and mobile speed of 1m/s.

4.4.1 Test Setup

Three users (represented by three laptops each connected to the IWF through its own SURF) try to download a large file. Each SURF shares the same RF connection to the Panda which means every user experiences the same fading environment. This condition is unrealistic but does not change the average sector throughput. Finally, while each user tries to download a large file—enough for a few minutes of transmission—their data rates are recorded. Also, our software logs the base station's assignment decisions and each assignment's E_c/I_{or}.

The following table lists the parameters of the scheduler tests.

Parameter	Value(s)
Radio Configuration	RC4
Band Class	PCS (1900 MHz)
Power Control	On, Target SCH FER = 5%, FCH FER = 1%, Mode 1
Coding	Convolutional, Turbo
Fading	1path, 10kmph
Geometry	4 to 12dB
Data Rates (effective)	(assigned by scheduler)
Frame Length	20ms, 40ms
Number of Users	3
Test Duration	2-3 minutes
MUX PDU	5
Burst Duration	32, 64, 128, 256 frames

Table 5: Table of parameters for scheduler tests.

Since each user is assigned the highest data rate the base station is able to support, three users are able to sufficiently load the system. Due to time constraints, only a one path slow fading scenario was explored.

4.4.2 Data

The following graph shows the average total sector throughput in thousands of bytes per second given the geometry. The burst assignment duration is 64 frames.

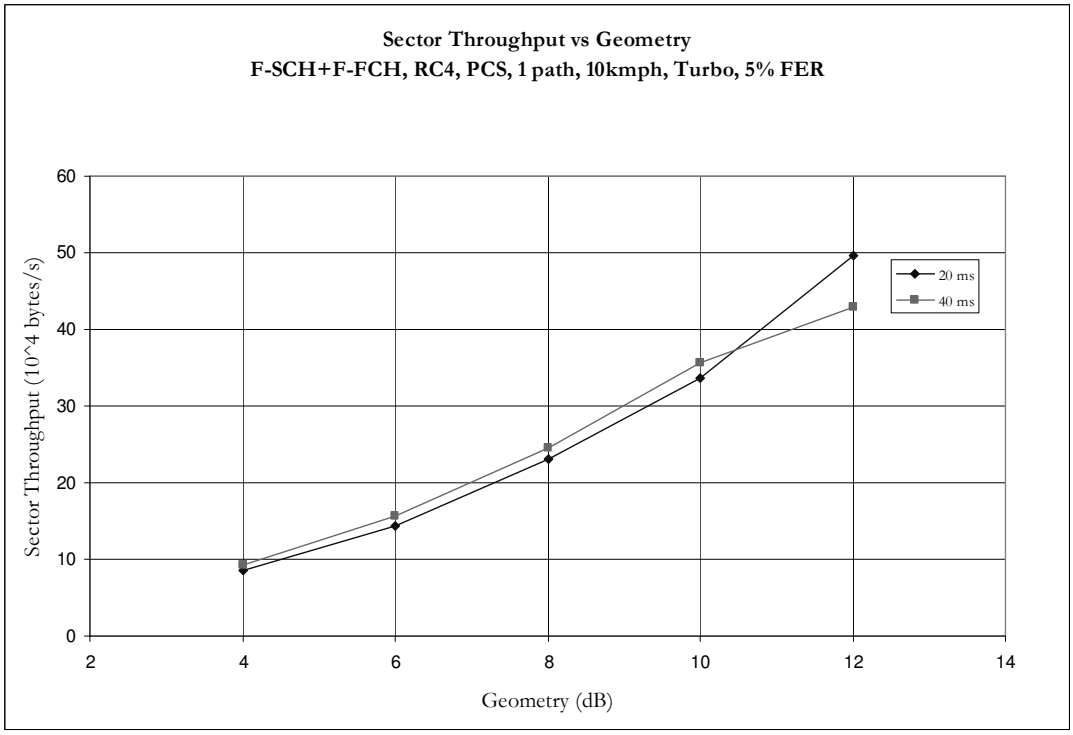

Figure 20: Graph of sector throughput for 3 users, 1path, 10kmph, and 64 frame bursts.

Table 6 shows the SCH rate assignments given the geometry of the channel. A "0x" indicates that the user was not given an assignment; a "~1x" means that the user was assigned a channel with effective data rate 9.6 kbps, and so on. For "~32x" there is no 40ms equivalent so the number of 40ms SCH assignments at that rate is always zero. For example, at 6dB geometry, a user was assigned a "~16x" (153.6kbps) SCH 4% of the time when using 20ms frames compared to 8% of the time when using 40ms.

		Average Assignments by Rate							Overall Average Assignment (x)
		0x	**~1x**	**~2x**	**~4x**	**~8x**	**~16x**	**~32x**	
G=4dB	20ms	0.39	0.16	0.15	0.22	0.08	0	0	1.99
	40ms	0.37	0.16	0.16	0.19	0.12	0	--	2.20
G=6dB	20ms	0.23	0.16	0.14	0.16	0.27	0.04	0	3.89
	40ms	0.22	0.16	0.14	0.17	0.24	0.08	--	4.32
G=8dB	20ms	0.17	0.12	0.14	0.15	0.14	0.27	0.01	6.80
	40ms	0.12	0.11	0.14	0.17	0.13	0.32	--	7.32
G=10dB	20ms	0.12	0.09	0.12	.014	0.17	0.22	0.14	10.27
	40ms	0.03	0.03	0.17	0.13	0.18	0.55	--	10.99
G=12dB	20ms	0.05	0.05	0.10	0.13	0.16	0.19	0.33	15.74
	40ms	0	0	0	0.01	0.32	0.67	--	13.25

Table 6: Table of SCH assignment history for 3 users, 1path, 10kmph, and 64 frame bursts.

Next, we consider the total throughput efficiency as a function of speed. The throughput efficiency of each user is computed in the same way as in the throughput tests, and the sum over each user is the total throughput efficiency of the sector. Figure 21 shows the percent difference between the 20 and 40ms values.

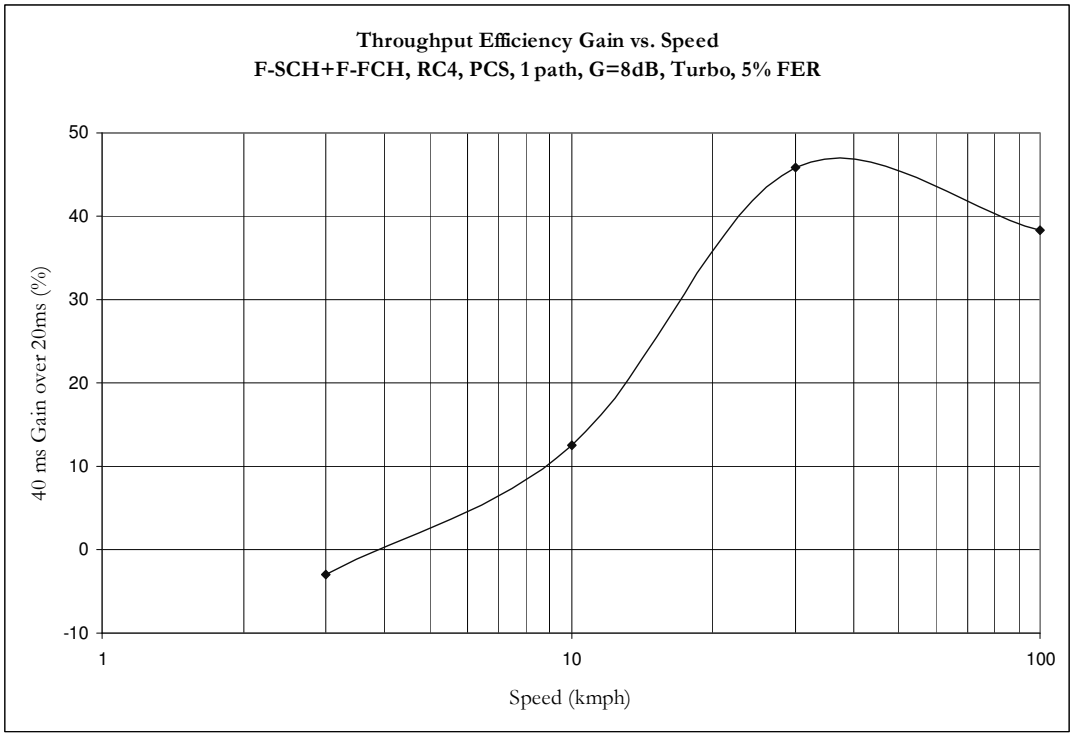

Figure 21: Graph of throughput efficiency improvements of 40ms frames vs. speed, with 3 users in 1path fading.

Finally, we look at the total throughput efficiency of the sector as a function of the burst duration. The burst durations selected were 32 frames (.64s), 64 frames (1.28s), 128 frames (2.56s), and 256 frames (5.12s). Figure 22 shows the sector throughput efficiency for 20 and 40ms frames.

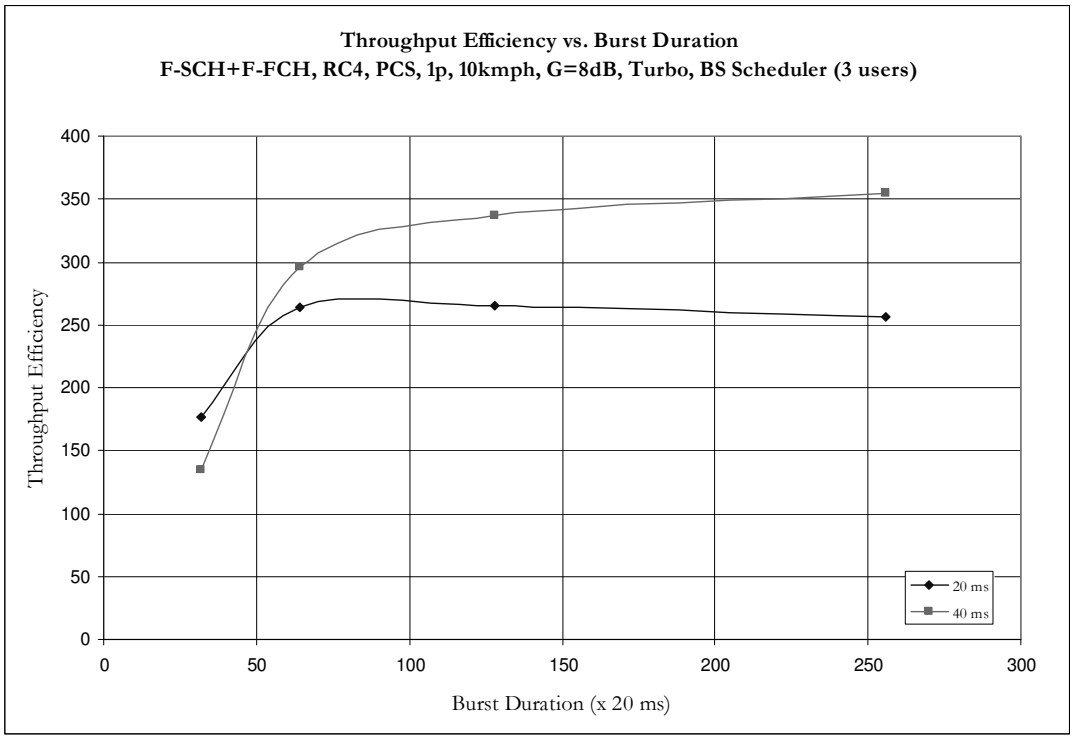

Figure 22: Graph of throughput efficiency vs. burst duration for 1path fading with 3 users.

4.4.3 Analysis

As mentioned in Section 2.1.5 about the forward link scheduler, power usage estimates are obtained by scaling the average power used for the fundamental channel. These scaling factors (or delta values when using units of dB) directly determine how aggressive the scheduler is in assigning channel rates, i.e. small scaling factors will lead the scheduler to be optimistic about future power requirements, whereas large scaling factors will do the opposite.

A couple limitations of the system make the scheduler's heavy dependence on the value of the scaling factors an issue. First, the values are static—one set of delta values must work for all types of channel conditions. Therefore, for any given assignment, the

scheduler's prediction of power requirements may be overly pessimistic or optimistic depending on the programmed delta values and the properties of the channel. This will lead to either underutilized system resources or, as we will see, due to the next limitation, unrealistic "over usage."

The second limitation is that in our test system, there is neither a software nor hardware limitation on the total power that is actually transmitted from the sector. Therefore, if the users experience a deep fade, forward power control could push each user's E_c/I_{or} to the maximum (usually -3dB), meaning the total E_c/I_{or} could exceed unity. This will cause excessive interference to neighboring sectors as well as provide unrealistic results since a real system would force such SCH channels to DTX.

The process to get around these two limitations is to use one type of channel condition at a time (as we have been doing) and to manually adjust the delta values. The rationalization for manually adjusting the delta values is that a more advanced scheduler could analyze its past predictions and dynamically change the scaling factors depending on whether power resources were being underutilized or if total E_c/I_{or} was exceeding its bounds.[9] The methodology behind manually adjusting the delta values was to target a certain aggressiveness which is the ratio of average total SCH power to average power available, where total power available is equal to maximum power minus total FCH power minus total overhead power. A good aggressiveness turned out to be 0.75. This value provided a good balance between using up available resources and leaving some margin for when power might spike up due to deep fades.

As we can see from Figure 20, sector throughput is higher for 40ms frames than for 20ms. However, the difference is not substantial—less than 10% improvement. While it seems that sector throughput is worse for 40ms at 12dB geometry, Table 6 shows that the system is actually Walsh limited with three users[10]. Introducing a fourth

[9] Dynamically changing the scale factors would only be plausible if the channel characteristics were not changing very fast, as is the case in our tests.

[10] "~16x" is the highest rate with 40ms SCH frames and the Walsh space limits us to at most two at a time. The next highest rate is "~8x". Since the scheduler assigns "~16x" two-thirds of the time and "~8x" the other third, and since power resources are not fully utilized, we conclude the system is Walsh-limited.

user would use up the available power and increase the total sector throughput for that scenario.

While the sector throughput provides some insight into the benefits of MFI, the SCH assignment history is perhaps more revealing. The data (Table 6) indeed shows that the average assignment is higher for 40ms than for 20ms. Moreover, users are denied an assignment less often at 40ms. This increase in availability is the result of what we said MFI would do: leverage power savings from other assignments to create room in the power pool for more users. At higher geometries, the assigned data rates are higher which means bigger savings that can be used to support lower data rates. However, at lower geometries, with mostly low data rates being assigned, the savings are usually too small to be applied to other users. This explains why users are denied an assignment much less at the higher geometries.

Figure 21 shows that throughput efficiency gains increase as speed increases and peak after 30kmph, with a maximum of about 50% better performance over 20ms frames. This data further supports our notion that MFI has maximum benefit for slow fading channels.

In Figure 22, there is substantial difference between the throughput efficiencies of 40ms frames when the burst duration is 32 frames versus 64 or more frames. This phenomenon is due to a slower outer loop when using longer frames[11]. Moreover, power control is stopping and starting every burst assignment, often for different data rates. This fluctuation in channel assignment causes power control to be less efficient, and, from an experimental standpoint, increases the variance in the measured average power over the test period. Nevertheless, the data suggests running 40ms frames with long burst durations to best utilize link level gains. However, long burst durations are undesirable unless the system can guarantee that data will be sent for the entirety of the assignment. Thus, there is a tradeoff between throughput efficiency (getting higher data rates for some power usage) and scheduling efficiency (assigning SCHs so that no bandwidth is wasted).

---

[11] The outer loop runs on frame boundaries; thus, the target E_b/N_t updates with a frequency of 50Hz for 20ms frames, 25Hz for 40ms frames, and 12.5Hz for 80ms frames.

Admittedly, the sector throughput results, even when using a reasonably long duration of 64 frames, are not that impressive. At 8dB geometry, the sector throughput gain is about 6%, while sector throughput efficiency gain is about 12%. This indicates that the system, for the 40ms case, is using less power than the 20ms case. Had the 40ms case been more aggressive with its assignments, perhaps its sector throughput would have been higher. This situation illustrates how hard scheduling for optimal power usage is. While results from the link level tests show very good gains in E_c/I_{or}, the scheduler is ultimately the system component that has to exploit the gains from each assignment into something beneficial. Invariably, the results of each assignment for 20 and 40ms should fall somewhere between no difference and 40ms have a higher total assignment. Therefore, we expect something in between and, as the results show, it is closer to no difference at all.

4.5 Latency

Here, we analyze Multiframe Interleaving from a different perspective: the added latency of 40 and 80ms SCH frames. Since the frames are longer, it should increase the link to link propagation delay. We compare theoretical delays with measured values, and the results are consistent for single frame transmissions. However, there is unexplained behavior when two frames are transmitted.

4.5.1 Test Setup

Using the standard network ping utility, we pinged an internal network site to measure round trip times. Given the header sizes of the various protocol layers, we can choose the size of the ping packets and data rate such that only one or two SCH frames are transmitted. The round trip times were then analyzed to determine the added delay introduced by using longer frames.

Table 7 shows the ping packet size and SDU type needed to achieve exactly one or two SCH frames.

# SCH Frames	20ms		40ms	
	Rate	**Size**	**Rate**	**Size**
1	4x (38.4/1, 20)	32 bytes	2x (19.2/3, 40)	32 bytes
2	4x (38.4/1, 20)	99 bytes	4x (38.4/2, 40)	99 bytes

Table 7: Table of ping packet sizes and SDU sizes for ping test.

4.5.2 Data

The average round trip times for 20ms frames were subtracted from the 40ms times to produce the relative latency.

# SCH Frames	40ms Latency Relative to 20ms	
	Expected	Experiment
1 (32 bytes)	+30 ms	+29.4 ms
2 (99 bytes)	+50 ms	+62.65 ms

Table 8: Table of latency of 40ms frames relative to 20ms

4.5.3 Analysis

We assume that the added latency of longer frames comes from the longer wait time in the MUX buffer and the longer transmission time, all on the base station side. We will call this the MFI delay. The wait time in the MUX buffer is approximately the time from which the packet enters the buffer to the next frame boundary. If we assume the packet arrival time at the MUX buffer layer is uniformly distributed along a frame, then the expected wait time is simply the frame length divided by two. The transmission time is the frame length times the number of SCH frames needed to transmit the ping packet.

Thus the expected latency for sending one 40ms SCH frame should be 30ms more than sending one 20ms frame. The experimental value supports this theory. For two 40ms frames, the expected increase in latency is 50ms. However, the experimental data is about 60ms—there is 10ms unaccounted for. One possible explanation is the additional time needed to process the extra frame.

The latency introduced by MFI is substantial when compared to internet propagation times and is a disadvantage of the feature. Such delays could cause TCP to presume the channel is worse than actuality, resulting in inefficient operation. While, the effect on upper layers would make an excellent subject of study, it is unfortunately outside the scope of this thesis.

5 Conclusions

5.1 Summary of Performance Analysis

Multiframe Interleaving shows the greatest improvement over 20ms frames at slow speeds and for single path fading. The benefits are observed in the link level curves, where there is as much as 3dB gain between 20ms and 80ms[1], and in the throughput efficiency tests, where there is about a 35% increase[2].

However, improvements diminish as mobile speed increases and as diversity is provided by other means, namely multipath. While realistic pedestrian speeds can be safely assumed to be slow, non-multipath is unlikely to be experienced in the field, especially in an urban environment where trees, buildings, and other structures provide ample surface area for reflecting waves. In any case, slow speed with multipath results show some link level gain but nowhere near the degree of non-multipath.

Even when using slow speed and single path fading for the scheduler tests, the overall gain in sector throughput improvement is less than 10%. SCH assignment histories revealed that the power savings for running at longer frames were, on average, being utilized to provide higher data rates to users. However, due to difficulty in predicting power requirements, the scheduler's aggressiveness was hard to control. Thus, at each assignment period, the scheduler was not always able to improve on the 20ms assignment, a task also made harder when power savings were very small. It will take a very smart scheduler with reliable predictive power to fully leverage MFI's gains as seen at the link level. Moreover, the choice of burst assignment period impacts the results: longer durations allow power control to stabilize and converge to the performance seen at the link level[3], while shorter durations allow the scheduler finer granularity in making assignments.

Hence, the limitations of lab testing and our lab testing methods in general—single user, long burst assignments, fixed channel environments—proved to be auspicious

---

[1] At 1path, 3kmph, 4dB geometry, 5% FER, with convolutional encoding.
[2] At 1path, 10kmph, 8dB geometry, 5%FER, with convolutional encoding.
[3] Recall that the link level tests were done with "infinite" burst assignments.

for MFI performance. However, none of those conditions hold in practice and barring perfect execution in our system, MFI fails to live up to its preliminary analysis.

While the final results were not the most encouraging, this study served its purpose in assessing the performance benefits for longer supplemental channel frames. It also provided a proof of concept through a working implementation and saved on possible field testing, which requires more valuable resources. Nevertheless, more study may be desired to totally discount Multiframe Interleaving's benefits in improving capacity and performance. There are some scenarios where MFI works great and some where its advantage is marginal; but in general, 40ms frames seems like it does no worse.

5.2 Future Work

Any immediate future work should focus on fixing RLP's delayed frame detection (see Appendix A.) The current implementation works for a non-MFI system but there is a problem when running with 80ms frames. Fixing the problem and conducting the experiments again with 80ms frames would almost complete the picture; the culmination would be exploring different fading scenarios, particularly multipath channels, with the scheduler. Despite less than spectacular results in single path cases, more data for other channel models will serve to broaden our understanding and perhaps give further insight about our current results or even ways to bring about improvements.

Second, this study was not able to run any type of application that would cause the forward SCH channel to DTX. The RDA is only programmed to categorize 20ms frames, as the algorithm relies on extensive data sets to do a nearest-neighbor type of classification. Collecting data for 40 and 80ms frames to program into the RDA would allow for DTX detection to be turned on and loosen restrictions on testing parameters.

Lastly, analysis of added-latency effects received little attention in this study but warrants further investigation. Applications that are described as bursty, such as HTTP, may suffer performance loss as SCH assignments take at least almost twice as long to get set up with longer frame durations. It would also be interesting to see the reaction of TCP to the increase in delay of the link.

6 Bibliography

Digital Communications

[1] Proakis, John G. *Digital Communications*. McGraw-Hill. 2nd Edition, 1989.

[2] Sklar, Bernard. *Digital Communications: Fundamentals and Applications*. Prentice Hall. 1988.

[3] Gallager, Robert. *Principles of Digital Communications*. MIT, 6.450 Lecture Notes. Fall 2002.

[4] Feramez, Michael. *Fundamentals of Radio Communications*. La Trobe University, ELE 52PMC Lecture 6 Notes. August 29, 2003.

[5] Feher, Kamilo. *Advanced Digital Communications*. Crestone Engineering. 1996.

[6] Rice, S.O. *Mathematical Analysis of Random Noise*. Bell Syst. Tech. J., Vol. 23, July 1944, p. 282, and Vol. 24, January, 1945, p. 46.

Standards

[7] Telecommunications Industry Association. *Physical Layer Standard for cdma2000 Spread Spectrum Systems, Addendum 2*. TIA/EIA/IS-2000.2-2. August 2001.

[8] Third Generation Partnership Project Two. *1xEV-DV Evaluation Methodology – Addendum (V6)*. WG5 Evaluation AHG. July 25, 2001.

Related Research

[9] Gutierrez, Alberto; Guo, Ning. "Sub-frames for cdma2000," Northern Telecom, TIA/EIA TR45.5 Subcommittee, Document TR45.5.3.1/98.08.18.37, August 1998.

[10] Gutierrez, Alberto; Li, Jun; Tong, Wen. "Multiframe Interleaver Performance Update and R-Rake Compatibility," Northern Telecom, TIA/EIA TR45.5 Subcommittee, Document TR45.5.3.1/98.11.17.33, November 1998.

[11] Kubota, Shuji; Kato, Shuzo; Feher, Kamilo. "A Time Diversity CDMA Scheme Employing Orthogonal Modulation for Time Varying Channels," *IEEE 43rd Vehicular Technology Conference*, pp. 444-447, 1993.

[12] Li, Jun; Gutierrez, Alberto. "Improvement of Link Performance for High Speed by Interleaving Over Multiple Frames," Northern Telecom, TIA/EIA TR45.5 Subcommittee, Document TR45.5.3.1/98.09.02.11, September 1998.

[13] Li, Jun; Gutierrez, Alberto. "Multiframe Interleaving – Performance & Scope," Northern Telecom, TIA/EIA TR45.5 Subcommittee, Document TR45.5.3.1/98.09.15.26, September 1998.

[14] Meyer, Michael. "Improvement of DS-CDMA Mobile Communications Systems by Symbol Splitting," *IEEE 45th Vehicular Technology Comference*, vol. 2, pp. 689-693, July 1995.

[15] Meyer, Michael. "Improvement of IS-95 by Symbol Splitting," *IEEE 4th International Symposium on Spread Spectrum Techniques and Applications*, vol. 3, pp. 1078-1081, September 1996

[16] Rowitch, Doug. "1X F-SCH Convolutional Code Performance," QUALCOMM, Inc.

[17] Rowitch, Doug. "1X F-SCH Turbo Code Performance," QUALCOMM, Inc.

7 Appendix A—MFI issues with RLP

7.1 Introduction

RLP or Radio Link Protocol is a NAK-based link layer protocol used with CDMA data services. RLP can significantly reduce the error rates of traffic channels as seen by upper layers such as TCP/IP. There are several RLP frame formats and depending on the channel rate, one or more RLP frames are transmitted in a single physical layer frame.

Since RLP frames are sequenced and transmitted (and hence received) in order, the receiver can detect missing frames as gaps in the received sequence numbers. To recover the missing frames, a NAK round is initiated by sending a NAK control frame to the transmitter. The round ends when the missing frame is retransmitted and successfully received or when a certain amount of time has elapsed. The RLP NAK scheme will specify the number of control frames per round and the number of rounds before giving up.

7.2 Issue

With longer frame durations, the current RLP implementation causes unwarranted NAKs to be sent from mobile to base station. The basic problem is that sequenced RLP frames, while transmitted in order at the BS, are received out of order at the MS because frames transmitted on the SCH have a longer latency than those transmitted on the FCH. When RLP detects out of order sequence numbers, it incorrectly assumes an erasure and sends a NAK for the corresponding frame. As it happens, this is only a problem with 80 ms frames; 40 ms frames do not cause unwarranted NAKs.

Figures 23 and 24 illustrate the problem. In Figure 23, at T = 0ms, one RLP frame is fitted into an FCH frame and four RLP frames are fitted into an SCH frame. These RLP frames are sequenced consecutively. At T = (20+D)ms, the mobile receives RLP frame 1. At T = 20ms, the base station sends the next RLP frame; this time, only the FCH is available for transmission. At T = (40+D)ms, RLP frame 6 is received. RLP does not know that frames 2 through 5 are still being transmitted and assumes they have been lost. Thus a NAK round begins. It just so happens that the 40ms case does

not have premature NAKs. This is due to the fact that the FCH frame gets filled before the SCH.

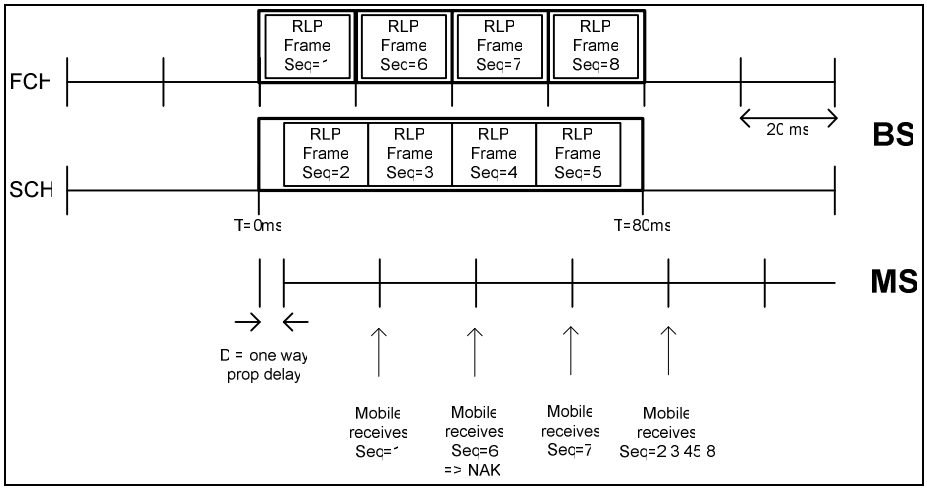

Figure 23: RLP Timeline for 80ms SCH Frames.

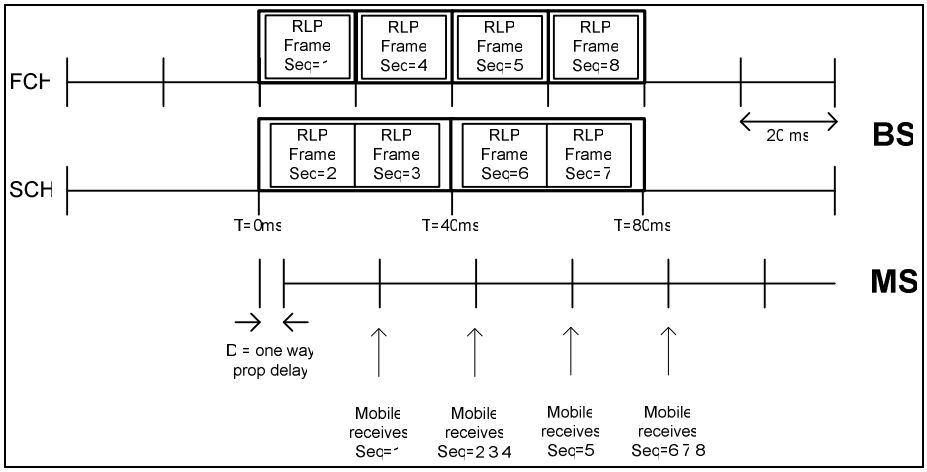

Figure 24: RLP Timeline for 40ms SCH Frames.

RLP 3 specifies a delayed frame detection algorithm that would eliminate the above problem. However, this detection was not implemented fully in the code and there are no current plans to complete it. Thus, any packet data testing using 80 ms frames is impossible.

8 Appendix B—Demodulator Performance Results

The following contains the complete set of data for section 4.1.

8.1 Summary

8.1.1 Convolutional, 20ms vs. 40ms vs. 80ms

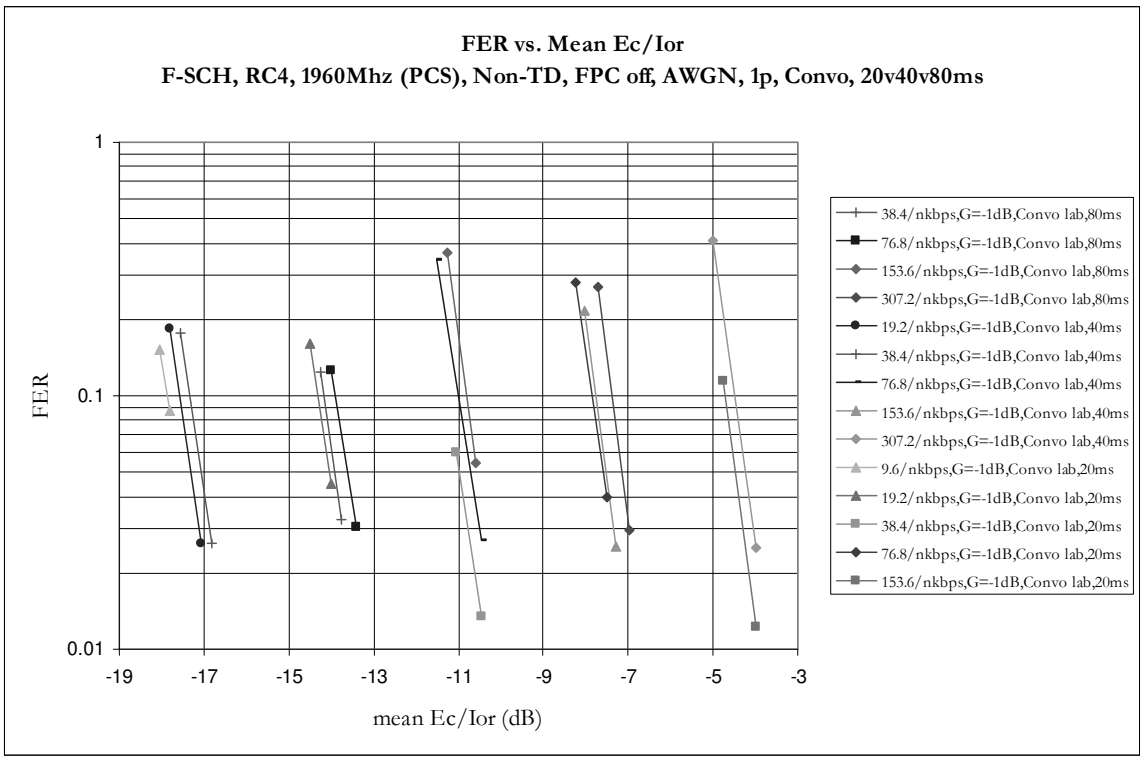

8.1.2 Turbo, 20ms vs. 40ms vs. 80ms

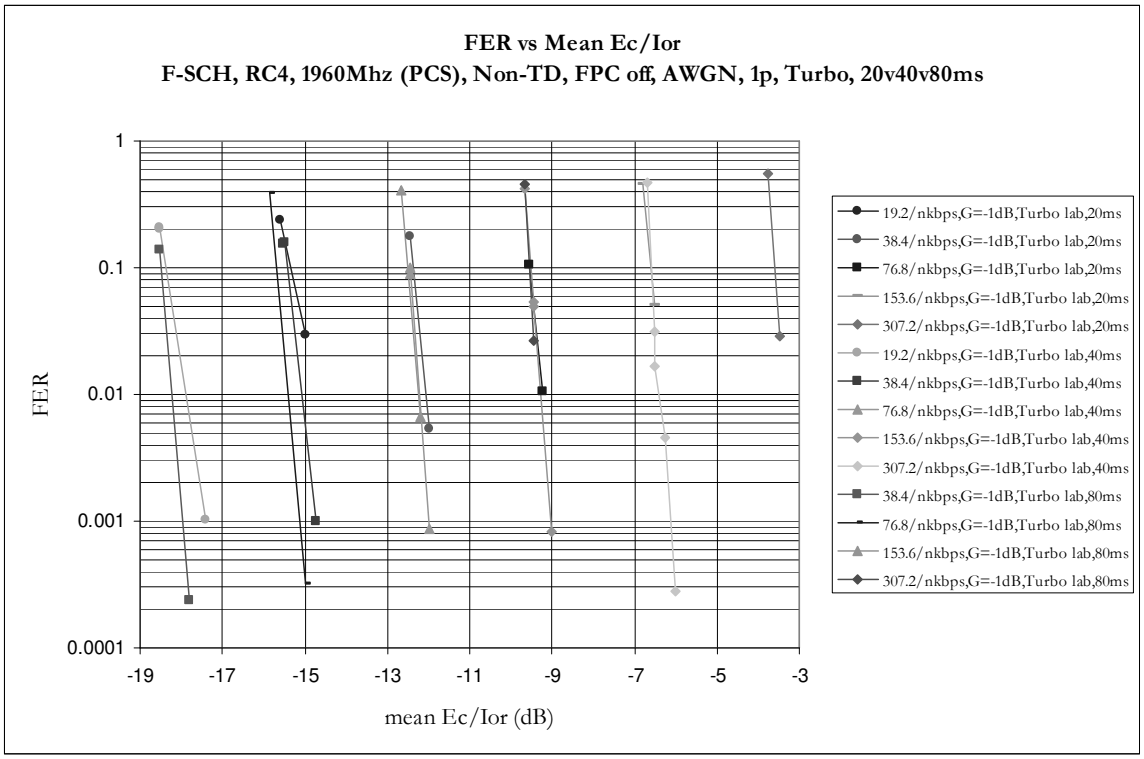

8.2 FER vs. Mean E_c/I_{or}

8.2.1 Convolutional Encoding

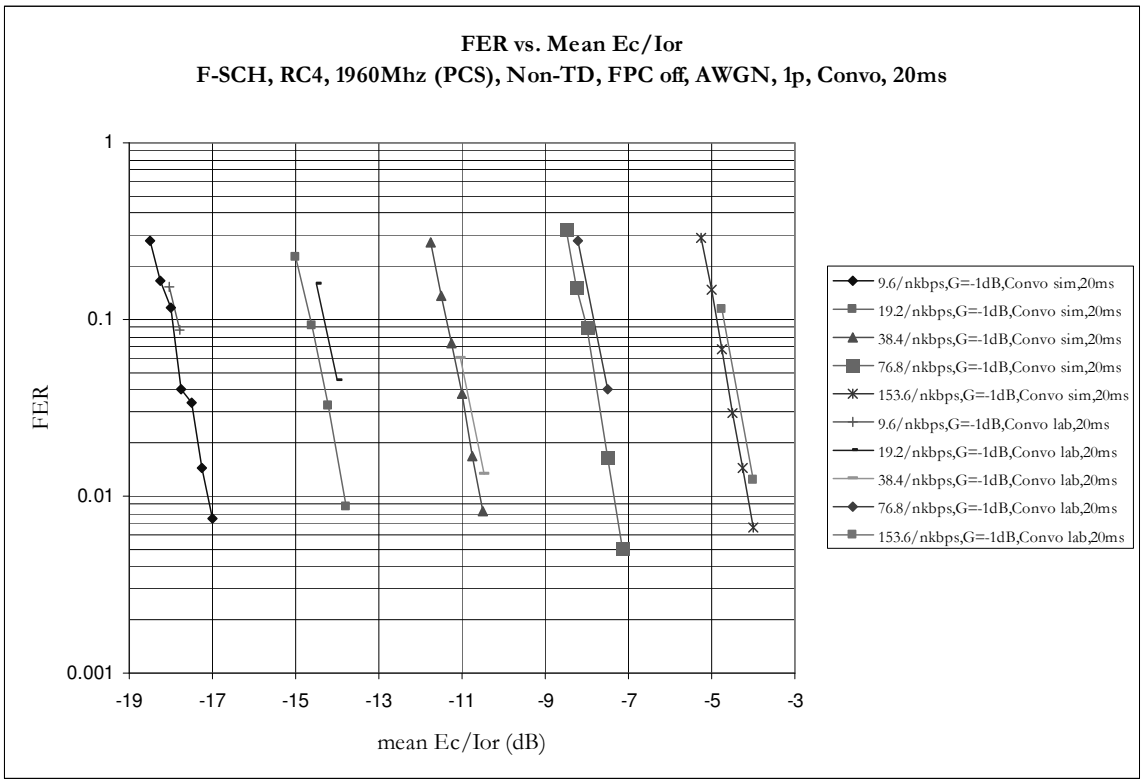

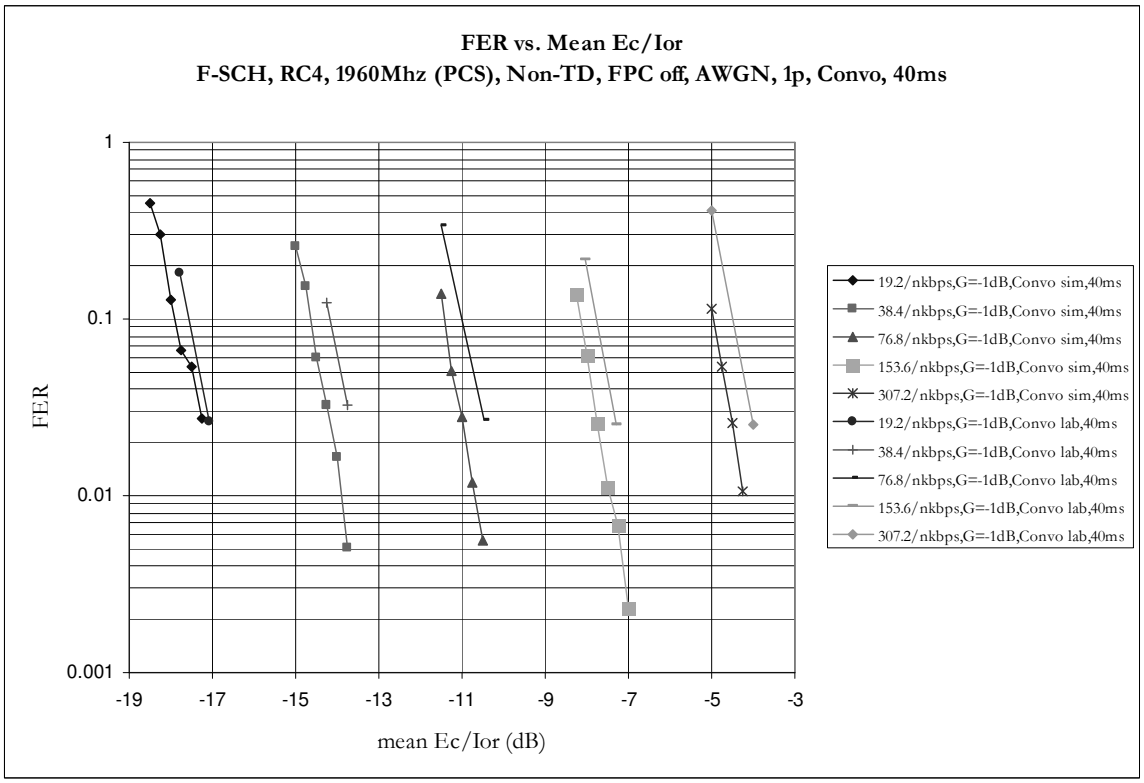

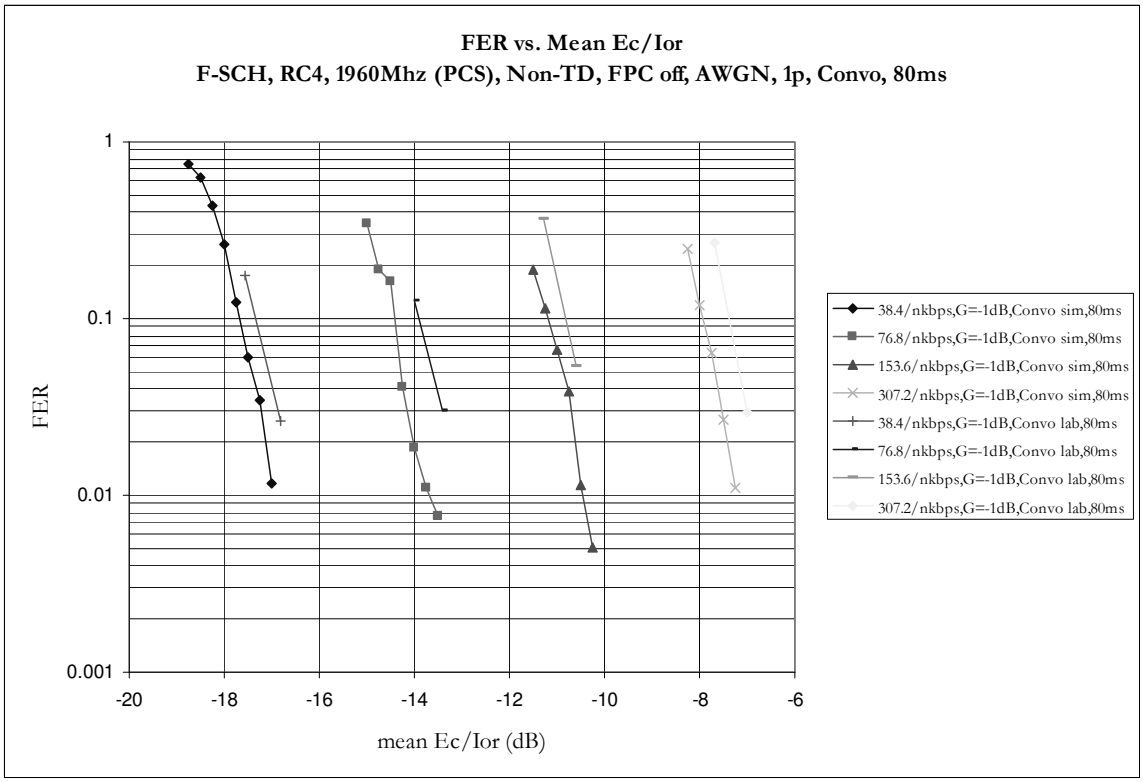

8.2.2 Turbo Encoding

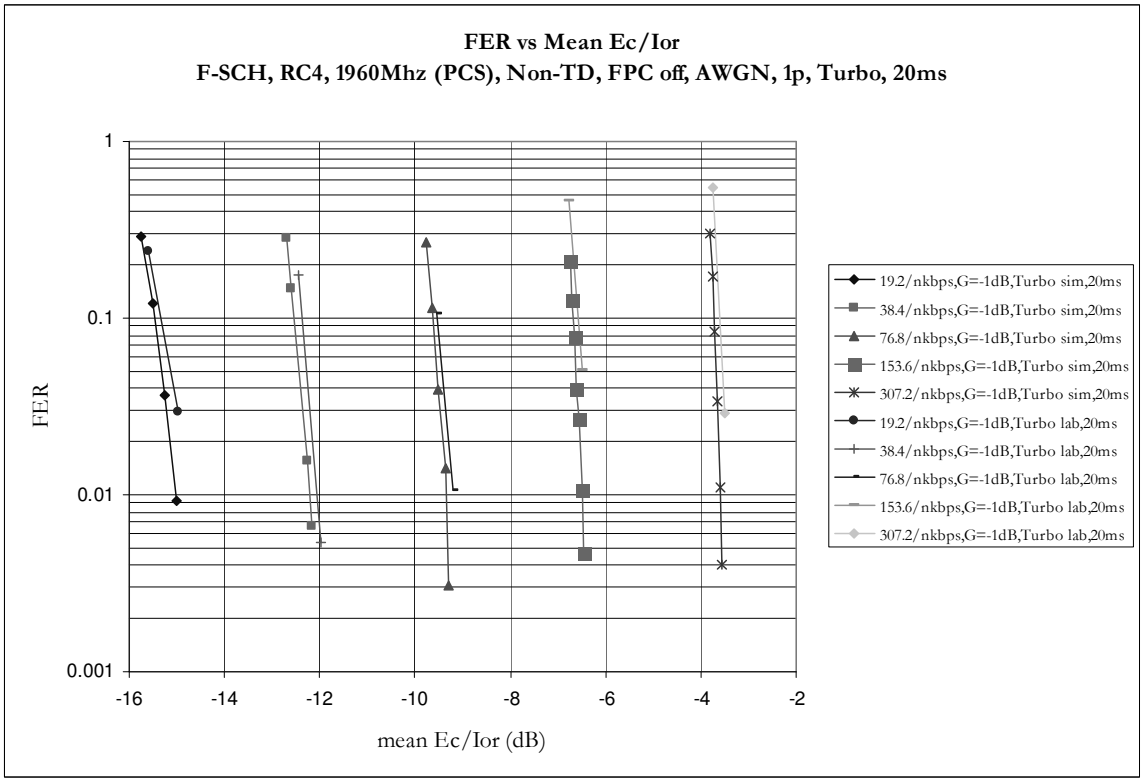

FER vs Mean Ec/Ior
F-SCH, RC4, 1960Mhz (PCS), Non-TD, FPC off, AWGN, 1p, Turbo, 40ms

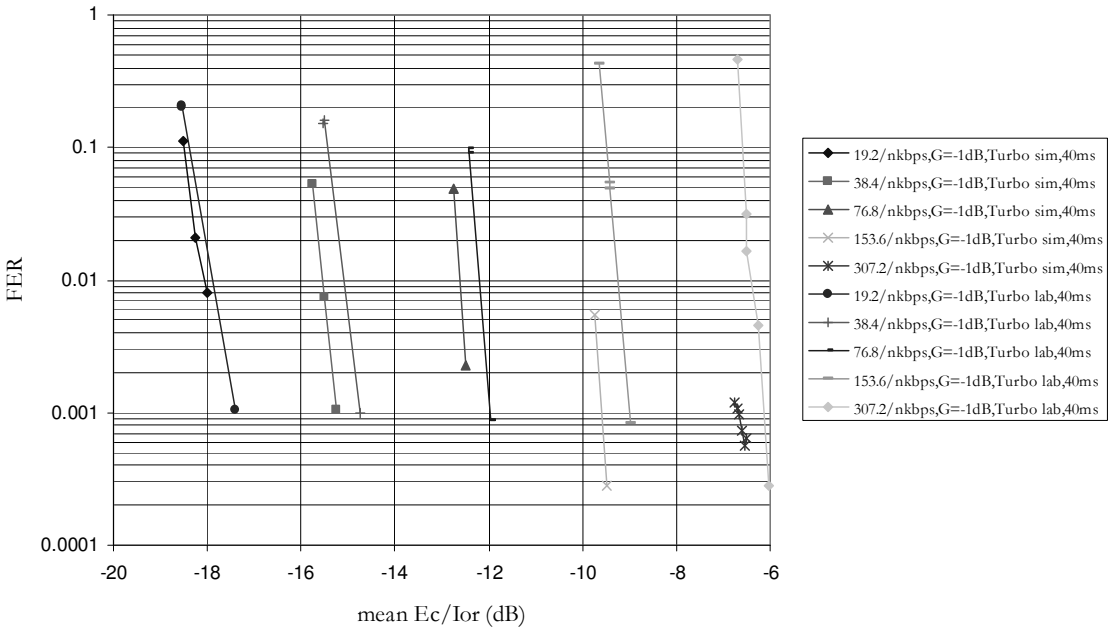

FER vs Mean Ec/Ior
F-SCH, RC4, 1960Mhz (PCS), Non-TD, FPC off, AWGN, 1p, Turbo, 80ms

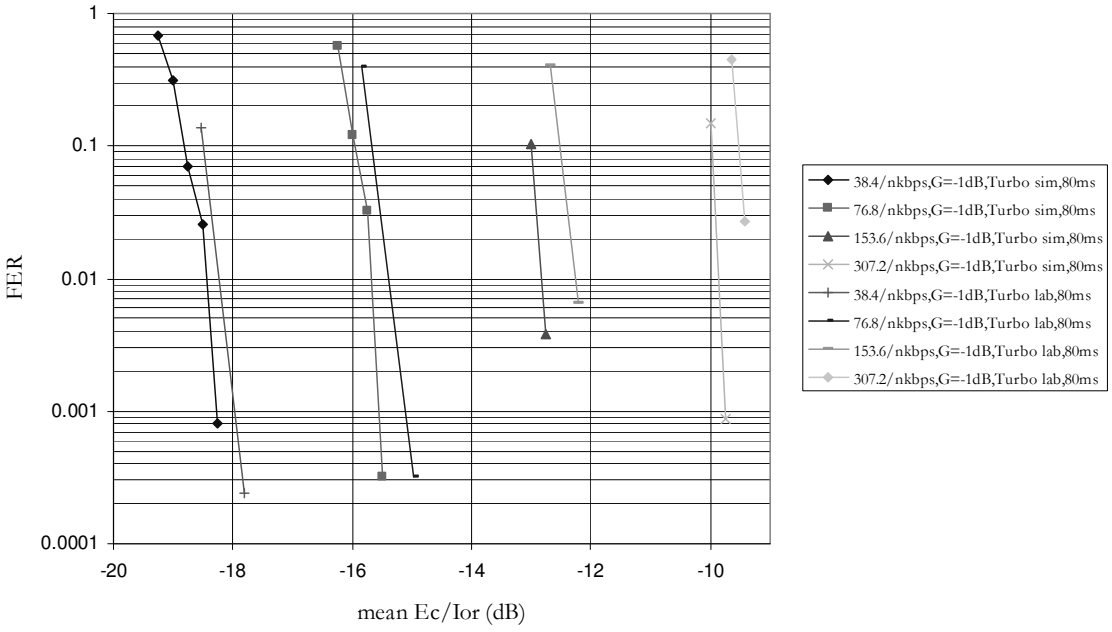

8.3 BER vs. Mean E_c/I_{or}

8.3.1 Convolutional Encoding

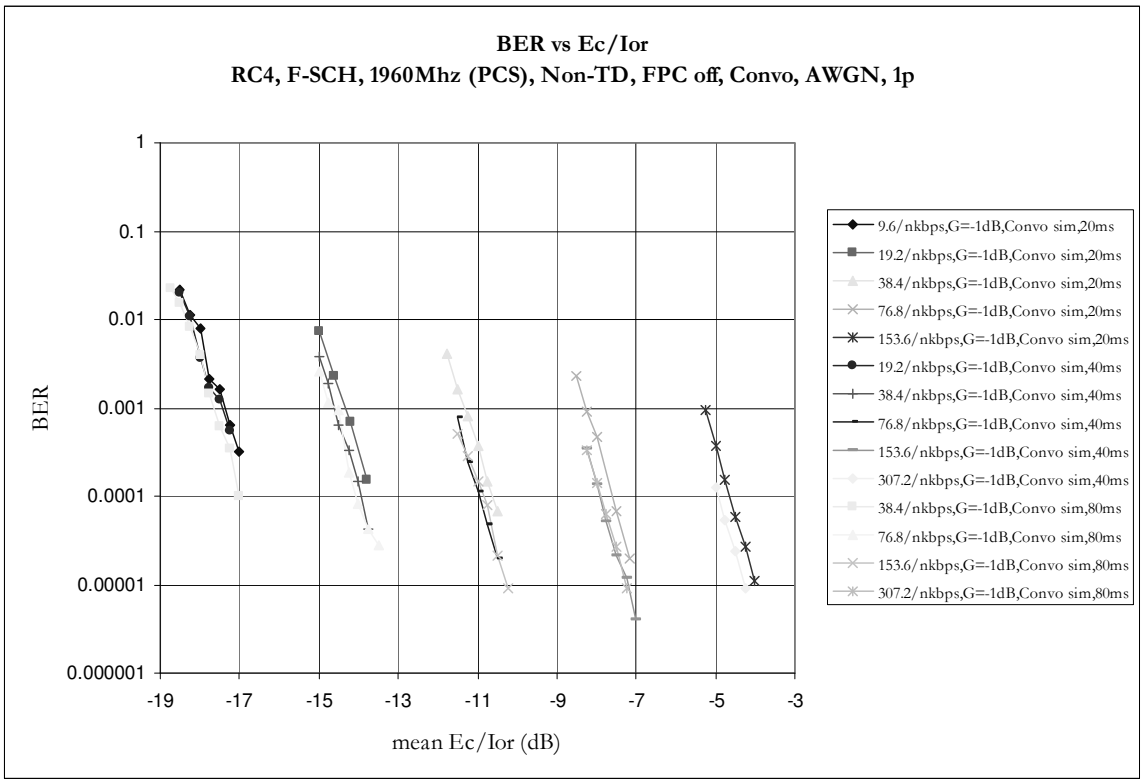

8.3.2 Turbo Encoding

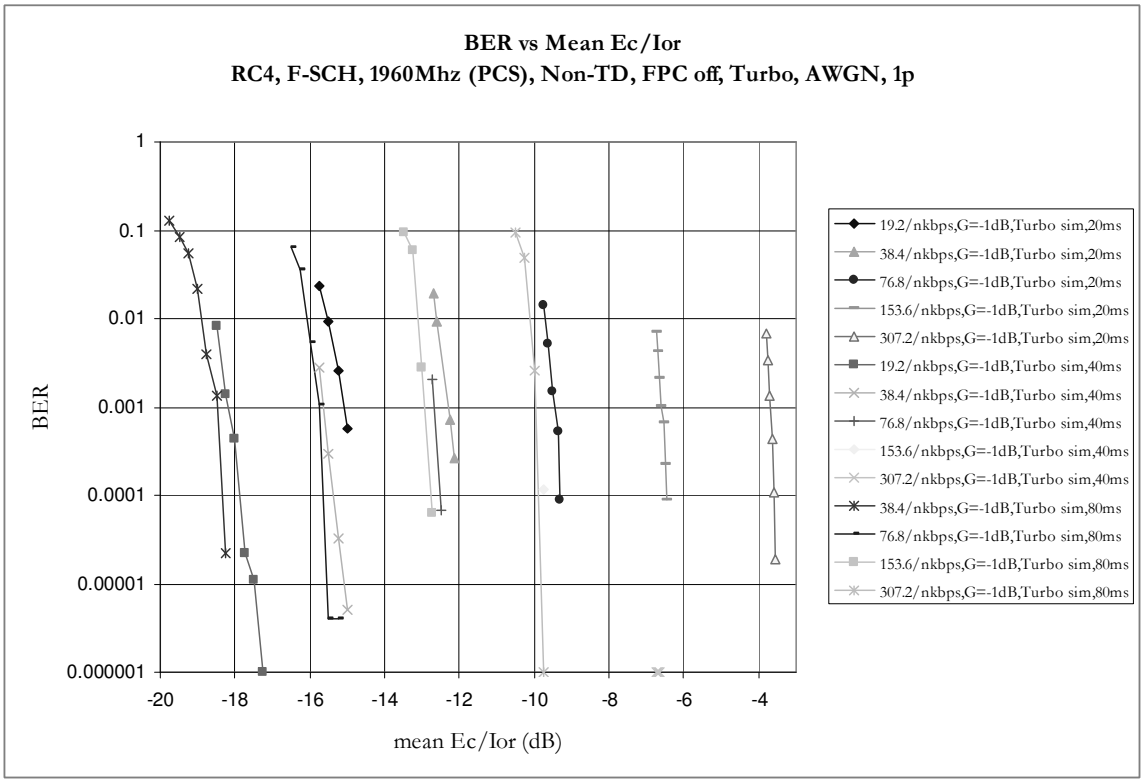

8.4 FER vs. BER

8.4.1 Convolutional Encoding

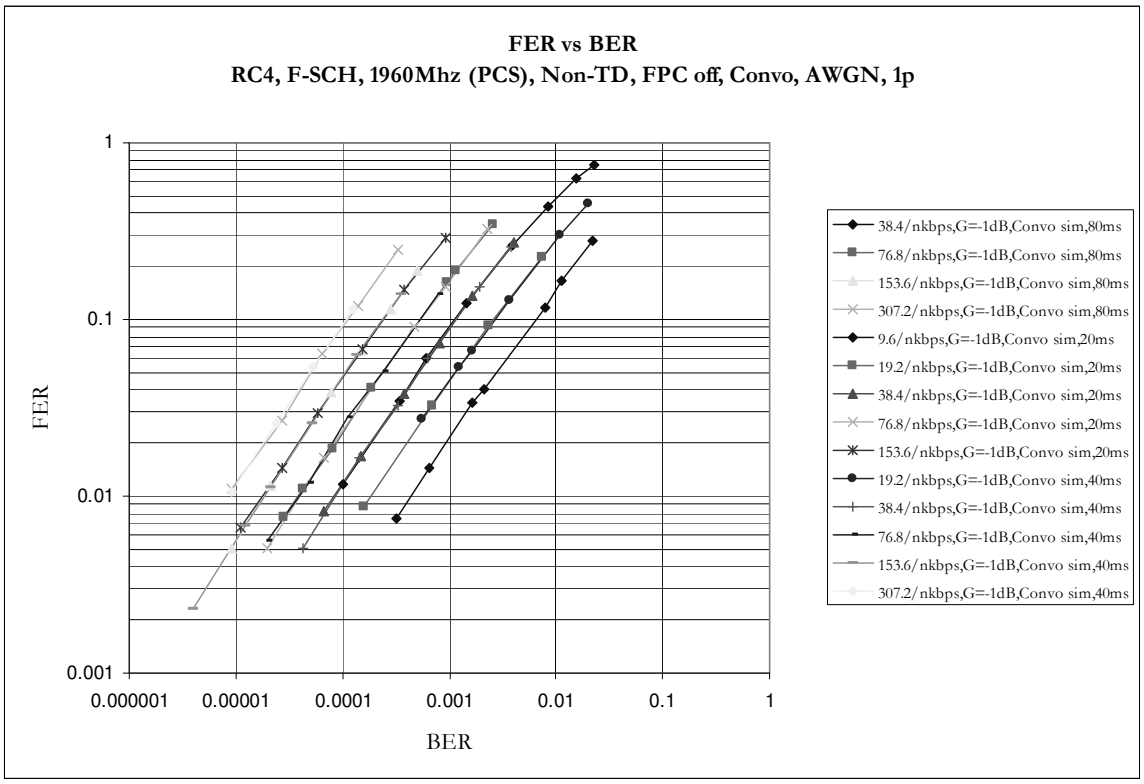

8.4.2 Turbo Encoding

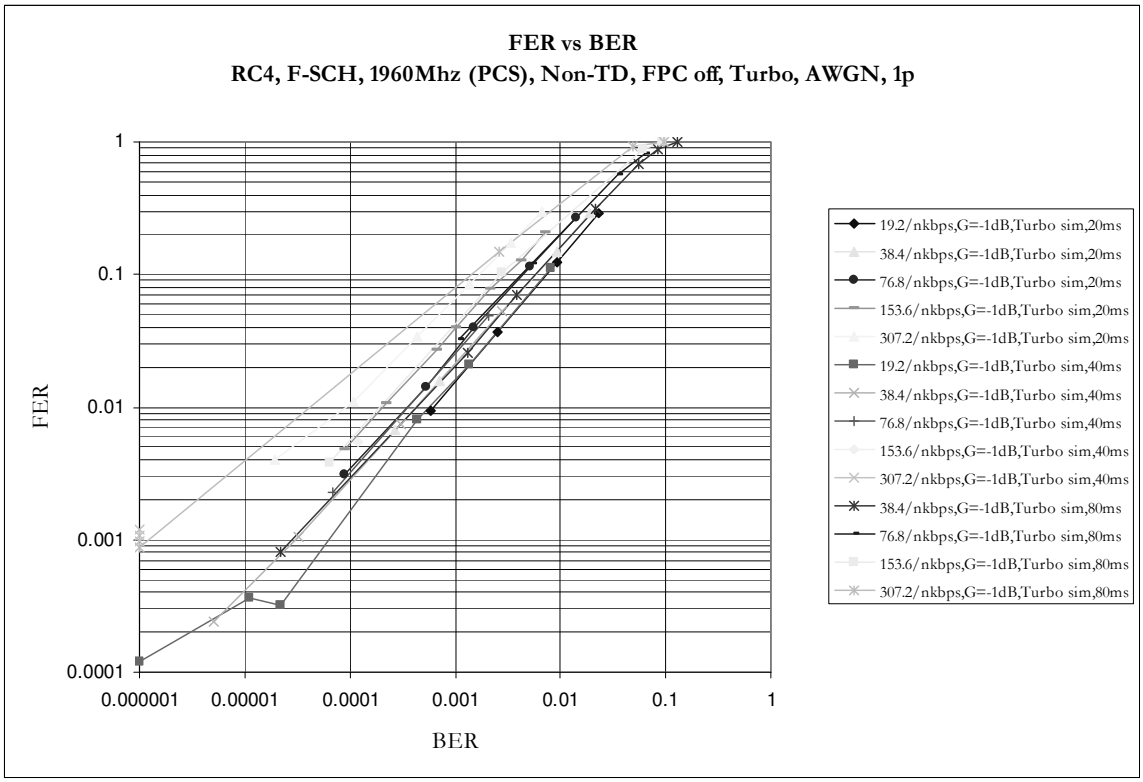

9 Appendix C—Link Level Performance Results

The following contains the complete set of data for section 4.2.

9.1 Summary

9.1.1 Convolutional Encoding

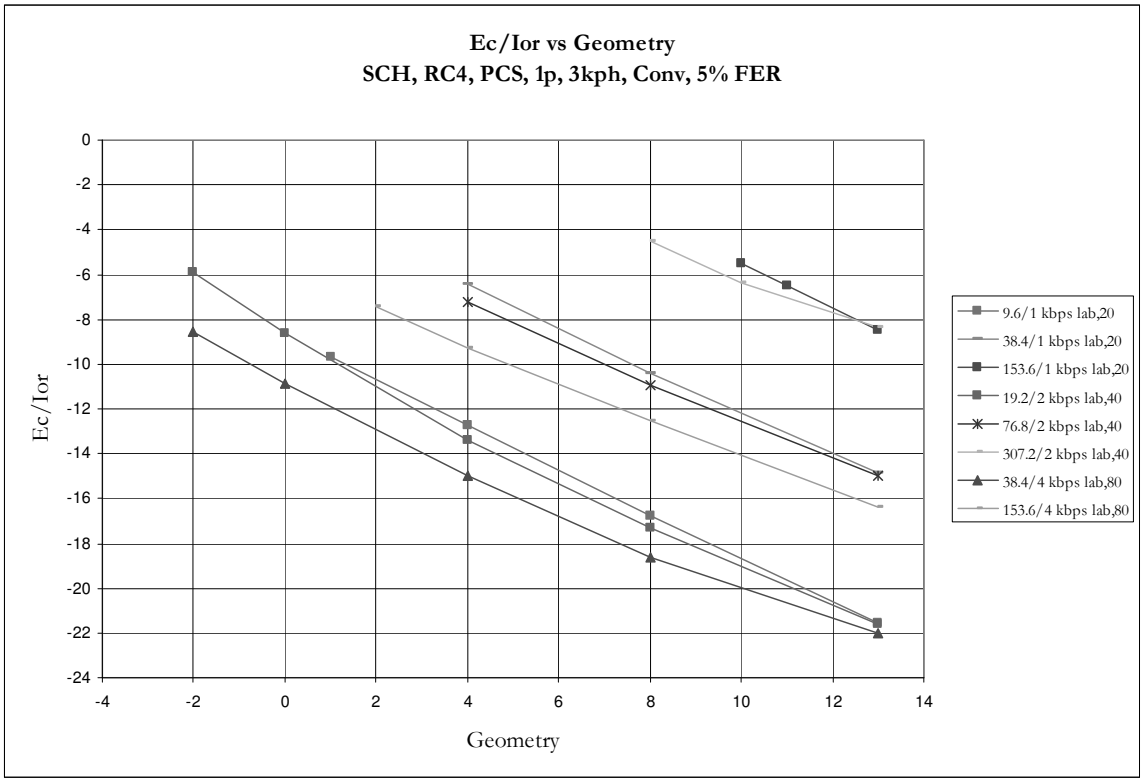

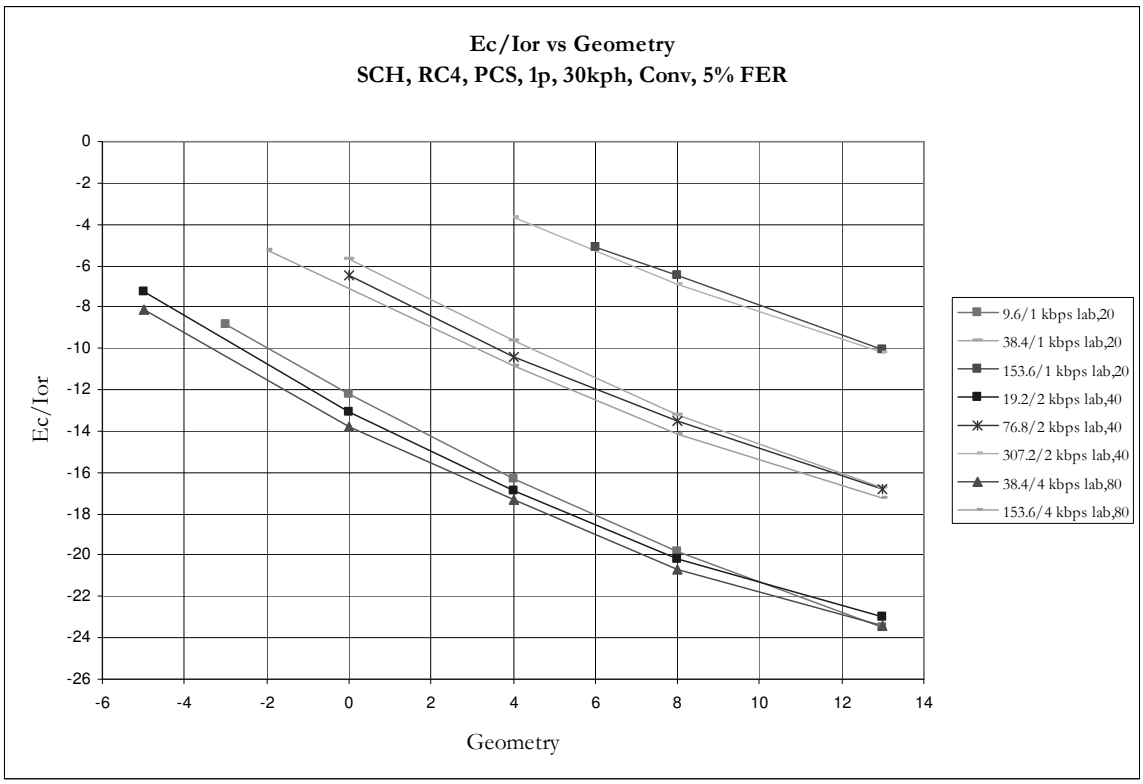

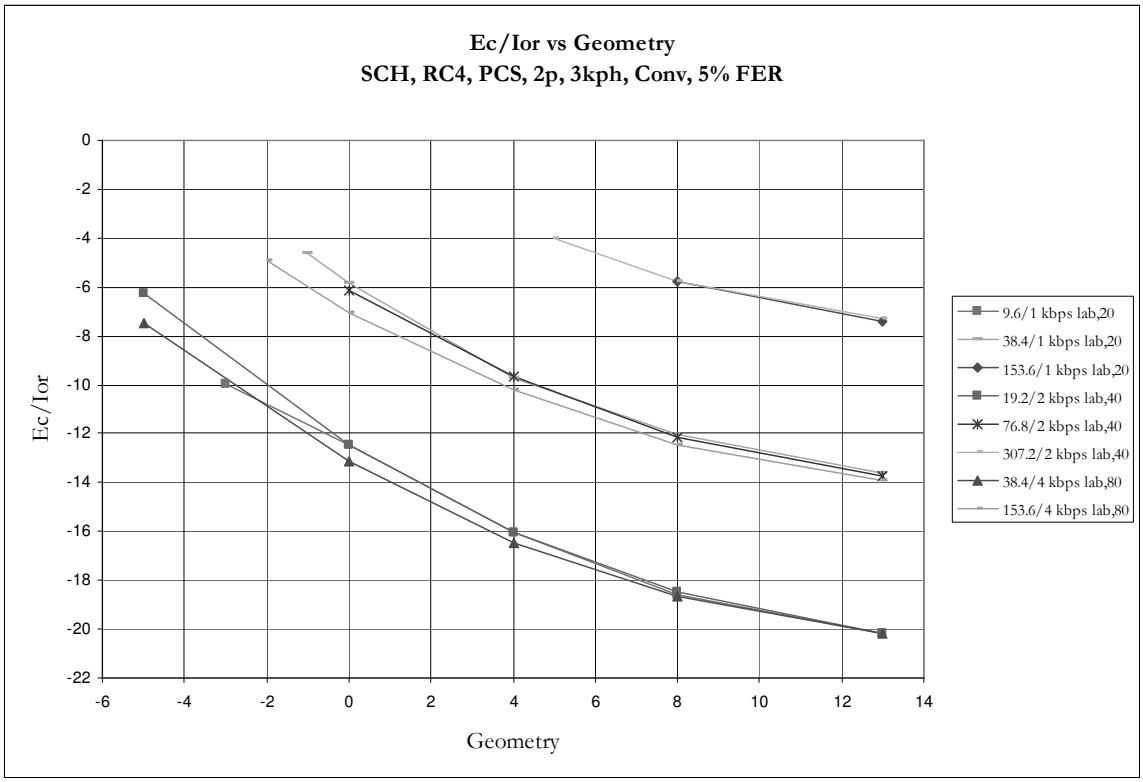
Ec/Ior vs Geometry
SCH, RC4, PCS, 2p, 3kph, Conv, 5% FER
Ec/Ior
Geometry
9.6/1 kbps lab,20
38.4/1 kbps lab,20
153.6/1 kbps lab,20
19.2/2 kbps lab,40
76.8/2 kbps lab,40
307.2/2 kbps lab,40
38.4/4 kbps lab,80
153.6/4 kbps lab,80

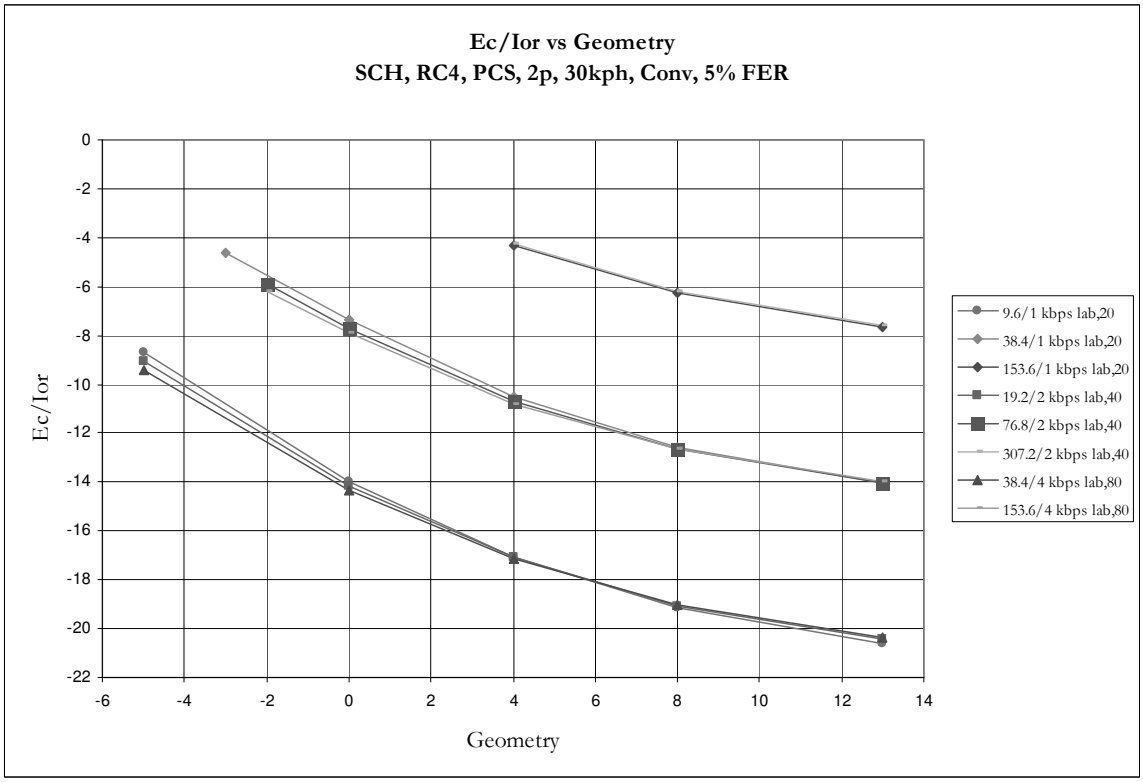
Ec/Ior vs Geometry
SCH, RC4, PCS, 2p, 30kph, Conv, 5% FER
Ec/Ior
Geometry
9.6/1 kbps lab,20
38.4/1 kbps lab,20
153.6/1 kbps lab,20
19.2/2 kbps lab,40
76.8/2 kbps lab,40
307.2/2 kbps lab,40
38.4/4 kbps lab,80
153.6/4 kbps lab,80

9.1.2 Turbo Encoding

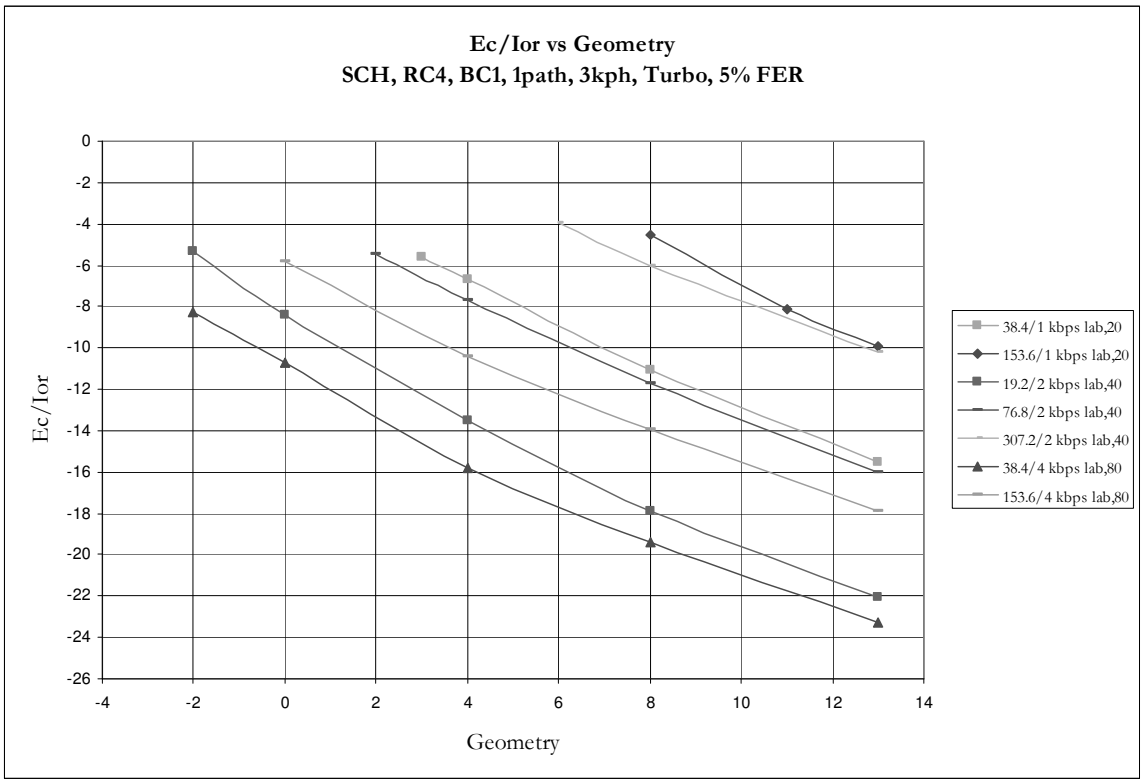

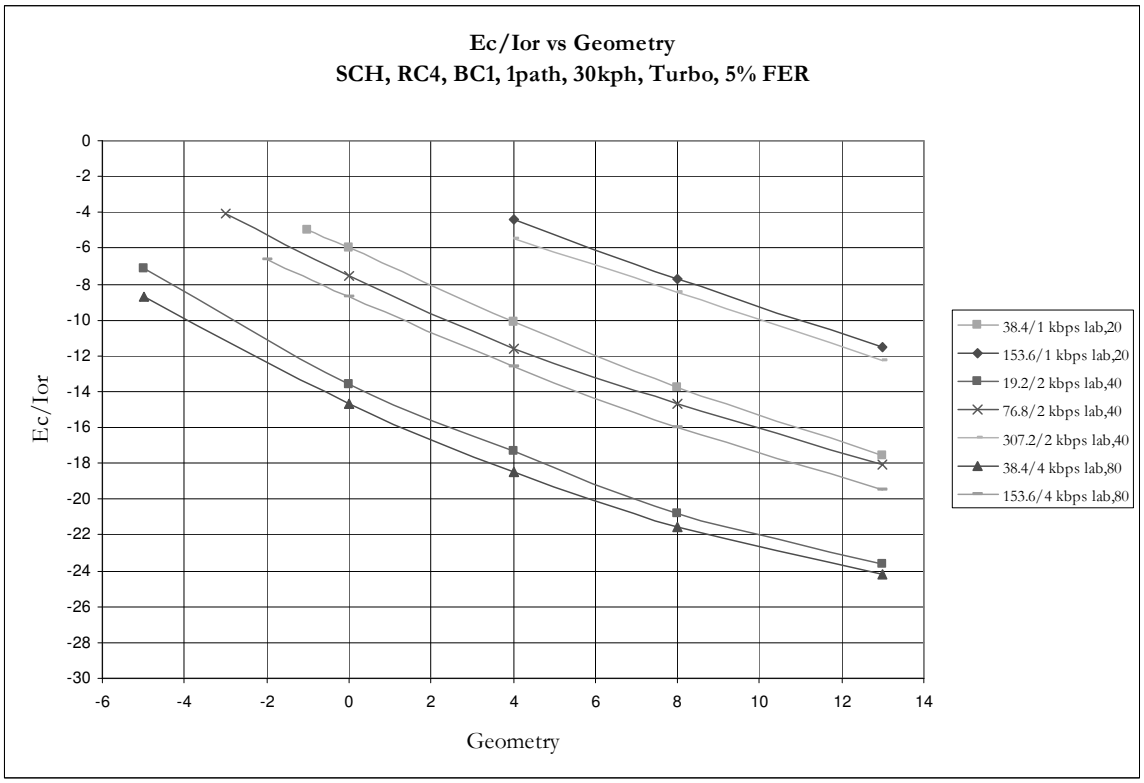

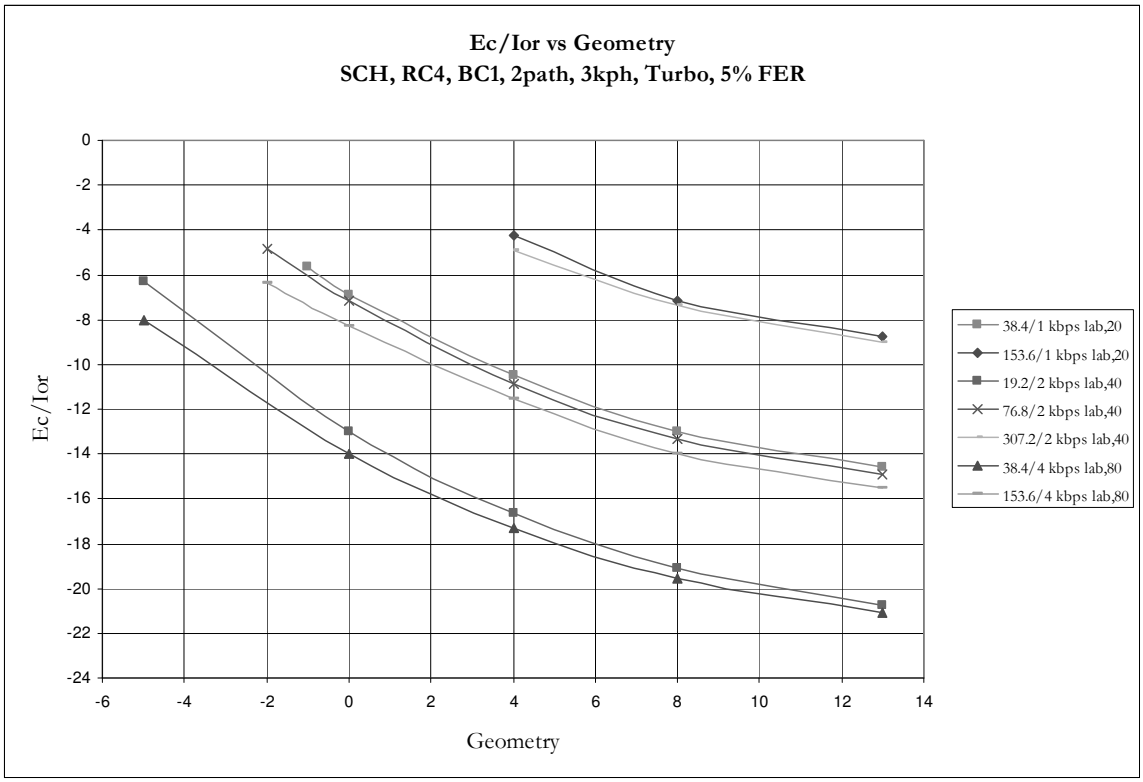
Ec/Ior vs Geometry
SCH, RC4, BC1, 2path, 3kph, Turbo, 5% FER
Ec/Ior
Geometry
38.4/1 kbps lab,20
153.6/1 kbps lab,20
19.2/2 kbps lab,40
76.8/2 kbps lab,40
307.2/2 kbps lab,40
38.4/4 kbps lab,80
153.6/4 kbps lab,80

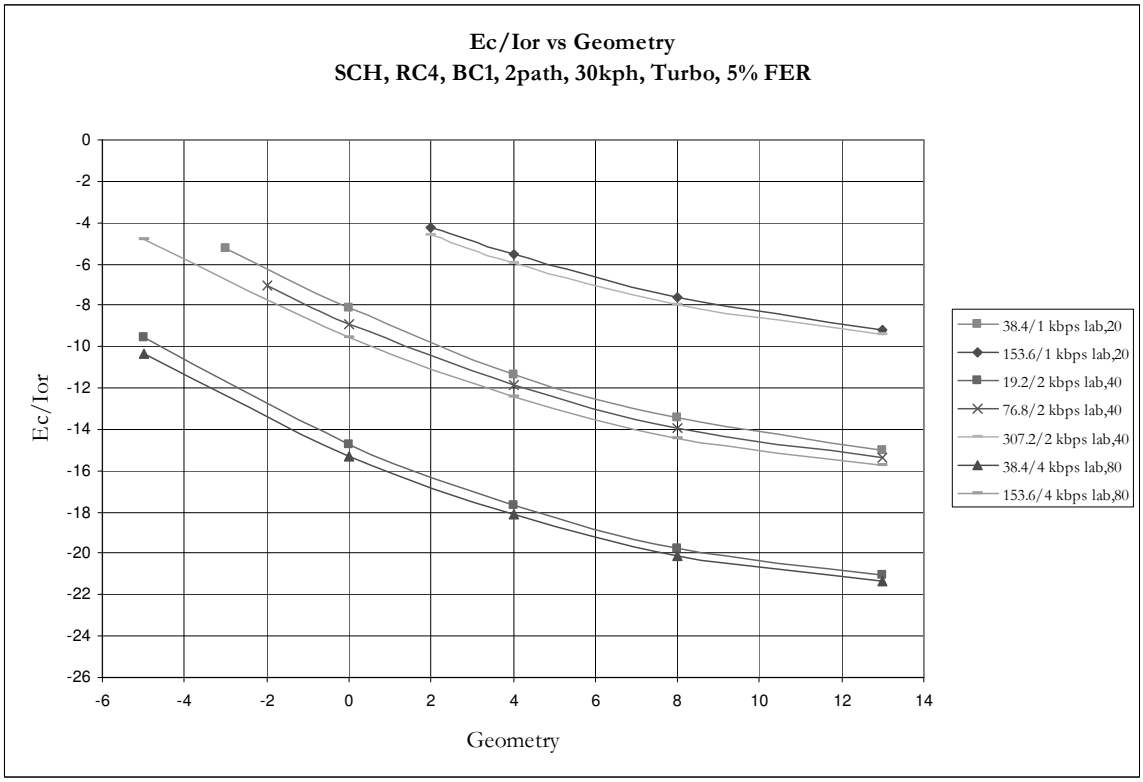
Ec/Ior vs Geometry
SCH, RC4, BC1, 2path, 30kph, Turbo, 5% FER
Ec/Ior
Geometry
38.4/1 kbps lab,20
153.6/1 kbps lab,20
19.2/2 kbps lab,40
76.8/2 kbps lab,40
307.2/2 kbps lab,40
38.4/4 kbps lab,80
153.6/4 kbps lab,80

9.2 E_c/I_{or} vs. Geometry

9.2.1 Convolutional Encoding

9.2.1.1 AWGN, 1 path

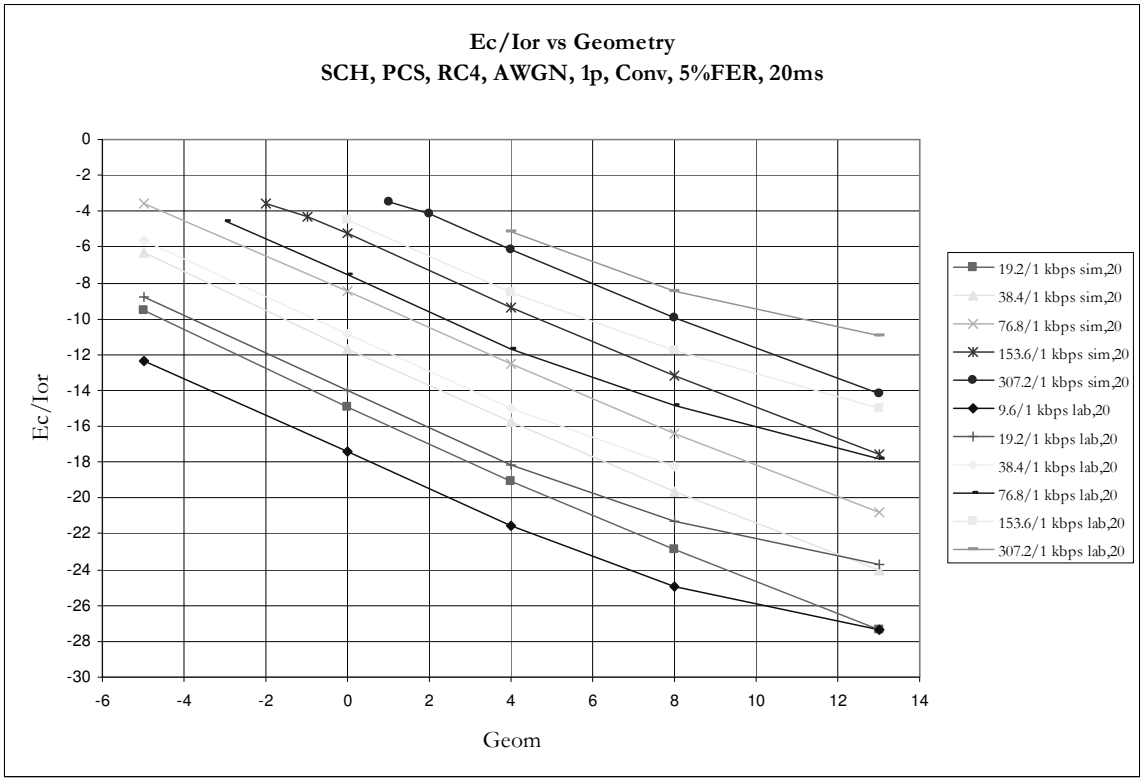

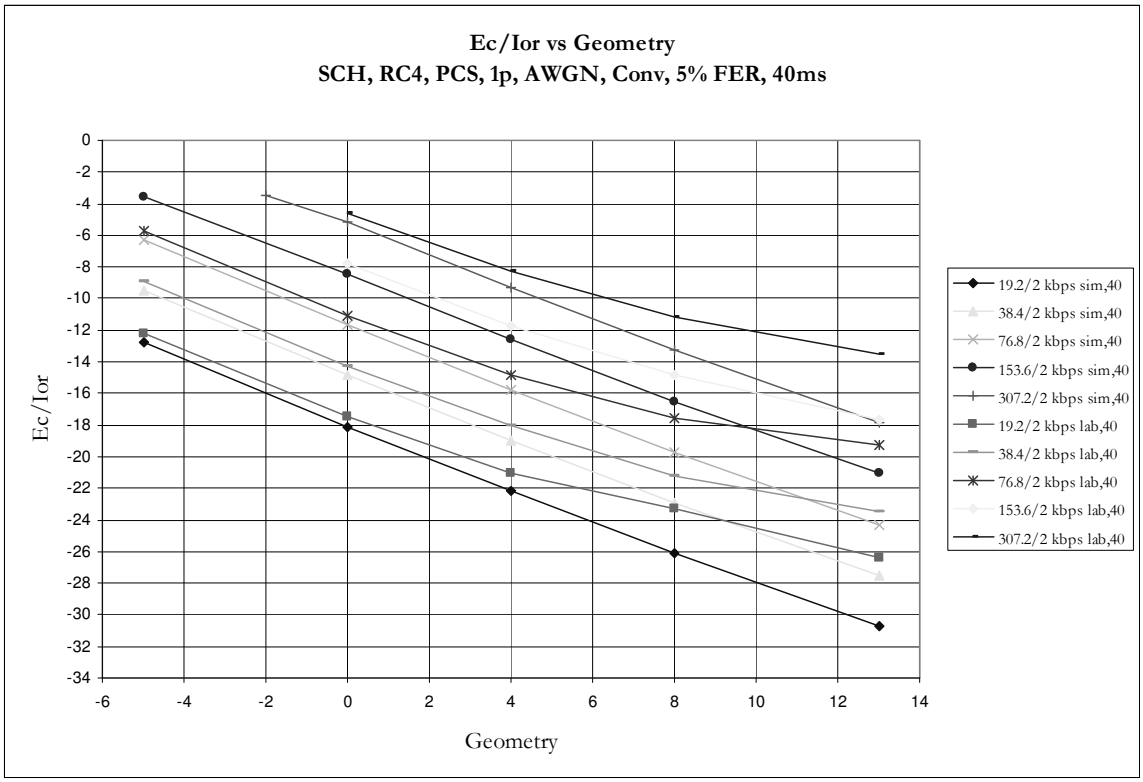
Ec/Ior vs Geometry
SCH, RC4, PCS, 1p, AWGN, Conv, 5% FER, 40ms
Ec/Ior
Geometry
0
-2
-4
-6
-8
-10
-12
-14
-16
-18
-20
-22
-24
-26
-28
-30
-32
-34
-6
-4
-2
0
2
4
6
8
10
12
14
19.2/2 kbps sim,40
38.4/2 kbps sim,40
76.8/2 kbps sim,40
153.6/2 kbps sim,40
307.2/2 kbps sim,40
19.2/2 kbps lab,40
38.4/2 kbps lab,40
76.8/2 kbps lab,40
153.6/2 kbps lab,40
307.2/2 kbps lab,40

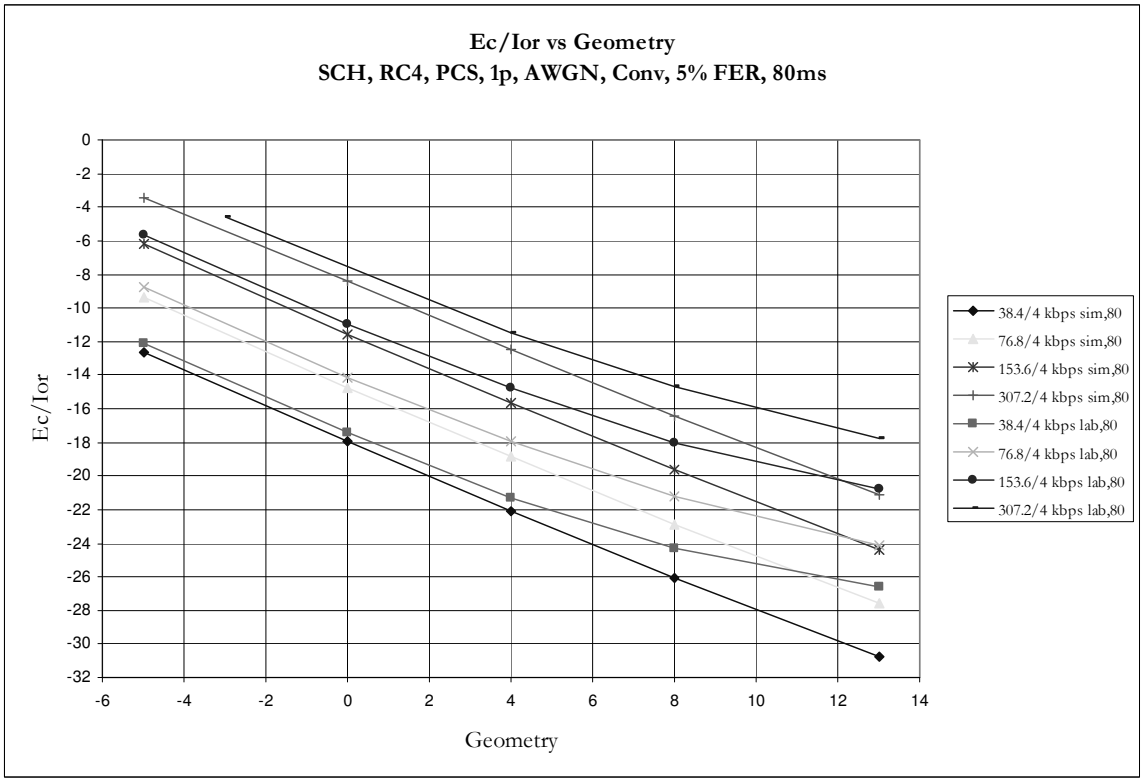
Ec/Ior vs Geometry
SCH, RC4, PCS, 1p, AWGN, Conv, 5% FER, 80ms
Ec/Ior
Geometry
0
-2
-4
-6
-8
-10
-12
-14
-16
-18
-20
-22
-24
-26
-28
-30
-32
-6
-4
-2
0
2
4
6
8
10
12
14
38.4/4 kbps sim,80
76.8/4 kbps sim,80
153.6/4 kbps sim,80
307.2/4 kbps sim,80
38.4/4 kbps lab,80
76.8/4 kbps lab,80
153.6/4 kbps lab,80
307.2/4 kbps lab,80

9.2.1.2 Rayleigh, 1 path, 3kmph

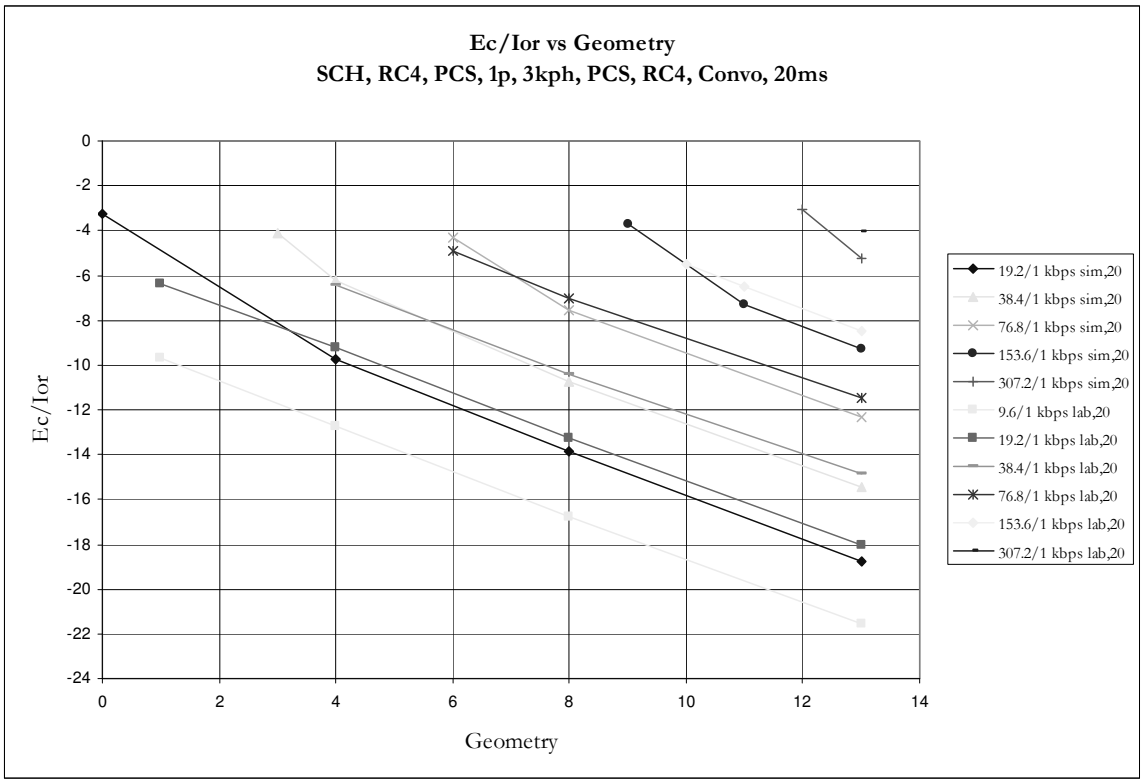

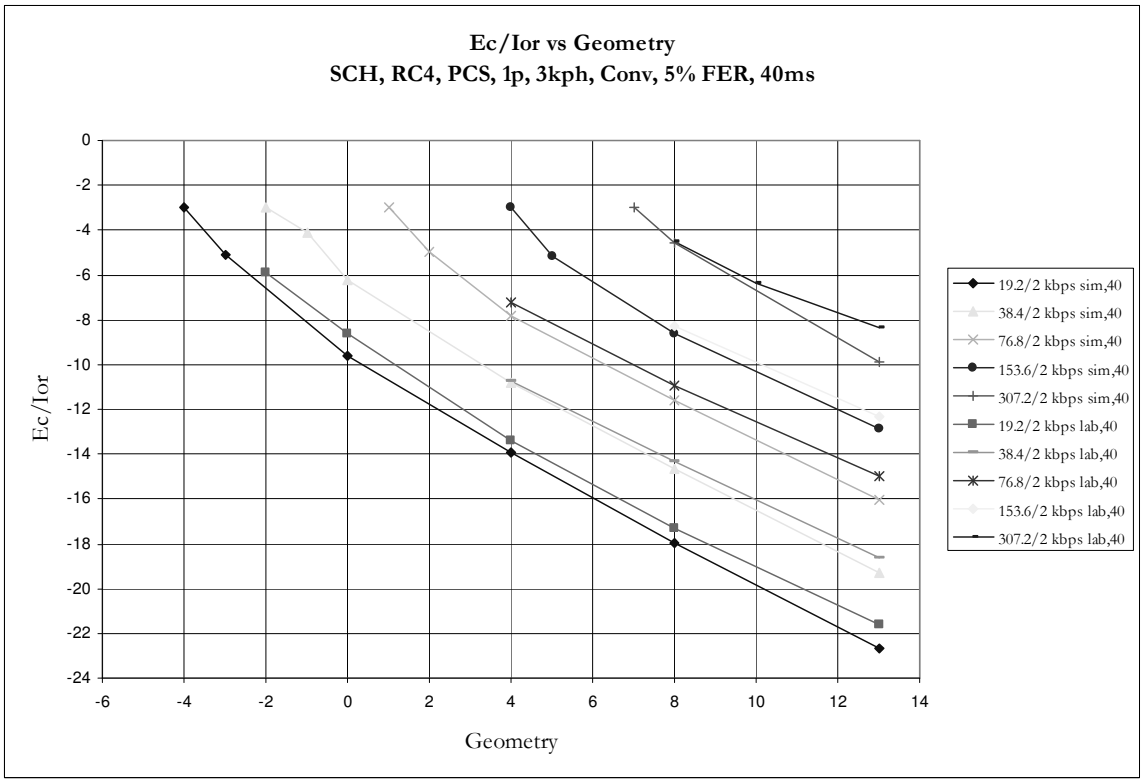

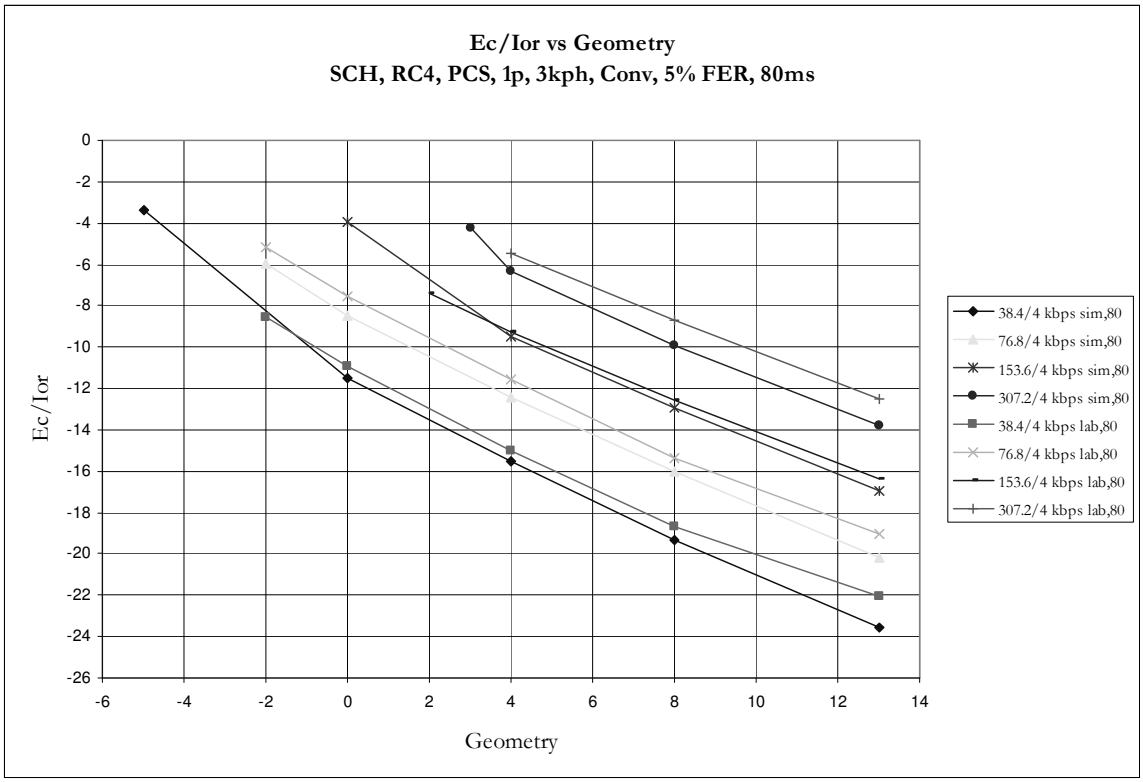

9.2.1.3 Rayleigh, 1 path, 30kmph

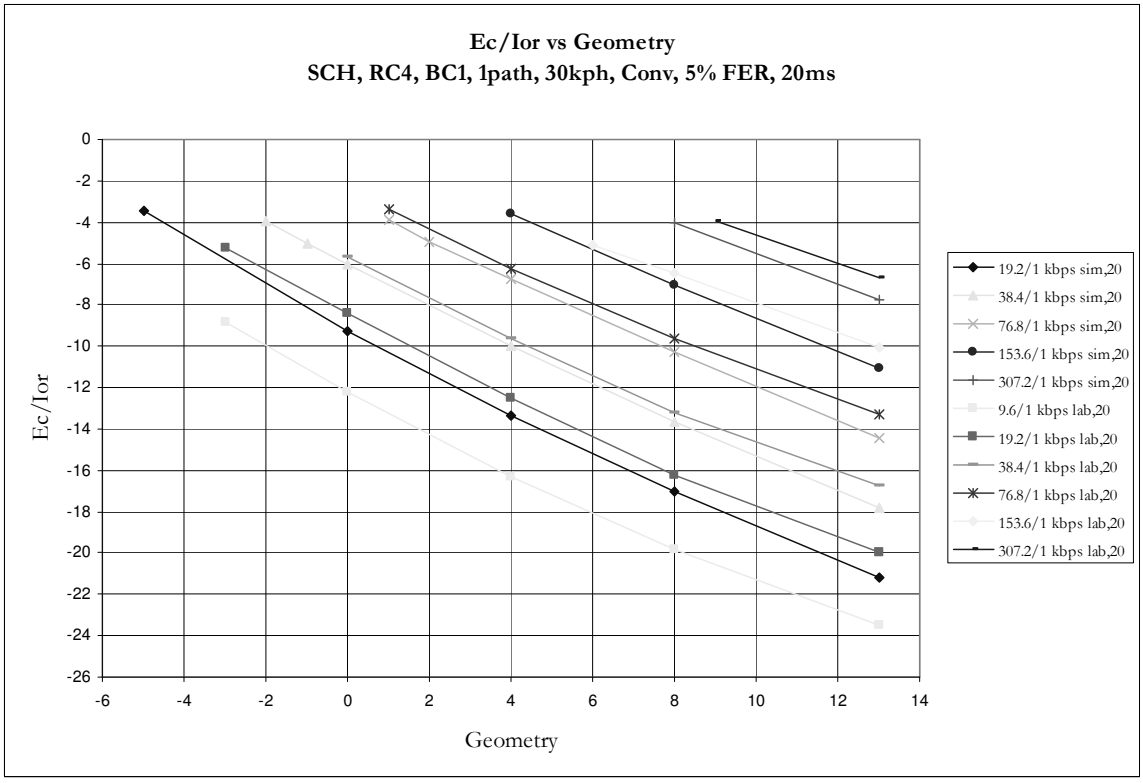

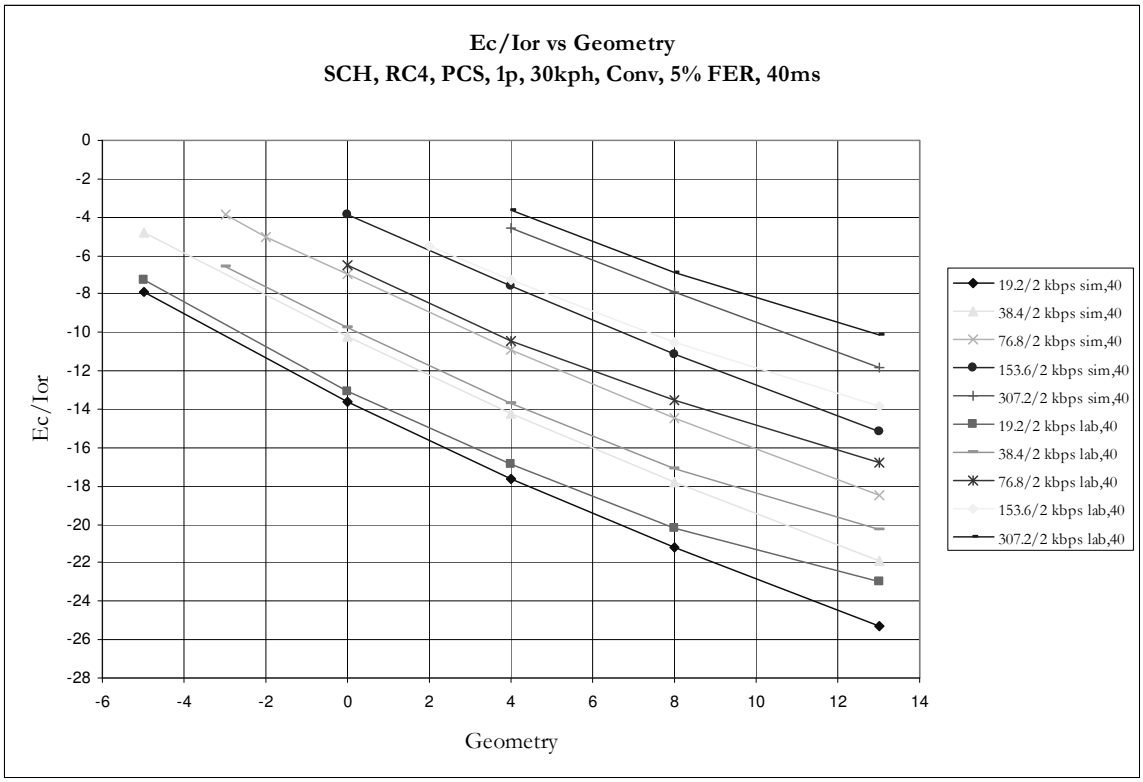
Ec/Ior vs Geometry
SCH, RC4, PCS, 1p, 30kph, Conv, 5% FER, 40ms
Ec/Ior
Geometry
0
-2
-4
-6
-8
-10
-12
-14
-16
-18
-20
-22
-24
-26
-28
-6
-4
-2
0
2
4
6
8
10
12
14
19.2/2 kbps sim,40
38.4/2 kbps sim,40
76.8/2 kbps sim,40
153.6/2 kbps sim,40
307.2/2 kbps sim,40
19.2/2 kbps lab,40
38.4/2 kbps lab,40
76.8/2 kbps lab,40
153.6/2 kbps lab,40
307.2/2 kbps lab,40

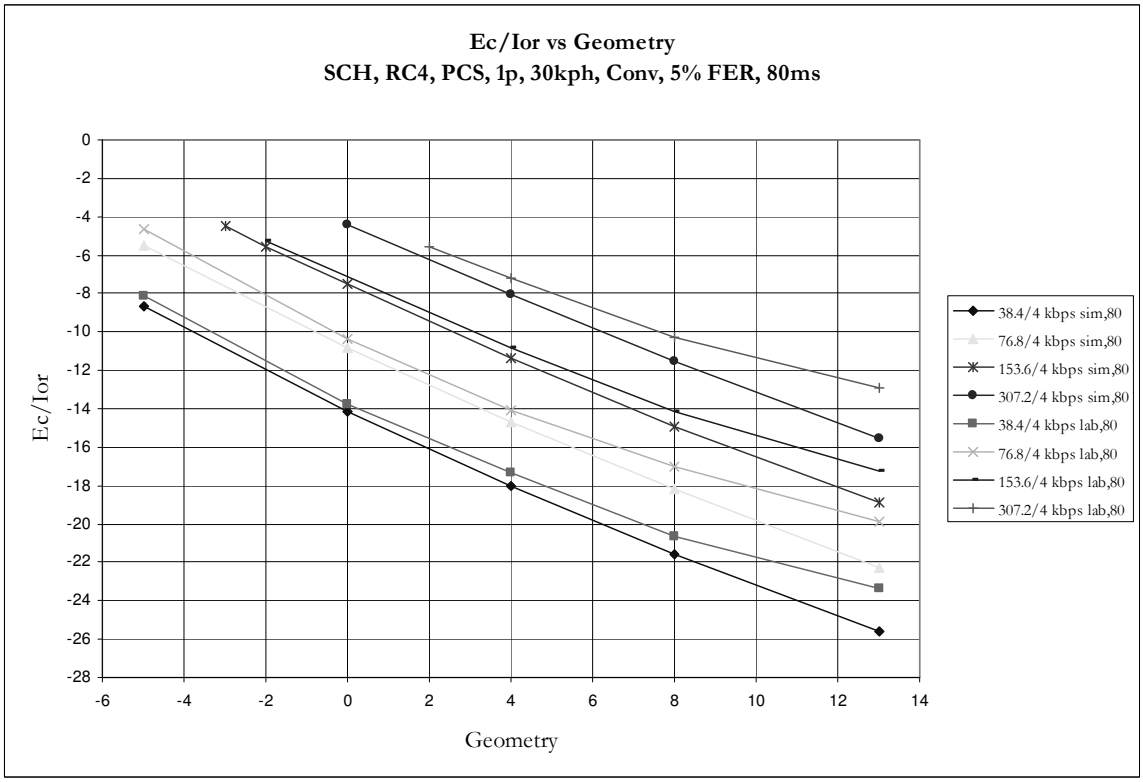
Ec/Ior vs Geometry
SCH, RC4, PCS, 1p, 30kph, Conv, 5% FER, 80ms
Ec/Ior
Geometry
0
-2
-4
-6
-8
-10
-12
-14
-16
-18
-20
-22
-24
-26
-28
-6
-4
-2
0
2
4
6
8
10
12
14
38.4/4 kbps sim,80
76.8/4 kbps sim,80
153.6/4 kbps sim,80
307.2/4 kbps sim,80
38.4/4 kbps lab,80
76.8/4 kbps lab,80
153.6/4 kbps lab,80
307.2/4 kbps lab,80

9.2.1.4 Rayleigh, 2 path, 3kmph

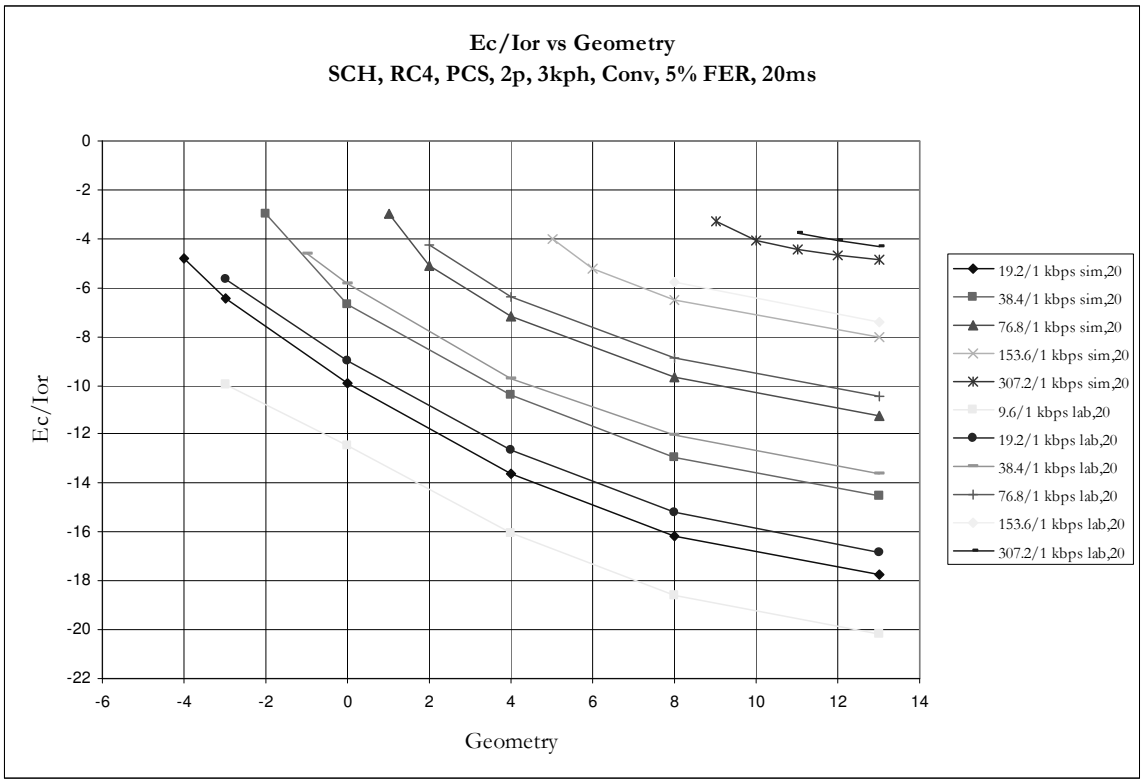

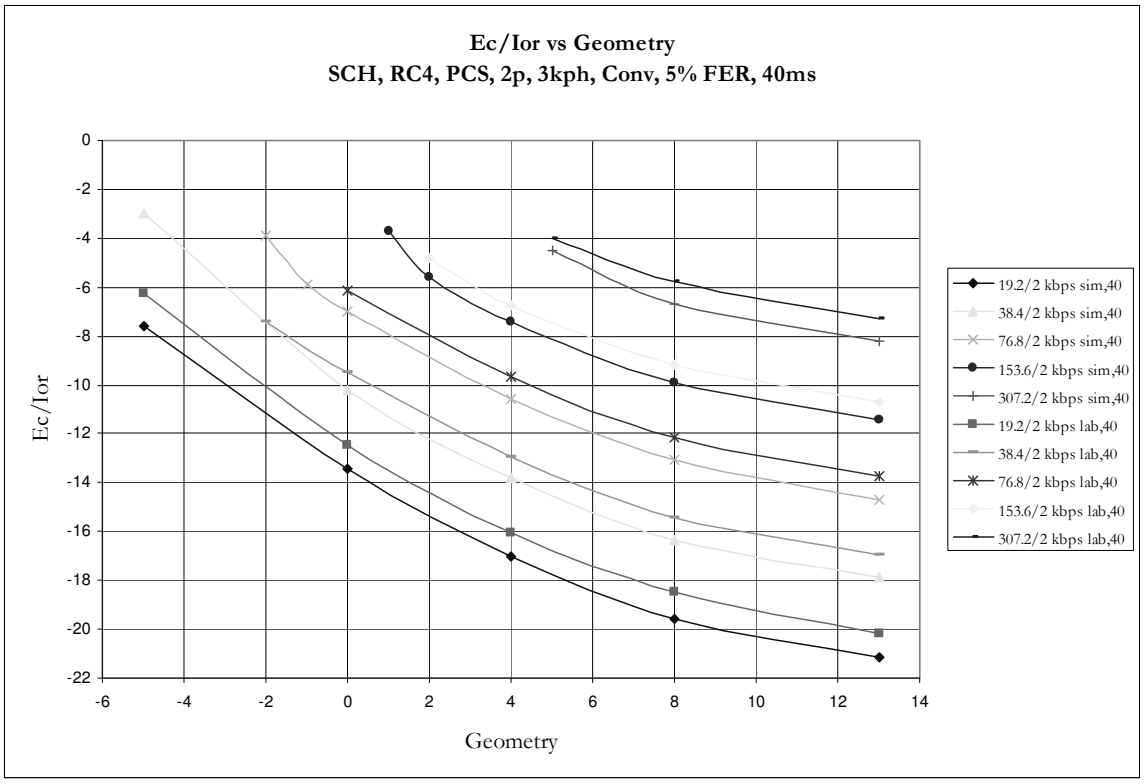

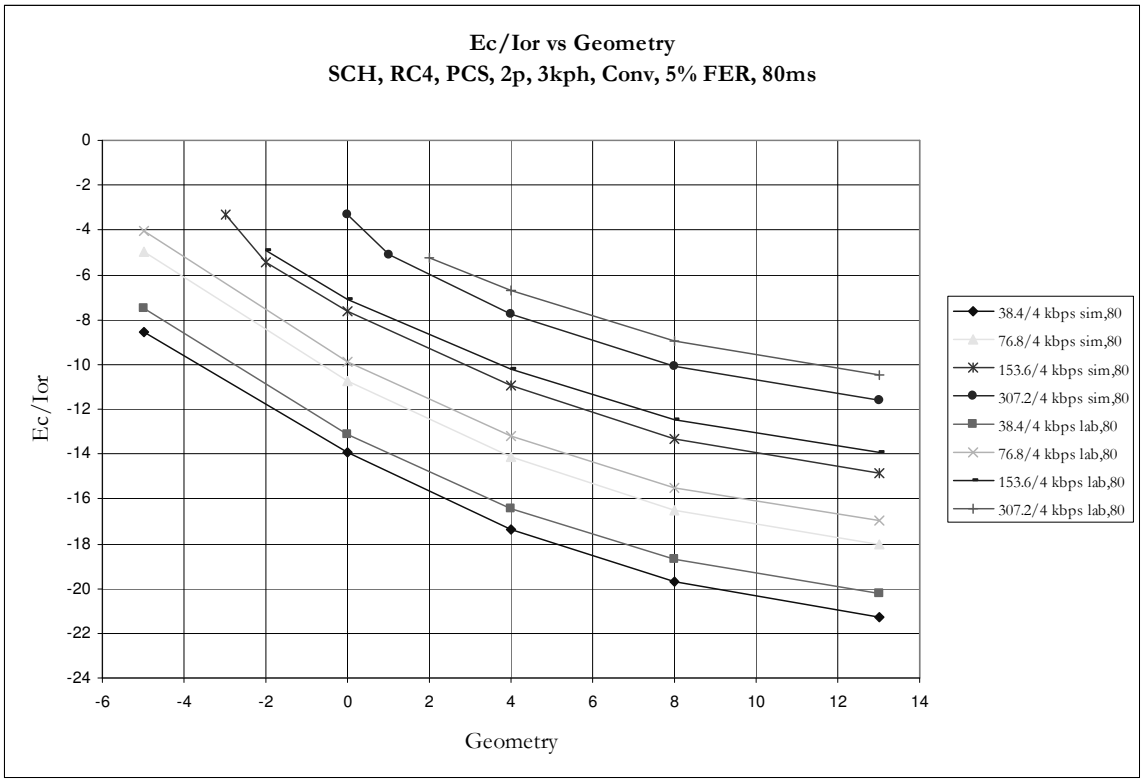

9.2.1.5 Rayleigh, 2 path, 30kmph

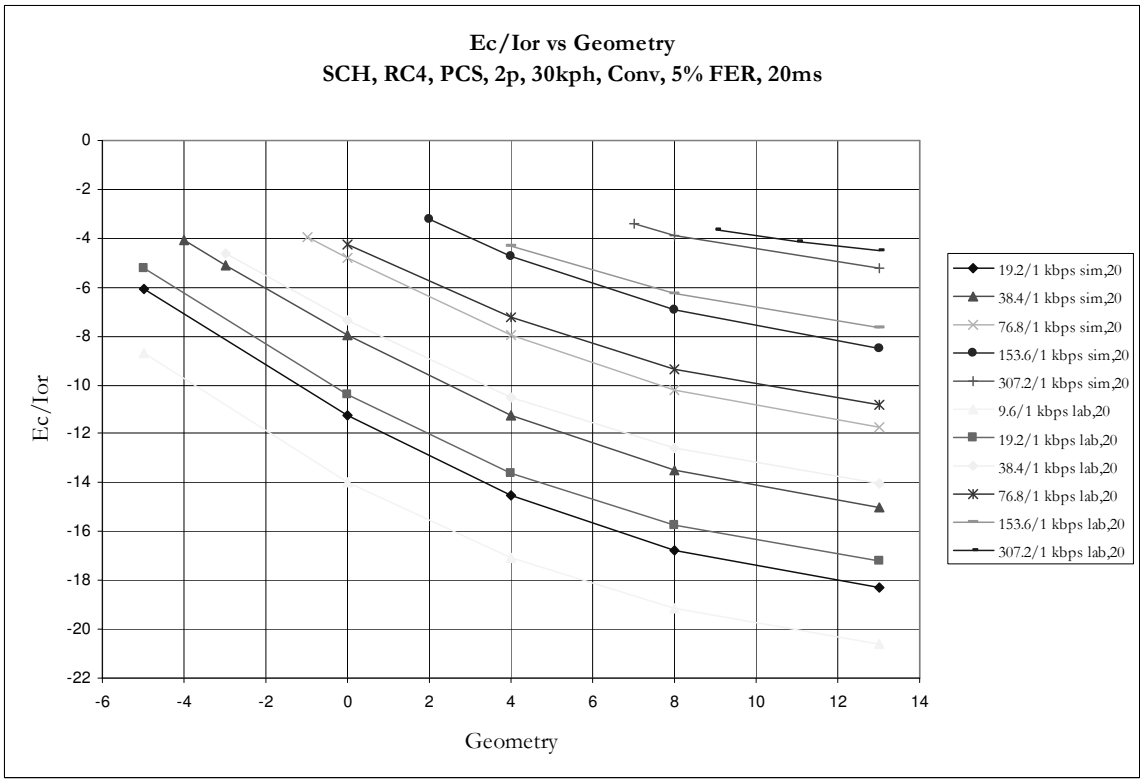

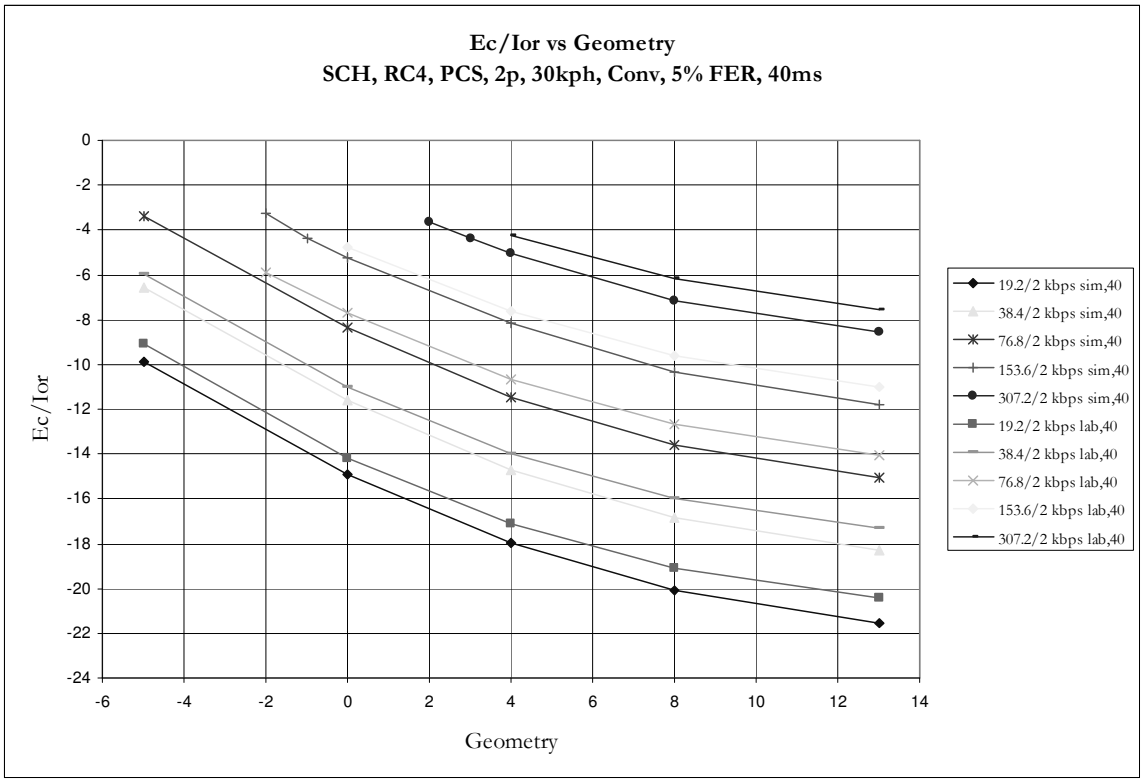
Ec/Ior vs Geometry
SCH, RC4, PCS, 2p, 30kph, Conv, 5% FER, 40ms
Ec/Ior
Geometry
0
-2
-4
-6
-8
-10
-12
-14
-16
-18
-20
-22
-24
-6
-4
-2
0
2
4
6
8
10
12
14
19.2/2 kbps sim,40
38.4/2 kbps sim,40
76.8/2 kbps sim,40
153.6/2 kbps sim,40
307.2/2 kbps sim,40
19.2/2 kbps lab,40
38.4/2 kbps lab,40
76.8/2 kbps lab,40
153.6/2 kbps lab,40
307.2/2 kbps lab,40

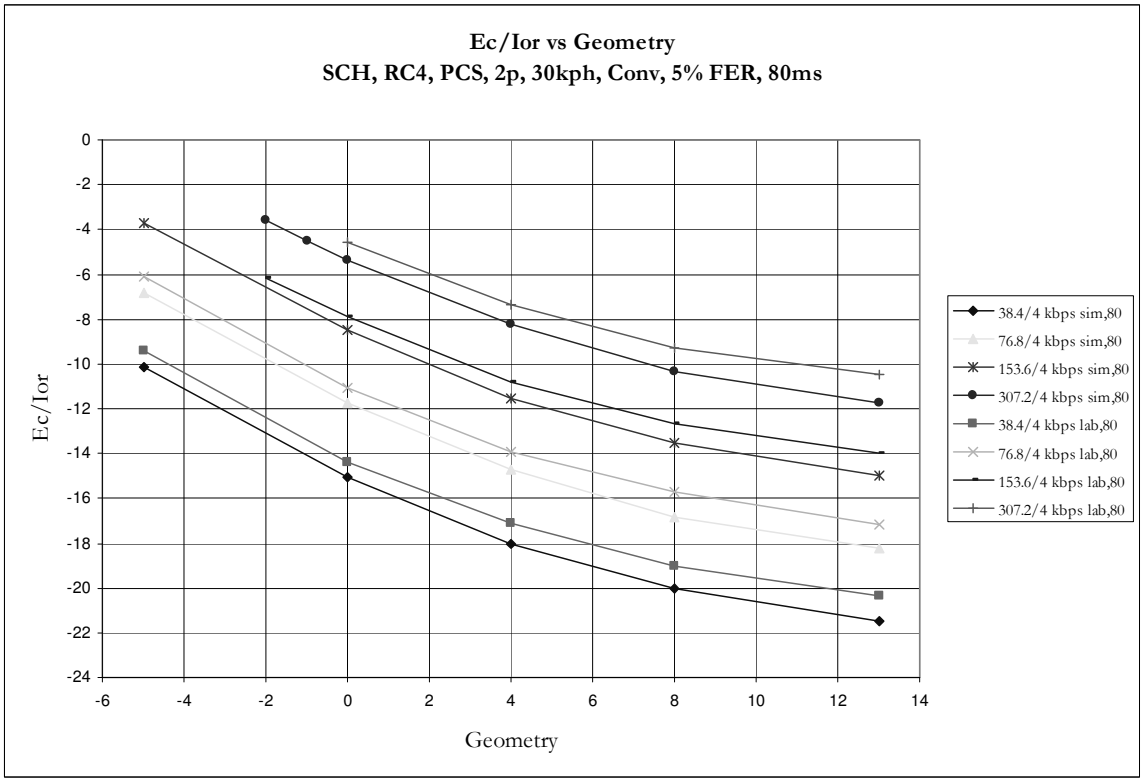
Ec/Ior vs Geometry
SCH, RC4, PCS, 2p, 30kph, Conv, 5% FER, 80ms
Ec/Ior
Geometry
0
-2
-4
-6
-8
-10
-12
-14
-16
-18
-20
-22
-24
-6
-4
-2
0
2
4
6
8
10
12
14
38.4/4 kbps sim,80
76.8/4 kbps sim,80
153.6/4 kbps sim,80
307.2/4 kbps sim,80
38.4/4 kbps lab,80
76.8/4 kbps lab,80
153.6/4 kbps lab,80
307.2/4 kbps lab,80

9.2.2 Turbo Encoding

9.2.2.1 AWGN, 1 path

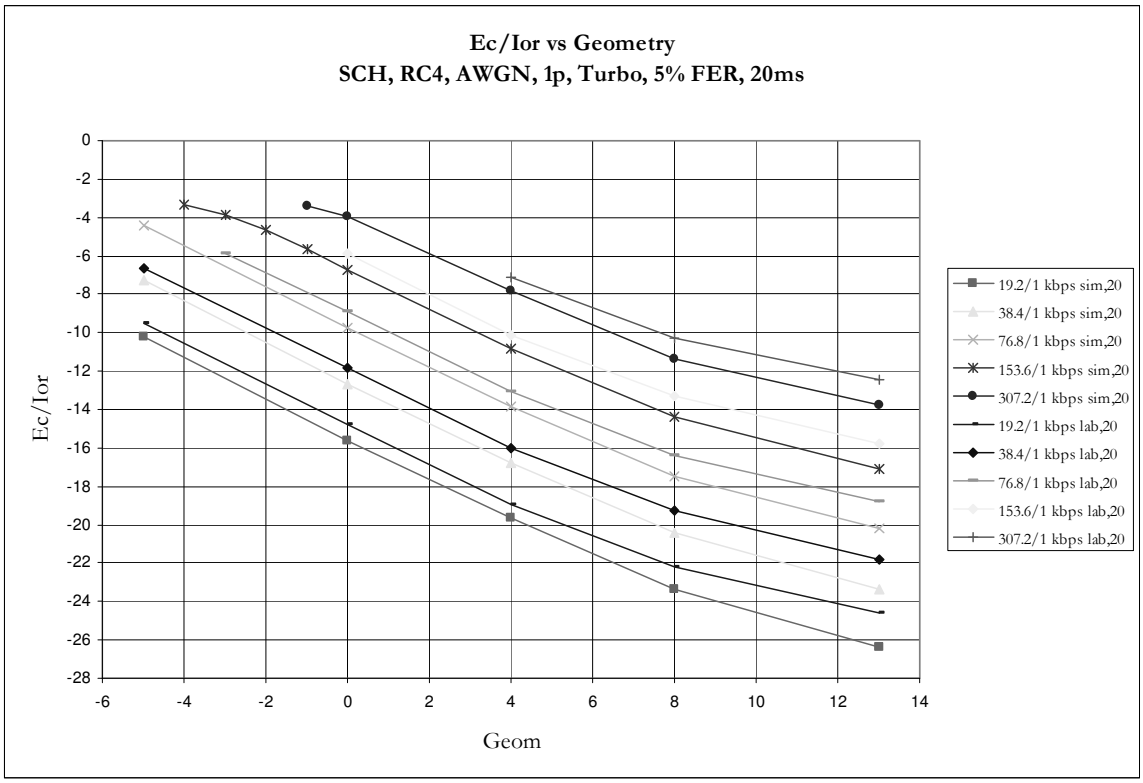

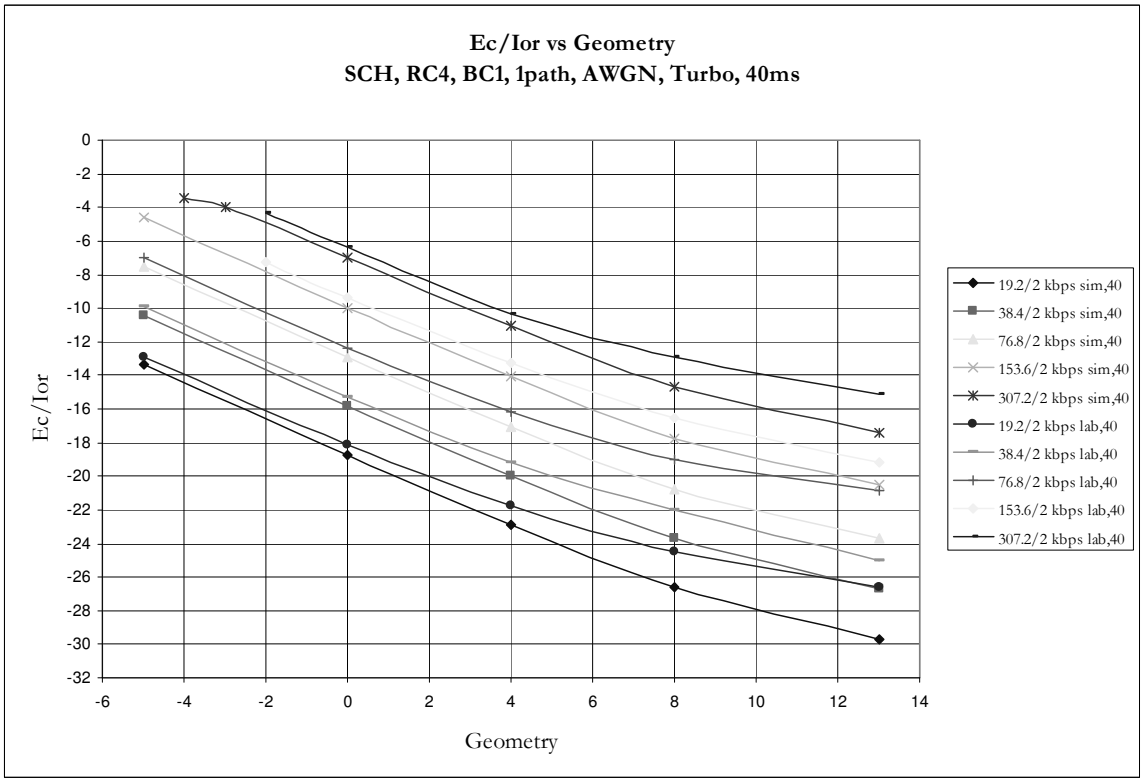

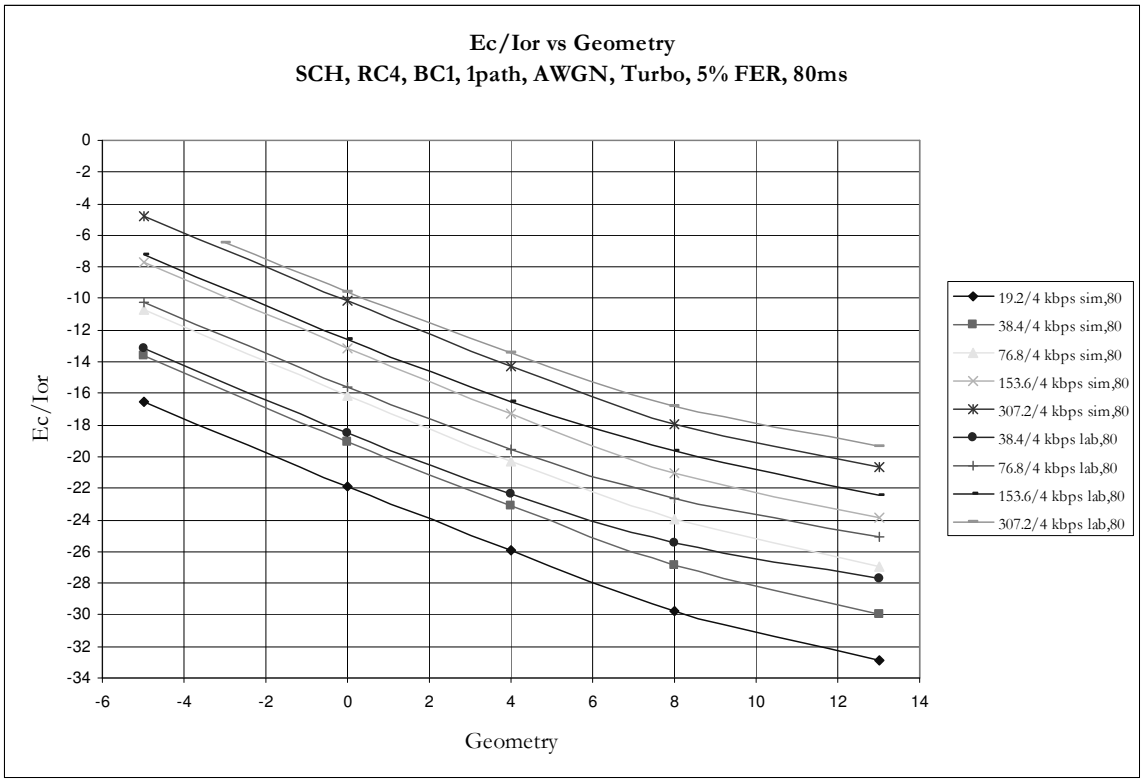

9.2.2.2 Rayleigh, 1 path, 3kmph

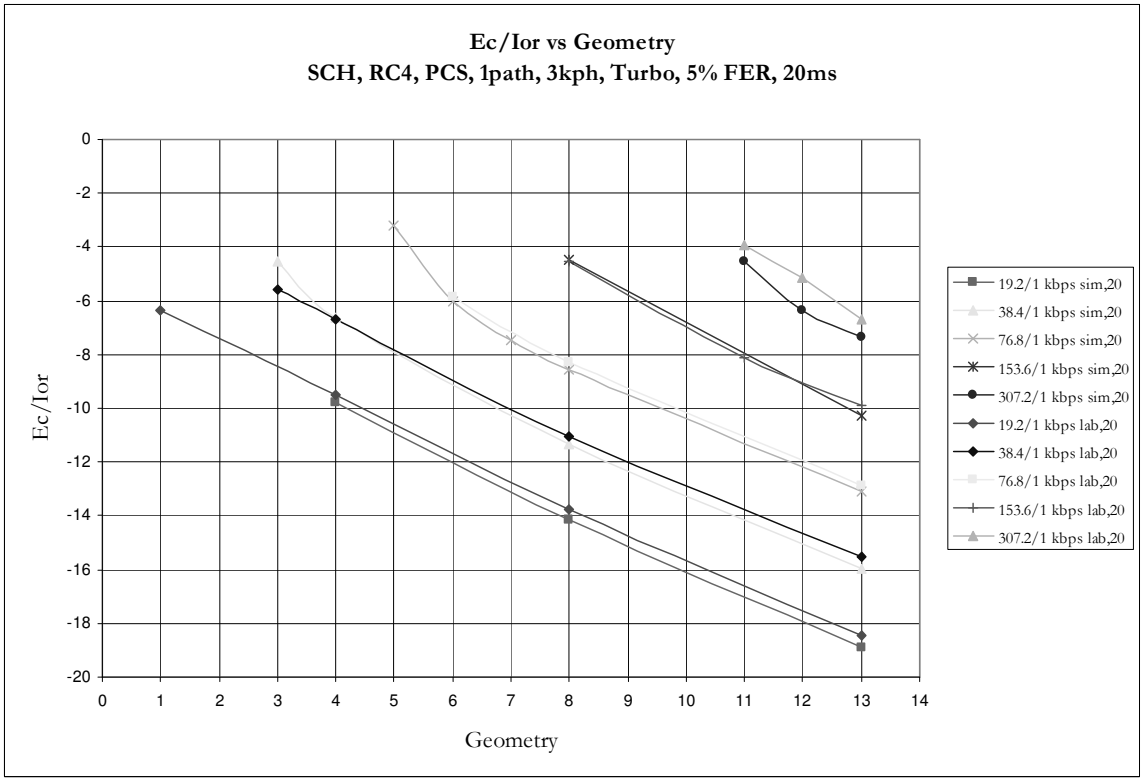

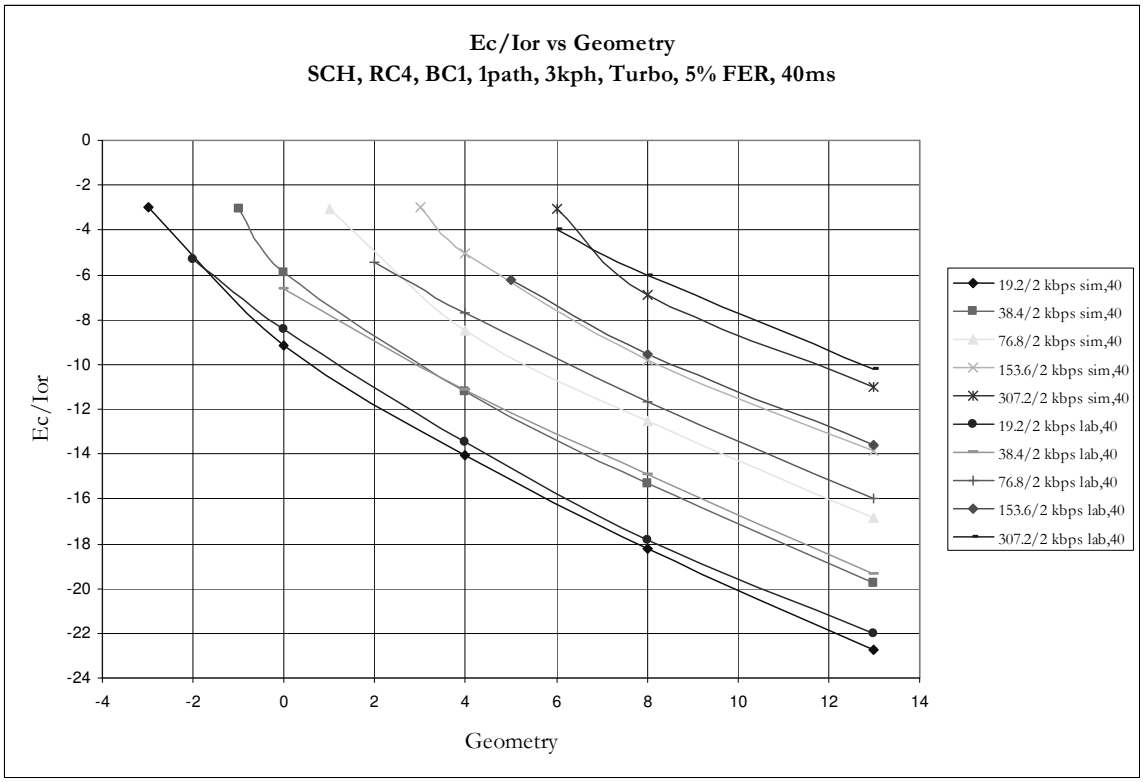
Ec/Ior vs Geometry
SCH, RC4, BC1, 1path, 3kph, Turbo, 5% FER, 40ms
Ec/Ior
Geometry
0
-2
-4
-6
-8
-10
-12
-14
-16
-18
-20
-22
-24
-4
-2
0
2
4
6
8
10
12
14
19.2/2 kbps sim,40
38.4/2 kbps sim,40
76.8/2 kbps sim,40
153.6/2 kbps sim,40
307.2/2 kbps sim,40
19.2/2 kbps lab,40
38.4/2 kbps lab,40
76.8/2 kbps lab,40
153.6/2 kbps lab,40
307.2/2 kbps lab,40

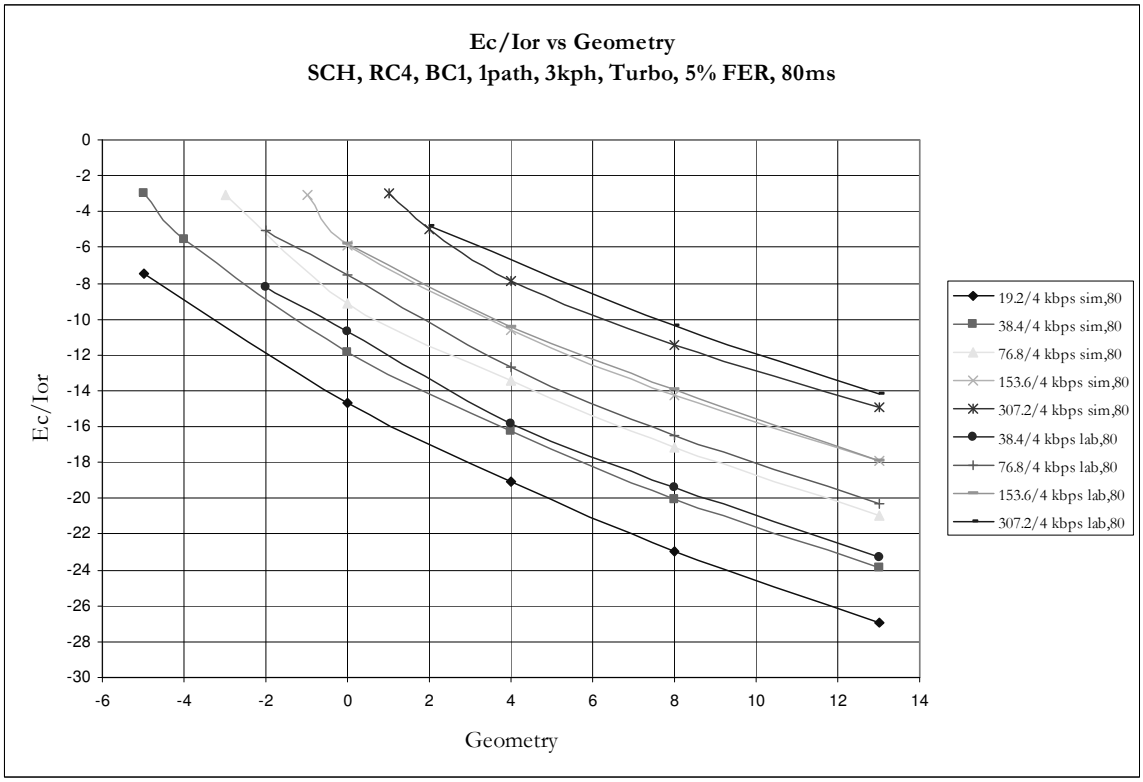
Ec/Ior vs Geometry
SCH, RC4, BC1, 1path, 3kph, Turbo, 5% FER, 80ms
Ec/Ior
Geometry
0
-2
-4
-6
-8
-10
-12
-14
-16
-18
-20
-22
-24
-26
-28
-30
-6
-4
-2
0
2
4
6
8
10
12
14
19.2/4 kbps sim,80
38.4/4 kbps sim,80
76.8/4 kbps sim,80
153.6/4 kbps sim,80
307.2/4 kbps sim,80
38.4/4 kbps lab,80
76.8/4 kbps lab,80
153.6/4 kbps lab,80
307.2/4 kbps lab,80

9.2.2.3 Rayleigh, 1 path, 30kmph

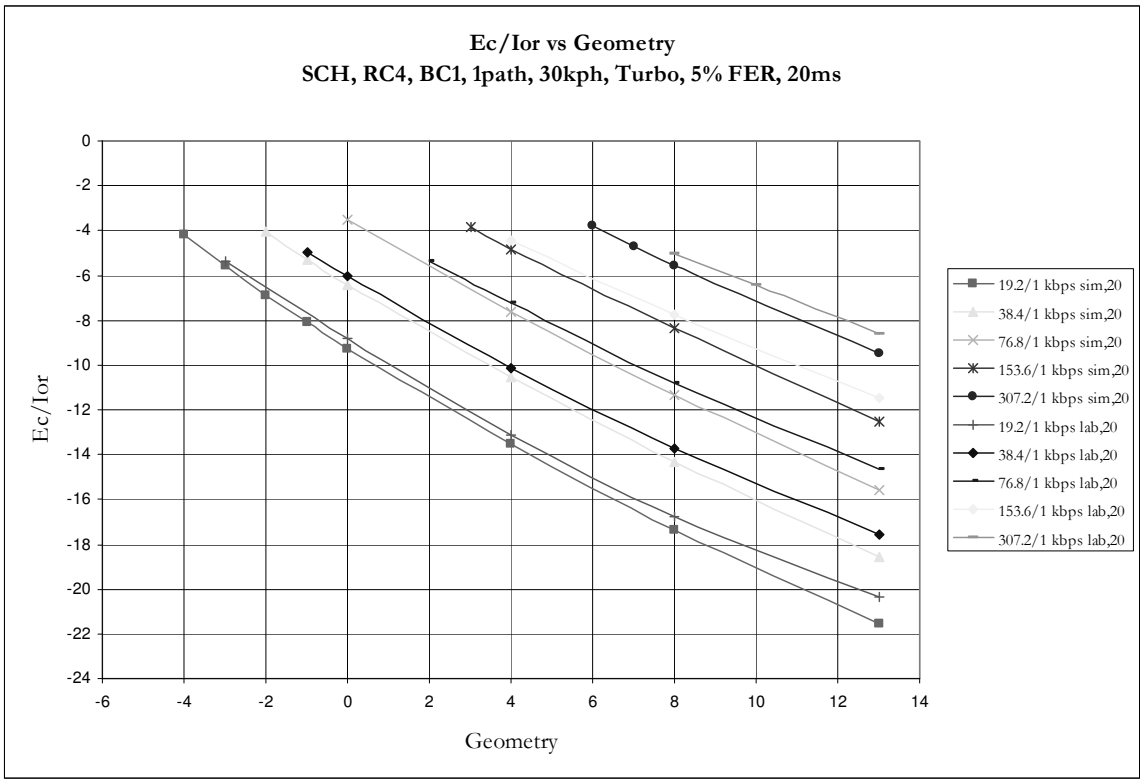

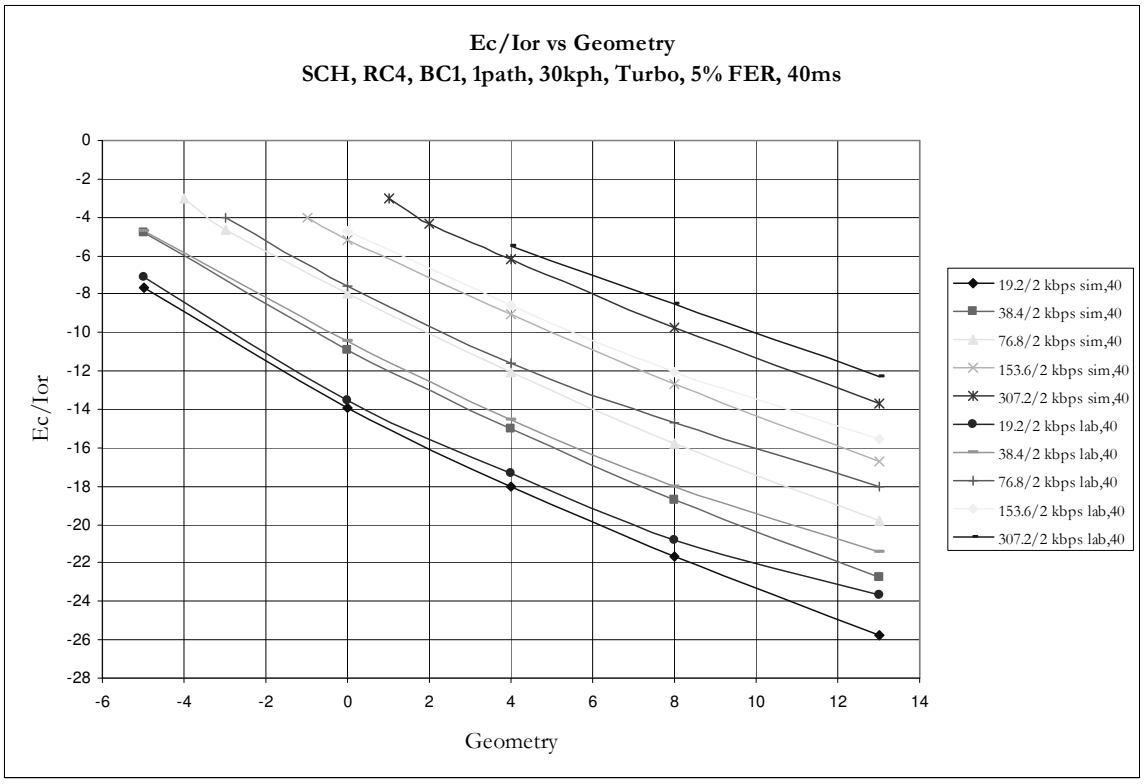

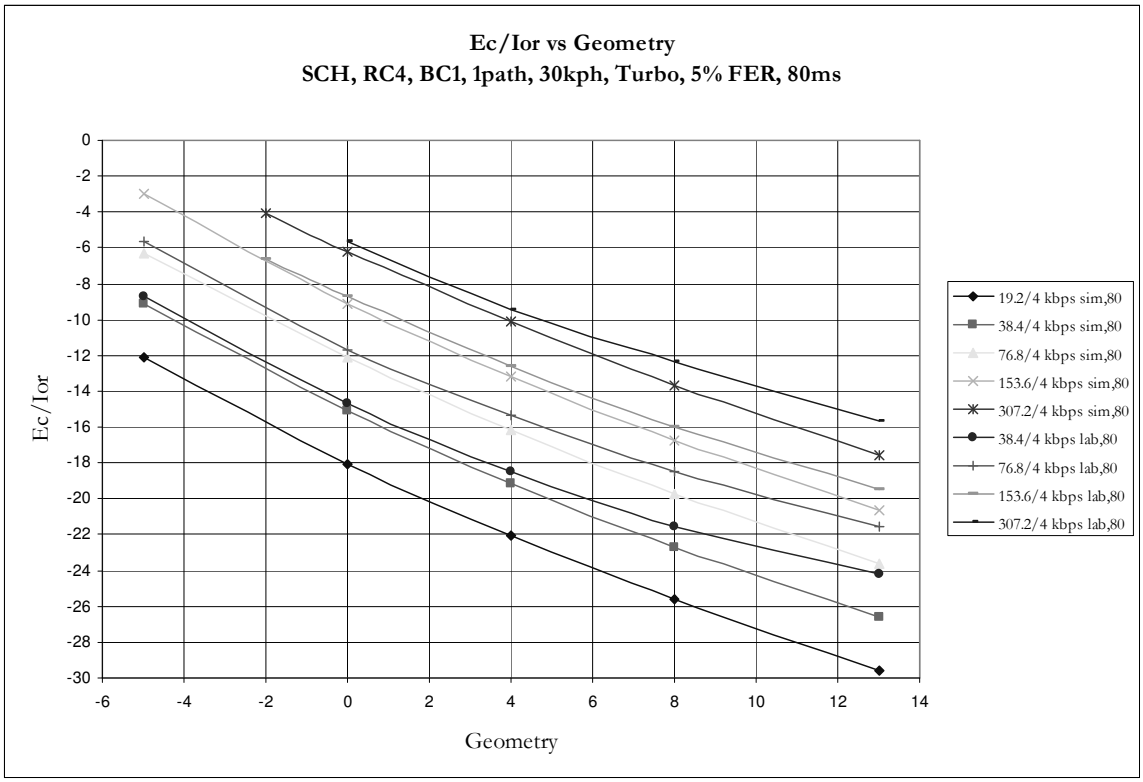

9.2.2.4 Rayleigh, 2 path, 3kmph

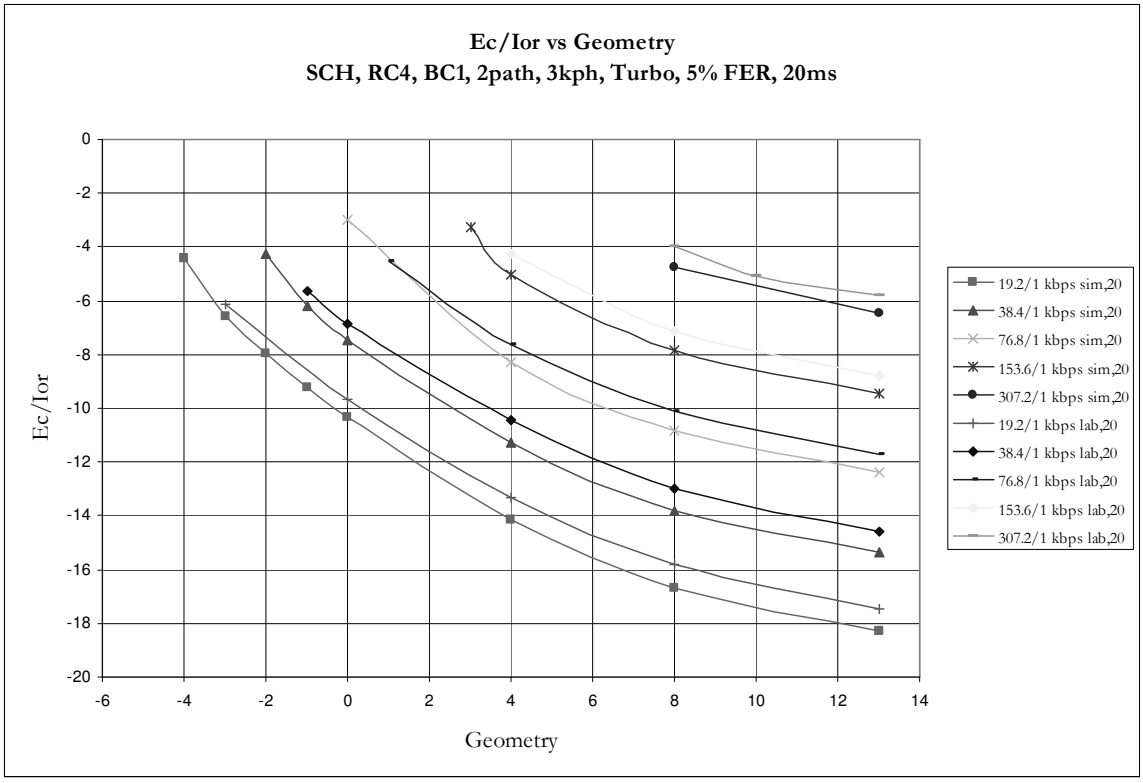

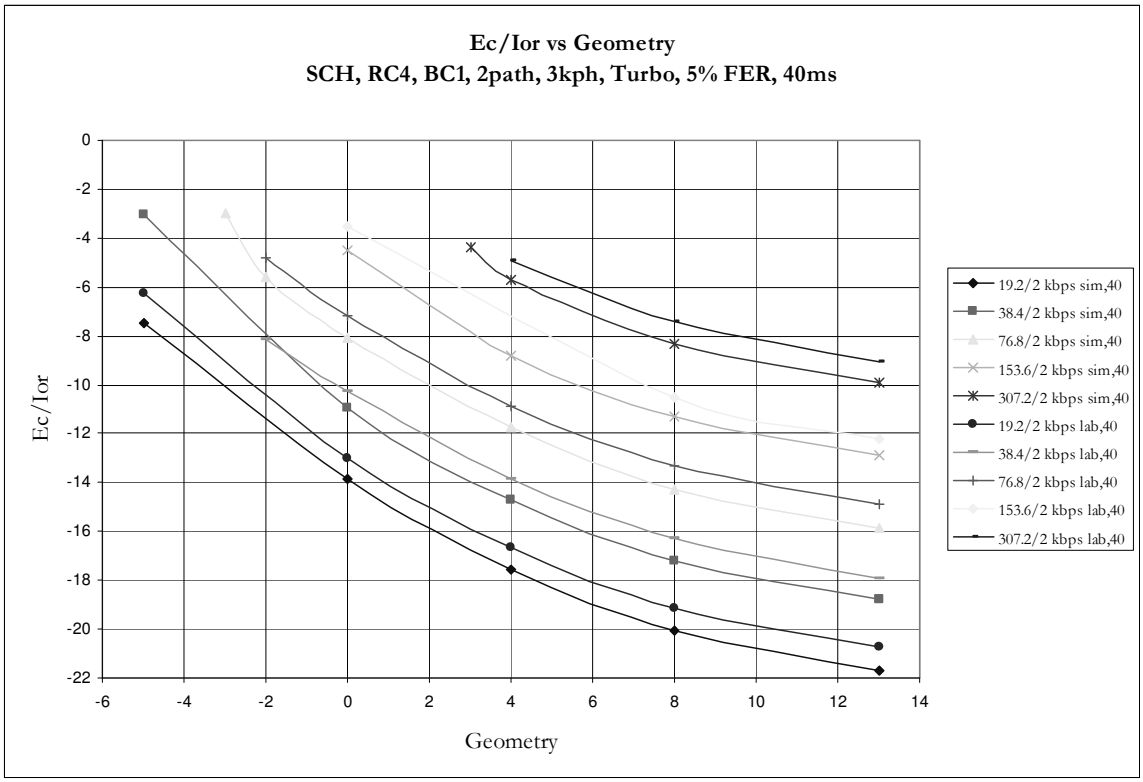
Ec/Ior vs Geometry
SCH, RC4, BC1, 2path, 3kph, Turbo, 5% FER, 40ms
Ec/Ior
Geometry
0
-2
-4
-6
-8
-10
-12
-14
-16
-18
-20
-22
-6
-4
-2
0
2
4
6
8
10
12
14
19.2/2 kbps sim,40
38.4/2 kbps sim,40
76.8/2 kbps sim,40
153.6/2 kbps sim,40
307.2/2 kbps sim,40
19.2/2 kbps lab,40
38.4/2 kbps lab,40
76.8/2 kbps lab,40
153.6/2 kbps lab,40
307.2/2 kbps lab,40

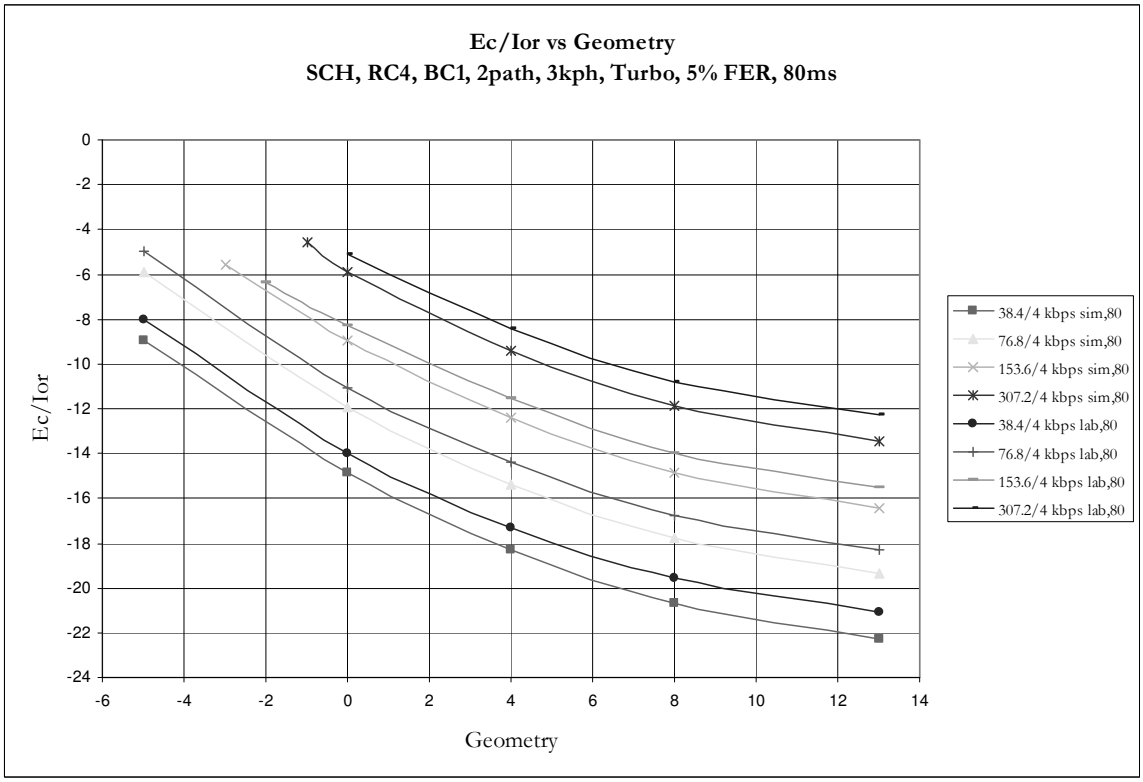
Ec/Ior vs Geometry
SCH, RC4, BC1, 2path, 3kph, Turbo, 5% FER, 80ms
Ec/Ior
Geometry
0
-2
-4
-6
-8
-10
-12
-14
-16
-18
-20
-22
-24
-6
-4
-2
0
2
4
6
8
10
12
14
38.4/4 kbps sim,80
76.8/4 kbps sim,80
153.6/4 kbps sim,80
307.2/4 kbps sim,80
38.4/4 kbps lab,80
76.8/4 kbps lab,80
153.6/4 kbps lab,80
307.2/4 kbps lab,80

9.2.2.5 Rayleigh, 2 path, 30kmph

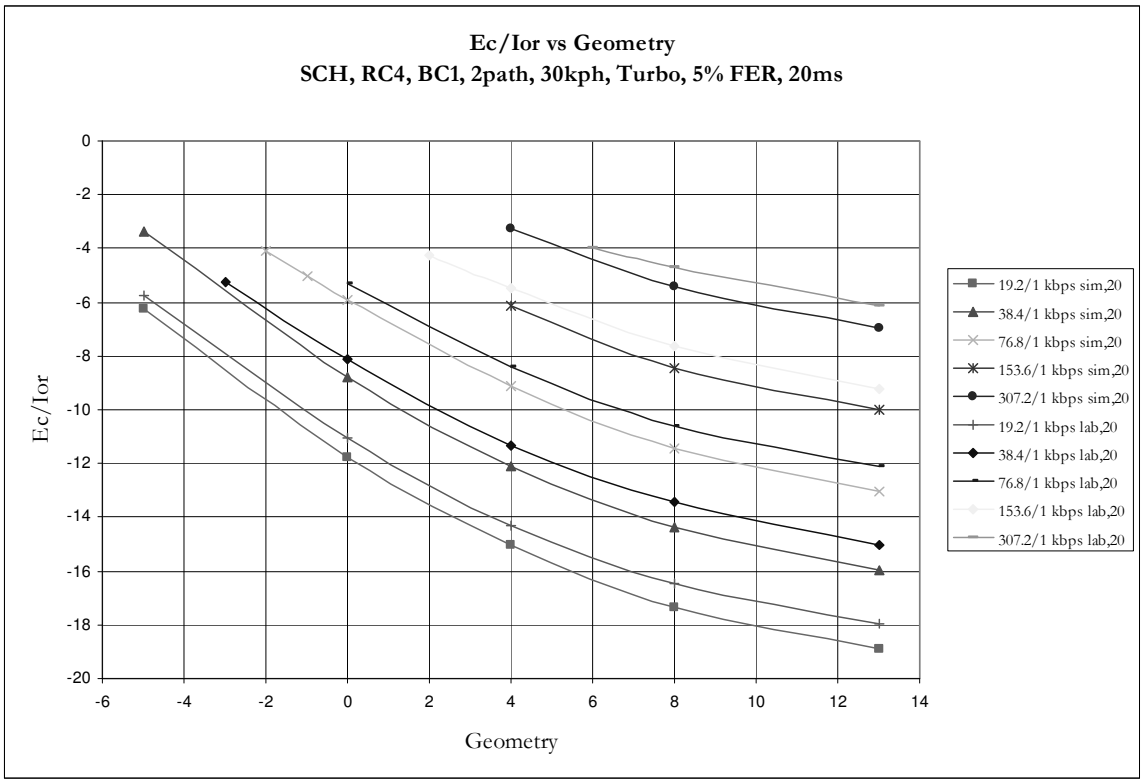

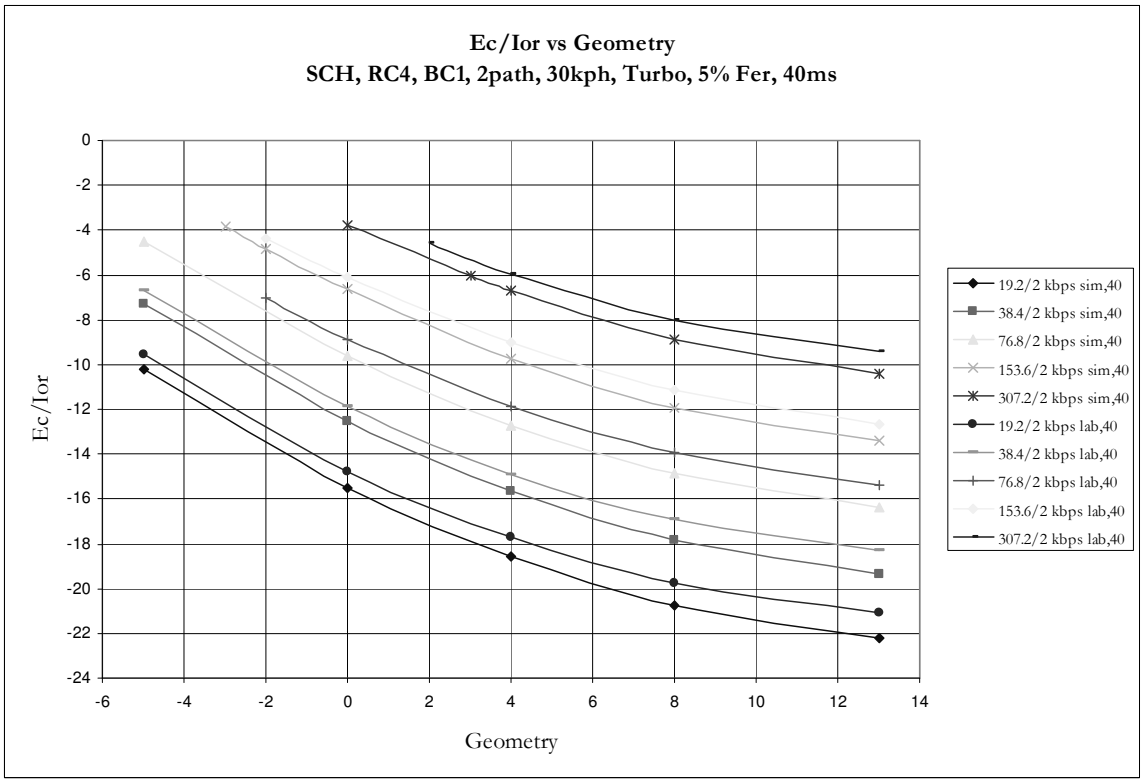

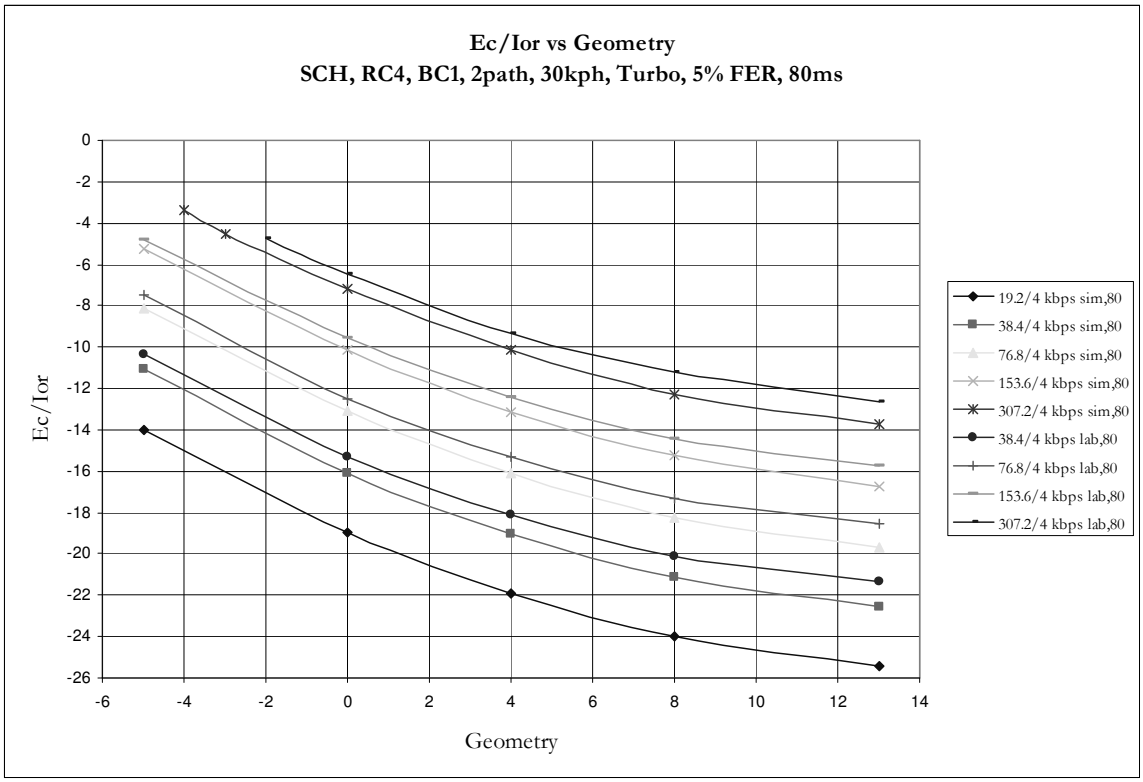
Ec/Ior vs Geometry
SCH, RC4, BC1, 2path, 30kph, Turbo, 5% FER, 80ms
Ec/Ior
Geometry
0
-2
-4
-6
-8
-10
-12
-14
-16
-18
-20
-22
-24
-26
-6
-4
-2
0
2
4
6
8
10
12
14
19.2/4 kbps sim,80
38.4/4 kbps sim,80
76.8/4 kbps sim,80
153.6/4 kbps sim,80
307.2/4 kbps sim,80
38.4/4 kbps lab,80
76.8/4 kbps lab,80
153.6/4 kbps lab,80
307.2/4 kbps lab,80

10 Appendix D—Throughput Performance Results

The following contains the complete set of data for section 4.3.

10.1 Maximum Throughputs

Conv/Turbo, Mux PDU 3	9.6/1, 20ms	19.2/2, 40ms	38.4/4, 80ms	19.2/1, 20ms	38.4/2, 40ms	76.8/4, 80ms
Max App Thruput (Bytes/sec)	960.4	2066.6	2113.5	3172.9	3266.6	3313.5
(Kbytes/sec)	**1.885**	**1.933**	**1.933**	**2.927**	**2.927**	**2.927**

Conv, Mux PDU 5	9.6/1, 20ms	19.2/2, 40ms	38.4/4, 80ms	19.2/1, 20ms	38.4/2, 40ms	76.8/4, 80ms
Max App Thruput (Bytes/sec)	1036.4	960.4	960.4	960.4	960.4	960.4
(Kbytes/sec)	**1.885**	**1.956**	**1.933**	**2.975**	**2.927**	**2.975**

Turbo, Mux PDU 5	9.6/1, 20ms	19.2/2, 40ms	38.4/4, 80ms	19.2/1, 20ms	38.4/2, 40ms	76.8/4, 80ms
Max App Thruput (Bytes/sec)		960.4	960.4	960.4	960.4	960.4
(Kbytes/sec)		**1.956**	**2.016**	**2.975**	**3.093**	**3.152**

Conv/Turbo, Mux PDU 3	38.4/1, 20ms	76.8/2, 40ms	153.6/4, 80ms	76.8/1, 20ms	153.6/2, 40ms	153.6/1, 20ms
Max App Thruput (Bytes/sec)	2160.4	960.4	960.4	960.4	960.4	960.4
(Kbytes/sec)	**4.917**	**4.917**	**4.917**	**8.896**	**8.896**	**16.854**

Conv, Mux PDU 5	38.4/1, 20ms	76.8/2, 40ms	153.6/4, 80ms	76.8/1, 20ms	153.6/2, 40ms	307.2/4, 80ms	153.6/1, 20ms	307.2/2, 40ms	307.2/1, 20ms
Max App Thruput (Bytes/sec)	962.3	962.4	962.3	963.4	963.3	963.4	965.3	965.4	965.4
(Kbytes/sec)	**4.917**	**5.012**	**5.012**	**9.085**	**9.085**	**9.559**	**17.233**	**18.180**	**35.422**

Turbo, Mux PDU 5	38.4/1, 20ms	76.8/2, 40ms	153.6/4, 80ms	76.8/1, 20ms	153.6/2, 40ms	307.2/4, 80ms	153.6/1, 20ms	307.2/2, 40ms	307.2/1, 20ms
Max App Thruput (Bytes/sec)	965.4	961.4	960.4	960.4	960.4	960.4	960.4	960.4	960.4
(Kbytes/sec)	**5.248**	**5.367**	**5.426**	**9.796**	**9.914**	**9.973**	**18.890**	**19.009**	**37.080**

10.2 Clean Channel, Infinite Burst

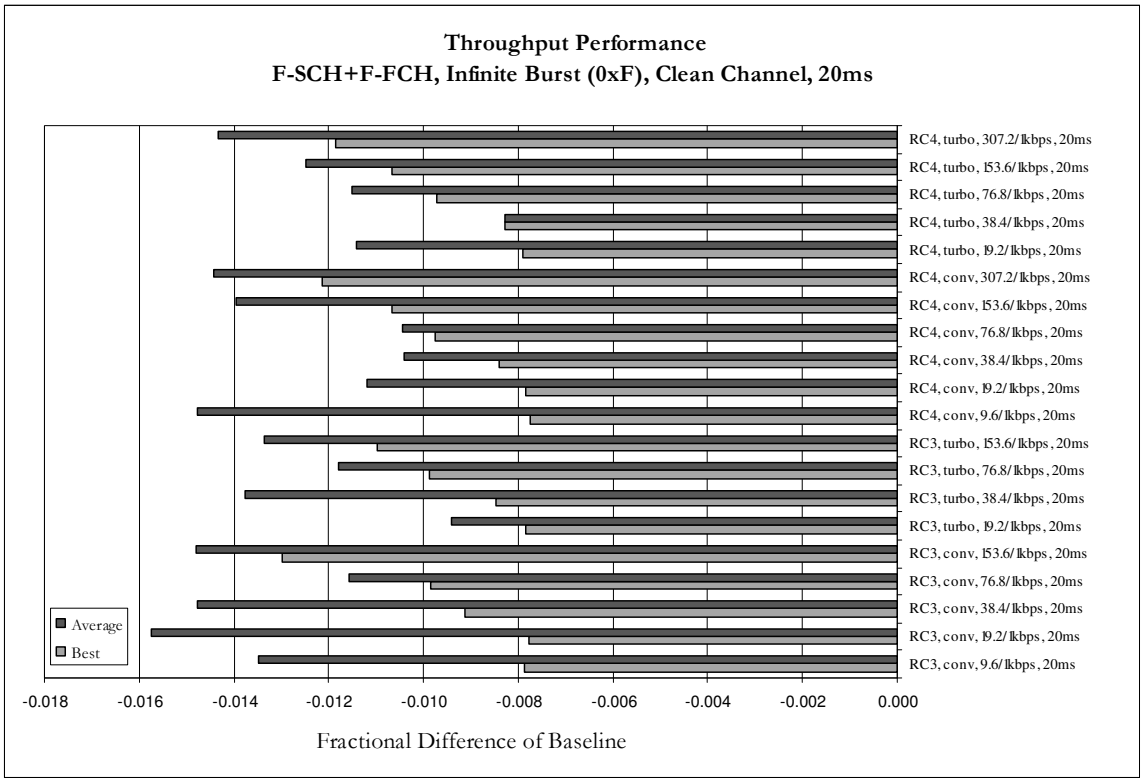

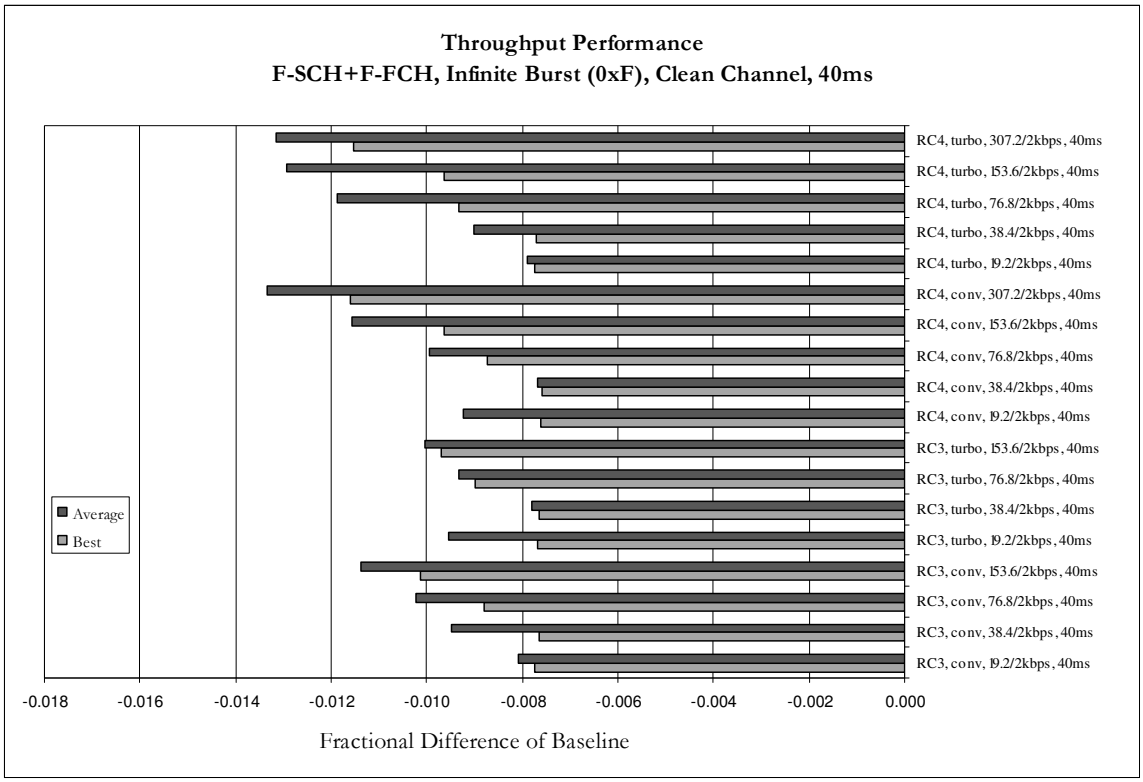
Throughput Performance
F-SCH+F-FCH, Infinite Burst (0xF), Clean Channel, 40ms
RC4, turbo, 307.2/2kbps, 40ms
RC4, turbo, 153.6/2kbps, 40ms
RC4, turbo, 76.8/2kbps, 40ms
RC4, turbo, 38.4/2kbps, 40ms
RC4, turbo, 19.2/2kbps, 40ms
RC4, conv, 307.2/2kbps, 40ms
RC4, conv, 153.6/2kbps, 40ms
RC4, conv, 76.8/2kbps, 40ms
RC4, conv, 38.4/2kbps, 40ms
RC4, conv, 19.2/2kbps, 40ms
RC3, turbo, 153.6/2kbps, 40ms
RC3, turbo, 76.8/2kbps, 40ms
RC3, turbo, 38.4/2kbps, 40ms
RC3, turbo, 19.2/2kbps, 40ms
RC3, conv, 153.6/2kbps, 40ms
RC3, conv, 76.8/2kbps, 40ms
RC3, conv, 38.4/2kbps, 40ms
RC3, conv, 19.2/2kbps, 40ms
Average
Best
-0.018
-0.016
-0.014
-0.012
-0.010
-0.008
-0.006
-0.004
-0.002
0.000
Fractional Difference of Baseline

10.3 Fading Channels, Infinite Burst

10.3.1 Model A, 1 path, 3kmph

10.3.1.1 RC3

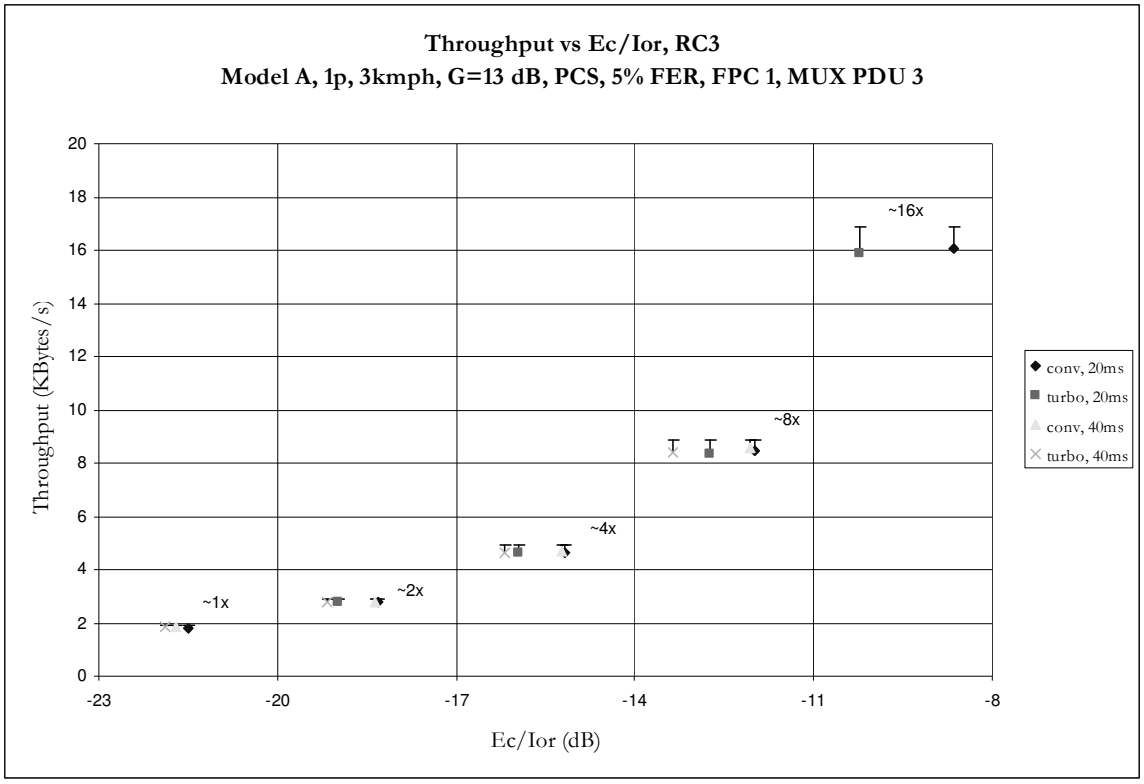

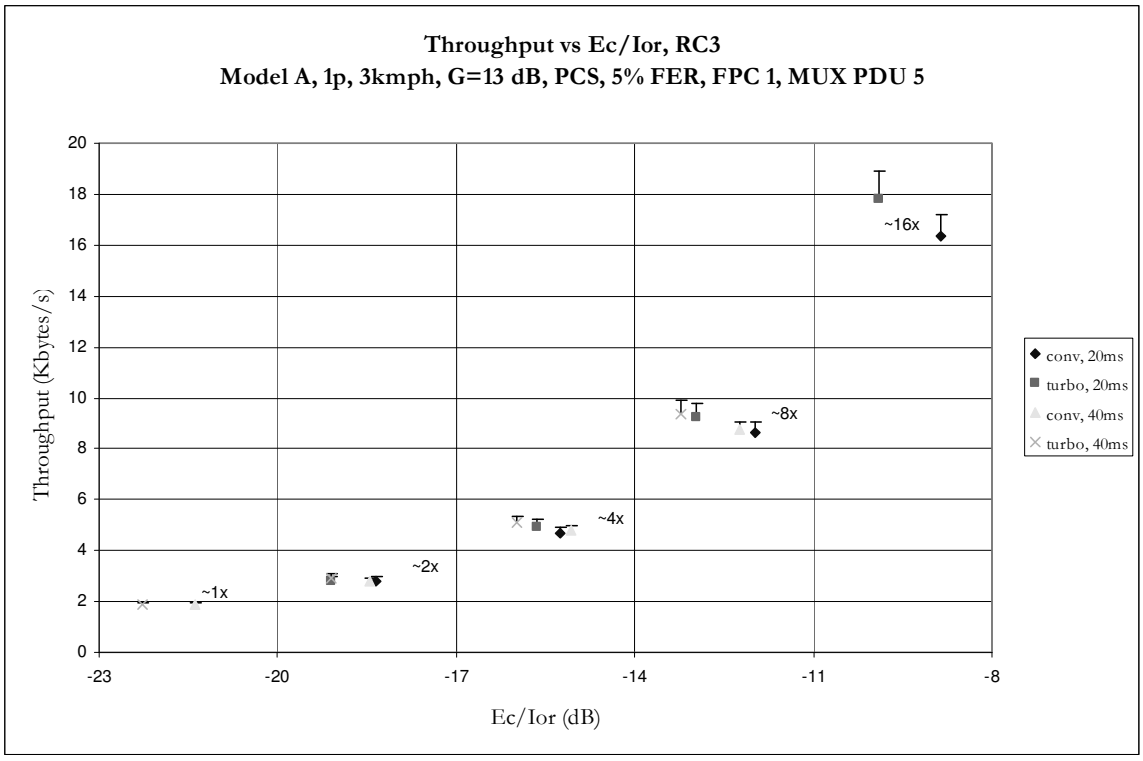
Throughput vs Ec/Ior, RC3
Model A, 1p, 3kmph, G=13 dB, PCS, 5% FER, FPC 1, MUX PDU 5
Throughput (Kbytes/s)
Ec/Ior (dB)
~1x
~2x
~4x
~8x
~16x
conv, 20ms
turbo, 20ms
conv, 40ms
turbo, 40ms

10.3.1.2 RC4

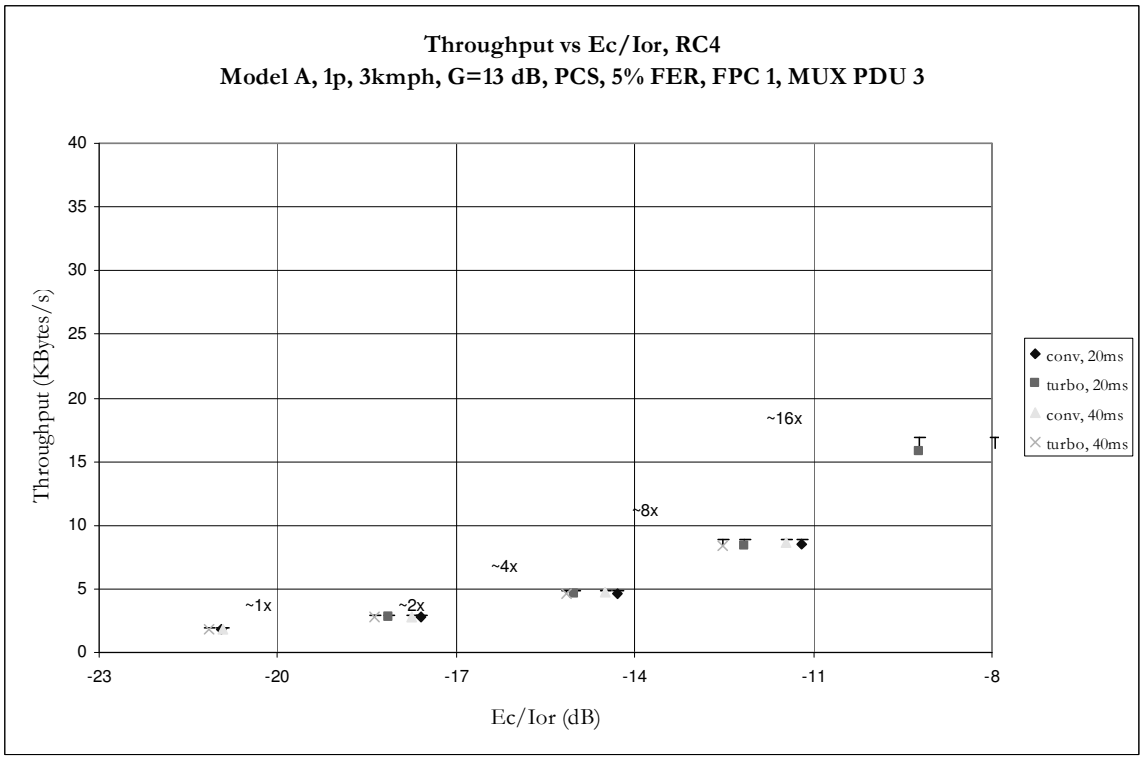
Throughput vs Ec/Ior, RC4
Model A, 1p, 3kmph, G=13 dB, PCS, 5% FER, FPC 1, MUX PDU 3
Throughput (KBytes/s)
Ec/Ior (dB)
~1x
~2x
~4x
~8x
~16x
conv, 20ms
turbo, 20ms
conv, 40ms
turbo, 40ms

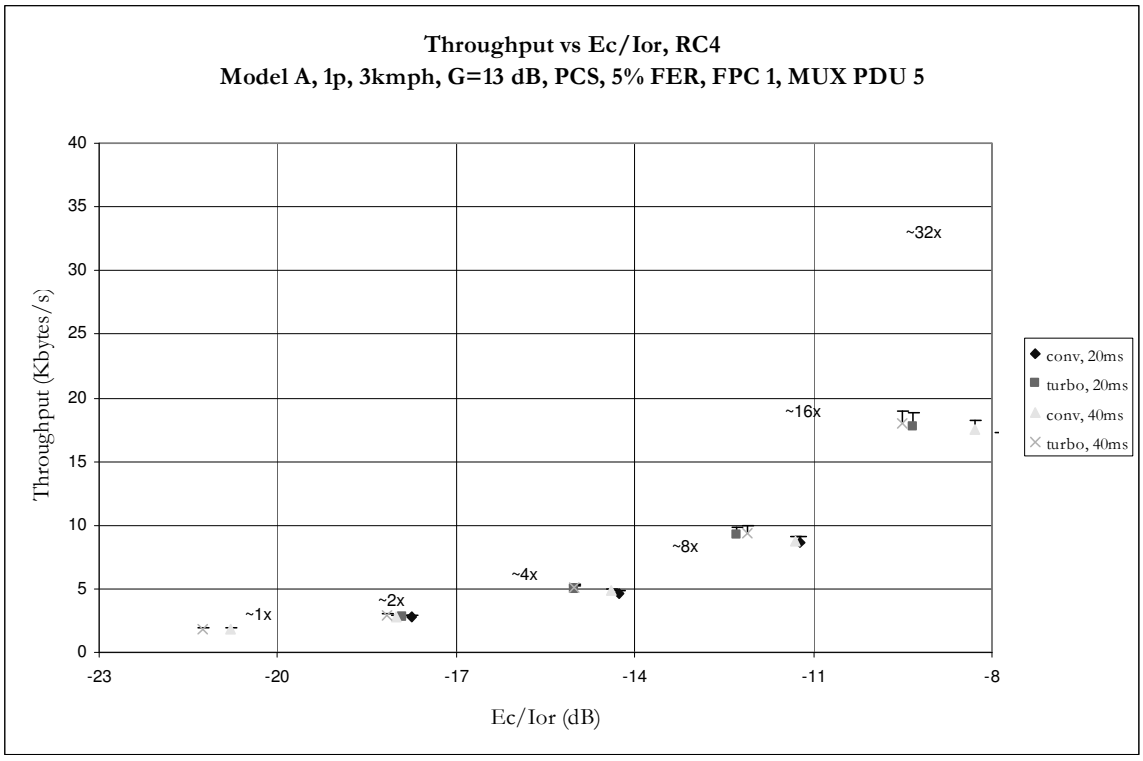

10.3.2 Model B, 3 paths, 10kmph

10.3.2.1 RC3

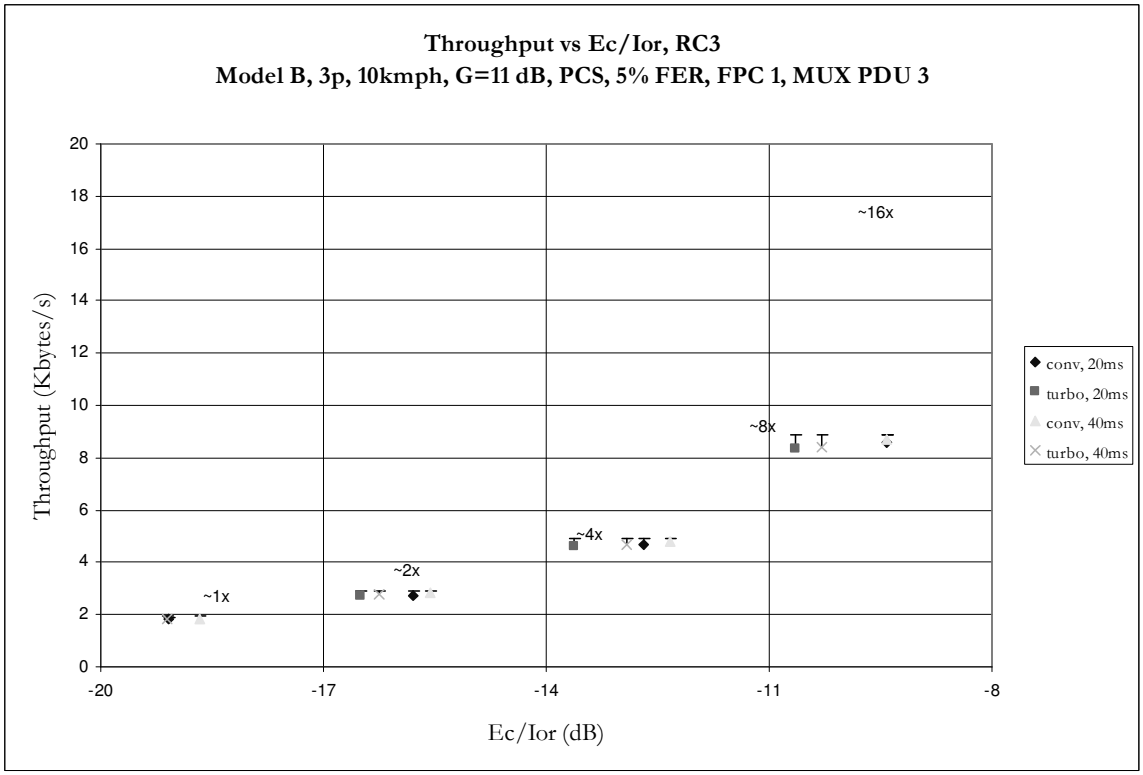

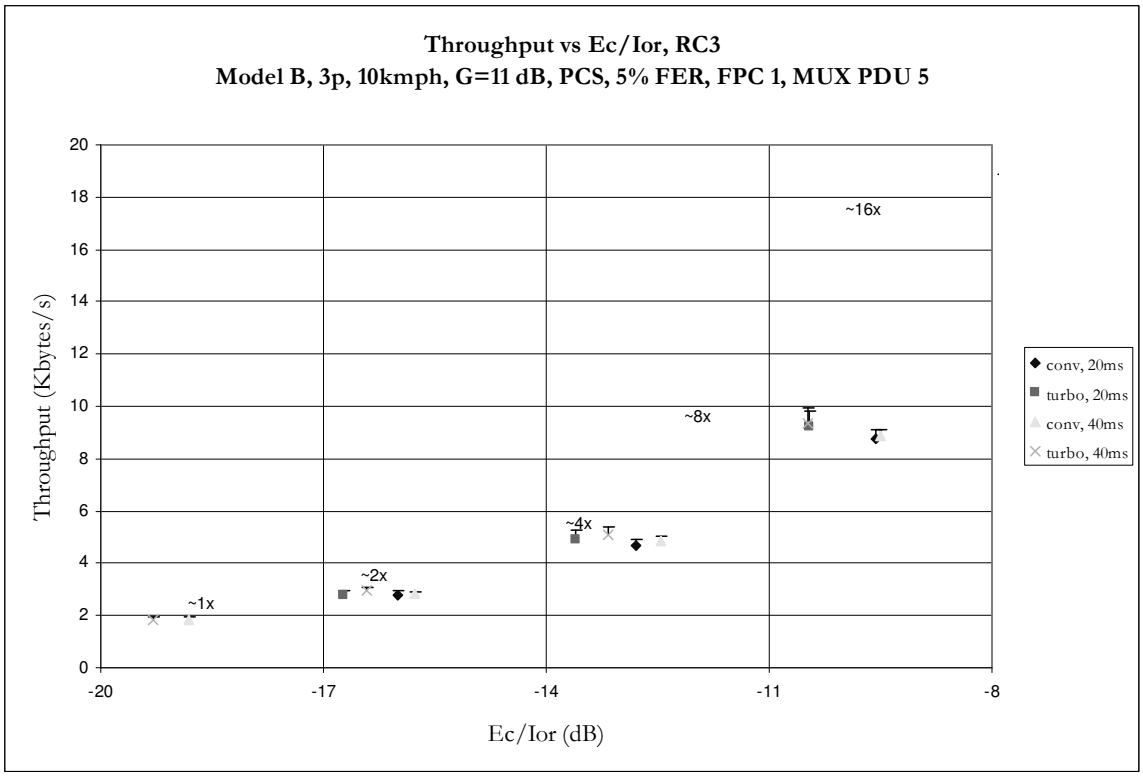
Throughput vs Ec/Ior, RC3
Model B, 3p, 10kmph, G=11 dB, PCS, 5% FER, FPC 1, MUX PDU 5
Throughput (Kbytes/s)
Ec/Ior (dB)
~1x
~2x
~4x
~8x
~16x
conv, 20ms
turbo, 20ms
conv, 40ms
turbo, 40ms

10.3.2.2 RC4

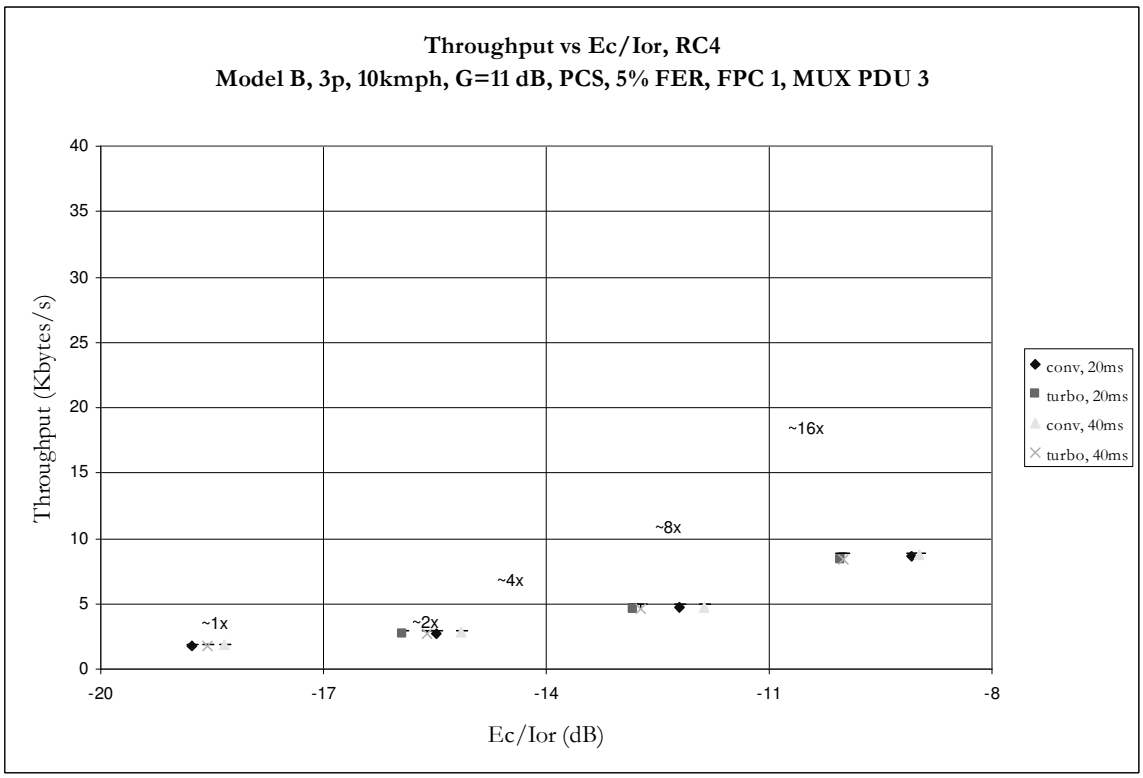
Throughput vs Ec/Ior, RC4
Model B, 3p, 10kmph, G=11 dB, PCS, 5% FER, FPC 1, MUX PDU 3
Throughput (Kbytes/s)
Ec/Ior (dB)
~1x
~2x
~4x
~8x
~16x
conv, 20ms
turbo, 20ms
conv, 40ms
turbo, 40ms

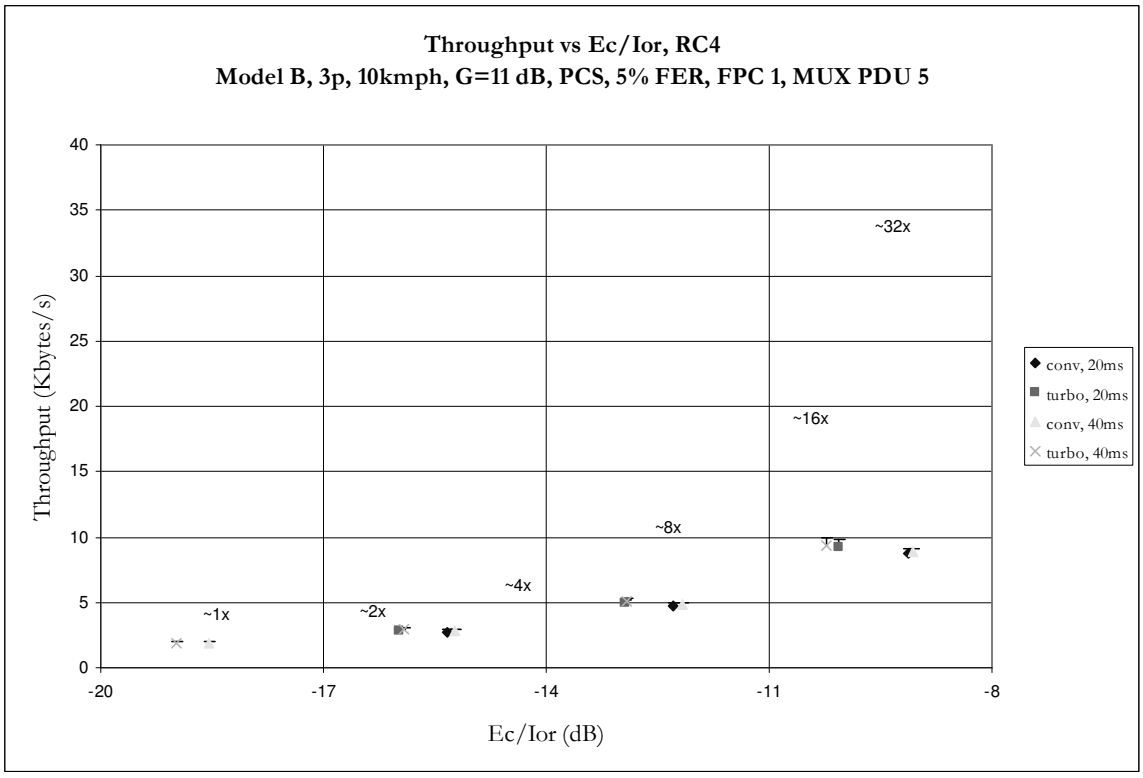

10.3.3 Model C, 2 paths, 30kmph

10.3.3.1 RC3

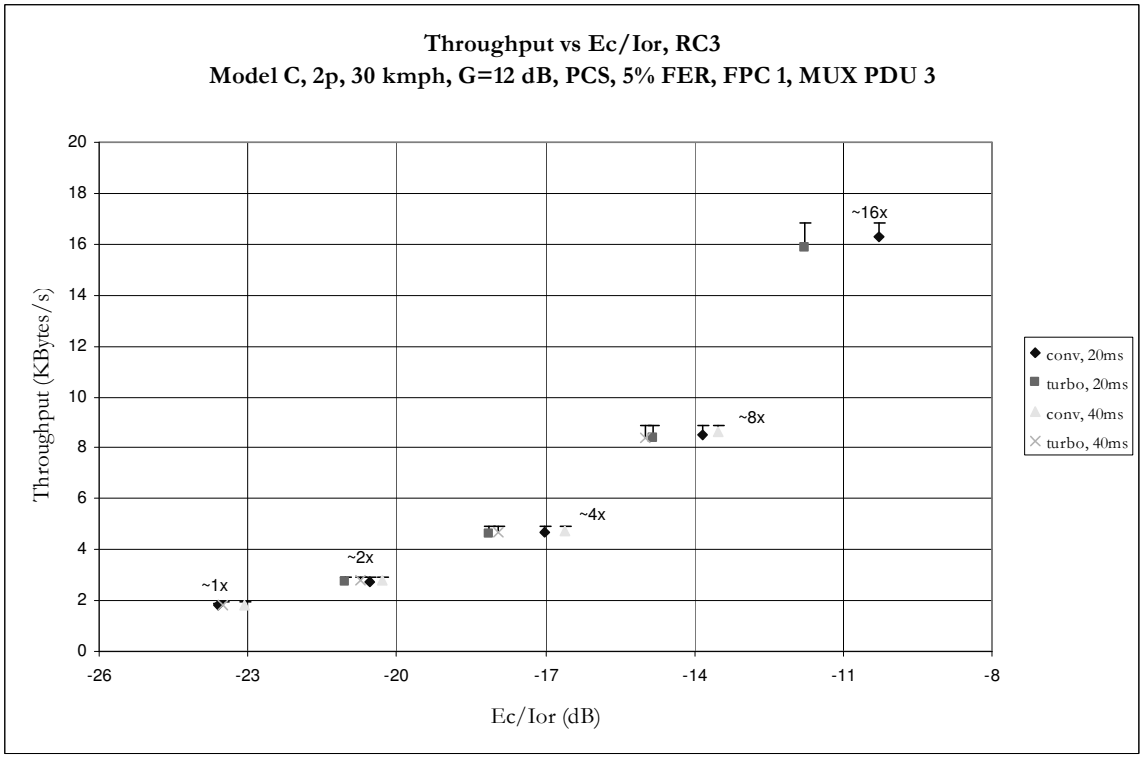

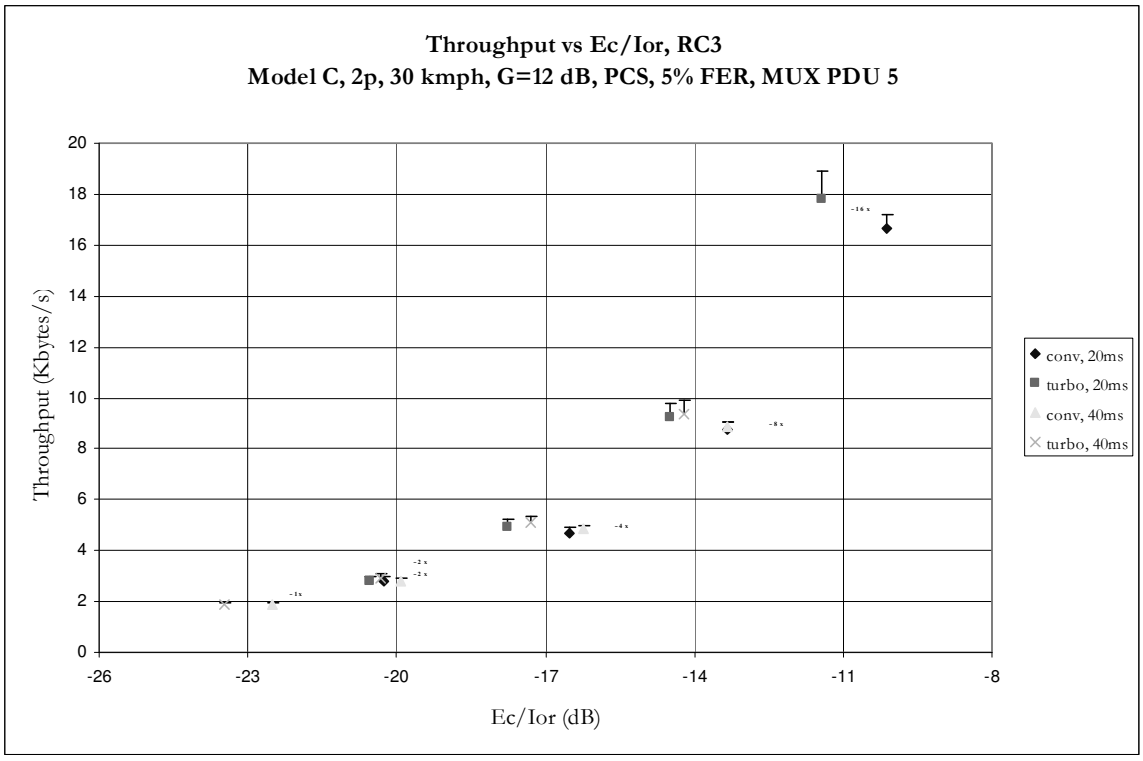
Throughput vs Ec/Ior, RC3
Model C, 2p, 30 kmph, G=12 dB, PCS, 5% FER, MUX PDU 5
Throughput (Kbytes/s)
Ec/Ior (dB)
conv, 20ms
turbo, 20ms
conv, 40ms
turbo, 40ms

10.3.3.2 RC4

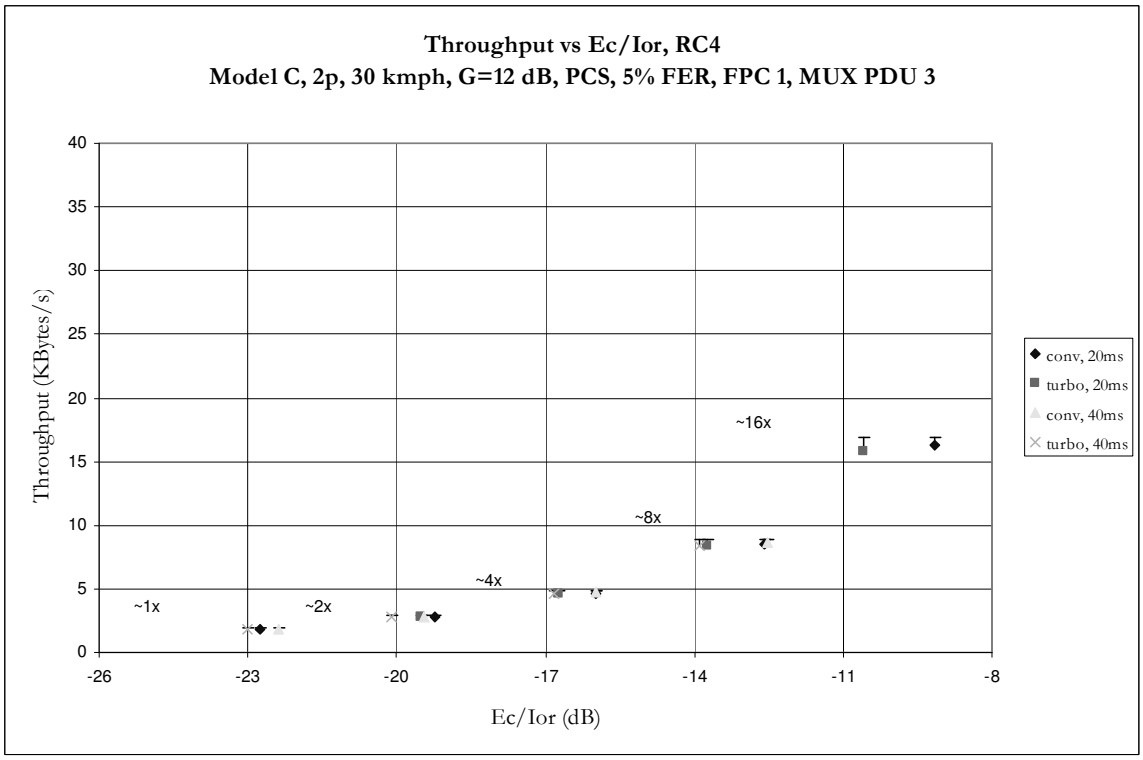
Throughput vs Ec/Ior, RC4
Model C, 2p, 30 kmph, G=12 dB, PCS, 5% FER, FPC 1, MUX PDU 3
Throughput (KBytes/s)
Ec/Ior (dB)
~1x
~2x
~4x
~8x
~16x
conv, 20ms
turbo, 20ms
conv, 40ms
turbo, 40ms

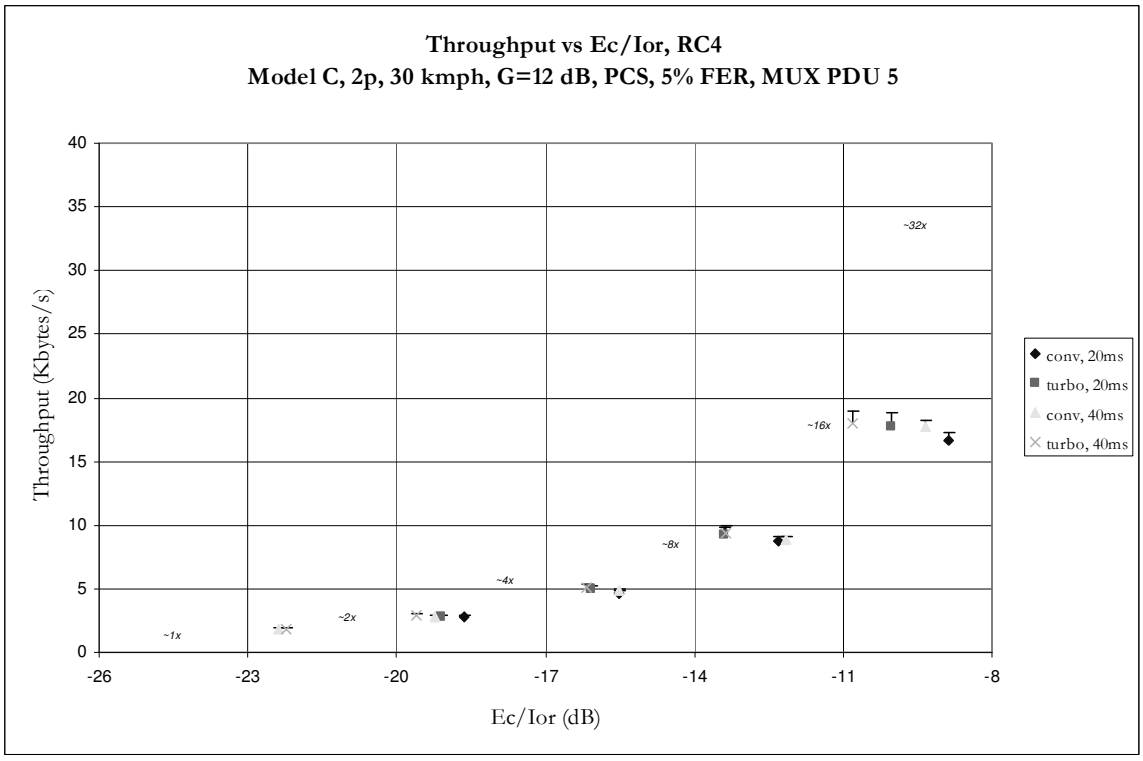

10.3.4 Model D, 1 path, 100kmph

10.3.4.1 RC3

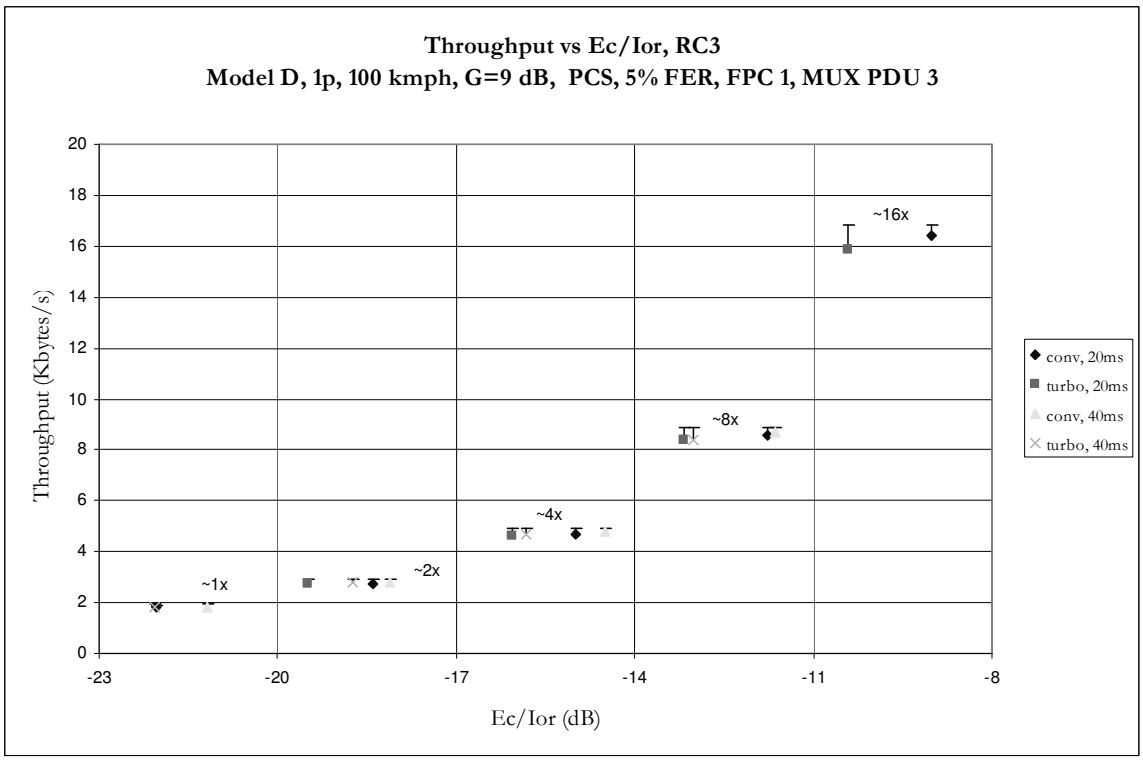

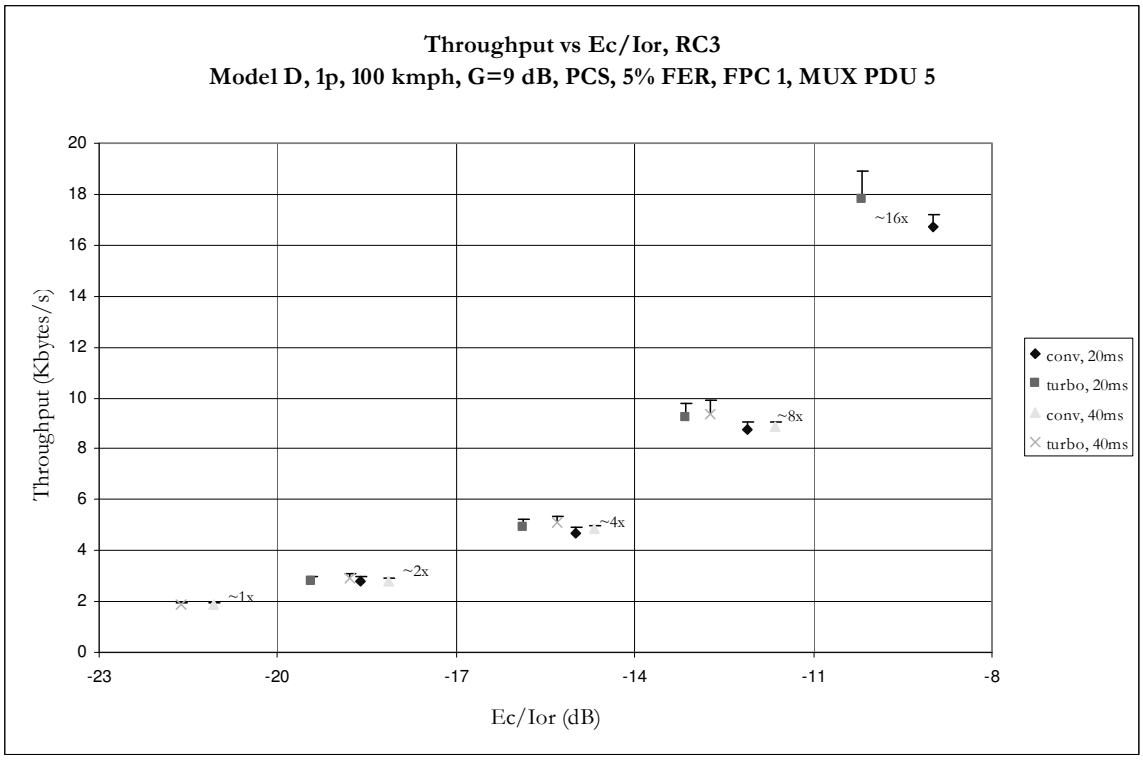
Throughput vs Ec/Ior, RC3
Model D, 1p, 100 kmph, G=9 dB, PCS, 5% FER, FPC 1, MUX PDU 5
Throughput (Kbytes/s)
Ec/Ior (dB)
conv, 20ms
turbo, 20ms
conv, 40ms
turbo, 40ms
~1x
~2x
~4x
~8x
~16x

10.3.4.2 RC4

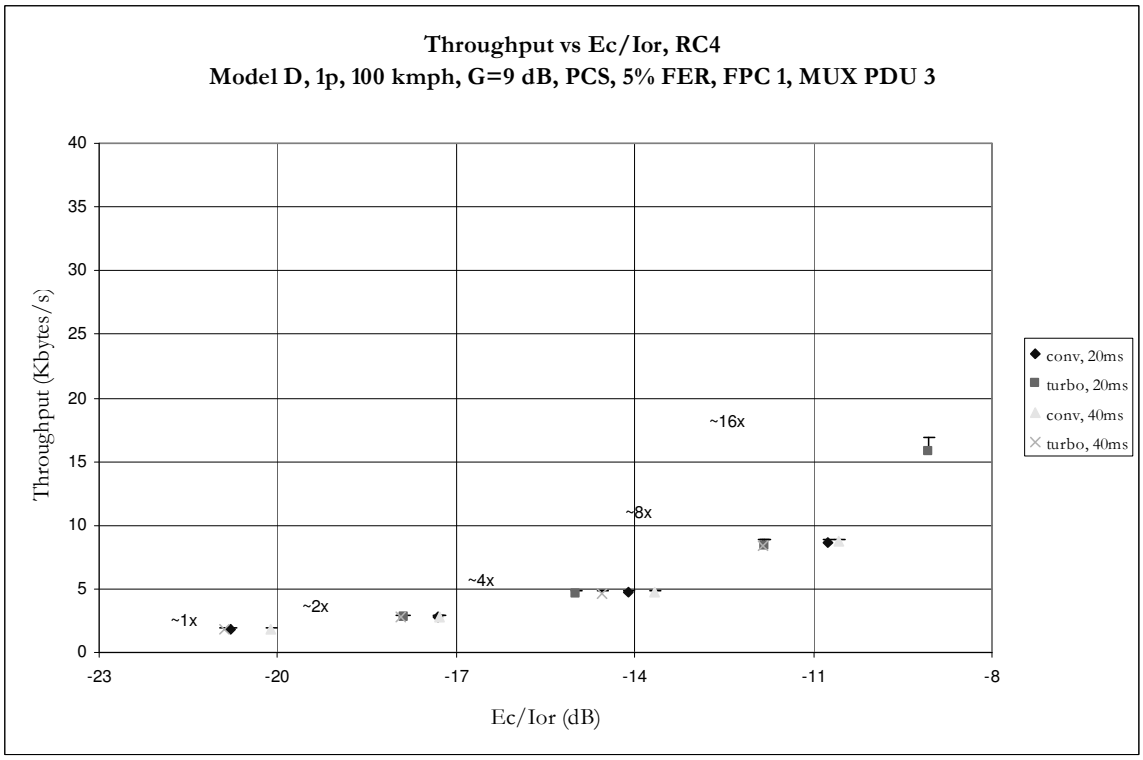
Throughput vs Ec/Ior, RC4
Model D, 1p, 100 kmph, G=9 dB, PCS, 5% FER, FPC 1, MUX PDU 3
Throughput (Kbytes/s)
Ec/Ior (dB)
conv, 20ms
turbo, 20ms
conv, 40ms
turbo, 40ms
~1x
~2x
~4x
~8x
~16x

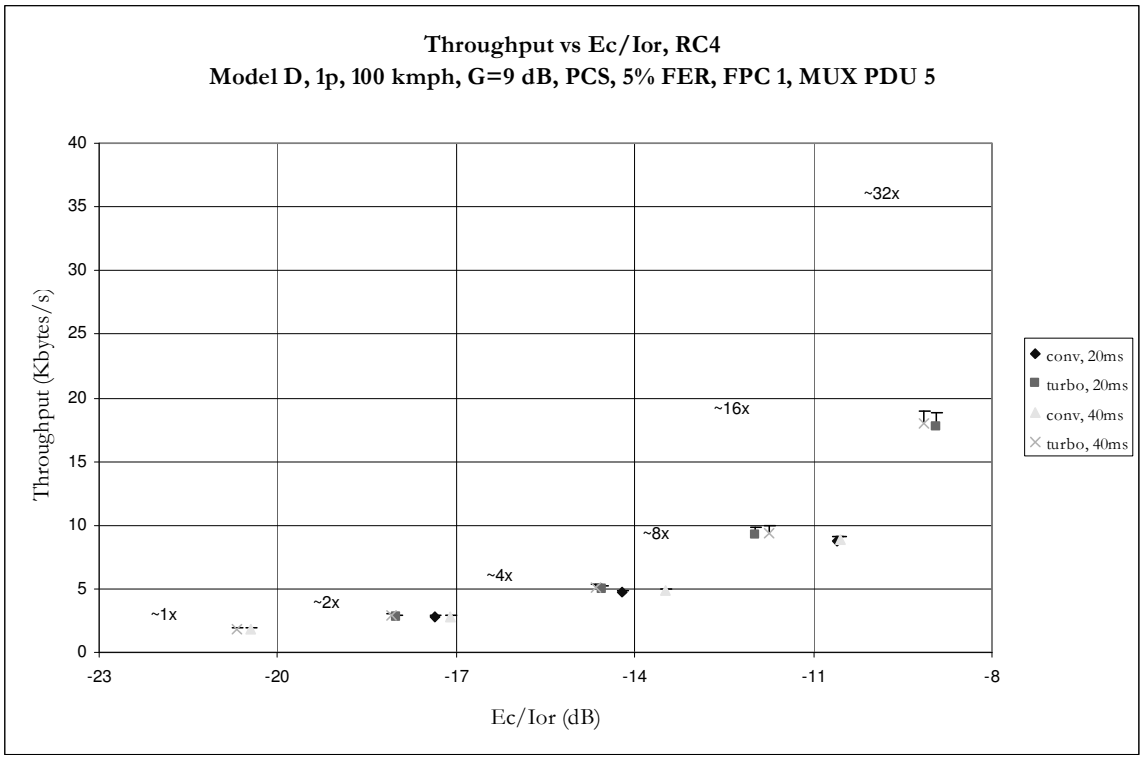

10.3.5 Model E, 1 path Rician, 1kmph

10.3.5.1 RC3

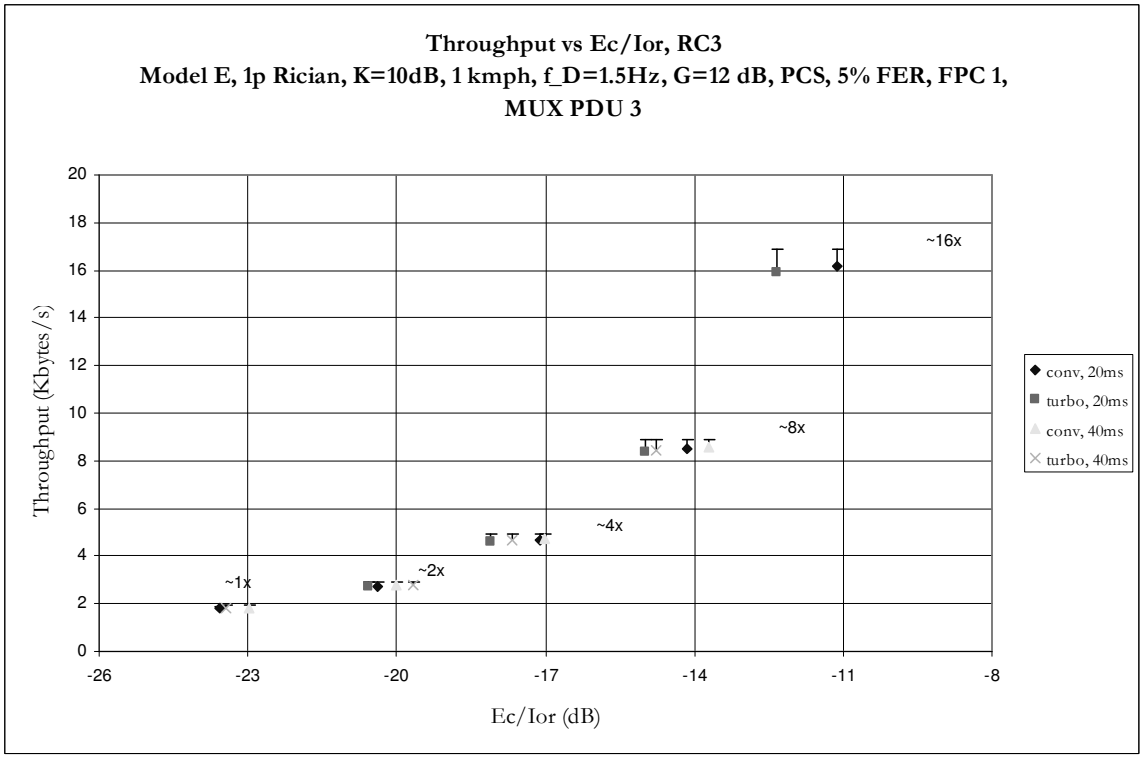

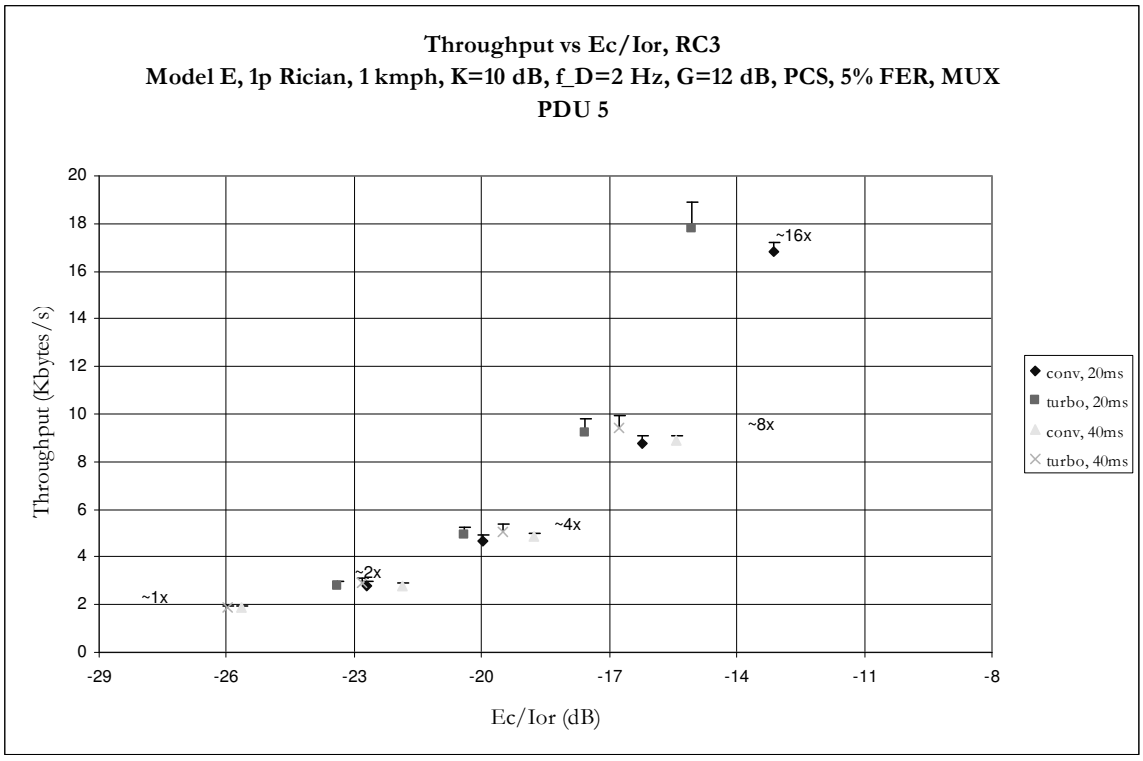
Throughput vs Ec/Ior, RC3
Model E, 1p Rician, 1 kmph, K=10 dB, f_D=2 Hz, G=12 dB, PCS, 5% FER, MUX PDU 5
Throughput (Kbytes/s)
Ec/Ior (dB)
~1x
~2x
~4x
~8x
~16x
conv, 20ms
turbo, 20ms
conv, 40ms
turbo, 40ms

10.3.5.2 RC4

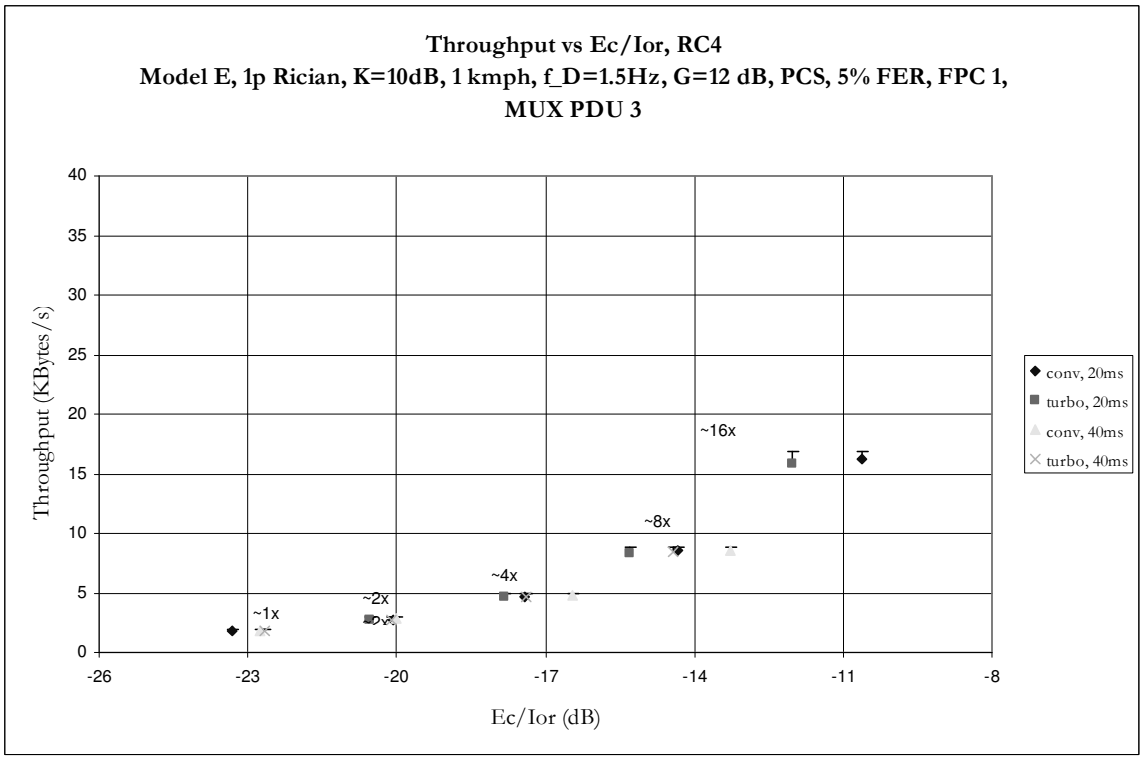
Throughput vs Ec/Ior, RC4
Model E, 1p Rician, K=10dB, 1 kmph, f_D=1.5Hz, G=12 dB, PCS, 5% FER, FPC 1, MUX PDU 3
Throughput (KBytes/s)
Ec/Ior (dB)
~1x
~2x
~4x
~8x
~16x
conv, 20ms
turbo, 20ms
conv, 40ms
turbo, 40ms

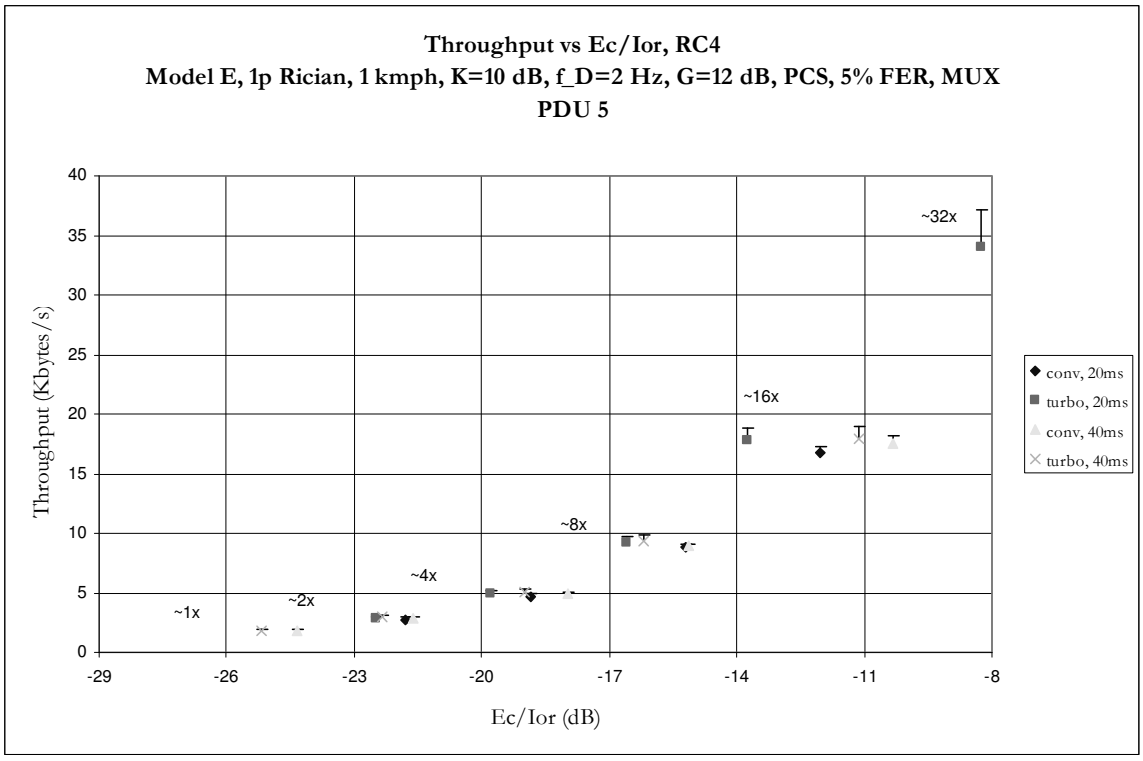

10.3.6 1 path, 10kmph

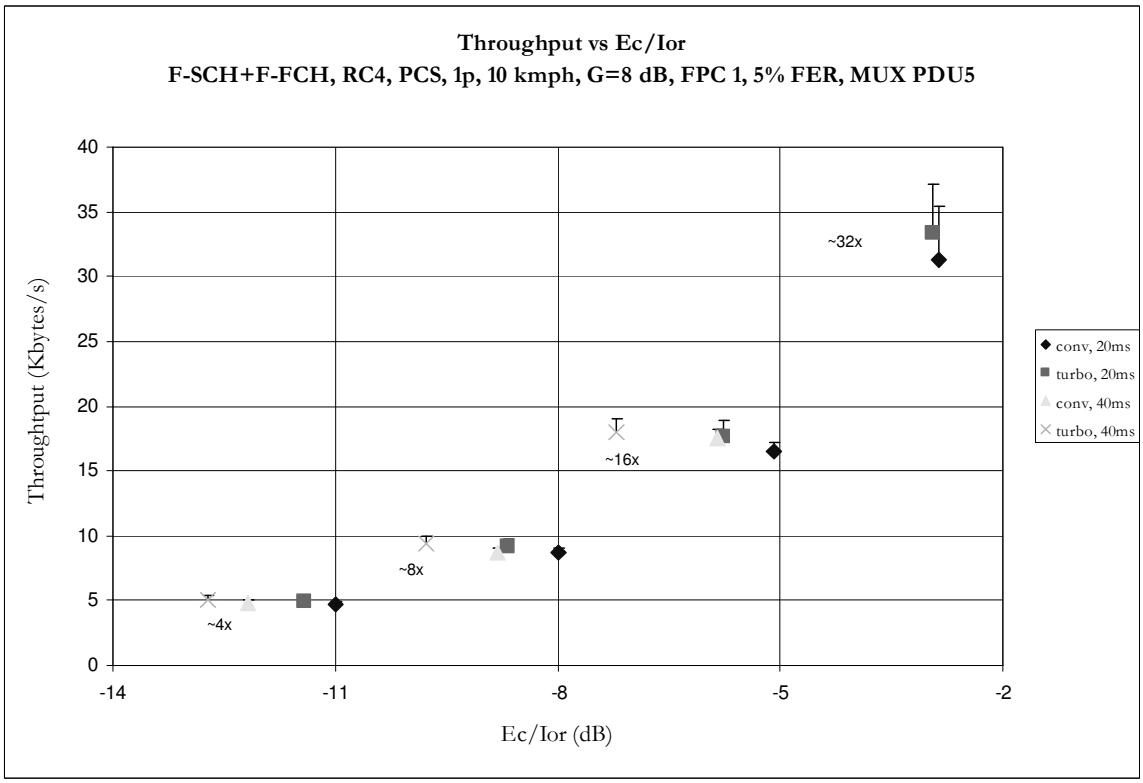

10.4 MFI Gain vs. Speed

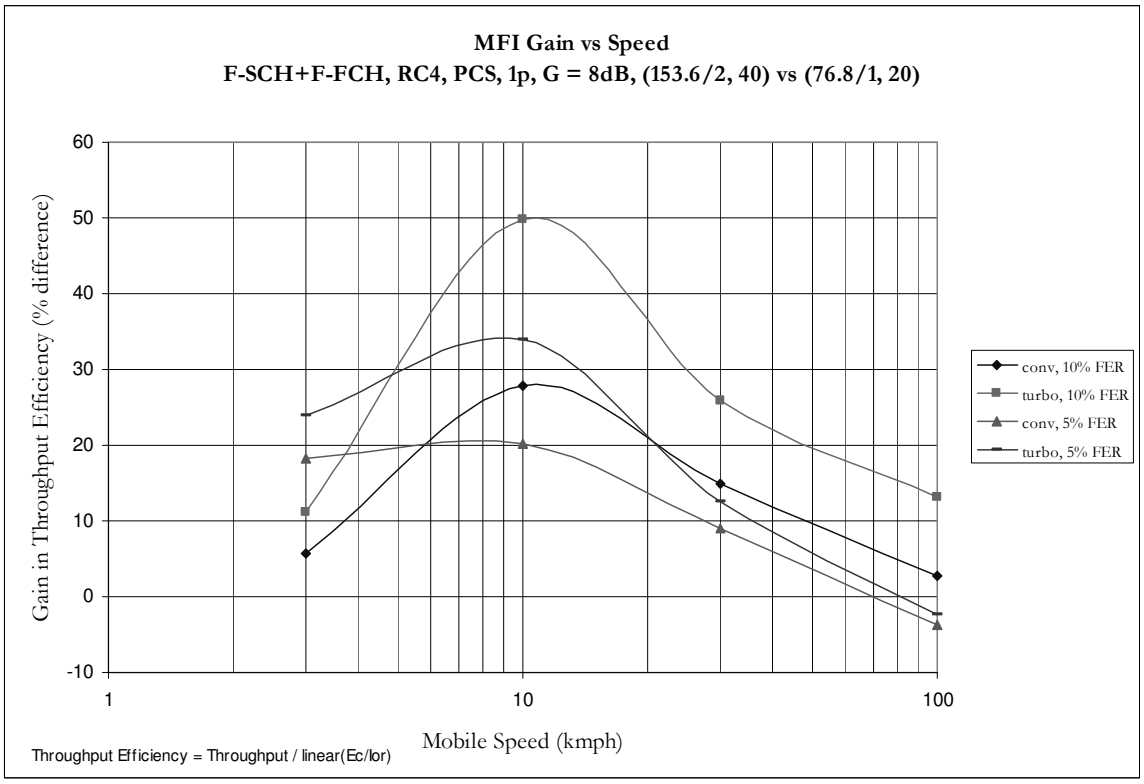

11 Appendix E—Scheduler Performance Results

The following contains the complete set of data for section 4.4.

11.1 Delta Values

	Delta Values for Forward Link Scheduler (dB)			
	RC3		RC4	
Rate	Convolutional	Turbo	Convolutional	Turbo
1x	0	0	0	0
2x	3.38	2.21	3.44	3.25
4x	6.39	5.06	6.44	6.05
8x	9.37	7.96	9.48	8.96
16x	12.38	10.89	12.47	11.94
32x	--	--	15.48	15.06

11.2 Turbo Encoding

11.2.1 Sector Throughput

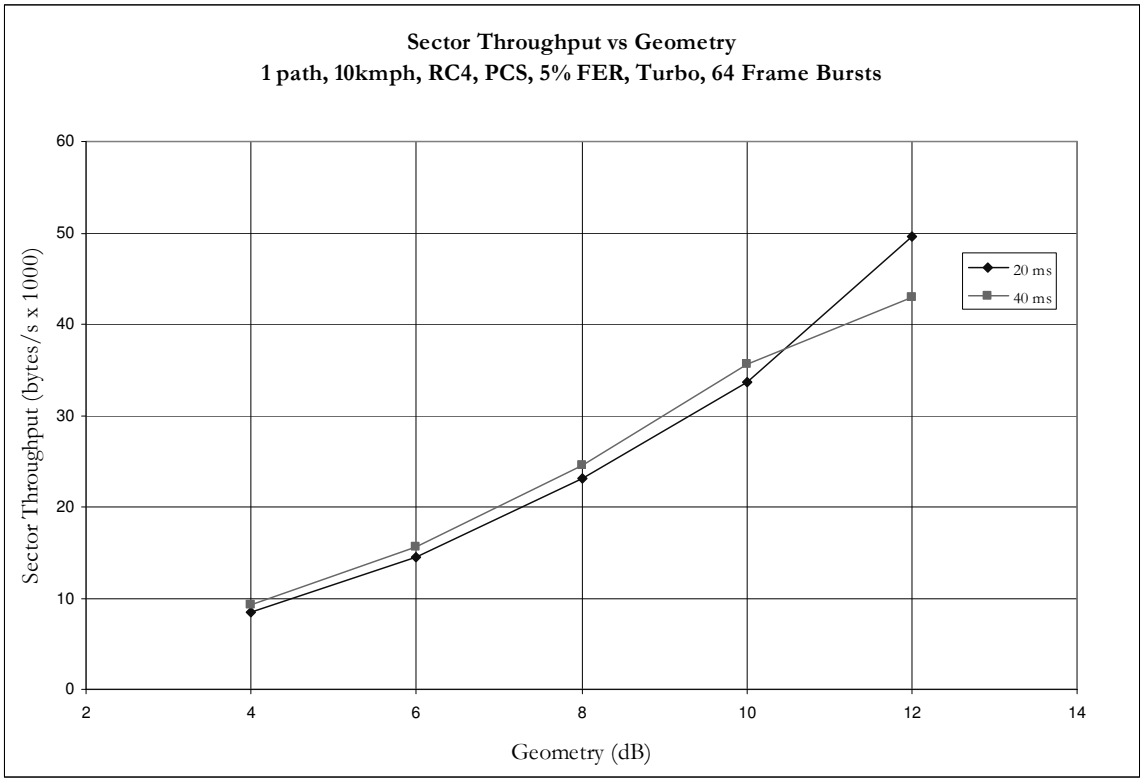

11.2.2 SCH Assignment History

		Average Assignments							Average Assignment
		0x	~1x	~2x	~4x	~8x	~16x	~32x	(x)
G=4dB	20 ms	0.39	0.16	0.15	0.22	0.08	0.00	0.00	**1.99**
	40 ms	0.37	0.16	0.16	0.19	0.12	0.00		**2.20**
G=6dB	20 ms	0.23	0.16	0.14	0.16	0.27	0.04	0.00	**3.89**
	40 ms	0.22	0.16	0.14	0.17	0.24	0.08		**4.32**
G=8dB	20 ms	0.17	0.12	0.14	0.15	0.14	0.27	0.01	**6.80**
	40 ms	0.12	0.11	0.14	0.17	0.13	0.32		**7.32**
G=10dB	20 ms	0.12	0.09	0.12	0.14	0.17	0.22	0.14	**10.27**
	40 ms	0.03	0.03	0.07	0.13	0.18	0.55		**10.99**
G=12dB	20 ms	0.05	0.05	0.10	0.13	0.16	0.19	0.33	**15.74**
	40 ms	0.00	0.00	0.00	0.01	0.32	0.67		**13.25**

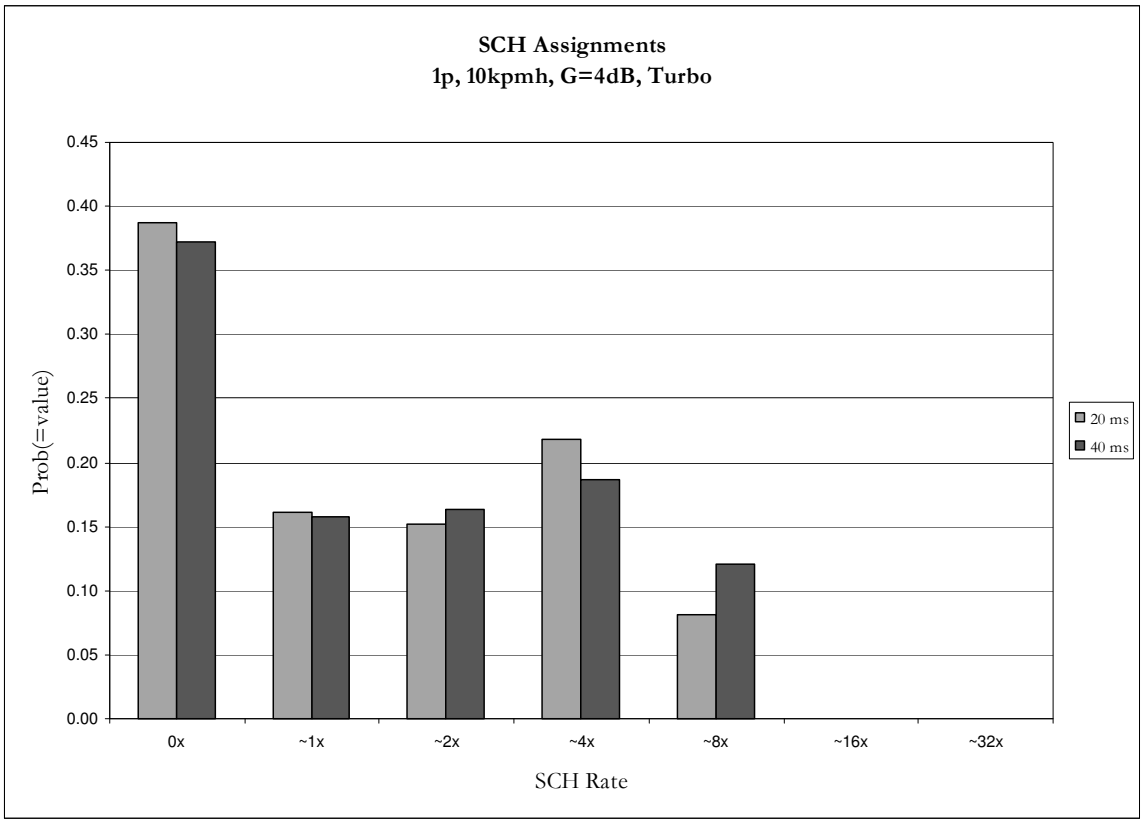

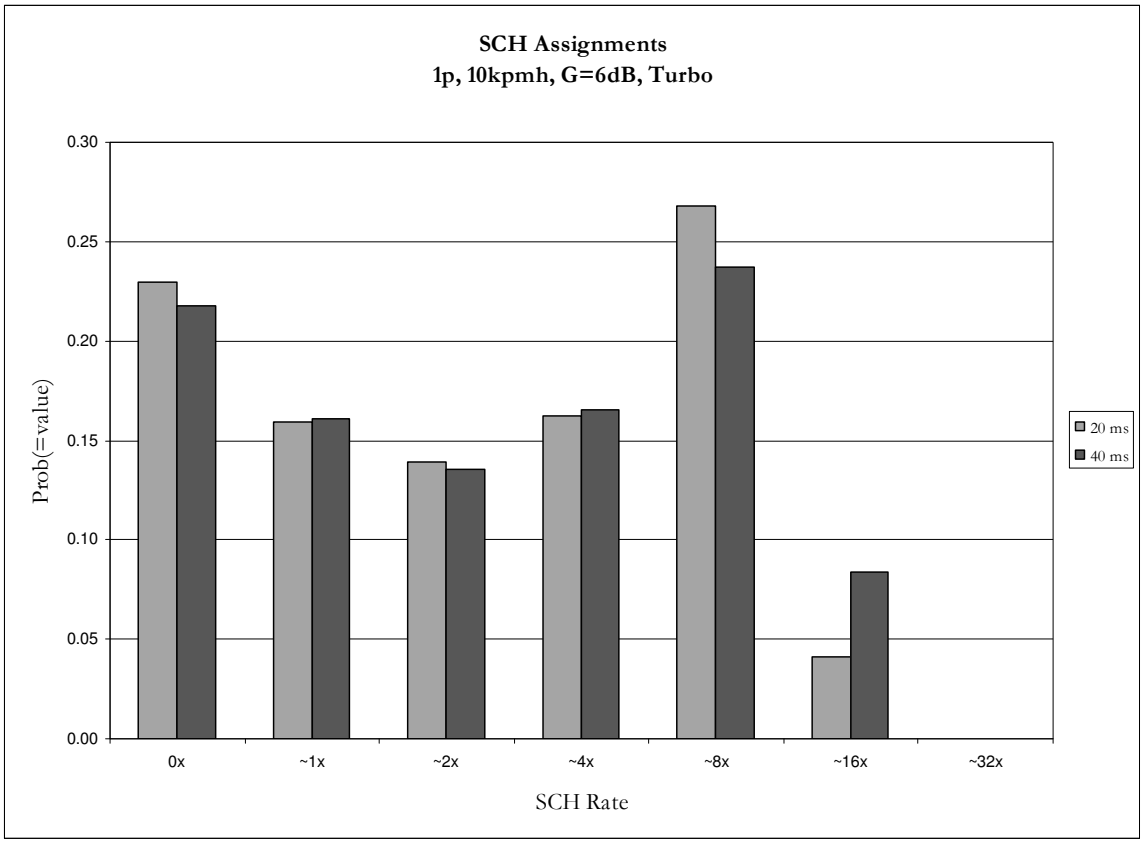
SCH Assignments
1p, 10kpmh, G=6dB, Turbo
0.30
0.25
0.20
0.15
0.10
0.05
0.00
Prob(=value)
0x
~1x
~2x
~4x
~8x
~16x
~32x
SCH Rate
20 ms
40 ms

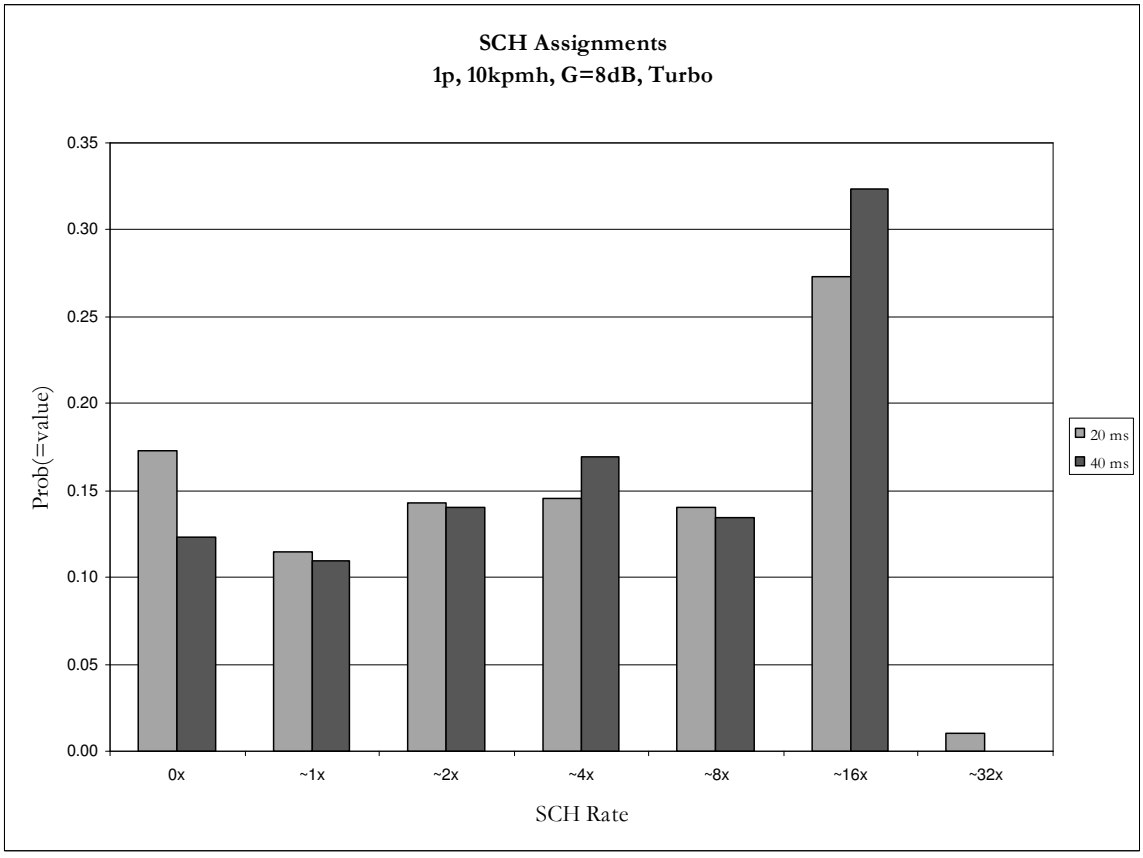
SCH Assignments
1p, 10kpmh, G=8dB, Turbo
0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
Prob(=value)
0x
~1x
~2x
~4x
~8x
~16x
~32x
SCH Rate
20 ms
40 ms

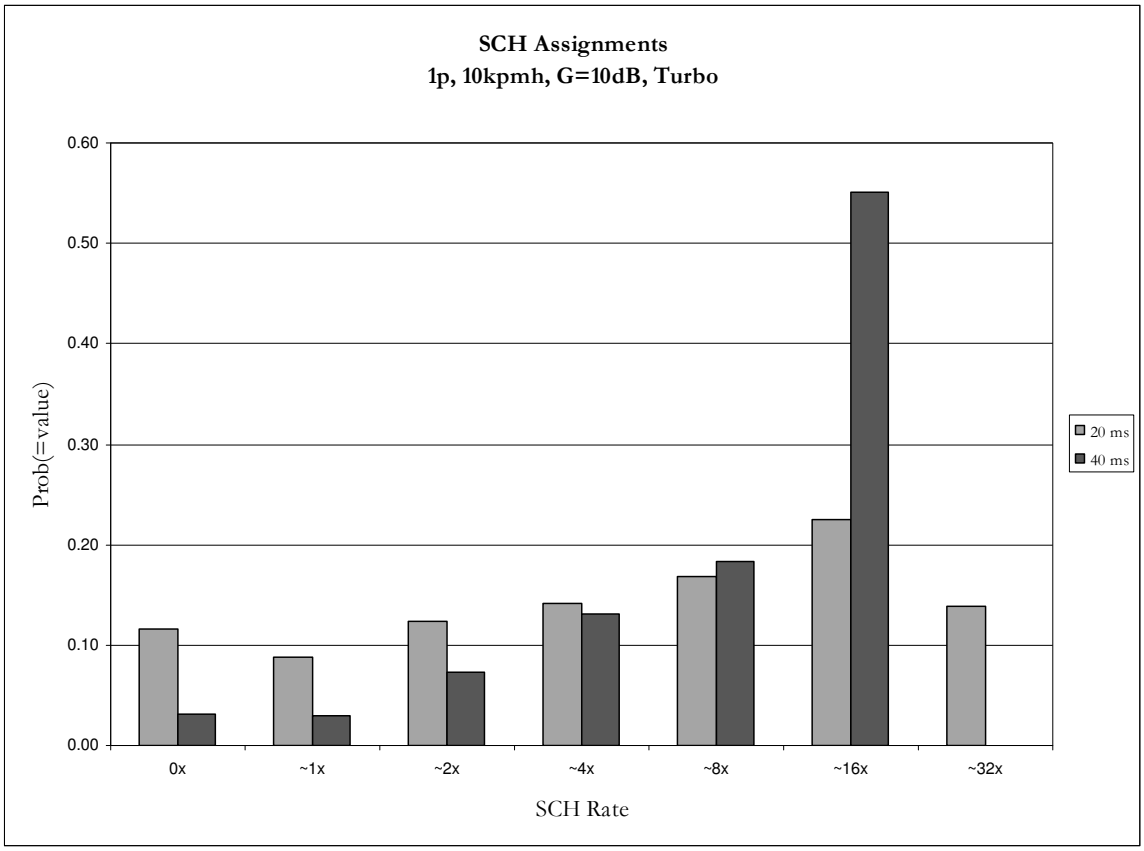
SCH Assignments
1p, 10kpmh, G=10dB, Turbo
0.60
0.50
0.40
0.30
0.20
0.10
0.00
Prob(=value)
20 ms
40 ms
0x
~1x
~2x
~4x
~8x
~16x
~32x
SCH Rate

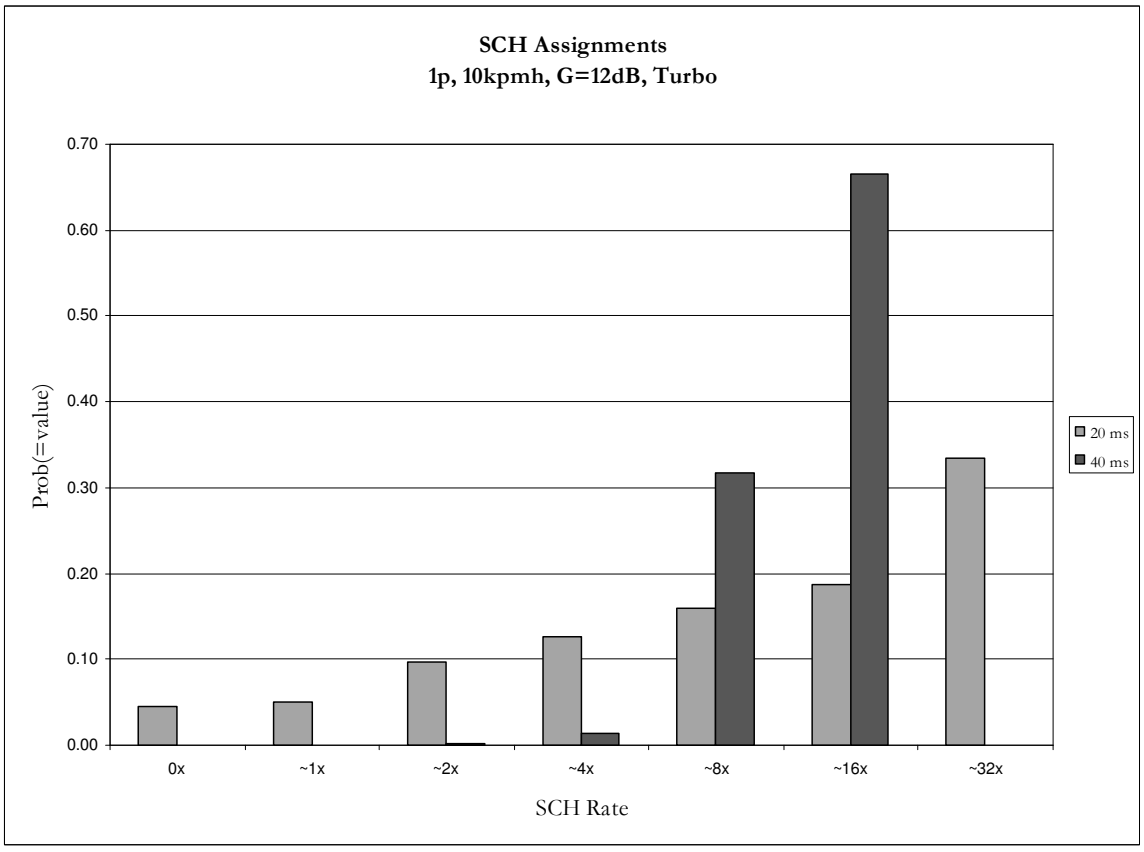
SCH Assignments
1p, 10kpmh, G=12dB, Turbo
0.70
0.60
0.50
0.40
0.30
0.20
0.10
0.00
Prob(=value)
20 ms
40 ms
0x
~1x
~2x
~4x
~8x
~16x
~32x
SCH Rate

11.2.3 Throughput Efficiency vs. Speed

G=8dB	Scheduler	Avg Throughputs (bytes/s x 1000)			Avg SCH Ec/Ior (dB)			Throughput Efficiency					
Speed		Call 0	Call 1	Call 2	Call 0	Call 1	Call 2	Call 0	Call 1	Call 2	Total	Diff	% Diff
3	20 ms	4.48	4.64	4.70	-10.04	-9.95	-10.26	45.13	45.91	49.91	141	**-4.23**	**-3.00**
	40 ms	4.75	4.89	4.81	-9.63	-9.56	-10.06	43.62	44.23	48.86	137		
10	20 ms	7.65	7.68	7.82	-10.56	-10.35	-10.77	87.10	83.20	93.32	264	**32.91**	**12.48**
	40 ms	8.35	7.94	8.19	-11.03	-10.62	-10.83	105.80	91.52	99.21	297		
30	20 ms	11.25	10.61	11.14	-10.08	-10.19	-10.27	114.66	110.83	118.64	344	**157.85**	**45.87**
	40 ms	12.07	12.30	12.29	-10.82	-11.64	-11.58	145.78	179.35	176.85	502		
100	20 ms	16.55	16.54	16.07	-10.25	-9.98	-10.15	175.42	164.68	166.19	506	**193.67**	**38.25**
	40 ms	14.34	14.29	14.62	-12.14	-11.74	-12.36	234.72	213.34	251.91	700		

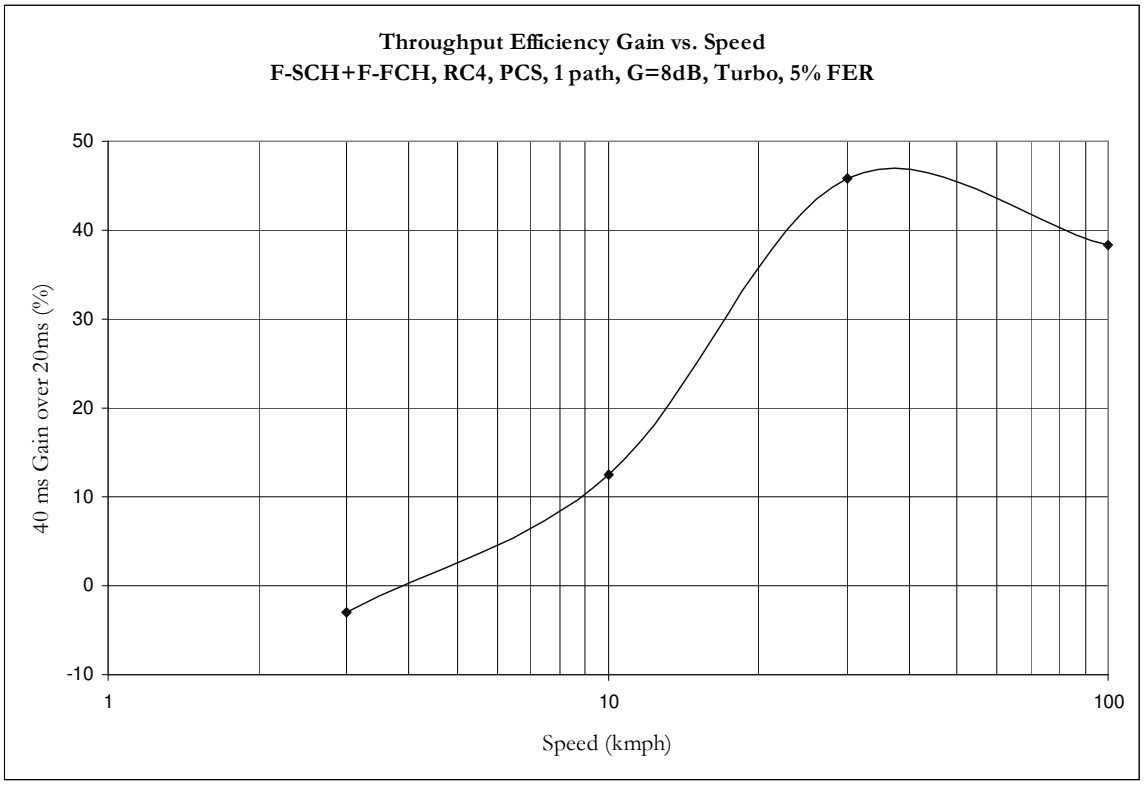

11.3 Convolutional Encoding

11.3.1 Sector Throughput

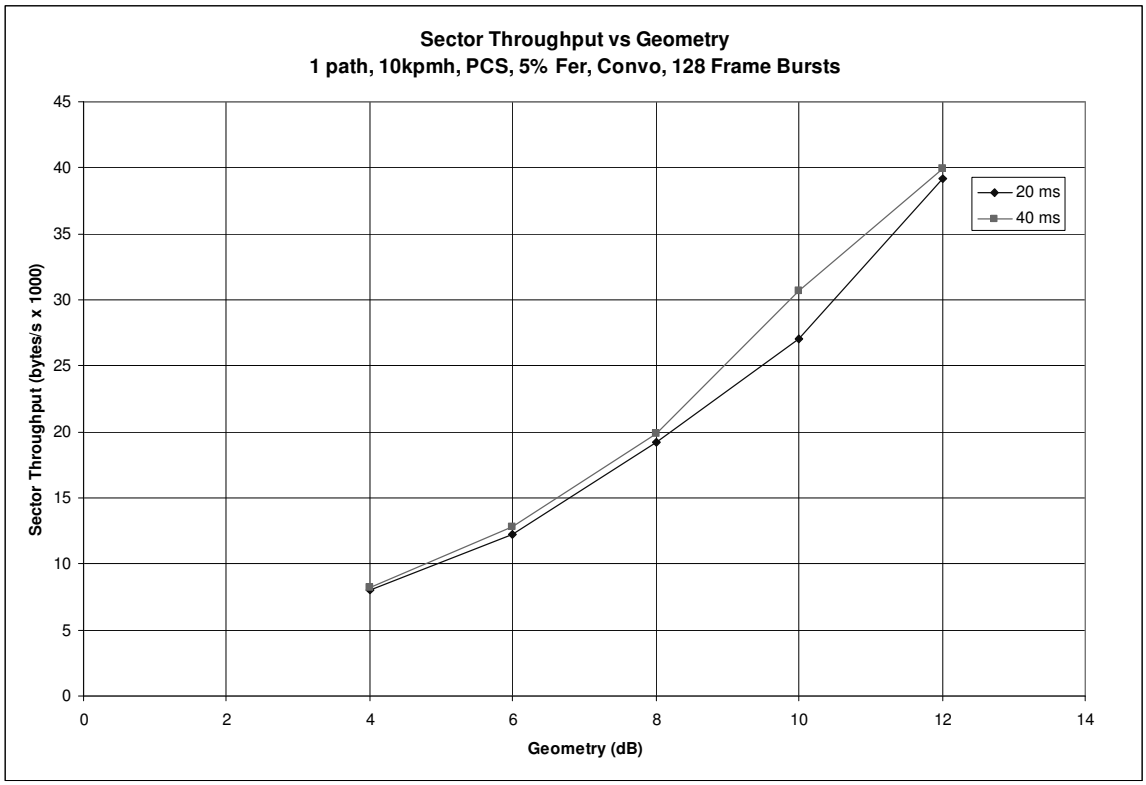

11.3.2 SCH Assignment History

		Average Assignments							Average Assignment
		0x	~1x	~2x	~4x	~8x	~16x	~32x	(x)
G=4dB	20 ms	0.35	0.20	0.14	0.26	0.06	0.00	0.00	**1.96**
	40 ms	0.34	0.20	0.16	0.24	0.07	0.00		**2.01**
G=6dB	20 ms	0.26	0.15	0.17	0.14	0.25	0.02	0.00	**3.35**
	40 ms	0.23	0.16	0.17	0.15	0.28	0.02		**3.63**
G=8dB	20 ms	0.19	0.13	0.16	0.15	0.12	0.24	0.00	**5.92**
	40 ms	0.21	0.11	0.16	0.16	0.15	0.21		**5.59**
G=10dB	20 ms	0.13	0.08	0.14	0.14	0.16	0.25	0.09	**9.16**
	40 ms	0.06	0.05	0.09	0.16	0.19	0.44		**9.52**
G=12dB	20 ms	0.08	0.07	0.11	0.15	0.18	0.16	0.26	**13.09**
	40 ms	0.02	0.02	0.03	0.11	0.19	0.63		**12.09**

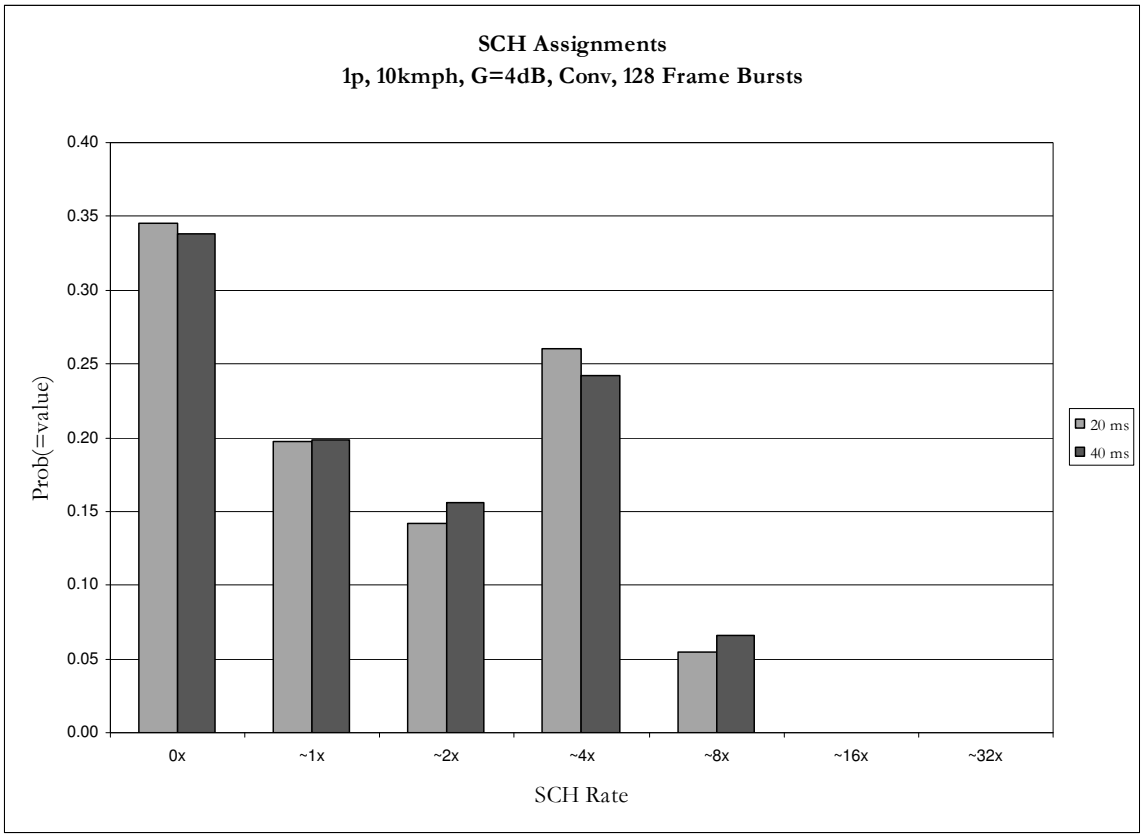
SCH Assignments
1p, 10kmph, G=4dB, Conv, 128 Frame Bursts
0.40
0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
Prob(=value)
0x
~1x
~2x
~4x
~8x
~16x
~32x
SCH Rate
20 ms
40 ms

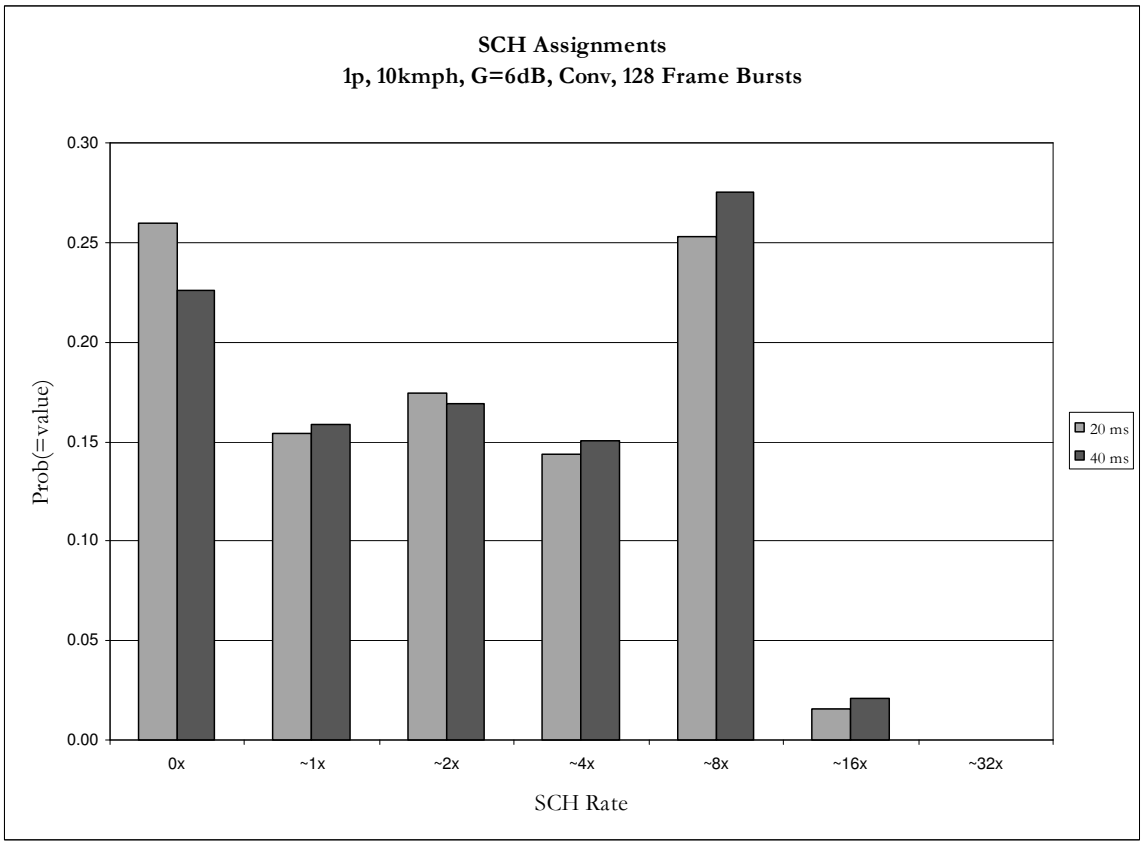
SCH Assignments
1p, 10kmph, G=6dB, Conv, 128 Frame Bursts
0.30
0.25
0.20
0.15
0.10
0.05
0.00
Prob(=value)
0x
~1x
~2x
~4x
~8x
~16x
~32x
SCH Rate
20 ms
40 ms

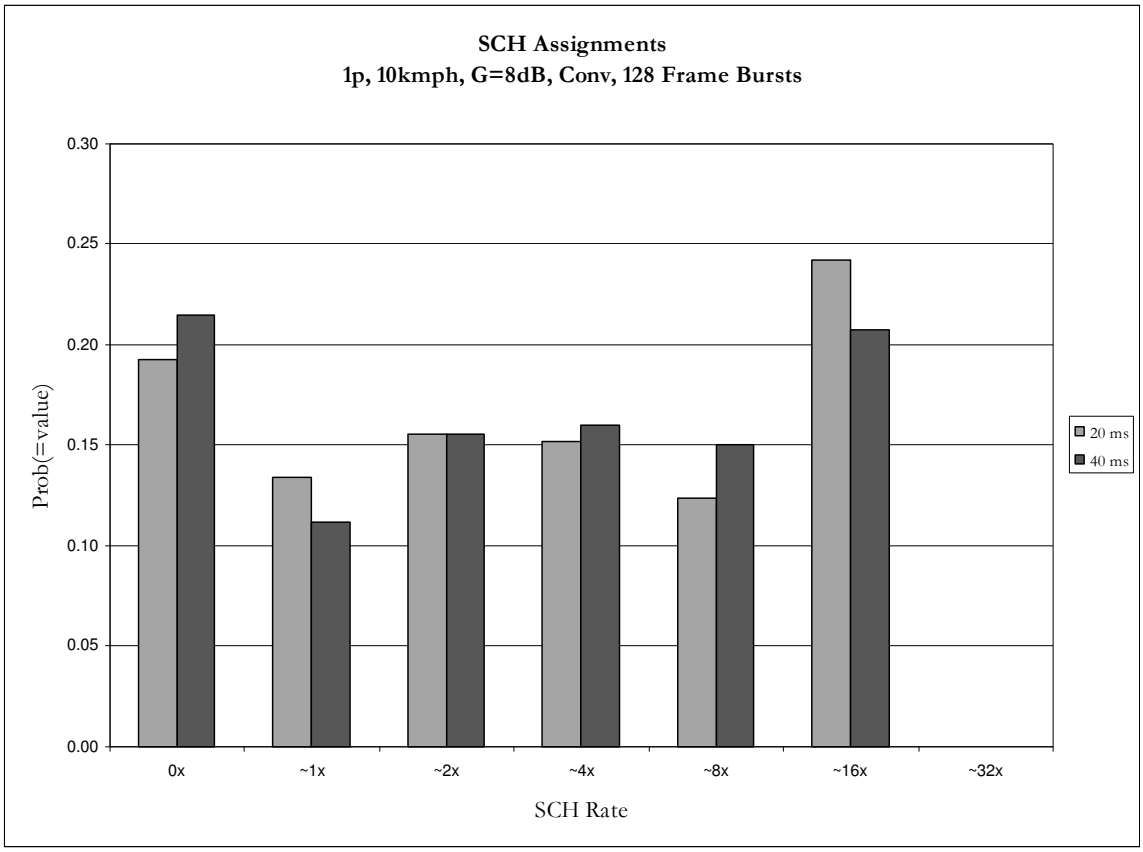
SCH Assignments
1p, 10kmph, G=8dB, Conv, 128 Frame Bursts
0.30
0.25
0.20
0.15
0.10
0.05
0.00
Prob(=value)
0x
~1x
~2x
~4x
~8x
~16x
~32x
SCH Rate
20 ms
40 ms

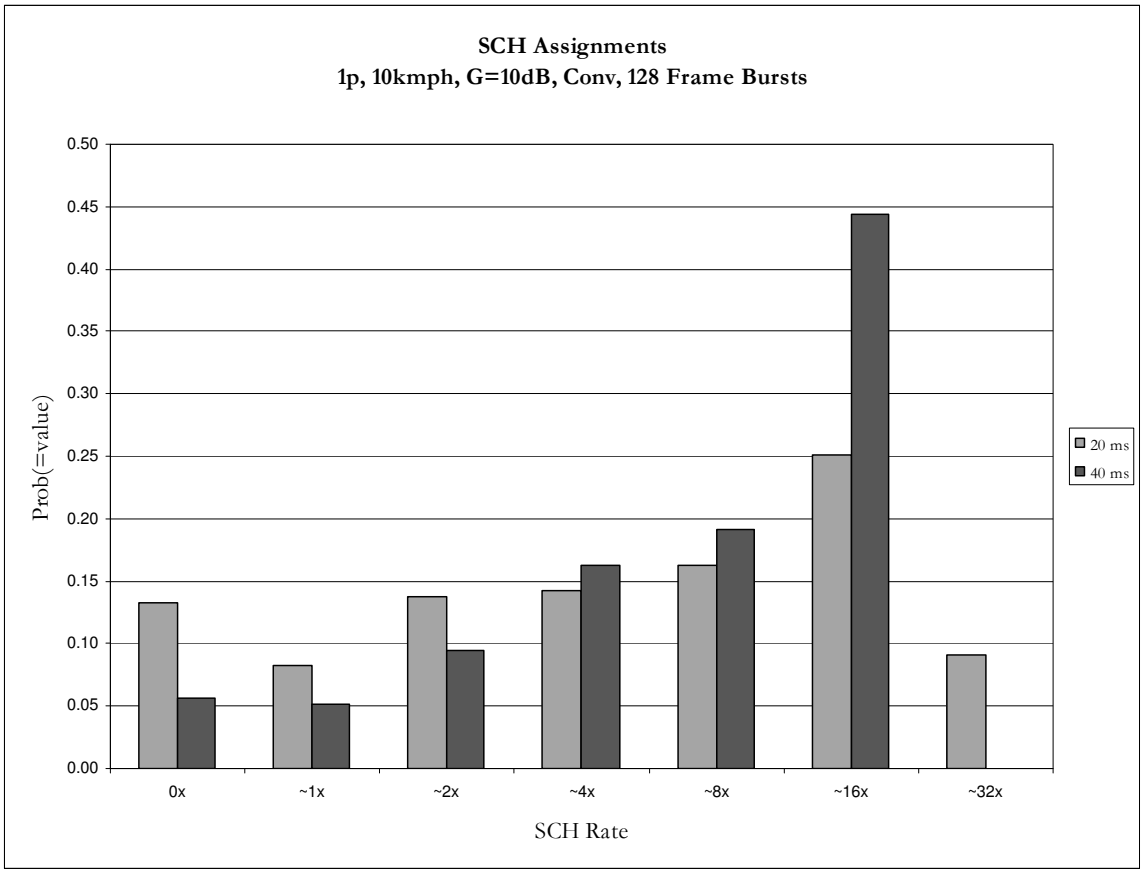
SCH Assignments
1p, 10kmph, G=10dB, Conv, 128 Frame Bursts
0.50
0.45
0.40
0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
Prob(=value)
0x
~1x
~2x
~4x
~8x
~16x
~32x
SCH Rate
20 ms
40 ms

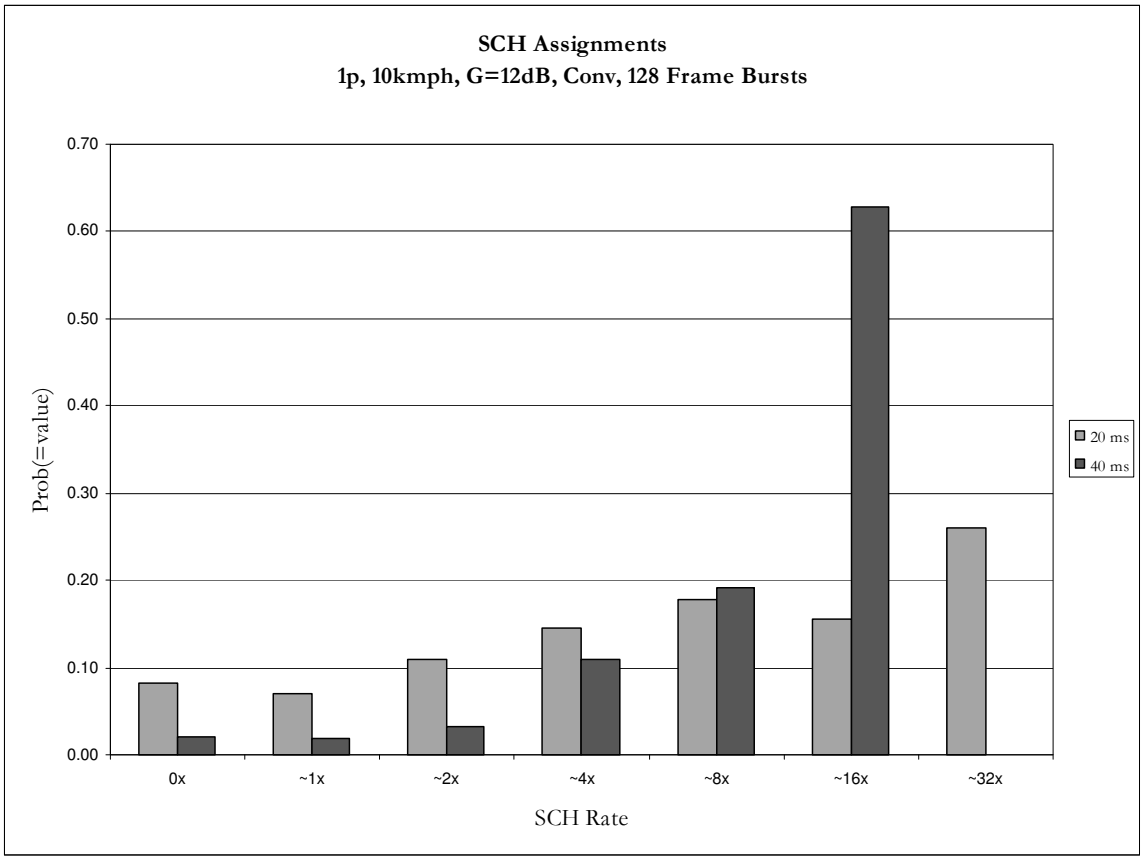

11.4 Burst Duration

	Burst Duration (frames x 20ms)	Avg Throughputs (bytes/s x 1000)			Avg SCH Ec/Ior (dB)			Throughput Efficiency			
		Call 0	Call 1	Call 2	Call 0	Call 1	Call 2	Call 0	Call 1	Call 2	Total
20 ms	32	6.50	6.03	6.48	-9.50	-9.48	-10.01	57.91	53.56	65.00	**176**
	64	7.65	7.68	7.82	-10.56	-10.35	-10.77	87.10	83.20	93.32	**264**
	128	7.96	7.79	7.89	-10.46	-10.45	-10.56	88.61	86.32	89.80	**265**
	256	7.74	7.15	8.15	-10.19	-10.91	-10.29	80.84	88.15	87.04	**256**
40 ms	32	7.30	7.34	7.32	-8.02	-7.78	-7.80	46.23	44.02	44.12	**134**
	64	8.35	7.94	8.19	-11.03	-10.62	-10.83	105.80	91.52	99.21	**297**
	128	7.49	7.61	7.61	-11.60	-11.83	-11.72	108.26	115.85	113.07	**337**
	256	7.84	8.42	8.01	-12.54	-11.26	-11.00	140.89	112.61	100.82	**354**

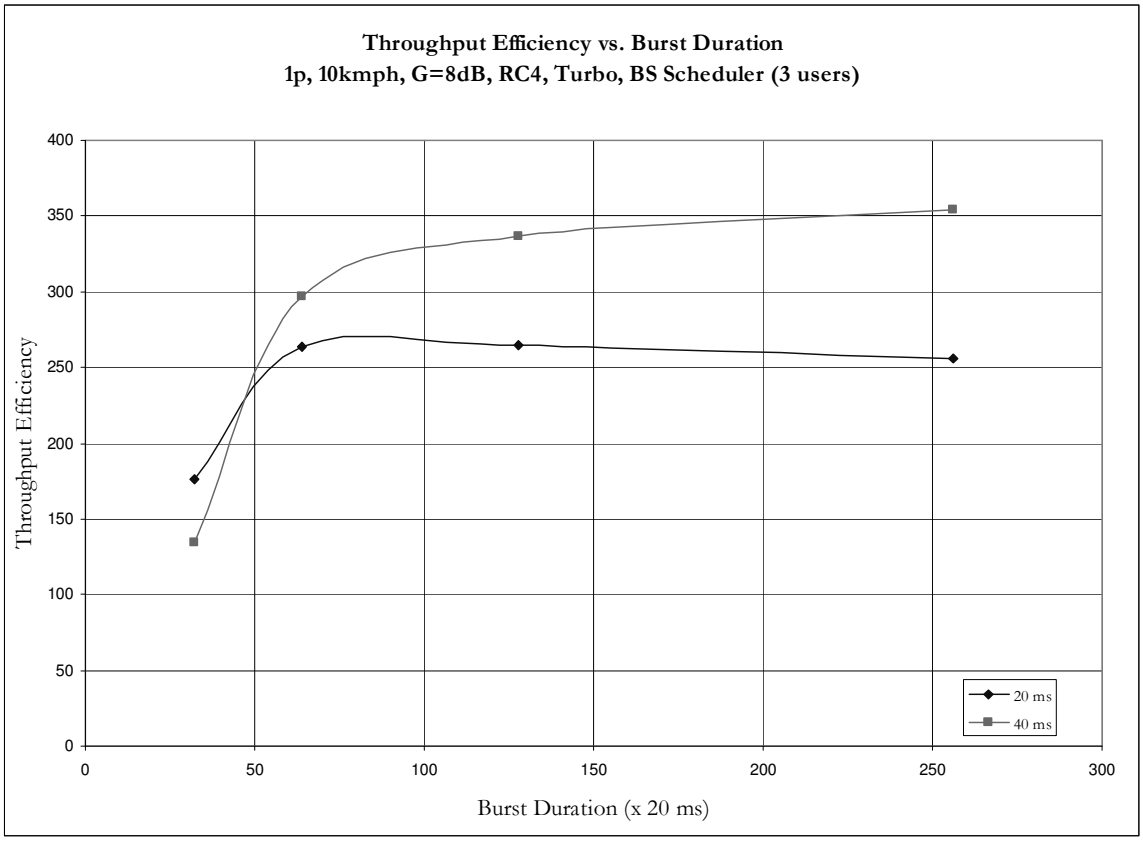
Throughput Efficiency vs. Burst Duration
1p, 10kmph, G=8dB, RC4, Turbo, BS Scheduler (3 users)
Throughput Efficiency
Burst Duration (x 20 ms)
0
50
100
150
200
250
300
350
400
20 ms
40 ms

12 Appendix F—Latency Results

The following contains the complete set of data for section 4.5.

12.1 Ping Packet Size and Data Rate

	20ms		40ms	
# SCH	Rate	Size	Rate	Size
1	4x (38.4/1,20)	32 bytes	2x (19.2/2,40)	32 bytes
2	4x (38.4/1,20)	99 bytes	4x (38.4/2,40)	99 bytes

12.2 Ping Data

	20ms, 32 bytes			40ms, 32bytes			20ms, 99bytes			40ms, 99bytes		
	min	max	avg	min	max	avg	min	max	avg	min	max	avg
	160	191	173	180	231	197	190	201	198	240	281	257
	170	191	176	181	230	204	190	211	199	240	281	257
	170	190	177	190	220	204	190	211	200	240	281	257
	170	191	178	190	221	205	190	201	199	240	281	259
	170	190	177	190	221	209	190	211	198	240	271	256
	170	181	174	190	221	209	190	211	199	240	291	259
	170	190	179	190	221	210	190	211	194	240	281	257
	170	190	177	190	221	202	190	211	197	240	281	259
	170	190	175	190	221	203	190	211	197	240	280	255
	170	191	175	190	211	200	190	211	194	240	281	260
	170	190	174	190	221	205	190	201	198	240	281	268
	170	191	175	190	221	205	190	201	198	240	281	263
	170	181	178	190	221	208	190	211	200	240	271	255
	170	190	176	190	221	209	190	211	197	240	281	262
	170	190	176	190	221	208	190	211	198	240	281	261
	170	190	176	190	221	204	190	201	196	240	280	263
	170	191	177	190	221	206	190	210	194	241	271	261
	170	181	178	190	221	207	190	211	200	241	271	261
	170	190	180	190	221	205	190	201	197	250	281	265
	170	190	173	191	221	212	190	211	195	250	280	266
avg	169.5	188.95	**176.2**	189.1	221.4	**205.6**	190	207.95	**197.4**	241.1	279.35	**260.05**
stdev	2.2	3.5	1.9	3.0	3.8	3.6	0.0	4.7	2.0	3.1	4.9	3.7

	Latency Relative to 20ms (ms)					
	20ms		40ms		80ms	
	Expected	Exper'mt	Expected	Exper'mt	Expected	Exper'mt
Data sent over 1 SCH frame (32 bytes)	0	-	30	29.4	90	-
Data sent over 2 SCH frames (99 bytes)	0	-	50	62.65	150	-

Notes: Each "min, max, avg" line above represents the average min, average max, and average average over 20 ping packets.