IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF PENNSYLVANIA

CARNEGIE MELLON UNIVERSITY, Plaintiff, v. MARVELL TECHNOLOGY GROUP, LTD. and MARVELL SEMICONDUCTOR, INC., Defendants.	Civil Action No. 2:09-cv-00290-NBF Hon. Nora B. Fischer

MARVELL'S MOTION FOR JUDGMENT
ON LACHES

Defendants Marvell Technology Group, Ltd. and Marvell Semiconductor, Inc. (collectively, "Marvell") hereby respectfully move the Court pursuant to Fed. R. Civ. P. 52(c) for judgment as a matter of law that pre-suit damages claimed by Plaintiff Carnegie Mellon University ("CMU") in this case are barred by laches. For the reasons stated in Marvell's Memorandum in Support of its Motion for Judgment on Laches, Marvell's Proposed Findings of Fact and Conclusions of Law on Laches, Affidavit of Sehat Sutardja, and Affidavit of Zining Wu, all filed simultaneously herewith, the trial record, all pleadings and papers on file in this action, such matters as are subject to judicial notice, and all other matters or arguments that may be presented with this motion, the Court should grant Marvell's motion and enter the proposed order attached hereto.

Dated: February 11, 2013

Respectfully submitted,

/s/ John E. Hall

John E. Hall
Timothy P. Ryan
ECKERT SEAMANS CHERIN & MELLOTT, LLC
U.S. Steel Tower
600 Grant Street, 44th Floor
Pittsburgh, PA 15219
Phone: (412) 566-6000
Fax: (412) 566-6099
jhall@eckertseamans.com
tryan@eckertseamans com

/s/ Edward J. DeFranco

Edward J. DeFranco (*pro hac vice*)
Kathleen M. Sullivan (*pro hac vice)*
Faith Gay (*pro hac vice*)
Raymond Nimrod (*pro hac vice*)
David Radulescu (*pro hac vice*)
Derek L. Shaffer (*pro hac vice*)
QUINN EMANUEL URQUHART & SULLIVAN, LLP
51 Madison Avenue, 22nd Floor
New York, New York 10010
Phone: (212) 849-7000
Fax: (212) 849-7100
eddefranco@quinnemanuel.com

Steven G. Madison (*pro hac vice*)
QUINN EMANUEL URQUHART & SULLIVAN, LLP
865 S. Figueroa St., 10th Floor
Los Angeles, California 90017
Phone: (213) 443-3000
Fax: (213) 443-3100
stevemadison@quinnemanuel.com

Kevin P.B. Johnson (*pro hac vice*)

Melissa Baily (*pro hac vice*)
QUINN EMANUEL URQUHART & SULLIVAN, LLP
555 Twin Dolphin Drive., 5th Floor
Redwood Shores, California 94065
Phone: (650) 801-5000
Fax: (650) 801-5100
kevinjohnson@quinnemanuel.com

Attorneys for Defendants, Marvell Technology Group, Ltd. and Marvell Semiconductor, Inc.

CERTIFICATE OF SERVICE

I hereby certify that on February 11, 2013, the foregoing was filed electronically on ECF. I also hereby certify that on February 11, 2013, this filing will also be served on counsel for CMU by electronic mail.

/s/ John E. Hall

John E. Hall
Timothy P. Ryan
ECKERT SEAMANS CHERIN & MELLOTT, LLC
U.S. Steel Tower
600 Grant Street, 44th Floor
Pittsburgh, PA 15219
Phone: (412) 566-6000
Fax: (412) 566-6099
jhall@eckertseamans.com
tryan@eckertseamans com

David C. Radulescu (*pro hac vice*)
QUINN EMANUEL URQUHART & SULLIVAN, LLP
51 Madison Avenue, 22nd Floor
New York, New York 10010
(212) 849-7000 Telephone
(212) 849-7100 Facsimile
davidradulescu@quinnemanuel.com

Attorneys for Defendants,
Marvell Technology Group, Ltd. and
Marvell Semiconductor, Inc

IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF PENNSYLVANIA

CARNEGIE MELLON UNIVERSITY, Plaintiff, v. MARVELL TECHNOLOGY GROUP, LTD. and MARVELL SEMICONDUCTOR, INC., Defendants.	Civil Action No. 2:09-cv-00290-NBF Hon. Nora B. Fischer

AFFIDAVIT OF SEHAT SUTARDJA IN SUPPORT OF MARVELL'S MOTION FOR JUDGMENT ON LACHES

I, Sehat Sutardja, declare as follows:

1. I am a Co-Founder of Marvell Technology Group, Ltd. and have served as the President and Chief Executive Officer of Marvell Technology Group, Ltd. since its inception in 1995. I submit this affidavit in support of Marvell Technology Group, Ltd. and Marvell Semiconductor, Inc.'s ("Marvell's") Motion for Judgment on Laches. Unless otherwise indicated, I have personal knowledge of the facts set forth in this affidavit and, if called as a witness, could and would competently testify to them.

Marvell's Investment in MNP.

2. Starting in 1998, Marvell invested in the development of its post processor design. In early 2001, Marvell further developed its post processor into its MNP design.

3. Marvell has continued to invest in the development and improvement of MNP, including its improved Enhanced Media Noise Processor, or EMNP. Marvell further improved EMNP with its Non-Linear Viterbi or Non-Linear Detector, or NLV/NLD.

4. Similarly, Marvell continued to invest in the development and improvement of other technologies included in its accused chips with MNP. For example, Marvell has

invested in developing run-length limited codes, iterative codes, and other data coding and error correction technology. This technology has been included in Marvell's accused chips.

5. To protect Marvell's numerous proprietary innovations, Marvell employs intellectual property attorneys to assist with the filing of patent applications. For example, Marvell has filed a number of patents related to its MNP technology.

6. As indicated by the following table, since 2001, Marvell's research and development expenses have substantially increased. This information is from Marvell's 10-K filings with the SEC, which I endorse as President and Chief Executive Officer of Marvell.

Year	**Marvell's Research and Development Expenses ($ thousands)**
2001	35,152
2002	93,422
2003	145,722
2004	213,740
2005	263,261
2006	311,498
2007	658,211
2008	988,996
2009	935,272

7. The below table summarizes Marvell's investments in research and development and selling, general and administrative expenses related to chips with MNPs, EMNPs, and NLV/NLDs, as calculated by CMU.

Fiscal Year	**Research & Development Expenditures ($ thousands)**	**Selling, General & Administrative Expenses ($ thousands)**
[redacted]		

8. As indicated by the following table, since 2001, Marvell's net revenue has substantially increased. This information is from Marvell's 10-K filings with the SEC, which I endorse as President and Chief Executive Officer of Marvell.

Year	Marvell's Net Revenue ($ thousands)
2001	143,894
2002	288,795
2003	505,285
2004	819,762
2005	1,224,580
2006	1,670,266
2007	2,237,553
2008	2,894,693
2009	2,950,563

9. The below table summarizes Marvell's unit sales and revenues from chips with MNPs, EMNPs, and NLV/NLDs, as asserted by CMU at trial.

Year	Unit Sales	Revenues
[redacted]		

10. Marvell's increasing sales since 2001 has led to a corresponding increase in the number of employees. As of March 31, 2001, Marvell had a total of 753 employees according to Marvell's 10-K form for the fiscal year ending on January 27, 2001. Marvell currently has more than 7,200 employees.

<u>Marvell's Anticipated Response Had CMU Provided Notification or Filed Suit in 2001</u>.

11. To my knowledge, prior to initiating this lawsuit, CMU never informed Marvell that it believed Marvell was infringing the asserted patents or that CMU intended to enforce its asserted patents against Marvell.

12. To my knowledge, prior to initiating this lawsuit, CMU did not request any additional information from Marvell about its activities, its simulations, its chips, or anything else relevant to CMU's instant claims of infringement.

13. If CMU had approached Marvell prior to initiating this lawsuit and requested information relevant to CMU's instant claims of infringement, Marvell would have provided sufficient information to allow CMU to evaluate its claims under a Non-Disclosure Agreement (NDA).

14. While engaging in licensing discussions with other companies, Marvell has presented information regarding the operation of Marvell's products. I have reviewed several sets of slide presentations that Marvell shared in 2009 with its competitor Freescale, labeled "Subject to Fed. R. Evid. 408" and "*CONTAINS CONFIDENTIAL MARVELL INFORMATION,*" because the slides include confidential information regarding the operation of Marvell's products and the particular accused circuitry. Of course, Marvell would have been more inclined to share confidential information with a university, as opposed to a competitor. Because of this, I would have approved the sharing of confidential information regarding our circuitry with CMU under an appropriate NDA.

15. Had CMU sued Marvell in 2001 - 2007 or notified Marvell of its intent to enforce its patents against Marvell, I would not have approved investing in EMNP and NLD in the manner that we did. Indeed, had CMU obtained a judgment against Marvell in 2003 – 2007, I would not have approved the continued investment in the development of enhancements to MNP, including NLD.

16. Further, had CMU sued Marvell in 2001 - 2007 or notified Marvell of its intent to enforce its patents against Marvell, I would have approved decisions to either: (1) develop chips with iterative coding that did not include NLD; (2) use approaches proposed by Dr.

Cioffi; or (3) re-prioritize our development schedule and accelerate development of 4 K sector, its high rate RLL, or zone servo technology.

17. Moreover, if CMU had approached Marvell in May 2001 and demanded 50 cents for every chip that Marvell sold with the accused technology, I would not have approved investment in MNP in the manner that we did.

18. Had CMU sued any time between 2001 and 2004, I would have approved the decision either to: (1) not add MNP, or at a minimum, to phase out MNP by the next product generation; (2) rely on alternative technologies such as 10-tap adaptive FIR, programmable target, sync mark improvement, high rate RLL code, 3-interleaved ECC, servo Gray code, 10-bit ECC, and permuted RLL and ECC to improve SNR; and/or (3) use another existing alternative such as Dr. Cioffi's approach.

19. If Marvell had to pay a royalty of 50 cents for every chip with the accused technology, I would have approved the decision to either not include MNP in 2001 or to phase it out in the next product generation in 2003-2004.

I declare under penalty of perjury that the foregoing is true and correct to the best of my knowledge. Executed this 11 th day of February, 2013, in Santa Clara, California.

Sehat Sutardja

IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF PENNSYLVANIA

CARNEGIE MELLON UNIVERSITY, Plaintiff, v. MARVELL TECHNOLOGY GROUP, LTD. and MARVELL SEMICONDUCTOR, INC., Defendants.	Civil Action No. 2:09-cv-00290-NBF Hon. Nora B. Fischer

AFFIDAVIT OF ZINING WU IN SUPPORT OF MARVELL'S
MOTION FOR JUDGMENT ON LACHES

I, Zining Wu, declare as follows:

1. I am the Vice President of Data Storage Technology at Marvell Semiconductor, Inc. I submit this affidavit in support of Marvell Technology Group, Ltd. and Marvell Semiconductor, Inc.'s ("Marvell's") Motion for Judgment on Laches. I have personal knowledge of the facts set forth in this affidavit and, if called as a witness, could and would competently testify to them.

2. I received a bachelor's degree in electrical engineering from Tsinghua University in Beijing, China in 1994. After receiving my bachelor's degree, I decided to come to the United States for my graduate studies because of the wide variety of quality programs available in electrical engineering. I received my Ph.D. in Electrical Engineering from Stanford University in 1999. I specialized in signal processing, digital communication technology, coding and detection for read channels. My Ph.D. thesis was on coding and iterative detection for magnetic recording channels.

<u>Background at Marvell</u>.

3. After being awarded my Ph.D., I started working at Marvell in 1999 as a Staff Design Engineer. As a Staff Design Engineer, I worked on signal processing and coding for magnetic disk drive read channels.

4. In late 2002, I was promoted to Manager at Marvell. As a manager, I continued engineering work on signal processing and coding for magnetic disk drive read channels and also managed other engineers in the read channel architecture group.

5. I continued to receive promotions approximately every two years at Marvell. I was promoted to Senior Manager, then Director, and then Senior Director. In 2008, I was promoted to my current role as Vice President of Data Storage Technology at Marvell. Throughout my tenure at Marvell, the majority of my work has involved signal processing and coding.

6. In my current role, I supervise the development of advanced technology for the read channel of hard disk drives. There are currently five teams, with a total of over 200 individuals that I supervise. Specifically, there are the following teams: architecture, digital design, application engineering, documentation, and optical data storage. The architecture team creates the high level design of the chip, similar to an architect who designs a blueprint plan for building a house. After the high level design of the chip is created by the architecture team, the digital design team builds the actual chip. The actual implementation of the circuit in a chip is a complex and expensive process requiring a large team. After a chip is laid out and a prototype is created, the application engineering team needs to confirm that the chip is operating properly. Since a chip is very complex, this is also a significant effort requiring a large team. The documentation team prepares specifications to describe to a customer how to

use a chip. The optical data storage team works on data storage technologies using optical, rather than magnetic, media.

7. Currently, more than 800 individuals at Marvell (excluding sales and marketing personnel) are involved in the development of systems on a chip, or SOCs, that include the accused technology.

8. When I began working at Marvell in 1999, there were approximately 130-150 employees at the company. There are currently over 7,200 employees at Marvell.

Marvell's Development and Commercialization of MNP and Related Technologies.

9. Throughout my tenure at Marvell, Marvell has consistently been working to develop and improve the performance of Marvell's chips. These improvements can be achieved through a large number of areas, including improvements to correction coding and data coding.

10. Starting in 1998, Marvell began to develop its post processor, initially developing a linear post processor. Marvell was awarded a patent on its linear post processor, U.S. Patent No. 6,427,220 to Vityaev. Between 1999 to 2001, I was also involved in the development of higher rate run-length limited (RLL) codes and 10-bit error correction codes.

11. By early 2001, I, along with Mr. Greg Burd and Dr. Toai Doan, were researching various options for addressing media noise. We reviewed literature and published papers by individuals in the field, such as by Dr. Cioffi, Dr. Zeng, Dr. Moon, and Dr. Kavcic. By mid-March 2001, Mr. Burd had simulated the method proposed in Dr. Kavcic's paper. Mr. Burd wrote to Dr. Doan on March 23, 2001 to inform him of "disappointing" results of the simulation. We concluded that Dr. Kavcic's approach could not be implemented because it was too complicated.

12. Marvell developed a media noise post processor, or MNP. I was involved in the design of the MNP in 2001, along with Mr. Burd and Dr. Doan. On January 3, 2002, we submitted a version of a report on our MNP design to the patent office in support of our provisional patent application directed to our MNP.

13. Marvell disclosed the asserted Kavcic patents to the patent office during this application process. In particular, in the application, Marvell described the differences between Kavcic's detector and Marvell's approach. The provisional application stated, "[e]ven though Kavcic's detector provides significant gains over conventional Viterbi detector in the presence of media noise, it is not very appealing due to implementation complexity."

14. In response to Marvell's patent application for its MNP design, the patent office issued U.S. Patent No. 6,931,585 ("the '585 patent"). I am identified as an inventor of the '585 patent, along with Mr. Burd. The asserted Kavcic patents are identified on the face of Marvell's '585 patent as "References Cited."

15. Marvell's MNP was included in Marvell's chips, with Marvell shipping its first read channel chips with MNP to its customers around late 2002.

16. Marvell continued to develop and improve its MNP after it shipped its first MNP read channel chips. Marvell researched and developed improvements to MNP, including its Enhanced Media Noise Processor, or EMNP. Engineers including Greg Burd and Hongxin Song were part of the team that worked on the improvement. I supervised the development of the improved EMNP.

17. Marvell further improved the MNP with its Non-Linear Viterbi or Non-Linear Detector, or NLV/NLD. Engineers including Greg Burd, Hongxin Song, Mike Madden, Ke

Han, and Mats Oberg were part of the team that worked on the improvement. I supervised the development of the improved NLV/NLD.

Marvell's Response Had CMU Filed Suit or Notified Marvell in 2001.

18. To my knowledge, prior to initiating this lawsuit, CMU never informed Marvell that it believed Marvell was infringing the asserted patents or that CMU intended to enforce its asserted patents against Marvell.

19. To my knowledge, prior to initiating this lawsuit, CMU did not request any additional information from Marvell about its activities, its simulations, its chips, or anything else relevant to CMU's instant claims of infringement.

20. Had CMU sued Marvell in 2001 - 2007 or notified Marvell of its intent to enforce its patents against Marvell, Marvell would not have invested in EMNP and NLD in the manner that it did. Indeed, had CMU obtained a judgment against Marvell in 2003 – 2007, Marvell would not have continued to invest in the development of its enhancements to MNP, namely its NLD (which has replaced MNP in the vast majority of its chips).

21. It is important to understand that one of the primary reasons Marvell developed its NLD enhancement was because MNP did not work well with iterative coding, which promised significant SNR gains. Marvell ultimately gained more than 3 dB from iterative coding, which it first included in chips shipped to customers in 2008 (chips with NLD were first shipped in 2007). Marvell expended years of research and development effort from 1999-2008 on its iterative coding technology. The chips with iterative coding also incorporated the accused NLD technology. However, neither NLD nor MNP were needed to obtain the gain from iterative coding. Had CMU sued Marvell in 2001 - 2007 or notified Marvell of its intent to enforce its patents against Marvell, Marvell would have developed chips with iterative coding that did not include NLD. With the large SNR gain achieved from

iterative coding, along with other previous SNR gains achieved from other technologies implemented in our chips, Marvell certainly did not need the minimal gain from NLD in its chips with iterative coding.

22. Marvell also did not need the SNR gain from NLD when it first shipped chips with NLD technology in 2007. By then, Marvell had already developed numerous other technologies from 2001 to 2005 that achieved several times more SNR gain than NLD, such as programmable target, sync mark improvement, high rate RLL code, 3-interleaved ECC, 10-bit ECC, and permuted RLL and ECC.

23. Consequently, had CMU timely filed suit, there would have been no reason for Marvell to include NLD in any of its chips. Rather, Marvell would have either (1) not included NLD; (2) used the approaches proposed by Dr. Cioffi and achieved substantially the same SNR gain as NLD; (3) or re-prioritized its development schedule and accelerated its development of 4K sector, its high rate RLL, or zone servo technology to get even more SNR gain than NLD.

24. Moreover, if CMU had approached Marvell in May 2001 and demanded 50 cents for every chip that Marvell sold with the accused technology, Marvell would not have invested in MNP in the manner that it did. For example, according to Marvell's customers, Marvell did not need to add MNP to its chips in 2001, or at any time thereafter. In fact, Marvell's largest customer, Western Digital, asked Marvell to remove the MNP. Moreover, around the time of the development of MNP, Marvell also developed numerous other technologies that achieved SNR gain, such as programmable target, sync mark improvement, high rate RLL code, 3-interleaved ECC, 10-bit ECC, disk synchronous write, and permuted RLL and ECC. In fact, Marvell was ahead of its competition with its C5575 chip, which did

not have MNP. And by the next production generation, Marvell was already seeing significant gains from 10-bit ECC, and knew that MNP was without question not needed.

25. Therefore, had CMU sued any time between 2001 and 2004, Marvell would have either: (1) not added MNP to its chips, or at a minimum, phased out MNP by the next product generation; (2) relied on alternative technologies such as 10-tap adaptive FIR, programmable target, sync mark improvement, high rate RLL code, 3-interleaved ECC, servo Gray code, 10-bit ECC, disk synchronous write and permuted RLL and ECC to improve SNR; and/or (3) used another existing alternative such as Dr. Cioffi's approach.

26. Once MNP was added to Marvell's chips, it became more difficult to remove from the current product generation in 2002 without re-allocating resources to modify the rest of the design. Therefore, it made little sense to remove MNP. Nevertheless, if Marvell had to pay a royalty of 50 cents for every chip with the accused technology, as I explain above, Marvell would have either not included MNP in 2001 or simply phased it out in the next product generation in 2003-2004. By the time of the next product generation, Western Digital had requested that Marvell remove MNP from the chip and other customers had reported lower than expected SNR gains from MNP. With knowledge of the higher gains from 10-bit ECC and other technologies implemented in that timeframe, Marvell would agreed to Western Digital's request and removed the MNP for the next generation chip in 2003-2004.

I declare under penalty of perjury that the foregoing is true and correct to the best of my knowledge. Executed this 11th day of February, 2013, in Santa Clara, California.

Dr. Zining Wu

PRODBEG	=	KAVCIC002256
PRODEND	=	KAVCIC002288
AUTHOR	:	Aleksandar Kavcic
FILENAME	:	memo.doc
EXTENSION	:	doc
FOLDERID	:	X:\CARNEGIE_MELLON_UNIV_ZZ_01\001\Kavcic, Alekdandar\Kavcic_Docs_072209\Laptop_02\Documents and Settings\kavcic\My Documents\Alek\reviews\tenure_review\
DATE_LASTMOD	=	09/21/2004
CUSTODIAN	:	Kavcic, Alekdandar
DOC_TYPE	:	doc
PAGECOUNT	:	33
TITLE	:	To: My review committee
CDVOL	:	CMU_030110_Kavcic
POD	:	Confidential
TEXTPATH	:	\KAVCIC000672\KAVCIC002256.txt
DOC_TEXT	:	ï»¿

To: My review committee

From: Aleksandar Kavcic

The purpose of this memorandum is to give a detailed report on the research, teaching and program building activities I have been undertaking since I joined Harvard. Our program in Electrical Engineering is indeed growing, but nevertheless it is very small in comparison to

To: My review committee

From: Aleksandar Kavcic

The purpose of this memorandum is to give a detailed report on the research, teaching and program building activities I have been undertaking since I joined Harvard. Our program in Electrical Engineering is indeed growing, but nevertheless it is very small in comparison to other institutions countrywide. For this reason, besides teaching and research, my job also comprises of efforts to achieve visibility as a researcher, as a research group, as an Electrical Engineering program and even as an Electrical Engineering faculty member within Harvard. What follows is a description of my visions, activities and results. I will point out where I believe I was successful and where I think I have failed (or rather where I believe the circumstances have forced me to delay and modify the execution of my plans).

Table of Contents

Building a Communications Program at Harvard University

After Computer Engineering, Communications is perhaps the largest sub-discipline of Electrical Engineering. When I started my job here, Harvard did not have a program in Communications at all. It seemed that my job had been very clearly laid out for me – build a Communications program at Harvard. Obviously, this task became easier as Harvard started adding more faculty members in this area (most notably professor Tarokh, whose hiring gave a large boost to the visibility of the program). Given the size of the area of Communications Engineering, where numerous academic departments around the country have teams of faculty members covering as many topics in this area as possible, Harvard faculty individually have to cover many sub-fields of the field of communications in order to compensate for the small size of our faculty. Obviously, my main field of expertise has been (and still is) Magnetic Recording. This has been the field that the scientific community mostly recognizes me for, and that accounts for most of my research funding. However, I have started branching out (either individually or through collaborative efforts) from this main research area into the wider field of communications, most notably in the direction of information theory, wireless communications, error correction codes and VLSI implementation.

My research group has grown over the years. It reached a steady state of 5 Ph.D. students, and 2-3 visitors (either post-docs, visiting students or faculty). To date, I have supervised 3 post-docs, who presently work as faculty and industry researchers, I have graduated 1 Ph.D. student (presently at Hitachi Research Labs) and 5 M.S. and M.E. students (who did research projects with me, and who now work either in academia or in the communications industry).

The following is a summary of what I have achieved thus far and where I have experienced setbacks:

Accomplishments:

1. The Magnetic Recording effort is fully operational at this point with 4 Ph.D. students and one research visitor working on magnetic recording related problems.
2. The students are (generally) a joy to work with and are producing what I (admittedly subjectively) believe are the best results in magnetic recording today. Their papers are frequently cited and they are invited to give talks at various institutions countrywide.
3. Harvard's program in Magnetic Recording has achieved international visibility. Evidence of this is that I am invited for talks within the United States and abroad, I have been repeatedly asked to organize conference sessions in this area and have served a 3-year term as an Associate Editor for the IEEE Transactions on Information Theory, I have been asked to consult for magnetic recording companies (Quantum, Agere, Link-a-Media) and even to consult for a research institute (the Data Storage Institute) in Singapore, a country that accounts for 40% of the world's disk drive production.

4. Despite nationwide cutbacks in the levels of funding of communications research, the program's funding is healthy, with funds from the NSF and from various industrial sources (IBM, Quantum/Maxtor, Agere, Seagate and the International Storage Industry Consortium). An indicator of the quality of the program is the fact that we still receive funding from the International Storage Industry Consortium, where several established researchers in the field have recently lost funding because of the inability of the industry to sponsor research as a result of a weaker economy.
5. Through the Magnetic Recording effort, I have initiated collaboration with Professor Michael Mitzenmacher (who is a Computer Scientist) and have started collaborating with Professor Gu Yeon Wei in order to implement my algorithmic solutions in VLSI technology.
6. In the area of wireless communications, I have an ongoing collaboration with Professor Tarokh. The task of the effort is to measure and model the indoor and outdoor wireless channel.
7. I have started several collaborative efforts outside of Harvard with Professor Marc Fossorier (University of Hawaii) in the area of error correction codes, with Professor Predrag Rapajic (University of New South Wales, Sydney, Australia) in the area of information theory and coding for wireless fading channels, with Dr. Emina Soljanin (Bell Labs) in the area of codes for wireless channels, with Professor Toby Berger (Cornell University) in the area of feedback capacity and codes, with Dr. Ravi Motwani (National University of Singapore) in the area of channels with synchronization errors, with Professor Hans-Andrea Loeliger (ETH Zurich) in the area of information theory for channels with memory, and with Professor Li Ping (City University of Hong Kong) in the area of multiuser coding and detection.

Shortcomings:

1. The pace of the group's growth was slow. The growth took one year longer than I had expected. In particular, it took 2 years to attract graduate students to the program and has taken the same length of time to attract a post-doctoral fellow. To date, I have graduated only one Ph.D. student. I attribute this slow start to the fact that our program in communications was non-existent prior to my arrival at Harvard, so it took several years for high-quality students to learn about the program and to start applying
2. While the funding situation in my group is healthy, most of the grants I received so far are small single-investigator grants that typically account for single-investigator research projects. Over the years, our program at Harvard has not been able to secure large team grants in the area of communications. In my view, this remains the biggest challenge for me and for our young communications program in the future. I believe that large grants will be instrumental in achieving visibility of Harvard's communications research *program* (and not just visibility as individual researchers).

Teaching activities

I have taught three Engineering Sciences courses at Harvard. **ES-50 – Introduction to Electrical Engineering** is a freshmen-level course whose content was radically changed two years ago in a collaborative effort between Professor Brockett and myself, and whose goal is to introduce the *entire* area of Electrical Engineering to a wider audience of freshmen through labs that cover all major modern sub-fields of Electrical Engineering. **ES-151 – Electromagnetic Communications** is a junior-level course whose goal is to expose the students to the basic principles of electromagnetic propagation. **ES-158/258 - (Advanced) Digital Communications** is a combined undergraduate/graduate course that teaches the basics of engineering the physical communications layer.

My view is that, like so many other things at Harvard, the topics that I teach are largely determined by the size of our faculty. Ideally, I would have liked to teach an undergraduate course in communications (which is something we need for the sake of the well-roundedness of our undergraduates) and a graduate-level course that would help the students build the foundations for conducting research in the areas of Signal Processing and Communications. Instead, the lack of faculty in Electrical Engineering (especially in the sub-area of Electromagnetics and Devices) has forced me to teach ES-151 for a few years (until we have hired Professor Donhee Ham), which only peripherally touched my research interests. I believe that this was not a course I was most qualified to teach and I believe the students sensed this (because my CUE Guide course rating is 3.8 for ES-151, while it is in the range 4.2-4.5 for the other courses I teach). Given that I cannot teach communications courses all year long, I have decided that the best way to cover the area of Communications is to teach a combined undergraduate/graduate course – ES-158/258. Recently, I have been teaching ES-50, which I view as one of the most important electrical engineering courses we teach at Harvard. It gives me joy and pride to see that the enrollment in the course has been steadily growing since Professor Brockett and I revamped the course two years ago. Our goals were to introduce the field of electrical engineering to freshmen in a manner that would be exciting for them, and with a goal to attract more freshmen to the field of electrical engineering. I am extremely proud of this effort and I believe that we are well on the way to achieving our goals.

Beyond the classroom, I have been involved in teaching through writings of book chapters and tutorial papers. Their goals were to familiarize a wider audience of engineers with my most recognized research papers.

I next give a slightly more detailed picture of my teaching efforts in the three courses and the book chapters and tutorials I wrote.

ES-50 Introduction to Electrical Engineering is a freshmen-level course in electrical engineering. Traditionally, introductory courses in electrical engineering have concentrated on introducing the fundamentals of electromagnetic field theory and circuit theory. When professor Brockett and I took over the course 2 years ago, our view was that the times have changed, that electrical engineering today spans a much wider horizon than a few decades ago, and that we need to introduce our freshmen to modern sub-fields (e.g., communications, robotics, signal processing, computer architecture, computer vision, etc.) very early on in order to show them the excitement that the field of electrical engineering entails. The challenge, however, was that we needed to do this without invoking the mathematical tools that are unavailable to freshmen. We decided to structure the course around weekly labs. Each week, we pick a different lab topic (e.g., operational amplifiers, analog-to-digital converters, data compression, data storage, data transmission, robotics, etc.) and teach the basic theory needed to understand and complete the lab. Thereby the emphasis is clearly on generating excitement during the lab. I view the course as a clear success. The CUE guide course grade went from the range of 2.5-3.5 two years ago to an average grade of 4.2 today. Admittedly, I would like the CUE guide grade for the course to be even higher, but objectively, this is a hard course to teach. Harvard freshmen come with diverse knowledge backgrounds, but they all are very demanding – they expect to learn, to be entertained and to acquire skills that will last them the rest of their lives. Admittedly, this is hard to achieve in a single course. However, I hope to get closer to this goal in the coming years as Professor Todd Zickler joins me in co-teaching the course and as we upgrade the labs to reflect some of his expertise in computer vision.

ES-151 Electromagnetic Communications is a classical course on Electromagnetic Theory. I have inherited the course from Professor Vic Jones, who has been teaching the course for several years. I found that the ABET requirements lowered my flexibility to teach the course. This course was envisioned as a fundamental classical course in Electrical Engineering (which until very recently has been a required course for all of our EECS concentrators). The students are exposed to the basic properties of magnetic materials, Maxwell's equations, electromagnetic waves, waveguides, transmission lines and antennas. I consider myself to have an above-average knowledge of this topic mainly because of my Ph.D. work in the area of Magnetic Recording. However, I do not consider myself to be an expert in this field. For example, I did not have the ability to draw on my research experiences and give illustrative examples to the students in my class. Without this valuable link to the research, I believe my lectures tended to be on the dry side, where I was only able to teach the theory without really being able to get the students interested in the practical aspects of engineering electromagnetic devices. My CUE Guide overall course evaluation grade is 3.8, which I do not believe reflects well on my ability to teach courses in the area of my research interest. For the past two years, the course has been taught by Professor Donhee Ham, who is undoubtedly doing a much better job given that the course is closely linked to his research interests.

ES-158/258 Digital Communications / Advanced Digital Communications is a combined undergraduate/graduate course. This course was not offered at Harvard before I joined. The material in the course is best suited for seniors and first-year graduate students. The prerequisites are ES-156 (Signals and Systems) and a course in probability theory, statistics or stochastic processes. The course introduces the students to the fundamental limits of communications, Shannon's theorems, orthogonal representations of signals, sampling theorem, digital modulation techniques, optimal receiver design, the Viterbi algorithm, and the basics on error correction codes. The graduate students are required to complete an engineering design project. The project topic changes yearly and is designed to reflect the state of the art in digital communications. For example, this year's project was to design an error correction code for a channel that suffers from deletion errors, which is a crude model for the Internet where packets are lost at random. Over the past 5 years, I have become very comfortable with this course. This is reflected by the CUE Guide reviews, where my grades average 4.5.

Book chapters and tutorials: Beyond the classroom, I have been involved in writing book chapters and tutorials. Their main goal was to make my better-recognized research papers known to a wider audience of engineers. The three book chapters [BT1, BT2, BT3] are about to appear in *Handbook of Coding and Signal Processing for Recording Systems*, CRC Press, 2004. These book chapters summarize my research on codes for partial response channels [J1,J9,C16,C20], capacity of finite-state machine channels [U2,JS3, JS5, C10, C11, C18, C19, C22, C26], and detectors for channels with pattern-dependent noise (or media noise in magnetic recording) [J12, J15, J16, J17, C33, C34], respectively. I also published a tutorial paper [BT4] in *IEEE Signal Processing Magazine, 2004*. The tutorial summarized my research on iterative timing recovery [J7, J11, C12].

Teaching summary: Generally, I think that my effectiveness as a teacher could improve. In particular, I don't think we are efficient in teaching Communications at Harvard. I am not sure how to improve this situation as long as our faculty roster in Electrical Engineering is so short. We have added Professor Vahid Tarokh as a communications faculty member, but frankly this is still not enough to cover all the courses in communications. Inevitably, we will still need to send our graduate students to MIT for 1-2 additional courses in the area of communications.

The following is a summary of what I have achieved thus far with the courses I teach:

Accomplishments:

1. I have successfully introduced a new course ES-158/258 – Digital Communications / Advanced Digital Communications. The course is now in a steady state. I am teaching it now for the 6th time, and I have reached a comfort level with the course much to the benefit of the students
2. With help from professor Brockett, I have revamped the introductory Electrical Engineering course ES-50. As far as I can tell, the course is a success. Further, it ties well into our ABET requirements, and it has been received well by the ABET committee. The course enrollment has been growing steadily since the course was revamped, which is yet another piece of evidence of its popularity among the students.
3. Every year I incorporate a lab project into ES-158/258. Labs are also integral components of ES-50. This gives the students an opportunity to learn state-of-the-art techniques in hands-on fashion.
4. I also try to set my graduate students up with internships. So far, all but one of my Ph.D. students has completed an internship in the industry. Though this mode of teaching is not the classical example of academic teaching, I believe it is

nonetheless valuable because it helps the students make a smooth transition into the workforce. Evidence of this is that my only Ph.D. graduate was aggressively recruited by the magnetic recording industry, and he eventually went to work for the company where he did his internship.

5. Through book chapters and tutorials, I have tried to widely disperse the knowledge I gathered by conducting research in magnetic recording. I do not have a method to measure how successful I have been in this endeavor, but occasionally I do receive encouraging words from people who read the chapters and tutorials.

Shortcomings:

1. My CUE Guide review grades in ES-151 (Electromagnetic Communications) were not as high as those of ES-50 and ES-158/258. I believe that this is a reflection of the fact that I do not conduct research in the area of Electromagnetics, so when I teach the course, I am not in a position to project enthusiasm onto the students.
2. In ES 158/258 (Digital Communications) I still only get at most 3-4 undergraduates per year. There is certainly interest among the students. However, the majority of undergraduates drop the course because they do not seem to be well prepared in Probability Theory. I am hopeful that this will change in the coming years because for the past 2 years Professor Navin Khaneja taught an excellent course in Probability Theory for Engineers (ES-150) that is perfect preparation for ES-158/258. The students who took the course should be seniors now, and will hopefully take ES-158/258.

Research

This is perhaps the most important section of this memorandum. I give a detailed description of my research activities. I consider my research areas to be Signal Processing, Information Theory, Communications and Magnetic Recording. Obviously, not all of these areas have received the same attention, but it is my goal to keep an active research effort in all four of these areas in years to come. Please refer to the section on Future Research Directions for my plans for the future. Here, I elaborate on the research that I conducted in the past 6 years. I divide my entire past research effort into 11 distinct topics.

1. Magnetic recording channel modeling
2. Detectors for magnetic recording
3. Information theory / magnetic recording channel capacity computations
4. Feedback capacity of Gaussian channels with memory
5. Coding for channels with memory
6. Optimization of LDPC codes
7. Timing recovery
8. Array signal processing

9. Feedback in detection and estimation
10. Wireless channel measurements
11. VLSI implementation of detectors for magnetic recording

Before I give a detailed description of my research efforts in these areas, I want to give a brief overview of the topics. I am probably best known in the Communications community for my work on topics 1 and 2 (Magnetic channel modeling and detectors) since this research is the continuation of my Ph.D. thesis. However, I consider my best results thus far to be in topics 3 and 4 (Information theory and channel capacity computation / Feedback capacity of Gaussian channels with memory) which I started about 3 years ago and where I have had success in cracking a 40-year old open problem. Topic 5 (Coding for channels with memory) was also started after I joined Harvard as a collaborative effort between Michael Mitzenmacher and myself. Topic 6 (Optimization of LDPC codes) grew out of the work on topic 5, but is now mostly concentrated on wire-line and wireless channels (not storage channels like topic 5). Another topic started after joining Harvard is topic 7 (Timing recovery), which was the topic of my NSF CAREER proposal (and award). Topic 8 (Array signal processing) was the subject of my undergraduate thesis, which produced 4 papers in 1994-1996, but recently I went back to that research area since it has applications in signal processing for multiple antenna communication systems. Topic 9 (Feedback in detection and estimation) is relatively new as it was started just a year ago with the aim to understand how feedback can aid in making more reliable statistical inference. Topics 10 (Wireless channel measurements) and 11 (VLSI implementation of detectors for magnetic recording) are both recently started collaborative efforts with other electrical engineering faculty in the Division, i.e., with professor Tarokh and professor Wei, respectively.

In the following I give a detailed account of my research activities in the past 6 years and cite my own papers. These papers are listed in my Curriculum Vitae, which is submitted with this memorandum. Citations that begin with the letter "P" refer to the patents, while those that begin with the letter "U" refer to unpublished work. Citations that begin with letters "BT" refer to book chapters and tutorials. Recent journal submissions and recent conference submissions are denoted by letters "JS" and "CS", respectively. Citations with the letter "J" refer to papers published in journals and those starting with the letter "C" refer to papers published as conference proceedings.

1) Magnetic Recording Channel Modeling

The driving force behind the magnetic recording industry is the constant push towards higher recording densities. Hitachi now makes disk drives barely bigger than a 25-cent coin. With densities increasing, the bit spacing between oppositely magnetized domains on a magnetic medium drops. At these very low bit distances (today at about 10-20nm apart), the non-uniform (zig-zag) boundary between the bits becomes observable. In fact, this non-uniformity was observed in the lab even at much lower densities (in the 1980s), but the recording densities were sufficiently low, so the effect did not affect the readback signal. Today, since the bits are much smaller, the sizes of the zig-zag domains are comparable to the bit sizes and the effect is certainly observable. These zig-zag domains give rise to what is known as media noise. In the 1990s, researchers started projecting that media noise will be the main limiting factor of future recording systems. In order to design signal detectors for magnetic recording devices, a good model was needed for the media noise and the readback signal. I experimented with several models. The first model I devised fitted a renewal stochastic process in the cross-track direction to describe the zig-zag boundary. I first published this idea in [C36]. Later versions expanded the model to include typical nonlinearities [J21], and to include read-head transfer functions in order to reconstruct waveforms [C35]. The model was validated in the lab and the results reported in [J19]. The paper that closed the topic was [J18], which is a comprehensive study of the underlying stochastic processes that I used to create the model.

Despite the success and the lab verification of the zig-zag media noise model, I soon found that the model was only good for creating waveforms via the computer, but it was not well suited for the design of read-back detectors. The reason for this is that the model was a media-level stochastic process. Though one can easily extract the structure of the stochastic process (as I have shown), by the time the noise passes through the read head, through the front-end amplifier and gets added to the stationary electronics noise (typically white noise) the nice structure of the media-level stochastic process gets lost. For this reason it became clear that if the model was to be useful to detector designers, a different model was needed. The main requirements for the model were that it be simple, yet accurate and realistic. The model needed to be simple enough to facilitate fast generation of waveforms to be used in Monte Carlo simulations of the receiver, it needed to be accurate enough to be able to predict the behavior of a real magnetic recording channel, and the knowledge of the model parameters needed to be linked to the parameters of the optimal detector. The model I proposed was a result of simultaneously working on the design of the optimal detector. It turned out that a data-dependent (i.e., time-variant) autoregressive model was the best choice because it had a simple description (a bank of discrete-time all-pole filters). Moreover, the same parameters that describe the noise, also describe the optimal detector – the maximum likelihood sequence detector (also known as the Viterbi algorithm). The model was first published in [J16]. Subsequently, I learned how to apply Rissanen's minimum description length (MDL) principle to determine the model size and the model parameters, and have published the work in [J13]. An interesting feature of the MDL principle applied to magnetic recording modeling is that the maximum likelihood equations are not tractable, so the work that appeared in [J13] used a minimum energy estimate as a substitute for the maximum likelihood estimate. The industrial community recognizes the model to be a very valuable tool. Several researchers in industry and in academic institutions use this as the standard model for the magnetic recording channel when simulating detectors and decoders for these channels. The National Storage Industry Consortium (which is made up of member companies that sponsor academic research) has also adopted this model as their standard model, which is evidence of the quality of the work. The National Industry Consortium has been funding research at Harvard for the past six years.

The last piece of work that resulted from the magnetic recording modeling effort was an elegant formulation of inversions of matrices whose inverses are known to be banded matrices [C32, J14]. The connection to the magnetic recording model is not very obvious as these papers concentrate on the fundamental properties of correlation structures and banded matrices. However, the connection becomes clear when we realize that the **inverse** of a noise correlation matrix (the **potential** matrix) is a banded matrix if the noise is a finite-memory autoregressive stochastic process. If such a process is Gaussian, then the autoregressive nature causes it to be a Gauss-Markov process. Moreover, the Cholesky matrix decomposition of the potential matrix reveals the autoregressive coefficients (both in the stationary and nonstationary case). These connections can now be exploited to find fast methods for matrix inversions and closed-form inversions of tridiagonal matrices as a Hadamard product of three matrices: a forward Wiener process covariance matrix, a backward Wiener process matrix and the covariance matrix of a unit-variance stochastic process – hence, the phrase "factorization of Gauss-Markov process" in the title of [J14].

I also had an attempt to model the magnetic recording media noise as a discretization of jitter noise. While this attempt proved to be a fun exercise for my coauthors and me [C29], its significance is far from the modeling work described above.

2) Detectors for Magnetic Recording Channels

The work that characterized my Ph.D. thesis is the formulation of the optimal signal detector for the magnetic recording channel that is sensitive to intersymbol interference, the data-dependent character of the noise and signal nonlinearities. Of all my work so far, this has made the biggest impact in the magnetic recording industry. All major read channel manufacturers (Agere, Marvell, ST-Microelectronics) utilize a form of the detector I proposed in their latest generations of read channel chips.

The presence of magnetic media noise was perceived in the 1990's as a hindrance for designing good detectors. In the early 1970's Forney showed that the Viterbi algorithm is the maximum likelihood sequence detector (i.e., the optimal detector) for a channel with intersymbol interference when the channel noise is white. Extensions to colored (but stationary) noise sources are straight forward through the use of whitening filters. However, when the statistics of the noise in the channel are dependent on the data transmitted through the medium, the design of the optimal detector was not obvious, and was the subject of research in the magnetic recording community during the 1990's.

My approach was to approximate the noise process as a data-dependent Markov process. When the Markov noise memory is sufficiently large, the Markov process is an accurate representation of the noise process. If the probability density function of the noise sample is Gaussian (as is the case in practice), then the noise is a Gauss-Markov process. Now, a Gauss-Markov process can be shown to be an autoregressive process, so this all ties very nicely to the autoregressive noise modeling work described in research topic 1). My initial attempts to formulate the optimal detector were clumsy. I first formulated the detector in a matrix form [C34, J17]. Later I realized that a more elegant (and computationally more efficient) formulation could be given in terms of filter banks [C33, J15]. This implementation was patented [P1]. Ties to Gauss-Markov processes, the Cholesky factorizations of covariance matrices, and autoregressive models were presented in [J14]. The symbol error rate analysis using union bound techniques was presented in [C30]. Extensions of this work towards soft-output detectors (which are needed for soft-output iterative decoding) were made in [C31], and the method was patented [P2]. In [J12], I have shown how to formulate the zero-forcing decision feedback equalizer (ZF-DFE) for channels with data-dependent noise, which is a suboptimal detector that has some implementation advantages if the system can handle error bursts. For example, error bursts are not tolerable in disk drives, but in magnetic tapes bursts are tolerable. This is because tape systems have other sources of error bursts that are not linked to the detector. Thus, the error bursts introduced by the detector are only a small fraction of the bursts otherwise present.

3) Information Theory / Magnetic Recording Channel Capacity Computation
Only a few years ago, while I was a Ph.D. student, the information capacity of a magnetic recording channel was not known. The simplest model for the magnetic recording channel is the *partial response channel* model. The partial response channel is a channel with **binary** inputs, corrupted by intersymbol interference and white Gaussian noise. Indeed, in an overview paper in 1998 (coinciding with the 50th anniversary of the birth of the field of Information Theory), Immink, Siegel and Wolf said that the biggest challenge in designing storage systems was to evaluate their capacities. In 1948, Shannon introduced the channel capacity as the supremum of the mutual information rate between the channel input and the channel output. Shannon showed that there exists a code that can approach any data rate arbitrarily close to the capacity with an arbitrarily small probability of error. Shannon also computed the capacity of the memoryless channel with additive Gaussian noise.

In the 1970's, for channels with discrete inputs (say, binary inputs), Arimoto, and independently, Blahut, found an iterative method to compute the capacity of an arbitrary discrete memoryless channel. For channels with memory and no noise (and discrete inputs), Shannon computed the capacity (this is actually the maximal entropy of the constrained inputs) in his 1948 paper. However, for channels with memory and constrained inputs (say binary inputs) the capacity was not known. Only rather loose upper and lower bounds by Verdu, Shamai, and Hirt were published in the 1980's and 1990's. Presented with this opportunity, I opened a new line of research when I joined Harvard – to compute the channel capacity of the partial response channel.

My approach to the problem has been as follows. Instead of optimizing the information rate over all input stochastic processes, I constrain the input process to be Markov with a fixed memory length M. I maximize the information rate for M fixed. Then I let M change and always maximize the information rate for the new M. In practice, M needs to be increased until further increase of M does not cause further increase in the maximized information rate. This then represents the numerical value of the channel capacity. Of course, I cannot claim that this is exactly the capacity, but rather a very tight lower bound. I can support this claim by computing an upper bound, and showing that it is very close to the lower bound.

The method I designed for computing the lower bound is a method that generalizes the Blahut-Arimoto method and Shannon's entropy maximization method. In fact, both of these methods, though seemingly unrelated, can now be shown to be special cases of the new method I devised. The method is an iterative Monte Carlo algorithm, with resemblances to the Expectation Maximization algorithm. The method resembles the Arimoto-Blahut algorithm in that it is an iterative procedure. On the other hand it generalizes Shannon's solution by defining the **noisy adjacency matrix** on the graph representing the channel. I describe the algorithm in a paper that was presented at the GLOBECOM conference in November 2001 [C26]. In this paper, I did not give the proof that the method converges (although this was evident from the numeric results). Today, we still do not have a proof of global convergence, but we can show that the method achieves a *local* maximum of the information rate [U2].

The algorithm I devised can be used to solve several open problems in information theory. For example, we can now evaluate the capacity of intersymbol interference channels if the inputs are constrained to a finite alphabet. Applied to magnetic recording, we can compute the capacity of the autoregressive channel model [C18]. Another application is the computation of the maximal information rate for constrained codes. Typical constrained codes are the run-length-limited codes used in communications systems. These codes impose a limit on the number of consecutive zeros to guarantee that the system clock is not lost due to a long stream of zeros (i.e., no signal) at the channel output. This is the application shown as the example in my paper [JS5].

A lower bound cannot be claimed to be tight unless an equally tight upper bound can be computed. I have given this task to Shaohua Yang, a Ph.D. student in my research group. In a very short time (surprisingly short for a first-year graduate student), Shaouhua has generated some remarkable results, working almost independently, with only conceptual guidance on my part. First, he used a known result that the **feedback** capacity of a channel with memory is higher than the regular capacity and computed preliminary lower bounds on the feedback channel capacity [C24]. Then he showed that the feedback capacity is achieved with a feedback-dependent Markov process, whose memory length equals the channel memory length, and showed that the optimal feedback is captured by the coefficients of the forward recursion of the Baum-Welch algorithm [C22, C19, JS5]. This is a remarkable result, which generalizes some known facts in information theory. For example, it is known that the capacity of a memoryless channel equals the feedback capacity of the channel, and that a memoryless random process achieves the capacity. Now we know that the feedback capacity is achieved by a process whose memory length equals the memory length of the channel (which is an intuitively pleasing result). Knowing that the optimal feedback is captured by a finite set of Baum-Welch coefficients, we can now solve the Bellman equation and compute this capacity (up to the quantization accuracy which can be made arbitrarily small). Preliminary results showed that the feedback capacity is a very tight upper bound to the capacity computed by the iterative Monte Carlo algorithm presented in [JS5]. Later, we found a way to further tighten the upper bounds by delaying the feedback [C11].

I consider this to be the best work I have done thus far. Not only was this a hard problem considered by many researches in the past, but also the solution now clearly shows what is achievable in terms of areal densities in magnetic recoding [J6]. Unlike the work on detectors, this work has no chance of being implemented in a chip, but it does have clear practical implications for the magnetic recording industry. It serves as a benchmark for future work on codes. For example, we now know how good our error correction codes are. We know that some (very long and impractical) codes are already very close to the capacity, and we know that practical codes still have a gap of 3-4dB to cover. This clearly sets the stage for future research on codes for magnetic recording.

The work on channel capacities of magnetic recording channels appeared in several dispersed papers (some of which have not yet been submitted or published) [U2, JS3, JS5, J6, C11, C18, C19, C22, C24, C26]. We have summarized the work in a book chapter [BT2], which I have listed as one of my most important papers.

4) Feedback Capacity of Gaussian Channels with Memory

The feed-forward capacity is the capacity of a channel if the receiver has no way of communicating back to the transmitter. Contrary to that scenario, if the receiver has a feedback link to the transmitter, we can talk abut the *feedback capacity*. Traditionally, with only a few exceptions, the computation of the feedback capacity of any channel has been much harder than the feed-forward capacity.

Unlike the work on the capacity of magnetic recording channels, the research on the feedback capacity of Gaussian channels with memory is of a much smaller practical importance. However, its theoretical importance is a lot higher. This has been an open problem ever since Shannon's groundbreaking work in 1948. Shannon showed that the feedback capacity of memoryless channels equals the feed-forward capacity. He also computed (in closed form) the capacity of a memoryless Gaussian channel. Since then, researchers have tried to compute the feedback capacity of Gaussian channels with memory. Most notably, in 1989, Cover and Pombra have expressed the feedback capacity as a solution to an optimization problem, but the solution to that problem is analytically and numerically intractable.

With only conceptual guidance from me, most of the work done on this topic was conducted by my Ph.D. student Shaohua Yang, who has already graduated and now works for Hitachi Research. Given what we have learned about feedback capacities of finite-state machine channels, we have decided to model the Gaussian channel with memory as a state space realization. This turned to be crucial for making a breakthrough because we were then able to utilize the vast knowledge available in the field of linear systems and control theory (specifically Kalman-Bucy filtering and dynamic programming). We have showed that Markov processes achieve the feedback capacities, and that the optimal processing of the feedback is achieved by the well-known Kalman-Bucy filter [J2, C4]. We have shown that the parameters of the optimal Kalman-Bucy filter are numerically computable using standard dynamic programming techniques such as value iteration. This marks the first method that is capable of numerically computing the feedback capacity of a Gaussian channel with memory. We also have an analytic solution that is valid only under the assumption that the optimal Kalman-Bucy filter is time-invariant in steady-state, but we were not able to prove that this assumption holds for a general Gaussian channel with memory, so the analytical solution is still a conjecture. However, this conjectured analytical solution numerically matches the numeric solution for all channels we tested.

The work on the feedback capacity of Gaussian channels with memory has been the major thrust of my Ph.D. student's thesis. The student, Shaohua Yang is the only Ph.D. student I have graduated so far, and I am tremendously proud of him. His work has been exceptionally well received by the information theory and magnetic recording communities, and he has had numerous academic and industrial offers. He decided to join Hitachi Research after graduation.

5) Coding for Channels with Memory

For about fifty years after Shannon introduced the concept of the channel capacity in 1948, researchers have been trying to design error correction codes to achieve the capacity. In 1993 Berrou and Glavieux introduced turbo codes and constructed an iterative decoder to decode them, showing a remarkable performance at only a fraction of a dB away from the Shannon capacity of the memoryless additive Gaussian noise channel. The best codes prior to turbo codes were not capable of getting closer than 3dB away from the capacity. This event sparked renewed interest in the design of error correction codes. Soon, LDPC codes (low-density parity-check codes, also known as Gallager codes) were rediscovered. Irregular LDPC codes were introduced by Luby and his coauthors in the late 1990's. By the early 2000's, Chung et al. demonstrated an irregular LDPC code that approached the channel capacity of the binary additive white Gaussian noise channel to within 0.005dB. Effectively this marked the closure of the field of capacity approaching error correction coding for binary *memoryless* channels.

For channels with memory, the story is slightly different. Until the very recent results from our group, the channel capacity for channels with memory (specifically, intersymbol interference memory) was not known. For this very reason, no capacity-approaching codes were known either. To try to design codes for channels with memory, Michael Mitzenmacher and I teamed up and submitted an NSF proposal for a project, which received funding. Our first task was to evaluate the performance of existing LDPC codes over binary-input channels with intersymbol interference memory. In [C27, J8], we proved a concentration result showing that every LDPC code over channels with memory has a noise tolerance threshold. If the noise variance is lower than the threshold, the iterative message-passing decoder can recover the codewords with an arbitrarily small probability of error, provided that the block length is large enough. We further proved that LDPC codes couldn't deliver arbitrarily low probabilities of error at code-rates higher than the i.u.d.-capacity [J8]. The i.u.d.-capacity is the mutual information rate between the channel input and the channel output, when the channel input sequence is a sequence of equally likely independent and uniformly distributed (i.u.d.) binary symbols. The i.u.d.-capacity is lower than the channel capacity for channels with memory. Therefore, it became clear that we could not hope to achieve the capacity of binary intersymbol interference channels with LDPC codes. For rates below the i.u.d.-capacity, Ned Varnica, a Ph.D. student of mine, has demonstrated that properly constructed LDPC codes can approach the i.u.d.-capacity arbitrarily closely [J9,C21].

In order to construct capacity approaching codes for binary intersymbol interference channels, we figured out a method to use the capacity computation algorithm [C18]. This algorithm computes an achievable (tight) lower bound on the capacity. A by-product of that computation is the characterization of a Markov process that achieves the capacity. The idea is now to construct a code that receives i.u.d. bits at its input and produces output bits that mimic the optimal Markov process. We used a trellis construction to produce a code with these characteristics. We call this trellis code a **matched information rate code**. The trellis capacity of the matched information rate code is very close to the information rate of the optimal Markov process, hence the name matched information rate code [J1]. To achieve the channel capacity (or more accurately, to achieve the trellis capacity of the matched information rate code), an outer random code can be utilized, as was proved in [J1]. However, random codes are unfeasibly complex to decode, so we use LDPC codes as outer codes [C25]. In a recent paper, we have shown how to construct outer LDPC codes and inner matched information rate codes that are less than 0.5dB away from the capacity and about 0.2dB above the i.u.d.-capacity [J1,C20]. These are the first codes for binary intersymbol interference channels to have been demonstrated to perform above the i.u.d.-capacity. We are now working on extending these ideas to insertion/deletion channels. All of this work is a joint research effort in collaboration with Michael Mitzenmacher, who is an Associate Professor of Computer Science in DEAS.

The research on capacity approaching codes for partial response channels (which are channels with memory specific to magnetic recording) has been summarized in a book chapter [BT1].

6) Optimization of LDPC codes and decoders

In the process of devising codes for partial response (magnetic recording) channels, we saw opportunities to utilize the acquired knowledge to optimize codes for wire-line and wireless communications. The challenge in these communications systems is no longer the construction of codes that achieve the capacity because there already exist low-density parity-check (LDPC) codes that can achieve (or rather approach) the capacity. However, these codes have a prohibitively large block length for practical applications. For moderately high block lengths, the codes experience error floors, which are undesirable in communication channels. On the other hand, if an LDPC code is constructed to have a short block length, the standard message passing decoder (implemented by localizing the Bayes rule) falls 1-2dB short of the performance of the maximum likelihood decoder.

In a collaborative effort with professor Marc Fossorier of the University of Hawaii, we devised a method to improve the performance of decoders of LDPC codes. We made an observation that when noise causes a parity check equation to fail, the most likely bit to be wrong is the one that configures in the largest number of other parity check equations. The modified algorithm is then very simple: for a parity check equation that failed, locate the bit that is also part of most other parity check equations, and set that bit to 1 or –1, arbitrarily, and continue decoding. Even though this is really an ad-hoc method, the results are quite remarkable – for short-length LDPC codes, we were able to achieve near maximum likelihood decoding performance, while for long-length codes we lowered the error floors (measured by the bit error rate) by 2-3 orders of magnitude. Remarkably, these results are not specific to any particular LDPC code, but can be applied to a variety of LDPC codes regardless of their structure, construction methods, size, regularity, etc. This work has recently been submitted for publication [JS2], after a patent has been filed [P4].

Other areas where we found that carefully optimized LDPC codes can achieve gains over non-optimized codes are in memoryless Gaussian channels, where we constructed constellation shaping codes [JS4,J3,C14,C16]], and in multiple access channels, where we showed that interleaving several LDPC codes can approach the multiple access channel capacity [C2]. The latter work is still in progress.

7) Timing Recovery

The afore mentioned improvements in the performances of error correction codes mean that future communication systems will have the ability to operate at very low signal-to-noise ratios. This means that these codes will have the ability to tolerate very high noise levels, and still recover all the data with extremely small probabilities of making an error in the transmission. However, all these statements hold if the timing recovery circuitry works well at these low signal-to-noise ratios, which may not hold. Indeed, most timing recovery circuitry used in communications systems today was designed for error correction codes less powerful than the recently introduced iteratively decodable turbo codes or low-density parity-check codes. Timing recovery is an essential component of a digital communication system since the incoming waveforms need to be sampled and digitized. If, due to the noise in the channel, the sampling instances are not accurately determined, the entire communication system will fail. To fix this problem, I have started a small research effort on timing recovery methods that would work in conjunction with iteratively decodable codes. This was the topic of my NSF CAREER proposal, which got funded four years ago, and has provided support for one Ph.D. student in my group.

The idea that we are trying to exploit is to use the power of the code to aid the timing recovery circuitry. To this end, we have identified the occurrence of cycle-slips as the main cause of timing recovery failure. Cycle-slips occur if the sampling clock skips a sample or inserts an extra sample. This error has a catastrophic effect on the performance of the receiver. The only way to recover from this error is to resample the waveform. However, in order to determine where the error occurred, we need to detect the presence of the cycle slip. We therefore made it our task to construct a cycle-slip detector. We used the soft decisions on the transmitted bits as inputs to the detector [C28]. (Soft decisions are a posteriori probabilities of the bit being either a zero or a one, where the term *a posteriori* means that the probabilities are determined after the waveform is observed.)

The next step was to design a cycle-slip detector for a channel with intersymbol interference. The problem with these channels is that the soft-decisions are obtained from the Baum-Welch algorithm (also known as the forward-backward algorithm). This means that the receiver needs to observe the entire waveform and then start a backward-going computation in order to provide the soft-output information – so the soft decisions cannot be obtained in real time. However, a timing recovery circuit cannot wait for the entire waveform to be received before it decides where to sample the data. For this reason, we conducted research on a class of forward-only soft output algorithms. This work was published in [J10], where we used graphical representation to describe the class of forward-only algorithms. These algorithms generalize the Viterbi algorithm, i.e., the classical Viterbi algorithm is an instantiation of a general class of forward-only trellis algorithms. The practical importance of such algorithms is that they can provide soft decisions with only a short delay, and can therefore be used in a real-time circuit that implements the cycle-slip detector for the intersymbol interference channel.

The research proceeded with formulating a joint iterative timing recovery and decoding system. We have made an empirical discovery that if we iterate between the decoder and the timing recovery algorithm, the decoder converges after a while. What this means is that we first sample the waveform and then run a decoding algorithm. The decoding algorithm does the best job of determining the transmitted bits. The determined bits (though they may be wrong) are then fed back to the timing recovery algorithm to resample the waveform and run the decoder again. The process is repeated until the decoding converges to a valid codeword (which usually is the transmitted codeword). This algorithm typically takes 30-40 iterations (i.e., 30-40 re-samplings of the waveform) to converge. We next discovered that if we use a cycle-slip detector to determine the point of cycle-slip occurrence, we reduce the number of iterations to 3, which is a substantial improvement over 30-40 [J11].

In 2002, Wei Zeng, a Ph.D. student of mine took over the project as it received funding from Seagate Technologies. Since then, we have made several improvements. First, we formulated the optimal soft detector (maximum a posteriori detector – MAP) detector for channels that suffer from intersymbol interference and synchronization errors (i.e., timing uncertainty) [J7,C12]. This work was also summarized in a tutorial paper [BT4], and a patent application was filed [P3].

Our next task was to optimize conventional timing recovery circuits, i.e., timing recovery mechanisms that cannot afford long delays to make timing offset decisions. The result was a trellis-based optimal timing recovery loop [CS2, U1]. Since Markov timing errors are used to model the timing error process, we also devised a method to estimate these parameters from channel waveforms [CS1]. The practical significance of this work is that we can now go to a disk drive, capture the waveforms, estimate the timing error parameters and use them to calibrate the timing recovery circuitry, hence the interest by Seagate Technology to fund the work.

Our next thrust in the direction of jointly combating detection, decoding and timing recovery is to evaluate the capacity of channels with synchronization errors. A small step in this direction has been a method to lower-bound the capacities of insertion/deletion channels [C6]. Further work in this direction is under way. We hope to characterize the capacity of channels with synchronization errors, and to devise error correction codes that will jointly solve the problems of timing recovery, symbol detection and error correction decoding at rates close to the capacity of such channels.

8) Array Signal Processing

I first started doing research in array signal processing as an undergraduate at Bochum University in Germany. The topic of my senior thesis was fast algorithms to track the signal and noise subspaces of a correlation matrix. These methods were used to track targets in a radar or sonar tracking problems. The basic method was based on tracking an approximate singular value decomposition of the covariance matrix, because the full decomposition was computationally too expensive to be affordable [C38]. The method went through several refinements in the direction of systolic array implementations and simultaneous subspace tracking and rank determination [C37, J20, J22]. Here, the rank defines the number of targets present, while the subspace determines the positions of the targets.

In itself, the publications that resulted from this undergraduate thesis have a made at best a transient impact on the community. Perhaps the best result of the work was that I received exposure to linear algebra and how basic matrix manipulation techniques can be used in array signal processing. These same techniques are used today for space-time signal processing of multiple antenna output signals, adaptive beam-forming, spectral shaping, etc. Recently, however, I came back to that topic.

The scenario I considered is a multiple antenna system, also known as the multiple-input multiple-output (MIMO) system. It is assumed that the channel is known both to the transmitter and the receiver, and that the receiver uses the back-cancellation detector to separate different users. It is also assumed that the transmitter utilizes a predetermined signal constellation (e.g., binary shift keying). Under this scenario, the bit error rate of an uncoded system depends on the precoder that is utilized for transmission. We show that the optimal precoder is a unitary matrix whose task is to turn the channel matrix into a matrix whose QR-decomposition (decomposition into a product of a unitary Q factor and an upper diagonal R factor) has an R factor whose diagonal entries are all equal to each other. Such a precoder gives rise to the lowest probability of block errors, and the resulting detector asymptotically approached the performance of the maximum likelihood detector as the SNR increases [J4,C23]. This precoder improves upon the back-cancellation detector used for the BLAST (Bell-Labs Space-Time) detector, whose equivalent precoder is a permutation matrix (where ours is a special unitary matrix). I consider the most important discovery of this work to be the special properties of the QR decomposition with equal entries the R factor, leading to statements of optimality of the precoder and detection method [J4].

Present work on the topic is concentrated on expanding the theory towards fading MIMO channels in scenarios where the transmitter does not know the channel realization.

9) Feedback in detection and estimation

Consider the problem of interrogating a subject. A good interrogator will approach the subject *not* with a list of questions, but with a *policy* on how to generate new questions based on the answers received. When it comes to utilizing feedback to control the information flow, humans present no exception. An example of an efficient interrogator in the animal kingdom is a bat. A bat will adapt the signal it sends at a target based on what it wants to know about it. It uses the feedback provided by the echo to generate new signals (questions) to refine the hypothesis tests. Despite these and other numerous examples of the usage of feedback in the natural world, humans have not yet figured out a way to understand information feedback even though we take great pride in our achievements that led to the "information age". While feed-forward methods for retrieving and generating information are well understood (e.g., the Neyman-Pearson lemma), we still lack basic mathematical tools to understand how to create optimal policies that utilize information feedback. Nevertheless, information feedback is enormously important and useful, as evidenced by the numerous examples of interrogation with feedback in the natural world that had millions of years of evolution to refine its methods.

In an engineering setting, the subject to be interrogated can be viewed as a communications channel, whose parameters (properties) we need to determine. We are allowed to send probing signals through the channel, and we receive the response (echo, or channel output). Based on successive channel inputs and outputs, we want to guess the channel properties in the most time-efficient manner. Thereby, our goal is to design a policy that would enable us to optimally utilize feedback in order to formulate the subsequent probing signals.

In the easiest setting, we (my Ph.D. student Patrick Mitran and I) assumed that the channel is a discrete memoryless channel (whose input and output alphabets are discrete) that could take one of finitely many realizations (hypotheses). Even in this simplest scenario, formulating the optimal finite-horizon policy is a very difficult problem. For this reason, we considered an infinite horizon problem, where our measure of optimality is the error exponent, i.e., the rate at which the probability of making an erroneous decision on the hypothesis drops with the number of transmitted pulses. So far, we have revealed several very interesting facts. First, we found that when we have only 2 hypotheses, the optimal open-loop policy consists of repeating a single input symbol for all time [C5], which is contrary to what the literature suggested. Indeed, we have shown that the Chernoff policy (which was thus far believed to be optimal) is indeed suboptimal [JS1]. Further, we have found a formula for evaluating the optimal open-loop policy, and for finding the policy that upper-bounds the error exponent in the feedback scenario [JS1]. Our recent work has shown that when the channel has memory (specifically Markov memory), the memory order of the optimal policy suffices to be only as large as the memory length of the channel [C1].

Obviously, this research effort has just started, but the results have been enlightening so far. Our goals are to formulate optimal feedback policies for a wider range of detection and estimation problems that reach beyond simple discrete memoryless channel identification. While we have only studied the detection problem thus far, a natural progression would be to consider estimation problems, which is where the project is heading next.

10) Wireless channel measurements and modeling
My involvement in wireless channel measurements has started two years ago when Professor Tarokh joined Harvard University. This particular research effort has been a collaboration between Professor Tarokh, Dr. Ghassemzadeh (AT&T Research), Professor Greenstein (Rutgers University) and myself.

Two years ago, the Federal Communication Commission (FCC) approved ultra-wideband (UWB) wireless transmission in the frequency range 3.1-10.6GHz. Never before has such a large bandwidth been approved for transmission, and clearly there was no known model for wireless (fading) channels with such a large frequency band. The most probable wireless applications in this band are likely to occur within the home or office environments, for which good channel models were not available either. Our goal in this effort was to make channel measurements to characterize the UWB channel in indoor environments. We picked 20 commercial and 20 residential buildings to conduct the measurements and compiled the largest available database of indoor UWB channel measurements today.

Or next steps were to utilize the measurements to extract a path loss model [J5, C8] and a delay profile model [C9]. Simulation studies were also conducted to verify the validity of the models. The repeatability of the model, combined with its simplicity, make it an almost ideal tool for engineers who wish to test their wireless algorithms in simulated environments.

Our next project was a study of interference between UWB signals and the already established wireless standard 802.11. The findings showed that UWB clearly interferes with 802.11, thus presenting a need to suppress this interference in the narrow band in which 802.11 operates. A recently accepted conference paper [C3], shows how this suppression can be achieved through time-hopping sequence designs.

11) VLSI implementation of detectors for magnetic recording
This project is a collaboration between Professor Wei and myself. It was initiated by the need of the magnetic recording industry to start replacing the traditional Viterbi detector with more advanced soft-output detectors. Soft output detectors are advantageous because their outputs could be used for soft decision decoding, which is on average 3dB better than hard decision decoding at the output of a Viterbi detector. However, replacing the Viterbi detector in magnetic recording channels is a difficult task for the industry because the detector is also used in timing recovery loops.

Recently, we actually showed that soft-decision detection (also known as BCJR detection) could be performed in architectures similar to the Viterbi detector [J10]. This soft-output detector is thus an ideal drop-in replacement for the Viterbi detector in magnetic recording read channels. The proposal was sent to Agere Corporation, and they decided to fund the project.

For the past year, we have been considering several VLSI design architectures, and only recently we decided on the architecture for the implementation [C7]. Presently, Ruwan

Ratnayake, the PhD student who works on this project is at Agere Corporation, where he is learning how to utilize the layout tools in order to map the architecture into silicon. We expect the chip to be fabricated by the end of the year, at which time we will start testing it. We also expect to put out a second-generation chip a year from now since we will inevitably learn how to optimize the preset design. This chip will mark the first time a forward-only soft-output algorithm has been implemented in VLSI.

Summary of research activities Over the past 6 years, I had two basic goals that I wanted to achieve:

1. Build a magnetic recording research program at Harvard University that would be known worldwide for its quality of research.
2. Branch out the research activities into other areas of communications (e.g., error correction codes, wireless communications, timing recovery, array signal processing, VLSI implementation, etc.) in order to achieve higher visibility of Harvard's communications program.

I believe that I have been successful on both fronts (although I should probably not be the one to make the judgment). Our program in magnetic recording is second to none worldwide and we are present at all major communications conferences and thus achieving visibility in all sub-fields of communications engineering (naturally, a lot of credit goes to Professor Tarokh whose presence gave a boost to our visibility). In my view, our future goal should be the formation of a research center that would attract large and sustainable funding. I touch on this topic in the following section on future research directions.

Future research directions

Future of research in communications and storage: The past decade has been one of the most eventful in the history of the telecommunications industry and communications theory. The industrial events that marked the past decade in communications are

1. the break-up of the AT&T monopoly on telecommunications,
2. the break-up of AT&T Bell Labs into Lucent Bell Labs and AT&T Labs,
3. the emergence of the Internet and dot-com industries,
4. the rapid spread of wireless technologies,
5. and finally, the overheating of the economy and the collapse of two major telecom companies (WorldCom and Global Crossing).

The scientific events that marked the past decade in the field of communications theory are

1. the emergence of turbo (and other iteratively decodable codes) that achieve the Shannon capacity of various communication channels – an accomplishment that evaded researchers for 5 decades,
2. and the emergence of novel solutions for wireless technologies (including space-time codes) that revolutionized the wireless industry.

The storage industry, which is a smaller sister of the communications industry, benefited from these scientific trends as well, but unlike the telecom industry, the data storage industry did not experience the turbulent roller-coaster ride in the past decade, but rather a slow and steady growth. This growth rate of the data storage industry is actually projected to increase in the next decade as ever smaller data storage devices make their way into consumer products such as i-pods, cell phones, camcorders, digital cameras, etc.

The biggest question today concerning communications researchers is whether communication theory has a future in terms of science and funding. First, it is already clear that the main scientific problem that drove physical-layer communication theory research over the past decade (i.e., the pursuit of the Shannon capacity) has now been solved. Second, the telecommunications industry is presently struggling, and it is very reluctant to invest into theoretical research when its main concern today is to stay profitable. Third, the government is no longer ready to heavily support research in telecommunications theory because it has done so for the past decade already, and because the overall research budgets are lower. Does research in this field then have a future?

My view is that there are enormous opportunities to make an impact in this telecom environment, but it requires a shift in research directions and goals. First, my view is that telecommunications and data storage are two of the basic needs of civilized societies (along with transportation, housing, health care, food, etc.). Telecommunications and data storage were present in very early civilizations (in the form of courier messengers and libraries), so it is very unlikely that these two industries will collapse in the future. The challenge for them, however, will be how to stay profitable and how to utilize scientific knowledge to their competitive advantages. It is true that the scientists today have provided theoretical solutions to many problems in communications theory, but it is also true that only a small fraction of these solutions have been implemented in products, even after these solutions have been available for a decade (or more). Take turbo codes for example. They were available since 1993, but no disk drive company is utilizing them in products. This is because it is still tremendously complex and costly to operate the BCJR detector (a key component of turbo codes) at speeds of 2-3GHz. The industry still utilizes Reed-Solomon codes because their decoders are simple to implement, even though the industry knows that these codes are suboptimal. Other examples are very old timing recovery circuits that date to the 1970s. Unless these are modified, the industry will not be able to move towards lower signal-to-noise ratios, and the inventions of the past decade will remain unutilized.

Because most researchers in telecommunications and data storage have focused on solving theoretical problems in the past, there is today a huge gap between what is theoretically possible and what is practically implementable. Filling this gap presents an enormous research opportunity. My opinion is that the main thrust of research in telecommunications and data storage in the next decade will be to bridge this gap. This is supported by the recent funding trends, where the funding for novel theories is very difficult to obtain, while funding from industrial sponsors for developing practical solutions is much easier to secure. Indeed, lately, more than half of my funding comes from industrial sponsors (Seagate, INSIC, Agere). Further, I am presently involved as a consultant to several companies in the data storage industry, so I have the opportunity to experience their desperate need for practically implementable solutions that will directly affect their product lines. These problems are no less challenging than theoretical/mathematical problems. Their solutions require knowledge of optimal solutions of theoretical problems, but also the flexibility to deviate form optimality (sometimes in ad hoc fashion) in order to strike a balance between what is theoretically possible and what is practically realizable.

Over the past several years at Harvard, my research was mostly focused on theoretical developments (derivation of optimal algorithms, capacity computation, asymptotic analysis, etc.). In the decade to come, the main thrust of my research will change towards practical implementations. I will, however, keep a small research effort in the theory of communications mainly because I believe that there are beautiful theories yet to be discovered, but my main focus will be to make an impact on products by constructing methods that are implementable.

I next elaborate on several research areas in data storage and communications where I intend to focus my attention. I also touch on my view of how to address Harvard's communications group's strategic goal to improve its visibility nationwide.

Detectors and decoders for channels with synchronization errors As recording densities in magnetic storage increase, the signal-to-noise ratios (SNRs) are dropping. This drop in SNR causes the timing recovery (sampling) units in disk drives to be less reliable. The industry has dealt with the problem by increasing the preamble length in each sector, but eventually this started cutting away from the usable disk area. In my group, we have already started work on devising novel timing recovery mechanisms. In the future, we plan to expand the effort to include error correction codes resilient to synchronization errors and methods to alter the preamble (say by dispersing the synchronization symbols within the data symbols), and save some synchronization symbols for the postamble. Obviously, any such effort will be preceded by a channel capacity study that will guide us in designing optimal methods for compensating for synchronization errors. Support for this research has been secured form Seagate Technologies.

Compensation of nonlinearities in magnetic recording The magnetic recording process is highly nonlinear. Most of the nonlinearities take place at the time of writing the data, while the read-back process can be considered to be linear. The readback process has been the subject of intensive studies over the past decade, but the writing process never received much attention because the readback had priority and because the writing process is harder to analyze due to the nonlinearities. All indications form the magnetic recording industry suggest that a comprehensive analysis of the writing process, and compensation of the most common nonlinearities (non-linear transition shift, non-linear magnetic percolation effects) are desperately needed. Early studies show that up to 10% of areal density can be saved by utilizing simple ad hoc methods to compensate for nonlinearities. My intention is to study this phenomenon more carefully and provide optimal and suboptimal solutions to the problem. Presently, INSIC has initiated funding of this project, which I expect to continue over the next few years.

Alternative decoders of Reed-Solomon codes Even decades after their invention, Reed-Solomon codes are still the only error correction codes used in the magnetic recording industry. The main reason for this is the tremendous simplicity of the Berlekamp-Welch decoder of Reed-Solomon codes. Alternative codes (such as turbo and LDPC codes) are capable of operating at much lower SNRs, but their complexities are far too big to implement in VLSI circuitry operating at 2-3GHz. Alternative decoders of Reed-Solomon codes (such as the Sudan algorithm, or the Koetter-Vardy algorithm) offer only small gains, yet are still too complex to implement in circuitry. Given that Reed-Solomon codes have found a strong footing in the storage industry, it is very unlikely that the industry will abandon these codes in favor of turbo codes. It is far more likely that the industry will look for alternative decoders for Reed-Solomon codes that are simpler than the Sudan or Koetter-Vardy decoders, yet more powerful than the Berlekamp-Welch decoder. Early results in my group suggest that such decoders are possible. This is obviously work that is still in its early stages, but there are indications that Link-A-Media Corporation would be interested in funding this work.

Feedback in inference and communications Empirical evidence suggests that there are advantages in utilizing feedback when retrieving or conveying information. A good interrogator will examine his subject by adaptively utilizing the feedback he gets in the form of answers to formulate new questions. Similarly a good teacher will utilize the feedback from her students to adaptively change the lecture or teaching methods. Despite this empirical evidence, science has not yet figured out how to utilize information feedback to formulate optimal policies for information retrieval (detection and estimation) or communication. In fact, the theoretical results in this area are still in their infancy. My research group has certainly pushed the state of the art when it comes to computing feedback capacities of channels with memory and formulating feedback policies for hypothesis testing. However, these advances are still in the realm of examples. My desire and ambition is to formulate a comprehensive theory that would unify how feedback should be optimally utilized in inference and communication. Early indications suggest that the solution will lie in the intersection of information theory, detection and estimation, and control theory, but the proper mathematical apparatus is still not known. Over the next several years, I plan to work on this topic in hope of developing this area from examples to a comprehensive theory. Unlike my other future projects, this one is a solely theoretical endeavor, for which I do not know at this point how I will secure the funding, but which I will nevertheless pursue because I believe that there are exciting discoveries to be made.

Merging wireless and recording systems In the past decade, data storage devices have experienced tremendous growths in recording densities. Today, a small disk drive the size of a 25-cent coin can hold 1Gbyte of data. This has made data storage devices (disk drives) attractive for integration into various hand-held devices, e.g., i-pods, cell phones, digital cameras, camcorders, etc. The pressure to reduce the cost of these devices will force the electronics of disk drives to be integrated with the other electronic components of the devices. For example, disk drive controllers and read channels will likely merge with base-band wireless electronics. The challenge will be to operate both the wireless channels and the disk drive channel by utilizing the same hardware. Today, these two electronics components are still on separate chips, and to date only one company (Samsung) has announced a wireless phone product that contains a disk drive. My prediction is that this will be an irreversible trend and that wireless companies will be making strategic alliances with data storage companies. Research on merging the two channel electronics into one is still non-existent. However, the Data Storage Institute in Singapore (with whom I have an ongoing collaboration) is seriously considering opening a new research track along these lines. I intend to pursue this research direction with them in the future – perhaps a year from now. Further, this trend of merging storage and wireless technologies presents an opportunity for Harvard University to form a well-funded high-visibility research center. We presently have on our faculty experts in all major areas needed for realizing such a center (data storage, wireless communications, networking, RF design and VLSI design). One of my goals in the future will be to see if such a center can start living at Harvard.

Summary of future research The focus of my research in communication theory in the next decade will shift from searching for what is theoretically possible towards what is practically implementable. Funding trends today indicate that the government is less likely to fund theoretical projects, while the communications and data storage industries are willing to fund research as long as it concentrates on solutions that are geared towards implementations. Over the past 6 years, the focus of my research has been mostly on the theoretical side (information theory, capacity computation, asymptotic analyses, etc.). However, over the past year or two, there has been a gradual shift in my research direction towards implementations, as evidenced by the newest projects I took on (VLSI implementation, timing recovery, etc.). I expect this trend to continue for the next 5 years, and I expect to utilize the theoretical knowledge I gathered over the past decade to make an impact on the practice of communication engineering. Another shift in my research philosophy will be towards the formation of a research center that will provide high visibility and sustained funding for our communications program.

University and community services

Over the past 6 year I served (and still serve) the division in several different ways

1. I have been on the Committee of Higher Degrees for six years,
2. I have served a 3-year term as the co-Director of undergraduate studies for Engineering Sciences,
3. I have been a member of the Information Technology Committee for 3 year,
4. I have served on three EE and one CS junior faculty search committees,
5. I presently serve on the ABET Committee, and
6. for the past 2 years I have been serving on the Graduate Admissions Committee.

I do not think I need to elaborate further on my services on these committees. This is probably best done by the chairs of the committees themselves.

I have served the electrical engineering (or, more specifically, the communications engineering community) through various activities in conjunction with the Institute of Electrical and Electronic Engineers (IEEE):

1. From 1999 to 2000, I was the vice chair of the Boston Chapter of the IEEE Communications Society.
2. From 2000 to 2001, I was the chair of the same chapter.
3. I have served as session chair and session organizer for numerous sessions of the IEEE International Conference on Communications, the IEEE Global Communications Conference, the IEEE International Symposium on Information Theory and the International Symposium on Mathematical Theory of Networks.
4. I served as the member of the technical program committees for of the same conferences for the past 3 years.
5. From 2001 until 2004, I served on the Editorial Board of the IEEE Transactions on Information Theory as the Associate Editor for Detection and Estimation.
6. In 2003/2004, I served as the guest editor for the January 2004 issue of the IEEE Signal Processing Magazine on Iterative Methods in Communications.
7. I am presently the vice chair of the Signal Processing for Storage Technical Committee of the IEEE Communications Society.

Students and post-docs

Over the past 6 years at Harvard, I have supervised 5 M.S. and M.E. students (I am counting only those masters students who did a research project with me), 6 Ph.D. students, and 3 post-docs:

M.S. and M.E. students:

- **Jiajun Gu** originally came to Harvard as a Ph.D. student to work on the DNA nanopore project. She received her M.S. degree in 2004 and is presently working towards a Ph.D. degree with Professor Liu at the Statistics Department here at Harvard.
- **Albert Huang** completed his M.S. degree in 2001 and stayed at Harvard to work on the DNA nanopore project for another year. He is presently working for Sony Labs in Taiwan.
- **Xiaowei Jin** came to Harvard as a Ph.D. student, but failed to pass the Ph.D. qualifying exam. He was working with me on the timing recovery project and received his M.S. degree in 2001. He is presently at Notre Dame University, working towards the Ph.D. degree.
- **Thorvardur Sveinsson** completed his M.E. degree in 2003, and wrote a thesis on measurements and modeling of the indoor ultra-wideband wireless channel. He is presently working for Kogun in Reykjavik, Iceland.
- **Zhiyu Yang** came to Harvard as a Ph.D. student, but after a year decided to transfer to Cornell, where he is presently working towards the Ph.D. degree. He received his M.S. degree in 2001.

Ph.D. students:

- **Jason Bellorado** started his Ph.D. studies in 2001. His thesis research is on decoders for Reed-Solomon codes. He will graduate within the next 2 years.

- **Patrick Mitran** started his Ph.D. studies in 2002. His thesis research is on collaborative methods for information flow management. He will most likely graduate within the next 2 years.
- **Ruwan Ratnayake** is jointly supervised by Professor Wei and me. He started his Ph.D. studies in 2002. His thesis research is on VLSI implementation of high-speed soft-output detectors for magnetic recording read channels. He is most likely within 2 years of his graduation.
- **Nedeljko Varnica** came to Harvard in 2000. His Ph.D. thesis research is on optimizing LDPC codes for magnetic recording, wireless and wire-line applications. His anticipated graduation date is June 2005.
- **Shaohua Yang** came to Harvard in 2000 and has earned his Ph.D. degree in 2004. His thesis topic was capacity computation for channels with memory. He is presently employed at Hitachi Research Labs in Almaden, CA.
- **Wei Zeng** started his Ph.D. studies in 2002. His thesis research topic is timing recovery methods, and detection and decoding algorithms for channels with synchronization errors. His anticipated graduation date is January 2006.

Post-doctoral fellows:

- **Xiao Ma** was a post-doctoral fellow in my group from 2000 to 2002. His research included iteratively decodable codes, information theory and detection and estimation. After spending two years as a research staff member at the City University of Hong Kong, he is now an Associate Professor at Sun Yat-Sen University in China.
- **Murari Srinivasan** was a post-doctoral fellow in my group from 1999 to 2000. His research with me spanned magnetic recording channel modeling and iteratively decodable codes (he also worked part-time with professor Brockett). He is presently working for Flarion Technologies in New Jersey.
- **Jian-Kang Zhang** spent a year in my group as a post-doctoral fellow in 2001/2002. After initially working on the nanopore project, he later worked on array signal processing methods. He is presently a research faculty member at McMaster University, Ontario, Canada.

Collaborative activities

I divide my collaborative research activities into those within the university, and those outside the university. My collaborations within the university include collaborations

- with Professor Michael Mitzenmacher on error correction codes for channels with memory (see publications [J8, C13, C25, C27]),
- with Professor Vahid Tarokh on ultra-wideband wireless channel measurements and modeling (see publications [J5, C3, C8, C9]), and
- with Professor Gu-Yeon Wei on VLSI implementations of high-speed soft-output detectors fro magnetic recording (see [C7]).

My collaborative activities outside of Harvard include collaborations

- with Professor Hans-Andrea Loeliger at ETH Zurich on channel capacity computation methods for channels with memory (see [U2, JS3, C10, C29]),
- with Dr. Saeed Ghassemzadeh at AT&T Research (who support our research) and Professor Larry Greenstein at Rutgers University on ultra-wideband wireless channel measurements and modeling (see [J3, J5, C3, C8, C9]),
- with Professor Marc Fossorier at the University of Hawaii on optimization of LDPC decoders (see [JS2, P3]),
- with Dr. Emina Soljanin at Lucent Bell Labs on LDPC codes for wireless communications (no publications to report to date),
- with Professor Li Ping at the City University of Hong Kong on codes for multiuser detection (see [C2]),
- with Professor Toby Berger at Cornell University on feedback methods in information theory (no publications to report to date),
- with Professor Predrag Rapajic at the University of New South Wales, Australia on channel capacity computation for fading channels (no publications to report to date),
- with Professor Sekhar Tatikonda at Yale University on feedback capacities of channels with memory (see [JS5, J2, C11]),
- with Dr. Ara Patapoutian at Maxtor Corporation (who supported our research) on modeling the magnetic recording channel (see [J16]),
- with Professor Willaim Ryan at University of Arizona (project funded by INSIC) on capacity computation of the magnetic recording channel (see [J6]),
- with Dr. Fatih Erden and Dr. Erozan Kurtas at Seagate Research (who support our research) on timing recovery methods, and on pattern-dependent detectors (see U1, CS2, BT3]),
- with Professor Brian Marcus, previously at IBM (who supported our research), and now at the University of British Columbia on codes for intersymbol interference channels ([see C27]), and
- with Dr. Ravi Motwani at the Data storage Institute in Singapore on timing recovery and detection methods for channels with synchronization errors (see [CS1, C6]).

Research support

In this section I list the sources of my research support over the past 6 years as well as presently pending support

- NSF GOALI Signal/Noise Modeling and Detector Design for Media Noise Dominated Magnetic Recording Channels.
  Duration: Aug 1999 – July 2002.
  Amount: $340,000 (sponsored research).
- IBM Partnership Award for developing a program in Magnetic Recording.
  Duration: one-time award in 1999.
  Amount: $37,000 (gift)

- Quantum Corporation Award for Electronic Servo Design in Multichannel Magnetic Recording.
  Duration: one-time award in 1999.
  Amount $30,000 (gift).
- INSIC (International Storage Industry Consortium) support for 5 years for work on various topics in data storage.
  Duration: 1999-pesent.
  Amount: $180,000 (gift).
- NSF CAREER Timing Recovery on Digital Magnetic Recording
  Duration: July 2000 – June 2004.
  Amount: $200,000 (sponsored research).
- NSF CCR (with Michael Mitzenmacher) Low Density Parity Check Codes for Channels with Memory.
  Duration: Sep 2001 – August 2005
  Amount: $515,000 (sponsored research).
- Seagate Technology Research Award: Receiver on a Graph – Robust Timing Recovery in Recording Systems.
  Duration: Jan 2002 – present.
  Amount: $149,000 (sponsored research)
- Agere Corporation Award (with Gu-Yeon Wei): Hardware Implementation of Soft-Output Detectors
  Duration: May 2004 – Apr 2006
  Amount: $80,000 (gift)
- John Loeb Associate Professorship
  Duration: 2002 – 2006
  Amount $25,000 (university funding)
- Link-a-Media (pending) Award: Alternative Decoders of Reed-Solomon Codes
  Duration: 2 years; start date to be determined
  Amount $80,000 (gift)

In addition to these awards, I have two proposals pending. They are submitted in the package with this memorandum.

Most important research papers

I list my most important research papers (with brief descriptions) in 8 research areas

- **Detectors for magnetic recording:** In the field of detectors for magnetic recording I derived the maximum likelihood detector for signals that suffer from signal-dependent noise. In magnetic recording such noise is also known as media noise. The paper that describes the optimal detector in these environments is [J15]. This is probably the paper that to date has made the biggest and the most direct impact in the magnetic recording industry. Virtually all chip manufacturers in the disk drive industry utilize a detector that is a variation of the one I proposed in that paper (either in their present products, or in their next generation products that will emerge on the market within 6 months).

- **Magnetic recording channel modeling:** Probably the most cited paper of mine is a 3-page paper [J16] on how to model the signal-dependent (media) noise in magnetic recording. The paper was so well received by the industry because it was written in very simple language and because it was made accessible to a wide audience. Further, understanding the model is a prerequisite to understanding the optimal detector, so in that sense it serves as a tool for designing other (sub-optimal) detectors. The model was adopted by most disk-drive companies as evidenced by their usage of the detectors that are derived from the model.
- **Capacity of partial response channels:** In my opinion, my best work thus far has been a collection of methods to compute the capacity of partial response channels. This has been an open problem in magnetic recording for 4 decades, and I am extremely proud that we solved it within my group. The tools combine a method to compute a tight lower bound [C26] and an equally tight upper bound – the delayed feedback capacity [JS5]. Although the nature of the work is theoretical, it has clear implications for the magnetic recording industry because it can now evaluate how close to the capacity its error correction codes perform. Since the work is divided into two papers [C26] and [JS5], the best single publication that covers the topic is a book chapter [BT2].
- **Feedback capacity of Gaussian channels with memory:** My best theoretical work so far has been the computation of the feedback capacity of Gaussian channels with memory [J2]. This work will probably not make much of an impact in the practice of engineering because it turns out that the feedback capacity is not much higher that the well-known feed-forward capacity. However, the significance of this work is that it was an open problem for 5 decades on which many well-known information theorists have worked. In our paper, my Ph.D. student, Shaohua Yang, and I were able to show that Gauss-Markov processes achieve the feedback capacity and that the Kalman-Bucy filter needs to be utilized in the feedback loop. We also derived a method to compute the parameters of the Kalman-Bucy filter, and thus have a method to numerically evaluate the feedback capacity. This marks the first time after 5 decades of research that anyone has evaluated the feedback capacity of the Gaussian channel with memory.
- **Codes for partial response channels:** My best paper in coding theory is [J8]. The paper establishes performance limits of random linear codes for partial response channels. It also establishes performance limits of random low-density parity-check (LDPC) codes for the same channels. It then goes on to show how to conduct density evolution for partial response channels and how to compute code performance thresholds for ensembles of LDPC coset codes. The importance of the paper is that it establishes the limits of performance (and capacity achievability) of codes for partial response channels, and as such is useful as a benchmark for any practical code constructed for these channels.
- **Feedback policies for system identification:** The proper utilization of feedback in hypothesis testing is still an open problem. The idea is to send probing signals to a target, and then utilize the echoes to design subsequent probing signals. This has proven to be a very difficult problem for which small improvements are made

only once or twice per decade. In our latest work [JS1], we were able to use the method of types to further advance knowledge in this area and to obtain some surprising results. The first surprising result is that when we only have two hypotheses to choose from, there is only one type of signal that needs to be transmitted toward the target, i.e., it does not help to mix signals. The second surprising result is that we showed that the Chernoff policy for choosing subsequent signals (which was widely believed to be optimal, but with no proof) is indeed suboptimal. In the paper, we also propose a method to derive the optimal policy. Obviously this work is still in progress, but I think that the mathematical tools developed in [JS1] and the surprising findings are already worthy of mentioning as one of my most important papers thus far.

- **Array signal processing:** My best result in the area of signal processing is a method to compute the optimal precoder for multiple-input multiple-output (MIMO) communications channels [J4]. For a number of years, it was noticed that different precoders may lead to different error rates, but it was not known how to choose the optimal precoder. We have now solved the problem. The interesting fact about the optimal solution is that it is the result of a special matrix decomposition. This is a special QR-decomposition that makes the elements along the diagonal of the upper triangular R-factor all equal to each other. I am not aware that this type of symmetry in the theory of matrix decompositions has been noticed before. The value of this work thus reaches beyond precoders alone, and adds a bit to the theory of matrix decompositions.
- **Matrices with banded inverses:** Several years back, Gilbert Strang (MIT), one of the leading experts on linear algebra, posed the inversion of a tridiagonal matrix as an open problem. I am actually proud to have solved this problem (though it came to me quite accidentally during my study of optimal detectors in magnetic recording). I have published this result in a paper in IEEE Transactions on Information Theory [J14]. The paper deals with methods to invert matrices when it is known that the inverse is a banded matrix. The inversion of a banded matrix is the reverse of this problem and it is much harder. I could track the problem analytically only if the matrix is tridiagonal. The work will probably not make an impact on the engineering community, but as an engineer I am actually proud that I was able to contribute a piece of theory in the field of mathematics, which is why I list this as one of my most important research papers.

List of experts in the field

John Barry	(Georgia Institute of Technology)
Toby Berger	(Cornell University)
Vijay Bhargava	(University of British Columbia)
Chong Tow Chong	(National University Singapore)
Tom Cover	(Stanford University)
Marc Fossorier	(University of Hawaii)
Costas Georghiades	(Texas A&M University)
Mark Kryder	(Seagate Research, Pittsburgh, PA)

Frank Kschischang	(University of Toronto)
Hans-Andrea Loeliger	(ETH Zurich)
Brian Marcus	(University of British Columbia)
Steven McLaughlin	(Georgia Institute of Technology)
Jose Moura	(Carnegie Mellon University)
Pierre Moulin	(University of Illinois)
Joseph O'Sullivan	(Washington University, St. Louis)
Vincent Poor	(Princeton University)
William Ryan	(University of Arizona)
Paul Siegel	(UC San Diego)
Hemant Thapar	(Link-a-Media, Santa Clara, CA)
David Thompson	(Hitachi Global Storage, San Jose, CA)
Sergio Verdu	(Princeton University)
Raymond Yeung	(Chinese University of Hong Kong)
Jack Wolf	(UC San Diego)

Summary

In this memorandum, I documented my activities from the day I joined Harvard as an Assistant Professor. My priorities over the past six years have been to build a teaching and research program in Communications Engineering. Thereby, I have tried to focus on building a world-class program in Magnetic Recording and on branching out the research activities in (for me) new directions of communications theory that helped Harvard's communications program gain visibility. I believe that we have established a very reputable program in Signal Processing and Communications at Harvard. I have also introduced a new course in Communications and revamped an old introductory course in Electrical Engineering. I see my research goals in the future turning towards implementations and towards formations of a communications research center at Harvard that would further increase our visibility among educational institutions in the area of communications. With this I conclude the memorandum as I eagerly await the feedback on my performance as a faculty member.

PRODBEG = CMU_00091748
PRODEND = CMU_00091750
AUTHOR_ORG : Katie Gramling
DATE_LASTMOD = 03/21/2006
CUSTODIAN : Moura, Jose
SUBJECT : Nomination Form
DESIGNATION : Confidential
DOC_TYPE : doc
CDVOL : CMU_030110
SOURCE : \Moura2\inbox-past\in-2006\mar-apr06\NAEMoura.doc
DOC_TEXT : ï»¿Nomination Form PLEASE NOTE STATEMENT ON CONFIDENTIALITY OF ELECTIONSON PAGE 1 OF INSTRUCTIONS

This two-page form will provide the only information used in the final balloting of your candidate. This form MUST be complete and typewritten (10 pt. font or larger).

1. Candidate Name; type in UPPERCASE (PREFIX, FIRST NAME, MIDDLE INITIAL, LAST NAME):
JosÃ© M F Moura

2. Engineering Category (Peer Committee assignment): Electronics Engineering

3. Type of nomination: Member

4. Primary employment position and institution (include complete address, telephone number, and e-mail, if available):
Title: Professor
Institution: Carnegie Mellon University Address: Department of Electrical and Computer Eng.
Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh PA 15213-3890 USA

Phone: +1-412-268-6341 Email: moura@ece.cmu.edu

5. Date and place of birth (city, state, country): 01/09/1946 Beira, Mozambique

6. Citizenship (name of country; if naturalized, must give date-month, day, year): United States, March 26, 1998

7. Education (list all degrees earned, dates, major fields, and institutions; honorific degrees

Nomination Form

PLEASE NOTE STATEMENT ON CONFIDENTIALITY OF ELECTIONSON PAGE 1 OF INSTRUCTIONS

This two-page form will provide the only information used in the final balloting of your candidate. This form **MUST** be complete and typewritten (10 pt. font or larger).

1. Candidate Name; type in UPPERCASE (PREFIX, FIRST NAME, MIDDLE INITIAL, LAST NAME):
José M F Moura

2. Engineering Category (Peer Committee assignment): Electronics Engineering

3. Type of nomination: Member

4. Primary employment position and institution (include complete address, telephone number, and e-mail, if available):

Title: Professor
Institution: Carnegie Mellon University
Phone: +1-412-268-6341

Address: Department of Electrical and Computer Eng.
Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh PA 15213-3890 USA
Email: moura@ece.cmu.edu

5. Date and place of birth (city, state, country): 01/09/1946 Beira, Mozambique

6. Citizenship (name of country; if naturalized, must give date-month, day, year): United States, March 26, 1998

7. Education (list all degrees earned, dates, major fields, and institutions; honorific degrees should be included under Item 15 – Professional Recognition):
Licenciado in El. Eng. 1969 Instituto Superior Técnico (IST) (Tech. Univ.), Lisbon, Portugal
MSc El. Eng. 1973 MIT; Electrical Engineer 1973 MIT; Sc.D. El. Eng. & Comp. Science 1975 MIT
Agregação in Electrical Engineering Sciences, 1978, IST, Lisbon, Portugal

8. Registered Engineer (list state/s): Portugal

9. Candidates work sector during accomplishments cited in Item 10 (select one): Academe

10. Specific outstanding technical and professional engineering accomplishments and contributions meriting election to the NAE:
a. Moura led a multidisciplinary, multi University team that developed SPIRAL, an automatic generator of optimized performance-tuned libraries of signal processing algorithms: SW codes (C, Fortran), HW IP cores, and SW-HW co-designed implementations. SPIRAL is highly versatile in terms of: applications, performance criteria (runtime, area, power, accuracy), platforms (desktops, PDAs, FPGAs, vector/parallel machines, multicore chip sets, clusters).
b. Moura and his student A. Kavcic designed an optimal detector for high density magnetic recording systems, versions of which are used by all disk drive chip manufacturers today.
c. Moura and his post-doc, and now colleague, M.Pueschel, developed a new algebraic theory of signal processing (SP) that provides a new framework for SP that *inter alia* allowed the discovery of many new fast algorithms for SP one- and higher- dimensional linear transforms.
d. With his students, Moura pioneered image/video content-based codecs, precursors of MPEG IV.
e. With his students, Moura developed high quality automatic segmentation/registration algorithms for MRI cardiac and kidney bio-imaging.
f. Moura has taken major leadership roles in several IEEE Boards and Societies.

11. Impact of work (technical, commercial, national, etc.) cited in Item 10. Please provide specifics in Item 13:
a. SPIRAL is a new breed of "intelligent" compilers that overcome the increasing gap between HW promise (Moore's law) and what human experts can deliver. SPIRAL automatically generates SW and HW implementations of signal processing algorithms that match or exceed the best manually tuned implementations. SPIRAL has been licensed by INTEL to automatically generate kernels for the MKL and IPP libraries in their new advanced architectures. SPIRAL automatically generated the first fast FFT codes to run on IBM's massively parallel BlueGene/L supercomputer.
b. In the 80's, it was recognized that the dominant noise source in magnetic recording is media noise and attempts were made to design an 'optimal' detector. By modeling the media noise as a Markov autoregressive(AR) process, Kavcic and Moura designed an optimal adaptive sequence detector, which was surprisingly simple to implement. A decade later, all major read-channel manufacturers are producing chips with versions of this detector. This represents the single most important invention in the development of magnetic recording read-channel electronics in the 90's. Their AR media noise model is today the standard modeling tool adopted by the entire magnetic storage industry.
c. The algebraic theory of signal processing is a foundational breakthrough with many new developments, including new FFT-like algorithms for DCT/DST and higher dimensional transforms.
d. His content-based video codecs (3 patents) precede by several years the MPEG4 standard.
e. His bio-imaging work is licensed and incorporated in Siemens' MRI cardiac analysis tools.

12. Record of professional work experience (list principal positions and describe primary responsibilities for each; include dates):
1984- :Professor CMU; 1984-86 and 1999-00: Visiting Professor, MIT; 1975-84: Professor at IST.
2005- :Cofounder and co-director of CMU's Center for Sensed Critical Infrastructures Research
1995- : Advisor to Minister of Science and Technology of Portugal; coordinated the international evaluation panels of all EECS research centers and labs. in Portugal.

PRIVILEGED INFORMATION

Candidate Name: José M. F. Moura

2002 :Member of Committee for Baseline Assessment of the Public Research System in Ireland.
2000-01: Member, NAS/NAE & Under Sec. Def. Sc. & Tech. "Adv. Res. and Exper. in Sensing" Panel
1987-95 :Co-founder and director LASIP (Lab. Automated Signal and Inf. Processing)at CMU
1981-84 :Founder and coordinator of MSc.program in ECE at Inst. Superior Técnico, Portugal.
1979 :Co-founded INESC (Nat. Inst. Sys. and Comp. Eng., funded by Portugal Telecom)
1975 :Co-founded CAPS (Center for Analysis and Proc. Of Signals, University Res. Center)

13. Contribution of record. List up to five contributions (publications and/or developed innovations/products of proven success) of most significant impact. Quantify where possible.

1. M. Püschel, J. M. F. Moura, J. Johnson, D. Padua, M. Veloso, B. Singer, J. Xiong, F. Franchetti, A. Gacic, Y. Voronenko, K. Chen, R. Johnson, and N. Rizzolo "*SPIRAL: Code Generation for DSP Transforms*," IEEE Proceedings, pp. 232-275, February 2005. Invited paper, *Special Issue on Program generation, Optimization, and Platform Adaptation.*
2. A.Kavčić and J. M. F. Moura, "The Viterbi Algorithm and Markov Noise Memory," *IEEE Transactions on Information Theory*, pp. 291-301, January 2000; also US Patent 6,201,839 and US Patent 6,438,180.
3. M. Pueschel and J. M. F. Moura, "*The Algebraic Approach to the Discrete Cosine and Sine Transforms and their Fast Algorithms.*" SIAM Journal of Computing, pp. 1280-1316, March 2003.
4. Y. Sun, M. Joly, and J. M. F. Moura, "*Integrated Registration of Dynamic Renal Perfusion MR Images*," disclosure filed with Siemens Corporation Intellectual Property Department,, March 11/2004; world wide patent filed in The Netherlands, March 2004; also "*Contrast-Invariant Registration of Cardiac and Renal Magnetic Resonance Perfusion Images*," Y. Sun, Marie-Pierre Joly, and J. M. F. Moura, filed U, March 11/2005(Serial No. 11/078,035).
5. José M. F. Moura, "Passive Systems Theory with Narrowband Constraints: Part III – Spatial/ Temporal Diversity," IEEE Journal of Oceanic Engineering, pp.113-119, July 1979; also Parts II OE-4(1), Jan 79 and Part I (w/ Baggeroer), OE-3(1), Jan 78; included in IEEE collection of selected papers in time delay estimation.

Summary: Refereed Publications:	Jour.: 73 Conf.:213	Patents:	Issued: 6 Pending: 2	Books:	2	other contributions:	Ch.Books:9 Editorials: 7

14. Principal technical society membership (including grades and offices) and activities, and other noteworthy pertinent accomplishments in engineering, business, and public service (boards, committees, NRC activities, technical leadership in business, and contributions to society):

- 2006-07: President Elect IEEE Signal Processing Society (SPS); 08-09 :President of SPS
- 2006: Member, Comm. of Visitors, Computing and Communications Foundations. Divn., CISE, NSF.
- 2005- :External Adv. Board, European Network of Excellence on Wireless Sens. Networks
- 2004- :Editor, Book Series on Signal Processing, Morgan & Clayton.
- 1995-99 :Editor in Chief, IEEE Trans. Signal Processing.
- 1999-01 :VP, Publications, SP Society; VP, Publications, IEEE Sensors Council.

- 2002-04 :Chair, IEEE Transactions Comm.(composed of the 80+ IEEE Trans. Editors in Chief).
- 1999-03 :Advisory Bd.,MATHSTAR,a fabless semiconductor company focused on programmable logic
- 1999- :Editorial Board of IEEE Proceedings, ACM Trans. on Sensor Networks.
- Steering com: IEEE Int.Symp. Bioimaging and ACM/IEEE Int. Symp.Infm. Proc. Sens. Networks
-.

15. Professional recognition (honors, awards, prizes, honorary degrees, etc.; give name of institution and date):

1. President-Elect, IEEE Signal Processing Society,2006-07; President, IEEE SPS (2008-09)
2. Fellow, American Association for the Advancement of Science, 2005.
3. Fellow IEEE, 1994.
4. Corresponding Member, Academia das Ciências, Portugal, 1992.
5. IEEE Signal Processing Society Meritorious Service Award, 2003.
6. IEEE Third Millennium Medal.
7. Genrad Chair (visiting Associate Professor)at MIT, 1984-86.
8. Plenary lecturer/ invited panel member at numerous international conferences.

16. Proposed citation (not more than 20 carefully edited words that reflect specific documented contributions):
For fundamental contributions to the theory and practice of Statistical and Algebraic Signal and Image Processing.

Nominator verification and reference identification

17. I have obtained personally from each individual listed as a reference an indication of willingness to serve in that capacity.
(Note: Nominator must be an active NAE member. References may be active or emeriti NAE members. No more than two of the nominator and references can be from the candidate's institution. Neither members of the Committee on Membership nor Foreign Associates can serve as nominators or primary references.)

Nominated by: Kailath, Thomas

Important Note: By Checking the box below, you are providing your electronic signature on this nomination form.
☐ Acknowledged by Nominator

References (no more than four may be listed):

1. Mark Kryder, Seagate

PRIVILEGED INFORMATION

Candidate Name: José M. F. Moura

Signature	Name Affiliation
	2. David Kuck, Intel Fellow
Name (typed) Thomas Kailath	Name Affiliation
	3. Sanjoy Mitter, MIT
Affiliation and telephone number	Name Affiliation
Stanford Univ, 650-494-9401/650-799-9145	4. Vincent Poor, Princeton
Date	Name Affiliation

Supplementary References (Optional; ONLY for completion by a foreign associate for a foreign associate candidate) optional:

1.	2.
Name Affiliation	Name Affiliation

IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF PENNSYLVANIA

CARNEGIE MELLON UNIVERSITY, Plaintiff, v. MARVELL TECHNOLOGY GROUP, LTD., and MARVELL SEMICONDUCTOR, INC., Defendants.	)))))))))	Civil Action No. 2:09-cv-00290-NBF

DECLARATION OF JOSEPH MILOWIC III

I, Joseph Milowic III, declare as follows:

1. I am a partner at the law firm of Quinn Emanuel Urquhart & Sullivan, LLP, counsel of record for Marvell Technology Group, Ltd., and Marvell Semiconductor, Inc. (collectively, "Marvell") in the above-captioned case. I submit this Declaration in support of Marvell's Motion for Judgment as a Matter of Law on Laches, Memorandum of Law in Support, and Proposed Findings of Fact and Conclusions of Law in Support of the Motion for Judgment as a Matter of Law on Laches.

2. Attached to this Declaration are true and correct copies of the documents cited in the accompanying Motion for Judgment as a Matter of Law on Laches, Memorandum of Law in Support, and Proposed Findings of Fact and Conclusions of Law in Support of the Motion for Judgment as a Matter of Law on Laches, as follows:

Exhibit A	Document identified at Bates Label KAVCIC 002256, including cover page that I inserted to show metadata of last updated September 21, 2004
Exhibit B	Document identified at Bates Label CMU 00091748, including cover page that I inserted to show metadata of last updated March 21, 2006

I declare under the penalty of perjury under the laws of the United States of America that the foregoing is true and correct.

Dated: February 11, 2013

Respectfully submitted,

By: /s/ Joseph Milowic III

Joseph Milowic III (*pro hac vice*)
QUINN EMANUEL URQUHART & SULLIVAN, LLP
51 Madison Avenue, 22nd Floor
New York, New York 10010
(212) 849-7000 Telephone
(212) 849-7100 Facsimile
josephmilowic@quinnemanuel.com

Attorneys for Defendants,
Marvell Technology Group, Ltd. and
Marvell Semiconductor, Inc.

IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF PENNSYLVANIA

CARNEGIE MELLON UNIVERSITY,	)	
	)	
Plaintiff,	)	
v.	)	
	)	Civil Action No. 2:09-cv-00290-NBF
MARVELL TECHNOLOGY GROUP, LTD.,	)	
and MARVELL SEMICONDUCTOR, INC.,	)	Hon. Nora B. Fischer
	)	
Defendants.	)	

[PROPOSED] ORDER

AND NOW, this ____ day of ____________, 2013, upon consideration of Marvell's Motion for Judgment on Laches, it is hereby ORDERED that Marvell’s motion is GRANTED and CMU is barred from recovering any pre-suit damages in this case.

____

Hon. Judge Nora Barry Fischer