

The
Spirit
of
Wi-Fi 

Cees Links

Wi-Fi Pioneer

THE **Spirit** OF **Wi-Fi**

Where it came from, where it is today and where it is going.

BY

Cees Links

Wi-Fi Pioneer of the First Hour

T H E
Spirit
O F
Wi-Fi 

Copyright © Information Links BV 2019

Written by Cees Links 2003 & 2019

All rights reserved

Published with Qorvo in 2019

Utrecht, the Netherlands

No parts of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the copyright owner.

This book is subject to the condition that it shall not, by way of trade or otherwise, be lent, resold, hired out, or otherwise circulated without the publisher's prior consent in any form of binding or cover other than that in which it is published and without a similar condition including this condition being imposed on the subsequent purchaser. Under no circumstances may any part of this book be photocopied for resale.

Cover by Qorvo

ISBN 9789090324326

NUR 986

Printed 2019, Belgium

The Bluetooth® word mark and logos are registered trademarks owned by Bluetooth SIG, Inc. and any use of such marks by Information Links BV and Qorvo US, Inc. is under license. Other trademarks and trade names are those of their respective owners.

TABLE OF CONTENTS

0.	PREFACE	1
1.	PROLOGUE	4
2.	INTRODUCTION	6
3.	THE ROOTS OF WI-FI: WIRELESS LANS	10
	3.1 The product and application background: networking.....	10
	3.2 The business background: convergence of telecom and computers?.....	14
	3.3 The application background: computers and communications.....	19
4.	THE ORIGINAL IDEA (1987–1991)	23
	4.1 The radio legislation in the US.....	23
	4.2 Pioneers.....	25
	4.3 Product definition.....	26
	4.4 Early marketing.....	29
	4.5 Early sales efforts.....	30
	4.6 Security.....	33
	4.7 Health risks.....	35
	4.8 Developing new technology in a large organization.....	36
	4.9 Summarizing the first period (1987–1991).....	37
5.	THE SYSTEM AND THE WORLD MARKET (1991–1994)	39
	5.1 An incomplete product.....	39
	5.2 The market outside of the US.....	41
	5.3 Incompleteness on the marketing side.....	44
	5.4 Reaching into the wired world.....	46
	5.5 OEM marketing.....	49
	5.6 Getting the product right.....	51
	5.7 Outdoor products.....	53
	5.8 Related technologies creating opportunities and confusion.....	54
	5.9 Summarizing the second period (1991–1994).....	57
6.	THE INITIAL STANDARDIZATION (1995–1998)	59
	6.1 The early days of IEEE 802.11.....	59
	6.2 The original IEEE 802.11 MAC standard.....	61
	6.3 The standardization process in IEEE 802.11.....	64

TABLE OF CONTENTS

6.4	The original IEEE 802.11 PHY standard.....	66
6.5	The Intellectual Property ownership in standards.....	68
6.6	The forming of IEEE 802.11a and IEEE 802.11b.....	70
6.7	European standardization: HIPERLAN/1 and HIPERLAN/2.....	74
6.8	What about IEEE 802.11g.....	77
6.9	Bluetooth®.....	79
6.10	Summarizing the third period (1995–1998).....	82
7.	THE BREAKTHROUGH (1999–2000).....	84
7.1	Wireless local loop.....	84
7.2	The Apple story.....	85
7.3	Others PC vendors following.....	92
7.4	Networking companies catching on.....	93
7.5	WECA.....	96
8.	THE HAY DAYS AND BEYOND (2001–2002).....	98
8.1	Managing growth.....	98
8.2	New markets: residential gateways.....	100
8.3	New opportunities: hotspots.....	101
8.4	Social consequences.....	102
8.5	Over the top: trouble looming.....	103
8.6	The solution.....	104
9.	THE FUTURE OF NETWORKING.....	106
9.1	The telecom networking industry today.....	106
9.2	The consumer networking industry today.....	111
9.3	POTS versus VoIP.....	116
9.4	The future of networking.....	118
9.5	Wireless networking all over.....	122
9.6	An additional but different viewpoint: telematics.....	125
10.	EPILOGUE.....	126
11.	ACKNOWLEDGEMENTS.....	127

P R E F A C E

2019 is quite a special year. We are celebrating 20 years of Wi-Fi this year, at the same time that we applaud the launch of the Apple iBook with wireless LAN included – without doubt, the breakthrough in the market. The IEEE has bestowed the NCR team in Nieuwegein with the technology breakthrough award for WaveLAN. And I have received the lifetime achievement award from Wi-Fi NOW for industry contributions leading to Wi-Fi dating back to 1989.

Looking back, and after rereading my own book, I can accept this award thankfully – and with a certain degree of amusement. I think about how the first 10 years were a struggle. How people kept on believing that there was something cool about what would eventually be Wi-Fi. But also about all the times that we had to make a “return on investment calculation” on the benefits of Wi-Fi. And about how Wi-Fi moved from a “nice-to-have” feature into a “basic requirement.”

From a personal level, I would say: if you would have known everything in advance, would you have gone on this journey? Honestly speaking? Probably not. The challenges and objections caused a lot of agony and pain. But at the same time. Looking back, I realize that it was more than worth it. What can an engineer expect as a greater reward than the success of the product in the market? I am sure I am saying this on behalf of everyone who participated in this.

My 18-year-old son has never lived without Wi-Fi. (Interestingly, Wi-Fi is considered synonymous for “internet access” for him.) When he has friends coming over, the first request is always to get the Wi-Fi password. For the young generation, connectivity is the first thing that comes to their mind. If I try to explain to him how difficult it was to get the Wi-Fi concept “sold,” he is just confused. It just doesn’t add up for him.

But of course, in hindsight, we can see that Wi-Fi was just waiting for the applications to become available that would be recognized as valuable. Wi-Fi was just waiting for the new generation of applications that would leverage continuous connectivity. The applications in the last decade of the previous century were satisfied with casual connectivity. Downloading was the key word. This all has changed now. Downloading has been replaced with continuous connectivity.

In 2016 I attended a conference in Washington, D.C., where a speaker from the FCC expressed his amazement about how homogeneous Wi-Fi is worldwide

compared to 3G, LTE (4G) and 5G. He received a copy of my book, which includes the background and description of how this all came together. In this respect, it is good to mention that the FCC in the USA always had a sort of mindset to match the amount of unlicensed frequency bands available to the amount of licensed bands. So, when the amount of licensed spectrum was extended for cellular networks, the amount of unlicensed spectrum was extended as well.

I will not go into the billions of dollars of economic value created by Wi-Fi in the unlicensed band. These numbers are highly speculative. That about 65% of internet traffic goes over Wi-Fi, and 35% goes over the cellular network, is telling enough. Of course, sooner or later all wireless communication ends up on a wire, but the freedom inherent in constant connectivity anywhere and at any moment is the beauty for which wireless is so highly recognized.

A lot has happened in the 20 years of world Wi-Fi market adoption. Even in the early days, Wi-Fi was under serious attack from both 3G and Bluetooth. Both planned to make Wi-Fi redundant and eliminate it. They were not successful. There maybe a deeper reason why, on the contrary, Wi-Fi, Bluetooth and cellular all three made it to our smart phone. Each of the three technologies represent how we experience life in the physical world.

Wi-Fi is the technology for our dwelling places – our homes and our offices. It is as if we own the spectrum in our dwelling places, and it is available license-free. Bluetooth is the technology for our personal bubble – when we move around, we bring our headset, or our fitness band, that stay in continuous contact with our phone. It also feels like we own the spectrum in our personal bubble. And then there is our “outdoor” connectivity, where a third party is providing connectivity, and for which we sign up and receive a token, called a SIM-card. This technology coexistence wasn’t planned, but it all worked out. And in a way, it seems to make sense.

What amazes me daily is how “alive and kicking” Wi-Fi is. Here in 2019, in the midst of the launch of Wi-Fi 6, the data rate has reached a factor 1000-times higher from the product that we released with Apple in 1999. There is something funny about Wi-Fi-6, and it shows how marketing works. The reason for the number 6 is just to be “one up” compared to 5G. The reality is that Wi-Fi 6 is actually the seventh IEEE standard, after 802.11, 802.11a, 802.11b, 802.11g,

802.11n, 802.11ac, and now 802.11ax. But I suppose there are only a few of us who know and keep track!

What is clear is that there always has been some (healthy?) rivalry between Wi-Fi and the cellular standards – 3G, LTE (4G) and now 5G. Personally, I have always found the train an interesting location in this respect. Is the train a hotspot, like a store, or a café? Or is the train a mobile environment? It looks like the latter is winning, and most people are satisfied with their 4G connection in the train. Or maybe they're very dissatisfied with the performance of the Wi-Fi internet access in the train?

The last 20 years also saw the emergence of the IoT, with the requirement for low power standard. Bluetooth saw the emergence of Bluetooth Low Energy (LE), and Wi-Fi saw the emergence of Zigbee or Thread (actually the emergence of IEEE 802.15.4). Market dynamics have not brought Zigbee/Thread and Wi-Fi close together, but with the distributed architecture of Wi-Fi 6, there is a new chance. Only the future will tell.

It is clear that the original Wi-Fi standard and market breakthrough in 1999 was not the end of the development. It was more the start of a very dynamic period that has not come to a close yet. With the recent notice by the FCC for allocating more bandwidth for Wi-Fi in the 6 GHz, higher capacity for Wi-Fi is becoming possible – that is to say, providing more users simultaneously with higher bandwidth. With the standard developments in the 60 GHz under what was known as WiGig (.11ad and .11ay), but which is now part of Wi-Fi, there are even more opportunities for capacity increases.

Yes, Wi-Fi is alive and kicking, and I would like to commend the standard and everyone who participated – and still participates on its further developments – with the success that has come out of it. When we started the Wi-Fi journey, we had the impression we were onto something interesting that probably would build good products. We had no idea how pervasive it was going to be, and we can only be genuinely and honestly amazed about it.

It was a long journey, more than 30 years long. But from my perspective, the results definitely made it worthwhile.

8-Jul-2019

PROLOGUE

This book is about Wi-Fi, and where it suddenly came from. At least that is the way it looks: that it suddenly appeared. This is despite the fact that it had existed in the market in its earlier forms since 1990. But it went unnoticed until the email culture started to bloom, notebooks became popular, and many of the telecom service providers started talking about wireless mobile Internet.

This book is also about the differences between the telecom industry and the PC industry, their different backgrounds, and the extensive misunderstandings they had about the other's market over the last three decades. Because of these misunderstandings, billions of dollars (euros) have been wasted on acquisitions and divestitures and, finally, on something that was called 3G. Quite a lot of these investments are still on the balance sheets of wireless telecom companies and will have to be written off, sooner or later.

This book should be for generalists, not for specialists. About many of the subjects in this book, the Internet can provide thorough and neatly structured information. This is a book that runs you through the subjects and events inside out, as I personally have worked through them, and as I have enjoyed them. I have tried to keep the technical content to a minimum, and at the conceptual level, to make this book accessible to a larger audience that wants to understand the general concepts of new technology development and marketing through the example of wireless LANs.

This book should also provide insight on the Internet and what can be reasonably expected. However, one warning, if history teaches us anything about technology, it is that the ways of technology are unpredictable, and even improbable, in foresight – although usually easy explainable in hindsight. Wasn't the early notion about computers that the whole world would only need a few? Didn't the success of the Internet take us totally by surprise? And weren't cellular phones an unprecedented success in the market, creating new consumer fundamentals, within just a few years? So, a humble warning is in place: the world will look different in ten years, and we just don't know what those differences will be.

This book is not about taking the market in a few heroic swipes; it is about stamina, endurance and staying power. I think my management had given up on me at least five times in the last ten years, believing that this whole notion of wireless LAN would go nowhere. Well, they were almost right... But as we all know: almost right is completely wrong at the same time.

This book can be read from front to back, from back to front, or be flipped through from anywhere in the middle. The last way is the way that I read this type of books myself, so I understand and will not take it as an insult!

Let me thank my bosses over the years for giving me the freedom to get on this exciting venture. Let me also thank my colleagues in NCR, AT&T, Lucent Technologies and Agere Systems for having lived through this with me, the bad times and the good times. I would not know how I could have lived through this without you all. If I tried to list names here, I could fill the next ten pages and still leave out important people, so let me not do this. But know that I have written this book with you in mind!

I probably should also thank our competitors – this is probably a more difficult one. But even early in our marketing efforts it became clear that it is very difficult, if not impossible, to build up a market without competition. Competitors create a reality check for customers, as well as provide a credibility mark for a market.

Most of all, let me thank the customers, who were patient with us, particularly in those early years of development. They were patient with us as we got the product right, got the price right and got the business model right to enable the real breakthrough in the market. In that last respect I want to thank Apple Computers, who despite their questionable supplier tactics, helped to create this wireless market.

Maarsse
The Netherlands Fall 2002

INTRODUCTION

1990–2000 will probably go down in history as the decade where the industry transitioned from computing to networking. It was clearly the decade that the Internet came to fruition, it included the rapid growth of “networking” companies like Microsoft, Cisco and Intel, and it became the decade of the explosion of the cellular phone industry. It was also the decade in which wireless LANs were born and Wi-Fi (Wireless Fidelity) was launched. But I believe that at some point in the future, wireless LANs in combination with Internet will be recognized as the real breakthrough of the 1990’s.

To quote Bill Gates on this, “After we woke up and realized that all the dotcom-mania was just smoke and mirrors, we will look back at the late 1990’s and realize that wireless LANs, at least, were a real innovation from that period!”

The wireless LAN development that I was responsible for during a period of about 15 years started in 1987 at NCR Corporation, a computer company that still exists and focuses on applications for the retail markets (cash registers, point-of-sales terminals for department stores, mass merchandisers, etc.) and the financial markets (financial transaction terminals, check readers for banks and financial institutions). After NCR was acquired by AT&T in 1991, the wireless LAN developments continued and moved internally from NCR to the Network Systems Division at AT&T, which turned into Lucent Technologies, the luster networking products company, when AT&T split into three companies in 1997. The other two were AT&T (the telephone service organization as we know it today) and NCR (computers again). So, we were not part of NCR anymore; we became part of Lucent Technologies. Lucent Technologies itself split up in three companies in 2000: Lucent Technologies targeting communication technology and products for service providers, Avaya Communications targeting networking for enterprise and businesses, and Agere Systems, a semiconductor company. That the wireless LAN organization became part of the last one was not because of any specific strategy, it more or less just happened. And considering the success that came from it, it was probably the right choice, although that was not known at the time.

Throughout these 15 years, upper management in these companies really didn’t know what to do with wireless LANs, with my organization, or for that

matter, with me personally. There was a general notion that what we were trying to accomplish made sense, but I think they didn't really relate to it. That may be hard to believe now, given the number of people today who enjoy their Wi-Fi connections and can hardly imagine how life was without it. Throughout all these years, I saw it as a large part of my job to ensure that these corporate organizational changes were kept as far away as possible from the rest my organization, so they could just keep working. Only the logo on our business cards changed, and that was it! Being able to keep the team in-tact over this pioneering decade was probably the key accomplishment, and I wonder where we and Wi-Fi would be today without it.

As I mentioned, the wireless LAN division ended up as part of Agere Systems, and it grew into a very profitable \$380M business by 2001. In October 2001, and after many discussions, I split the total business into two roughly equal parts: a semiconductor business unit (developing chipsets), and an infrastructure end-product business unit (developing products: ORiNOCO). In 2002, Agere Systems divested from the infrastructure business unit when it was acquired by Proxim Corporation, and the semiconductor business unit has been integrated with other Agere Systems business units – at this time wireless LANs had become a sound and solid mainstream chipset business, and companies no longer could be vertically integrated from core technology chipsets to branded end-products.

Since 1988, I have been personally involved in driving the wireless LAN market and product development at NCR and its successors, and as I lucky to be able to keep all my business notes over this long period. Originally this was thanks to Time/System, a comprehensive paper organizer system and method, and later by keeping notes on my computer notebook (which I was used to carrying with me all the time, thanks to “Wireless!”) I thought that writing a book about this “real innovation” of the late 1990s might be fun and entertaining, but also useful from a variety of angles: what is it to develop a new and innovative technology and bring it to market, and what it means to do this as part of a large corporation – which is probably not that different than doing this as a start-up.

Making a new market is one of the most intriguing things in business. Although a lot is known about it, I think nobody fully comprehends it. Sometimes it succeeds, but sometimes it completely flops and is quickly forgotten, like the

pen tablet computers of the mid 1990s or a wireless data networking system like Metricom's Ricochet in early 2000. Making sure new technologies get adopted is crucial, particularly for tech businesses, since they thrive on new ideas and innovations, and failures can be very expensive. Over the years I have learnt a lot about the theories that are meant to make sure that new technology proceeds. Many of the theories make a lot of sense, are relevant and need to be applied. But in reality, as far as I can see, there is no golden bullet. Some ideas just catch on and other ideas just don't. The common theme is usefulness and convenience at a reasonable price, but this is not a guarantee, as not all good ideas at a reasonable price make it in the market. As with so many things, hindsight is pretty clear. It's trying to predict the future that is the tricky part. Still I believe over time, we will get better at it. As I read through my notes, I am simultaneously appalled and amused about the mistakes that I made, and I think this review would improve the chances of success a next time around.

However, of all the new business start-ups that I have seen succeed, there is a common theme – an idea, hard work, persistence, risk-taking and a never-give-up attitude. Resilience and the will to survive and succeed are the keys. This was certainly true with our efforts to create this new wireless market, since being part of a large corporation created a whole extra set of challenges. I remember a comment from Bob Holder, a Senior Vice President in AT&T, "I know that we are not very friendly for internal start-ups." In reality, "very unfriendly" might have been a better description. What I learned was that large corporations have a preference for quick, instantaneous success, just like many start-ups. They are continuously looking for the next \$100M or \$1B opportunity, without any appreciation for the \$1M and the \$10M phase one has to go through to get it right.

What makes it more challenging is that large companies are also very particular about the technologies they want to embrace, and the way they serve their markets. The marketing of a new technology can be out of alignment with the existing company model, and therefore it can create major havoc when trying to fit new products into existing sales and distribution models. Our wireless LAN technology did not really fit NCR Computers, because NCR was not really a networking company. And it didn't really fit AT&T. Bell Labs was used to lots of innovation but to not risk-taking. It also didn't really fit Lucent Technologies. Lucent was dial-tone networking and no data-networking. And finally, even

when wireless LANs had developed to a \$1B market, it only partially fit with Agere Systems (the semi-conductor part of the business). Consequently, Agere spun-off half the business to Proxim in 2002 (the products and systems part).

In the mean time we had gone through a lot of agonizing questions about “how to get it right,” something that even led to a serious attempt to take the business private via a management buyout early 1996. Unfortunately, the Lucent Technologies management at that time did not want to cooperate, and the business stayed part of Lucent, which was a great disappointment for the wireless LAN management team. The feeling with Lucent’s management was that there was “something” there, but they were never sure what.

This same Lucent management then tried to sell the business to various companies (3Com, Intel and Cisco in 1998 and 1999), but they were unsuccessful. At this point, companies (even these three that pride themselves in data networking leadership) had not really “seen the light” or fully underestimating the potential of what would become Wi-Fi technology.

In spite of these distracting and sometimes disturbing events, it was an exciting period, and I cannot describe the satisfaction when in 1999 we saw our revenue growing from \$30M the year before to over \$70M. This was nine years(!) after the first product release at Networkd in Dallas in 1990. This was clearly just the start. We saw the revenue grow further to \$210M in 2000, and, despite the big bust happening in the telecom industry, we continued growing to \$380M in 2001, with an estimated market share of 40 to 50%. This was the ultimate proof that the concept had been right from the early days, but it also showed how big of a challenge it was to get the product right and the market ready.

THE ROOTS OF WI-FI: WIRELESS LANS

This chapter is in three sections. First there is the product and application background of wireless Local Area Networking (LANs). Then there is the business background that, as far as I can see, is still not fully evaluated and understood. This section is about the war between the PC industry and the telecom industry, with “data networking” as the battle field. Interestingly enough, the consumer electronics industry has been recently pulled into this battle, which will probably give it a completely new direction.

The last section is a description of the breakthrough of the key applications that were necessary to make this technology come to fruition. After all, without any applications, (i.e., “what you can do with it”), hardware is pretty much a futile effort.

3.1 THE PRODUCT AND APPLICATION BACKGROUND: NETWORKING

To understand what wireless LANs are, it is probably good to understand what LANs are in general. To loosely define it, Local Area Networks are the way a set of computers are connected in a company or business. Connecting the computers allows them to communicate with each other – sending and receiving email, direct messaging, and sharing files, for example. Another reason for networking is to be able to share common and costly resources like printers, mass storage devices, or a connection to the Internet.

These LANs were more-or-less the natural successor in the 1980s of the (mini) computer networks. Up to the later 1980s, “computer networks” existed out of a large (mini)computer in the computer room, connected with a set of computer screens, the so called “dumb” terminals, and the (mini)computer in the computer room itself did all the work. In the 1980s we saw the dumb terminals being replaced by “smart” terminals, in essence personal computers, on which you could run local applications like word-processing (WordStar) or spreadsheets (VisiCalc). These smart terminals could usually operate as “dumb” terminals at the same time for the application that ran on the (mini)computer. This was the start of the migration to the network as we know it today. Now, to a large extent, the central (mini)computers are replaced with “servers,” in essence multiple large personal computers. The reality is also that the “dumb” terminal mode does not exist anymore on the personal computer; all the PC computing has become PC network computing.

Interestingly, a PC can operate when it is “on-line,” on the network, or “off-line,” when it is stand-alone and off the network. After a PC has been off-line and is then connected to the network again, it needs to synchronize with the other computers and servers that are on the network. This makes sure that all the functions that were suspended during off-line mode are caught up on-line.

To initially connect or to “network” PCs, these computers required a plug-in card, also called a Network Interface Card (NIC). A NIC had a connector, and a special cable was required to connect computers to other computers, usually via central boxes (ethernet hubs). In the very early networking days, there were two types of cards, and two popular systems were the de-facto standard: ARCNet and Omnet. However, the industry recognized the need for formal standardization, and the industry body that took ownership for this effort was the IEEE (Institute of Electrical and Electronics Engineers). Interestingly the IEEE is a standardization body based on a mixture of democratic majority and consensus, so it took quite a while and compromise/coalition-forming before these de-facto standards were replaced with three formal industry standards: Ethernet (IEEE 802.3), Token Bus (IEEE 802.4) and Token Ring (IEEE 802.5). Of these three, the Token Bus standard quickly “died” in the early 1990s, and Token Ring is today almost completely abandoned in favor of Ethernet, maybe with the exception of some small areas. The IEEE is important to mention here, as they would eventually play a crucial role in the definition of a wireless LAN standard.

Around the same time that networking started, another trend emerged – computers became smaller. When I studied at the Twente University in Enschede, the Netherlands, the computer room was incredibly large, housing a sizable DEC-1020 mainframe. The reason that this room was so large was the expectation that computers in the future would be even larger, and all that extra space would be needed. The reality of today is that the current computer room is about one tenth of its original size.

This trend toward even smaller computers continued during the 1990s, and PCs became so small that portable and/or mobile computing became popular. Originally nicknamed “luggables,” because of their hefty weight, these small computers jumpstarted the transition to real “personal computers” – computers that are not bound to a desk, but on the contrary are very personal to someone and go wherever s/he goes. I believe that today about half of the personal

computer market is represented by laptops, notebooks and/or any other form of mobile (and therefore truly personal) computing.

Interestingly, these two trends (computers networked via a LAN and computers becoming mobile) actually conflict with each other. The cable required for networking destroyed mobility. This conflict was clearly recognized in the late 1980s in our facility, where we did all kind of interesting networking “stuff” in NCR, but where we also were looking for a new product and market opportunity. We decided to explore the opportunity of a wireless LAN, a Network Interface Card with a radio transceiver and antenna at the end, instead of a cable transceiver and connector, as we felt that the two trends could converge as wireless networking. In those days, though, the market for portable computers was not very well developed, so new marketing ideas focused more on cable replacement than on mobility.

Initially we called the concept “Radio-LAN.” Actually, it was a “radio NIC”, a radio Network Interface Card, with an antenna connected to it that could be slid into a computer to provide a radio-based networking connection.

Originally this idea was ridiculed as “completely impossible.” It would be way too slow, way too sensitive to interference and data would be garbled, far too risky from a security perspective, and just too expensive. As Darrell Clark, Vice President of NCR Corporation put it at that time, “Companies have money, they can pay for wiring.” By the way, to put some date perspective on this, at this time electronic mail was only sporadically used by some U.S. west coast companies, and definitely not by NCR, which was an east coast company. This was also the time that NCR’s popular (mini)computer, the Tower series, was one of its main cash-producing product lines.

But to the credit of NCR, its President and CEO at that time was Chuck Exley, who later managed to sell NCR to AT&T, and he showed himself to be a visionary leader who not only supported the program, but was also successfully navigated company politics so that the development program could continue. It is also important to mention that Wiek Schellings, my General

Manager at that time, also really believed in this program and gave me the freedom to thoroughly work it out and bring the product to market.

As I mentioned, although the vision of wireless LANs involved networking for mobile computing, the positioning during the early years of the first products more about replacing the network cable with radio waves. The original theme for wireless LANs in the early 1990s was “cut the cable,” and the market positioning was about ease of installation, ease of relocation, and reducing total cost of ownership for the network. There were a variety of reasons for this “cut the cable” market positioning. One of NCR Corporation’s key markets was cash registers and Point-of-Sales terminals. Installing these terminals was quite an expensive effort, not so much because of the power cables, but because of the specialty network cables. Also, these network cables were error prone and sensitive to damage. Although it took quite some time, we have today seen a significant shift with retailers from wired to wireless terminals. And still today, NCR is selling wireless LANs as a cable replacement for their terminal installations, as do their competitors in this field, including IBM, Symbol Technologies, Intermec, LXE and Teklogix. These were niche markets that stayed niche markets for a long time, even when the product was considerably reduced in size. The consequence of this was that the “cut the cable” marketing campaign became a problem, as it was too hostile to the cable and networking industry, which made a lot of money on wiring and wiring products. We had to find a way to be complementary.

This original wireless LAN NIC was only the first phase of what is included in the concept of wireless LANs today, which are now so prevalent for company facilities that they are equivalent to cellular phones on the global scale.

By the early 1990s, we had already started to work on size and power reduction to make a wireless LAN card fit the size and power requirements of a notebook computer, our original goal. We also realized that we had to work on a system that would support “seamless roaming” through a larger office building, where the radio network fully takes care of the connectivity in every location in that building.

To that end, the concept of access points was developed. Actually, access points would become for the wireless LAN industry the equivalent of cellular base stations for the cellular phone industry, and a pretty sizable building can be completely covered (lit up) with a reasonable number of access points. One access point covers about 5,000 square feet (500 square meters) on one floor of a building. Access points can be connected together (networked again!) with

ethernet and software in a notebook computer. The access points take care of a seamless and unnoticeable handoff for the end-user, once he is getting out of range of one access point and into range of another.

3.2 THE BUSINESS BACKGROUND: CONVERGENCE OF TELECOM AND COMPUTERS?

This war actually started in the '70s with a first fence-off between AT&T and IBM. AT&T, before spinning out Lucent Technologies, was clearly the worldwide representative of everything that was telecommunication. It was a telephone service provide to millions, it owned a large part of the world's telecom infrastructure, it developed all communication switching products itself, it was a very prominent player in the telephone switching industry for enterprises, and on top of all that, it owned Bell Labs, its renowned research arm, that seemed to be able to generate an unlimited stream of new technologies. AT&T was synonymous with telecommunications.

In the same way, IBM was computing and computers. If there was any company that had managed to leverage the changeover of manual administration to electronic administration, it was IBM. They covered the world of computing from mainframe down to the smaller machines that were available at that time, including the hardware, the operating system, the programming software and compilers, the application software, etc. The slogan in the IT world was, "Nobody in IT will ever be fired over having selected IBM instead of one of its competitors." IBM was the safe choice. At NCR, we definitely feared IBM.

An interesting side note to mention here is the fact that IBM always had to deal with competition. On its main turf, large computers, companies like Control Data and Sperry were challengers; at the lower end it was originally DEC (later Digital); on the PC side it was Apple, and several others; for operating systems, it was Microsoft. IBM was an incredibly strong company. I remember the comments at NCR during the 1980s that IBM was growing every year by the size of a complete NCR. As NCR employees, we looked at each other in amazement, because we thought NCR was pretty big, and growing NCR was one of our bigger challenges.

The first battle in the 1970s started was started with the realization that telephone switching technology could be computerized, and that a telephone call could essentially be digitized and run over a computer. This led AT&T to

decide to go into “computing,” and IBM decided to go into “switching.” AT&T started its Computer Systems Division, and IBM bought Rolm, a switching division from Siemens AG. Apparently, Siemens had already decided not to compete any longer with AT&T. Cut to ten years later, and there was a lot of red ink flowing. IBM was no success in switching and divested from Rolm, and AT&T was no success in computing and was looking for ways to get out and keep its customer base without losing credibility. The result was the acquisition and the divestiture of NCR during the 1990s.

And even today, it looks quite logical to combine the two: voice networking and data networking. AT&T did quite well during this period. Their digital switches (based on computer technology) made us all say goodbye to rotary phones. However, owning and developing these computing capabilities did not make AT&T a computer company.

IBM, on the other hand, became a big networking company. In the early 1990s with Token Ring, IBM was actually the largest network vendor in the market, at least for enterprise networking. But apparently there was a big gap between enterprise computer networking and telephone networking, and the technology that AT&T owned was just too large, because IBM never made it there.

The underlying reason of why these “convergence” efforts, as they were called, never came to fruition was the fact that voice networking and data networking were too different. In the first place from a marketing perspective, the company decision makers for telephones were different people, usually a real estate department in the company, responsible for buildings, gas, water and electricity. The decision makers for computers were usually made in the IT department, where teams were overworked and in catch-up mode continuously. Despite that this changed in the late 1980s, when IT departments began to take control over the PBX and telephone equipment, this still did not create enough synergies for success.

Probably even more important were the very deep product differences between telephones and computers and the ways those two industries were originally built. Phones were essentially “dumb” and “cheap,” and all the “brains” of a telephone system were in PBX or in the switch. But at the same time, the trend in computing was to move away from dumb terminals and move into smart

PCs, as mentioned earlier. This created a large divergence in the industry. (Of course, the current trend in the cellular phone industry is “smart” phones, but more about that later.) There were other attempts. The telecommunications industry made a serious run on the general enterprise and home networking via ISDN (Integrated Services Digital Networking), which was developed in the 1970s. ISDN would bring integrated voice and data networking at a speed of 128 Kbit per second. At the time, this sounded attractively fast, as “dumb” terminals were connected to mainframes at much lower speeds.

Unfortunately for ISDN, these were also the years that the first rumblings could be heard about the privatization of the telecom industry, and different telecom companies began to jockey for position, along with some “loyal suppliers” (like an AT&T equipment division in the USA, and Siemens in Germany) who started to dream up proprietary ISDN implementations. The consequence was that the rollout of an ISDN standard definition was seriously delayed, and so different variations were implemented in various countries. The real momentum started only in the late 1980s and early 1990s, but by then, ISDN was already being overshadowed by much faster data networking technologies emerging from the computing world, like Token Ring and Ethernet. On top of it all, PC companies were not enthusiastic about ISDN because of its different variations around the globe. This made them look for alternatives, because if computer companies need one thing to proceed with the adoption of something new, it’s one worldwide standard. Despite all of this, serious efforts were made to create products that brought “ISDN to the desktop.” They all failed.

ISDN had another major drawback that related to the lack of understanding of the difference in architecture between voice networks and data networks. ISDN is switch-based technology, which means building a connection between two points – a sort of physical connection is built. Token Ring and Ethernet however are packet-based, which involves transmitting a packet from one address to another address. In this case, only a temporary (virtual) connection is created. Packet-based technology is significantly cheaper and more effective for data, where “quality” is secondary to speed. This made ISDN a losing proposition for the longer term, and when Token Ring and (especially) Ethernet developed next levels of speed, ISDN dropped out of the market before it had even reached a reasonable penetration.

Here is an often-overlooked example of how different telephones and computers are: compare the numeric keypads. A telephone starts with the “1” in the top left corner and works its way right and down. Not so with the computer (or even the pocket calculator) numeric keypad. There the “1” is to be found in the bottom left corner, from where it works its way right and up. There are logical and historical reasons for this. Telephones follow the normal Western world script style (i.e., left to right, top to bottom), but computers came from the original mechanical calculators, where the location of the numbers is ergonomically optimized for number usage. In accounting, the lower digits are used more often than the higher ones, so this layout places the higher numbers farther away. The result, though, is that I and almost everyone else have to always look at a numeric keypad to make sure to enter the right numbers – especially phone numbers. Surely this is the source of many errors in numeric key entries! There was a point that I seriously considered changing the numeric pad of my telephone to be the same as my calculator, by switching and re-soldering the wires in the telephone. Why aren’t there phones with computer-based numeric keypads? I think the best solution for this confusion would be that both computer vendors and telephone equipment vendors offer a choice between computing or telephone pads on their equipment.

On top of all of this, the real underlying factor is the difference in technology – or actually in the quality of the technology. I do not mean to say that the quality of data networking is low, but the reality is that the quality and reliability of switched voice networking is just a few factors higher.

Voice technology needs to be synchronous. Many of us probably remember the lack of “quality” that manifested itself in the delays in turnaround time that initially plagued transatlantic calling. Giving a response to a question of someone at the other side of the Atlantic quite regularly interfered and collided with an extra explanation of the question, and when both parties realized the situation, both parties would be quiet, and then to make things really awkward, both parties would start talking at the same time again, starting the communication difficulties all over again.

The underlying architecture and focus of research and development in the telephone industry is to make the communication experience instantaneously present. To be able to hear the subtleties in someone’s voice. This is realized

when the voice technology is characterized by very high-quality, instantaneously synchronous, and two-way (so you can interrupt each other!) networking. Interestingly enough, it actually does not need a lot of bandwidth. The reality of today is that the bandwidth of a voice call is 64 Kbit per second or less; it can be even as little as 16 Kbit per second.

Computer technology, however, does not look at networking this way at all. For computer technology, bandwidth is very important. Ethernet was offering 10 Mbit per second early in the 1990s, however with relatively low quality with regard to timeliness/synchronization. A few hundred milliseconds earlier or later did not really matter, so long as the data packet came over correctly, or could be requested to be retransmitted in case it was garbled. This is “low quality” compared to voice, but more than good enough for data. So, data networking needs “fat” pipes, the bigger the better, the faster the better. But the timeliness issue was more or less a “don’t care.”

Given this, it makes sense that the technologies have never merged. A compromise between “the high timeliness” (quality of service) of the voice networking and “the bigger the better” of the data networking would indeed compromise them. Either “quality” is too low to have a real intuitive voice communication, or it is too slow for a real larger data communication like a transfer of a large file. And if it could do both, it would be altogether too expensive.

The industry has tried to merge these needs. ATM (Asynchronous Transfer Mode) technology was a typical example of trying to find a compromise, but it became too expensive. There is today a fair amount of ATM technology implemented, including its many variants. But ATM’s goal to reach the desktop turned out to be a non-starter – this compromise was too expensive.

By the 1990s, IBM’s and AT&T’s endeavors to invade each other’s turf had unraveled. IBM had never had any real success or made any real money with Rolm, so they could divest themselves relatively easily. AT&T was in a more difficult spot. They had made progress with their computer division and won quite a lot of customers, thanks to their technology base and the AT&T brand name. Unfortunately, they were literally losing money on every computer sold. For them to get out of the computer business was not as simple as closing it down. They had their reputation to preserve.

By this time AT&T had already come under serious competitive pressure in their various divisions, not least from MCI (now WorldCom) in their core service business. So, Bob Allen, the CEO of AT&T at that time, came up with a smart way to resolve this. Instead of retrenching and divesting like IBM did, he opted for “fleeing forward,” and acquired an experienced computer company, NCR. I believe the deal was \$6B, in those days a pretty large number. His strategy originally was to have NCR take over and run AT&T Computer Systems under the name of AT&T GIS (Global Information Solutions). However, the reality was that AT&T started to run NCR, which only made the original AT&T problem worse. The newly integrated company started to lose even more money than AT&T Computer Division had been on its own. Eventually, AT&T divested itself from its computer division and put NCR back on the market, now for less than \$1B, a considerable cut from the \$6B that it had originally paid for NCR.

It is important to grasp this decade-long misunderstanding between the computer and the telecom worlds, in order to understand how separated the voice and the data industries already were in the wired world – and the billions of dollars that were lost because of this. In the wireless world this separation would be even sharper and even less understood, costing companies tens of billions of dollars, as we will see later.

3.3 THE APPLICATION BACKGROUND: COMPUTERS AND COMMUNICATIONS

As the word itself indicates, computers were originally designed for computing, and there are still a small number of computers today designed specifically for doing real computing. However, although most computers today can run some computing applications – popularly called “spreadsheets” – most computers are now developed for communications and entertainment. And whether we are talking about email, direct messaging, or writing documents, the communication (and therefore the networking) is essential. If the telecom had invented the computer, it would surely have been called the “communicator.”

I remember many conversations during the 1990s trying to come up with “the ‘killer’ application for wireless LAN.” Its equivalent in the telecom world was clearly the cordless phone. A computer was a piece of office equipment, why wouldn’t people want a cordless computer in the office? This was the original thought behind wireless LANs, and to be honest it did not go very far. Interestingly, cordless phones had become very popular in the consumer market but had

never made it from the home into the office. This was a signal; cordless phones for the office, based on the DECT standard never really got great acceptance. Even though there were some wireless DECT PBX phone installations and wireless LAN sales in the original years, they did not make even a little dent in the ever-growing volume of Ethernet and Token Ring networking.

Then the killer application showed up. It was the Internet, or more specifically, email. Not that anybody realized that this was the killer application initially. Email had existed since the early-and mid-1980s on mini computers, largely on internal corporate networks. But this communication application that would transform the computer into a communicator, sneaked in slowly but surely. And everyone who started using it realized that there would be no way back. Technology makes it ways along many unexpected roads.

The Internet and wireless LANs are not directly connected, but the usage of the Internet and email fiercely increased the need for networked computers, and for networking in general. The Internet stimulated the usage of notebook computers, and with that the need for wireless LANs to support the networking of notebook computers.

Email in itself is an interesting thing, and I have spent some serious time thinking about its meaning. Before the “email society,” there were generally two means of indirect communication: the memo (the “letter” for the romanticists) and the phone. The memo was very indirect and usually had a delay of one or more days. Of course, it had its advantages and, one might argue, it was a higher quality communication! No rash “reply-all,” but instead a thorough and considered form of exchanging viewpoints.

If there was something urgent, you picked up the phone and had a conversation. It may have interrupted the day or task of the person you called, and it was much more expensive than the memo. But compared to travel and a face-to-face meeting, it saved a lot of time. I think it’s fair to say that nobody developing a computer would have thought about a very popular application that would generate a new form of communication in between these two extremes. It would be fast and inexpensive, though still indirect. I would like to categorize “chatting” or “messaging” over the Internet in this same category, but maybe as a simplified form of email – you see someone online, you send a message or

two, fast and cheap, and not as interrupting as a telephone call.

A quick personal story on this... Before NCR, I had worked (until 1986) MAI Basic/Four out of Tustin, California. MAI Basic/Four was a mini-computer development and manufacturing company, and they had implemented a company-wide email system, even to the point that it pushed every employee (white and blue collar) to have an account, and to have access to a terminal to read and write email. This was a challenge for the factory workers, but it happened. When I moved on and joined NCR, another mini-computer company, email was an unknown concept for them. So now I had to “unlearn” email and start writing memos again. I cannot overstate how difficult and painful this was, as I was used to the instantaneous delivery and typically quick responses, which the memo system definitely did not offer. My desperation grew deeper when the Vice President of the Division, Darrell Clark, asked for ideas – he had some budget left for the year that he needed to spend. I pleaded that implementing an email system would be an excellent investment, but apparently, I was not convincing enough. He dismissed the idea, and it took NCR another three years before they adopted email.

The 1990s showed a few technology applications that together paved the way for the wireless LAN industry at the end of that decade. These applications were: computers becoming small and mobile, computers getting networked, and (finally) email and the Internet becoming pervasive applications. As I will discuss below, these developments created the environment for wireless LANs to break through, although it still required a catalytic event to make it really happen. Steve Jobs from Apple Computer, a person known for his technology views, started to promote wireless LANs.

Getting back to email, I have another interesting observation to share here. The buzzword in the industry in the mid-1990s was “the paperless office.” But things certainly did not look like they were going that way. On the contrary: it looked as if computers created more paper rather than less. There were piles of computer paper output that we had to read through, with maybe one (literally one) interesting bit of information to be found on each page. It certainly seemed like the computer generated more paper, and the paperless office was just a hoax. In retrospect, I think it was a transitional phase, and by the late 1990s, we saw a turning point. With the increased usage of notebooks, print-outs became less

necessary – why make a print-out if you carry the electronic version around anyway? Also, the flood of emails that ensued made it virtually impossible to effectively sustain regular print-outs. And with the reduction in print-outs came a reduction in photocopies. I think that by 2002, we made only a fraction of the number of paper copies compared to 1992.

This transitional phase did lead to some pretty hilarious, even appalling, situations. In 1998, my Lucent Technologies Vice President was clearly from the previous generation (“dial tone country and memo land”). She visited our office, and while she was very nice, email was just a bridge too far for her. She clearly thought that “typing was for secretaries” and managers need to manage. So, while visiting our facility, she had her secretary printing out her emails, and faxing them over to our facility, where my assistant sorted them and handed them to her. She then marked up with pen the faxed and printed emails and gave them back to my assistant, who then faxed them back to her own secretary, who took the comments and typed them as replies to the emails. The hilarious part was that at some point I had written an email to her during her visit, and that afternoon my assistant gave me her reply via the printed/faxed/sorted/typed version. And then an hour or so later, I got the email version of the printed/faxed/sorted/typed reply.

As I mentioned, this was 1998. Unfortunately, she was not the exception to the rule at Lucent. I think a lesson here is that new technologies do not only have technological requirements that should be satisfied. The social context also needs to be upgraded appropriately. In the example above, “keyboard fear” was a clear roadblock to using a computer, or even a notebook computer, to create some efficiency.

In the following chapters I will describe how wireless LANs moved through the 1990s from the early idea phase to a real product, and from its initial niche markets in retail and via education into the “mass” volume where it is today. I would estimate the number of wireless NIC cards being shipped at about 1 million per month. These shipments include standalone cards (PCI or PCMCIA-bus) sold through distribution to end-users, as well as integrated cards (mini-PCI) sold via the computer industry.

THE ORIGINAL IDEA (1987 – 1991)

Few people realize that Wi-Fi goes back to the late 1980s, that the technology really developed in the heart of the computer world, and that it had quite a rough time getting accepted in the industry. Everything around what Wi-Fi is today needed to be invented: from spectrum to standardization, from implementation to customer acceptance.

4.1 THE RADIO LEGISLATION IN THE US

It most likely will be impossible to trace back and identify the first person that created the idea of wireless LANs. I think the original concept came from an idea in the mid-1980s by the FCC (the Federal Communications Commission) in the USA to enable some frequencies for data communications in the 915 MHz, 2.4 GHz and 5.8 GHz bands, probably for warehouse data collection terminals. These three bands existed already for the so-called ISM usage, where the three letters stand for Industrial, Scientific and Medical. Examples of applications that have been and are running in this band are microwave ovens, garage door openers, and cordless telephones(!). This band was also used by a popular retail security system from a company called Sensormatic (now part of Tyco Corporation), that was running in the 915 MHz band. Considering that retail was one of the target markets, it is probably not surprising that there was immediately a serious interference concern. Fortunately, Sensormatic and we at NCR could quickly sort out how to use the 915 MHz spectrum (actually ranging from 902 MHz to 928 MHz), so mutual interference was avoided.

In the late 1990s Bluetooth also started to use the ISM band, in the 2.4 GHz band this time, causing interference. More about this later.

These bands were called “unlicensed,” since one did not require a license to have radio equipment using this band, as long as certain restrictions were met. The usage of these bands was restricted to low transmission power (less than 1 Watt), and in secondary status. This last requirement meant that the primary user of the band (amateur radio) could request to stop transmission in case they received harmful interference from a secondary user – something that I never ran into during my 15-years of work in the ISM band. However, this and the fact that the band was “unlicensed” (also meaning “free for all users”) had many major companies and service providers concerned in the early days, which limited the acceptance of the technology.

“Unlicensed” was also a major marketing concern. The typical customer question was, “What happens if I have setup and invested in a wireless LAN configuration, and a radio amateur asks me to close down my system?” The answer and assurances were usually not good enough, and the customer was left with FUD (the standard marketing term for Fear, Uncertainty and Doubt). Promises like, “We will immediately replace your wireless LAN with Ethernet” usually did not help and left customers wondering why they should bother going through all the hassle of wireless in the first place.

In our development we have always targeted low-power transmission and high receiver sensitivity. In better marketing terms, we focused on “being very good and sensitive listeners, so our cards do not have to shout so loudly or have to transmit a lot of power.” But at this stage in the selling cycle, this story amounted to not much more than an interesting fact.

This customer uncertainty would only fade away with enough success stories over the long term – and hoping that harmful interference of a primary user would not take place.

None of this meant that the FCC regulators had any particular product idea in sight, or a clear idea how such a product would have a broad benefit – or that it would even work. As happens often in this type of situation, product ideas started to pop-up in various places. Besides NCR’s location in Nieuwegein, Utrecht in the Netherlands, there were at least four other places in the late 1980s where work on wireless LANs was happening – probably all driven by the thought that cordless phones would lead to an equivalent of cordless computers.

By the way, the “cordless” term never really caught on, except perhaps in the UK, even though it was probably a better description of what we were doing than terms like “wireless” or “radio.” So, it’s probably not surprising that it usually required quite some time to explain what we were really trying to achieve. We called our project initially Radio-LAN, until we chose the product name WaveLAN though an internal contest that provided a free dinner for the two winning employees, Martin Jansen and Ton Wormgoor and their spouses.

4.2 PIONEERS

Of the companies and locations where wireless LAN ideas were initially developed, the Motorola location of Schaumburg, IL in the USA was probably one of the more important ones. If there was any company known for radio technology, it was Motorola. Tom Freeburg and his team had worked out and mastered an incredibly difficult technology, which ultimately proved to also be incredibly expensive. Interestingly, it was not based in any of the ISM bands, but instead on an 18 GHz radio frequency, using all kind of sophisticated algorithms to work properly. The second important one was relatively close by in Toronto, Canada, where a start-up called Telesystems tried an approach that was fairly common to ones tried in other locations, including us at NCR. The Telesystems approach based on the so-called ISM band in the 915 MHz led to a product called Arlan. On the West Coast, in Spokane, Washington, there was a company called “RLAN” (probably for Radio-LAN), and then on the East Coast there were two companies – LAWN and O’Neill Communications. These companies also used 915 MHz ISM technology and even tried some products in the market, but all dropped out, sooner or later. I think Proxim also existed at this time, but they were focusing on wireless meter reading. And then there were some efforts coming from Symbol Technologies (San Jose, CA) in the late 1980s. This was more or less the initial competitive playing field.

Looking back at what is left of these initial efforts by the large companies, Motorola dropped out quickly (in late 1992 or early 1993). Their 18 GHz product line, named Altair, turned out to be based on a too-complicated technology architecture, although it had a higher speed (3-4 Mb/s was advertised). Overall the product turned out to be way too expensive; most likely every product was sold at a loss.

NCR survived following a long and winding road via AT&T and Lucent Technologies into Agere Systems, where it split into a chip division and an infra-structure product division, that merged with Proxim. Proxim itself has lived through the years, facing good times and bad, but surviving and eventually merging in 2002 both with Western Multiplex and the infrastructure products division “ORiNOCO” that Agere divested itself from in that year.

Telesystems developed a very interesting history. It was originally acquired by Telxon in the mid-1990s, to allow Telxon to compete with Symbol

Technologies, who by that time had built up their internal radio knowledge and effectively competed with Telxon in the wireless data-collection market. But later in the 1990s, both Symbol Technologies and Telxon concluded, that wireless technology needed to be independently competitive in the market. Both companies had unsuccessfully tried to compete in the OEM market by selling their radio technology to third parties, but neither of them had made real progress. So, they tried to divest their wireless LAN divisions. Telxon was successful in initially putting Aironet in the market as a separate division in 1997 and then as completely independent in 1999. As of mid-2002, Symbol still owned their own wireless LAN division, probably simultaneously as an asset for the key knowledge it provides to their business, and as a liability because of the struggle to meet effective price points due to the relatively low volumes in their retail market segments. In early 2000, Aironet was taken off the market, as it was acquired by Cisco. Considering that they were a late 1980s start-up, this can be considered a great success story, especially given the future opportunities that Wi-Fi has in store.

4.3 PRODUCT DEFINITION

But let's not get ahead of ourselves here, as in the early 1990s the plans were big, but the challenges were probably even bigger. At NCR, we had the typical problems of a new start-up finding itself within the boundaries of a larger organization. In the first place the definition of the product in the context of NCR was a big problem, and even some fundamental technology choices had to be ironed out. My division happened to be part of the NCR Financial Systems Division.

Unfortunately, wireless LANs did not have much affinity with ATMs, check readers and sorters, and this sort of equipment. The other two important NCR divisions to mention were RSD (Retail Systems Division) and PCD (the PC Division.) The Retail Systems Division was the most interested in wireless LANs, despite the fact that they did not have a handheld or data collection product line like Symbol Technologies had. Their main interest was the ease of installation and relocation of cash registers or POSs (point of sale terminals) for department stores like JC Penney's and SEARS, or mass merchandisers like Wal-Mart. The PC Division was not interested at all. In the early 1990s, the PC industry was already as extremely competitive and fragmented as it is today. A relatively small player like NCR could not afford to spend the time and

bandwidth to bring wireless networking into the market. Despite the fact that we repeatedly approached this division, they were not really interested in even trying. They also thought the price of the product was prohibitive. Ultimately it would take many efforts over many years (until 1999, in fact), before the PC industry started to embrace wireless LAN technology.

In these early years at NCR, the relationship between the Financial Systems Division, that we were part of, and the Retail Systems Division, where the first applications would be possible, did not develop very well. It was a typical turf war between divisions in a large company, that for us started with different understandings of the technology and then moved from a factual phase into an emotional phase. The battlefield itself was not that interesting; it was the issue of what protocol would be used. It was Token Bus (RSD's preference) versus CSMA/CA (do not worry about what it means, but it was the name for our preference – Carrier Sense Multiple Access with Collision Avoidance). This difference in opinion was uncertainty about the actual performance of either protocol in real circumstances and a lack of facts to really support either view. But it blew so out of proportion that RSD decided to go their own path and subcontracted their development to a company called SSS at that time, later changing its name to Omnipoint.

But this was not the only product decision that needed to be made. At this time, the PC market was quite diverse in terms of hardware interfaces (“busses”) and software interfaces (“network operating systems”). Additionally, local area networks were starting to become commonplace, while discussions were still ongoing about whether “client-server” computing would be viable at all (compared to a “dumb” terminal-based architecture). Looking back, this may seem weird, but the reality at that time was that IT departments were struggling with the somewhat renegade concept of PCs in general – and probably rightly so. The management, maintenance and support costs of client-server-based systems quickly showed to be significant, offsetting the gains made because of the lower costs of the equipment.

As I mentioned, the PC market was still quite fragmented at this time, and the choices that needed to be made were far from obvious. Our original choices were AT (8-bit) for the bus and Novell Netware as the operating system – probably the right choices at the moment. Very quickly, however, we had to support AT (16-bit)

and Microchannel, the new hardware interface bus that was defined by IBM and that had some limited success for a while. At the operating system level, our initial choice was quickly challenged when Microsoft started their efforts to get into the network operating system market with “LAN Manager.”

But there was also another layer of difficult architectural choices to be made. For instance, should a radio-based network be built on a LAN architecture as was standardized by IEEE, or should it only be a transmitter and receiver connected to an Ethernet card itself? Or would it be better to base it on ISDN, or even just on DECT (the Digital European Cordless Standard)?

The new DECT standard was trying to harmonize the many proprietary cordless phone technologies that were in use around the globe, and at NCR we seriously looked at using this. It would have had a lot of advantages, including taking away the FUD about interference. In the end we rejected this as a solution for reasons of throughput. Combining all the available DECT voice channels would give us a product running at 1 Mb/s (Megabit per second), but at that moment not allowing any voice conversation. This was when we were looking at 3 channels, each of them running at 2 Mb/s.

Interestingly, Olivetti, at that time one of the larger IT equipment conglomerates located in Italy and one of the key promoters of DECT, went all the way to implementing a product that used wireless LANs running over DECT. Andrew Budd was the great inspirer of this project, and he was also involved in a variety of European standardization activities around it. The product however, just like the Motorola Altair product, proved to be very expensive and probably was not in the market for longer than a year.

And last but not least, there were these “other” technologies: implementing a LAN with infrared technology or via transmitting signals over the power cables. Both technologies had several runs at the market with the implementation of real products, but none were really successful. Of course, this is all hindsight. At the time that we made the choices, the total picture was showing many options.

During the first years of development, the number of people that worked on wireless LANs slowly grew to around 40 or 50 people by 1991. They were a mix of radio and digital hardware development engineers, manufacturing and

test system development engineers, and software and system test engineers. It was always amazing to see how many software engineers were required for what looked like a hardware product.

The target cost for the initial product was around \$300. We actually hit this target pretty well, after about a year of manufacturing. The initial products cost around \$600 for the initial low volumes. However, this was still significantly lower than Motorola's Altair product that we estimated to cost over \$1,000.

4.4 EARLY MARKETING

After the separation from the Retail division, we continued on our own path. We continued our development based on the CSMA/CA protocol, and this decision proved to be a good one. This protocol was later adopted, with many changes and upgrades, as the core of the 802.11 MAC protocol as we know it today. Retail's subcontracted development with SSS was a failure; there was a large discrepancy between what Retail thought that SSS could offer, and what they actually did. This meant that in the years after, Retail retraced their steps and followed our approach by using our products, although some special adaptations were required.

Actually, this temporary separation turned out to be a blessing in disguise, as it helped us to focus development on something concrete and start translating the wireless ideas into a tangible product. Our interest from the beginning had been in the horizontal market, the wireless notebook computer, instead of a vertical segment like retail.

So, we started with a lot of "horizontal" market research. I remember that we hired a bureau to do focus groups, and we watched from behind a two-way mirror as a group of about 10 IT employees discussed the wireless LAN concept. We did three of these focus groups in the US (New York, Chicago and Los Angeles), and we learned about what people liked and disliked in general. The outcome was generally positive, although the reactions varied. I think these sessions were useful, as they were also followed up with a quantitative study that included phone interviews of about 150 companies. In a third phase I personally had in-depth follow-up interviews with about 10 of these 150 companies, and so we got a picture what a wireless LAN product should look like from a user perspective.

The general opinion was that wiring in general was difficult and quite expensive – significantly more than one would expect at first sight. The actual numbers ranged from a \$200 to \$1,500 per “drop” (connection). Not too many buildings at that time had a structured wiring plan in place. Another problem was the expertise required for installing and maintaining a cabled network. In particular, relocations were a “royal pain,” as I remember one of the participants saying. It was clear that notebooks had not really entered the market, or the minds of people, and so we slowly moved from the original concept of mobility, to the concept of “cutting the cable,” and defining lower cost of installation and relocation as the key selling feature. We had to get back on the “mobility” track later.

The product launch created major excitement at Networld in September 1990 in Dallas. It also got us in major trouble, because despite the fact that we had working prototypes at the show, the productizing, radio certification and manufacturing startup took significantly longer than expected. Radio products are clearly a different “breed” than the digital products that we were used to at the time. In December 1990, we started to ship products in very limited quantities, but the general product release took until May 1991. The main reasons for the delay were around product stability, in particular getting the Novell drivers to pass certification, as well as the availability of the product diagnostics.

4.5 EARLY SALES EFFORTS

On the sales side, we came into some challenging situations as well, separate from the fact that we initially had only a limited number of products available for pilots. Customers generally tried to understand the technology, and were fascinated about its capabilities, but the benefits were seen as too marginal and the price as too high, so relatively few customers really started buying and implementing it. The launch price of the product probably was too high at \$1,390 per card, including the Novell Netware driver, as Novell Netware was the reigning NOS (Network Operating System) at that time. In these days, an ARCNet card was selling at \$300, an Ethernet card at \$495, and a Token Ring Card at \$645.

We quickly decided to lower the price to \$995. This was a step forward, but the number of products we sold stayed relatively low. The initial question was, “How do we sell the product?” NCR’s background was clearly in the direct sales environment – selling mini-computer systems to large accounts in horizontal and

vertical markets. That was not exactly the sales environment that was appropriate for something as small as a wireless LAN card. The product was offered to many retail customers, to many financial customers and to several governmental agencies. We did a lot of trials and testing with JC Penney's, with SEARS, and from the early days we could cite some success with Chemical Bank in Raleigh, NC and with Revenue Canada. Actually, Revenue Canada was our first large order – about 1,000 cards, quite a phenomenal size order at that time.

These days, however, NCR was also trying to move into the indirect distribution channel, working with distributors, system integrators and value-added resellers. We did do some initial business with companies like Computerland, Softsel, Ingram, and with Micro/D. (The last two merged into one company, Ingram Micro, and is still selling ORiNOCO, the successor of WaveLAN today.)

While this was a more appropriate distribution environment, selling a wireless LAN card in these early days was a complex sale, and despite many efforts we undertook on training and providing sales material, we had only limited success.

Interestingly, NCR also had an OEM salesforce that was related to their semiconductor division. They had relationships with many of the computer companies and networking companies, and with their help we made direct calls on quite a lot of them.

Several of them were quite interested – IBM, Apple, Compaq, Toshiba and Digital, to be specific. Unfortunately, that's as far as it went. What I learned from this is that here is quite a bit of herd mentality amongst PC vendors, and few are taking the step into something new, unless everybody else does as well. This is quite amusing, as many of them claim to be innovative and provide thought leadership. Clearly wireless LANs at this stage were a bridge too far.

The only exception I must make is for Apple. Unfortunately, Apple was already working with Motorola – though not with their Altair product line. They were trying to do something with a technology called FHSS (Frequency Hopping Spread Spectrum). Our product was using DSSS (Direct Sequence Spread Spectrum), and for one reason or another, Apple had an “internal hang-up” about DSSS and had declared that FHSS should be the technology of choice, because of its perceived better interference resistance.

This was the first skirmish of a technology battle that raged through the mid-1990s. This war between DSSS and FHSS delayed the uniform standardization and caused the market to stall for several years. Looking back, the FCC had done something very useful by allowing data communication in the ISM bands. However, by also allowing two radio technologies in the band, they had also put in there a poison pill that the industry failed to cope with for a period of at least five years.

In addition to PC companies, we also approached networking companies with our technology. 3Com was not interested and said to come back when there was an IEEE standard. Xircom, a networking company that was eventually completely taken over by Intel, seemed uninterested – until the next year, when they launched their own internal development program.

We had some success with Puredata, at that time a successful Ethernet card supplier that is still around as a communication company. It was a mixed bag, though, because once they were signed up, the first customer they went after was Revenue Canada, trying to “steal” them away from the direct sales channel, where we had an earlier win. This was our first exposure to what is known in distribution channel management as a “channel conflict.” It was an enlightening experience, but not a very pleasant one.

One definite surprise in the first year was the need for outdoor point-to-point connections. More than half(!) of the leads coming in, were for outdoor applications, despite the fact that the product was completely designed for indoor networking – in an office, a store, a warehouse or factory. But what customers really liked to use the product for was to connect the building on the other side of the street, or further downtown. Fortunately, our design was modular, and the omni-directional antenna that was specifically designed for indoor could easily be replaced with a directional antenna – and, if necessary, antenna boosters, whether meeting FCC requirements or not. Several VARs, Persoft and Hiperlink are two to mention, did a lot of pioneering work and started businesses based on this long-range value proposition and achieved amazing results. Anecdotal stories talked about 50-mile ranges in Kansas to connect buildings.

Most of the initial customer enthusiasm clearly came from this type of application, as it was not only significantly cheaper than leasing a T1 line, it was also under the customer's own control and quickly installed.

4.6 SECURITY

Another big marketing problem was (and still is today) the security of wireless networks: eavesdropping and/or breaking-in. This also added seriously to the FUD factor that I mentioned before. Radio waves go everywhere – even when the range is rather limited. And there is no way to avoid the fact that networking radio waves can be picked up outside of a building, for instance from the parking lot, allowing unauthorized people to listen in. At the same time, and from the same parking lot, radio signals can be transmitted into the building and into the data network that is used.

Now in general this was not as easy as it appeared. There were, and are, all kinds of security levels built in, ranging from physical security to network security, and from what radio technology was used at what frequency to password protection on network access. But as with all security measures, given enough time and money, they can be broken.

Initially we had identified a good solution for this. We provided as an option in WaveLAN a DES (Data Encryption Security) encryption chip. DES was a publicly known and described algorithm, and it looked like a reasonable choice. Unfortunately, it created a whole set of complications that we became painfully aware of later, when our legal department was informed that the implementation and distribution of this algorithm was controlled by the government via the NSA (National Security Agency).

During 1990 and 1991, I was told several times to come to Washington to the NSA and The Hague in the Netherlands, to the Dutch equivalent of the NSA (NBVB, the Netherlands Bureau for Connection Security) to explain what we are doing with DES, as well as with Spread Spectrum, the special radio technology that we used for the first out product and that was approved by the FCC for the 915 MHz band. Both technologies were considered under control of the NSA, and under regulation for export. The problems with Spread Spectrum were quickly resolved. We were using Direct Sequence Spread Spectrum with a low spreading code 11 bits per chip, so this was OK. I never got a real explanation

why this was OK, probably they realized that Spread Spectrum with such a low spreading code was very easy to unscramble.

We ran into many more difficulties with DES itself, the data encryption methodology that we used. The authorities seemed to be in a panic about it, though the reasons for that were very confusing. In the first place, the DES algorithm itself was fully published and available in any standard cryptography text book – so that was not new. Secondly, we implemented DES on a separate chip, developed by a Dutch chip development company in the Netherlands and manufactured in a foundry in France. So, I wondered why the NSA would even have jurisdiction over this. But they exercised their power as an American institution over the American firm (NCR) that we were part of, and they appeared to be very strict and nervous about this. This was actually when I learned the expression “red tape.” It was also the NSA that involved the NBVB in the Netherlands.

Anyhow, the bottom line was that we were allowed to ship this DES security feature with our product on the following conditions: it was free distribution in the United States and Canada, and it was a controlled export to only financial institutions (mainly banks) in other countries in the world, as long as these countries were not blacklisted (like Libya, North Korea or out South Yemen). The DES-chip also had to be on a socket on the board, it was not allowed to be soldered on directly. This last requirement was a very strange one. If the chip is soldered on the board, it is much more difficult to take it off and do improper things with it. Also, removing a soldered-on chip from a board by de-soldering it can easily lead to its destruction or at least reduction of its lifetime. So, let's make it easy and put it in a socket?

This controlled export forced us to implement a thorough administration, up to and including the end-user of the product. For non-financial customers in the non-US market we found a workaround. We developed a simpler security algorithm ourselves, based on DES, but with some reduced complexities. This algorithm had to be disclosed to the NSA for screening and approval, and after that we got it implemented – again by a Dutch chip development company, with manufacturing in France. For marketing reasons, we called the chip AES (Advanced Encryption Security), and we shipped it over several years. This all came to an end when the IEEE came up with an encryption standard as part of

the wireless LAN standard this was called WEP (Wiring Equivalent Protection). That was the moment that we de-released both DES and AES.

Even as of today, I do not understand the deeper meaning behind all of this. What I know is that Microsoft ran into similar problems when they wanted to freely ship DES encryption as part of a security feature in one of their earlier versions of the Windows Operating System.

Actually, just a few years ago, the restrictions on DES were largely lifted. But by that time, it was a moot point for us, because of the implementation of the IEEE WEP algorithm: RSA instead of DES.

4.7 HEALTH RISKS

The last FUD factor to be mentioned here was the health risk. These concerns came and went quite regularly, almost with a predictable repetition of once every other year. And all understandably, as our society is seeing a growing number of manmade radio emissions, in particular over the last decades. The real problem is also that there has been no research that really proves that radio waves are not dangerous. The usual situation is the other way around – research has not been able to show that radio waves are dangerous. This last statement is clearly not enough to be true assurance for customers.

What helped was being able to show that usually only one computer was transmitting at one time, and that the radiation of our product (around 100 mW) was significantly lower than the radiation transmitted for the early generations of cell phones (around 2-3 Watt), and also that the cell phone was kept very close to one's head, but that the antenna of a wireless LAN card was usually some distance away. This was a helpful explanation until there was a case of someone in the U.S. who got a brain tumor and sued the cellular phone company and the cellular phone equipment provider.

This event led a whole set of initial customers to start raising questions again, and the concept of “accumulated” radiation became a hot topic. What about the effects when one had multiples, for instance ten wireless LAN cards in a small area, what would be the impact? It got to the point where customers were asking for a guarantee or a health certificate – something that was not possible for us to provide. I think we definitely lost sales in these days. It was going to

take years to get the public to a level of comfort about radio waves at the low power levels that are used in the cellular phone industry and wireless LAN industry. Over the years, additional scientific information has been collected – none of it being able to show that there are negative effects from radio waves on the public health at the power levels that we were using.

4.8 DEVELOPING NEW TECHNOLOGY IN A LARGE ORGANIZATION

Actually, it was somewhat strange that NCR developed this new technology internally. This was clearly the old way, in which large corporations had internal “advance development” teams that tried to come up with ideas and then tried to translate them into profitable business. The more “modern” way was looking at companies who had launched something in the market and received a positive business response. Later in the 1990s, this difference in approach was very clear between Cisco and Lucent Technologies. Cisco grew through acquisitions of “proven” companies and technologies; Lucent Technologies originally tried the same through developing technologies in Bell Labs and bringing them to market. From the position of both companies today, it may be clear which of the two ways is preferable. But that may be a too-quick conclusion – besides their internal development, Lucent also acquired companies, with Ascend as probably the largest example.

Looking back, the Bell Labs model truly produced Nobel prize winners, and internally we had jokes about this, as it seemed this was creating PR value, but not necessary any revenue. The efficiency of a model where many small start-ups try new ideas and technologies is unmatched. Many of these start-ups die quickly if the idea does not have enough commercial merit. The few that survive are proven in the market and may have a high price tag – but this price tag may actually be lower than maintaining an advance development internal team. Although in the late 1990s, there was also a frenzy of buying start-up companies before any commercial viability was shown – just to make sure that the competition did not acquire this company first. Not surprisingly, then, the price of these companies was driven up considerably.

Wireless LANs were an internal NCR development, initiated by NCR’s Corporate Advance Development in Dayton, Ohio, and championed by Gary Spencer and Don Johnson. Don retired from his efforts in wireless LANs in 2000, working tirelessly until that time on standards and legislation. He certainly qualifies as one of the true wireless LAN pioneers of the very early days.

The Corporate Advance Development organization at NCR was not solely focused on maintaining their own momentum. Part of their model was to fund Advance Development programs in the different Business Units. The Business Unit in the Netherlands had a set of skills that seemed to match the original ideas of what would be necessary to successfully develop wireless LANs – a blend of knowledge in three areas: analog radio, digital and software. Plus, the team had a few really brilliant experienced engineers, who later have been recognized as key developers of the standard. In this respect I want to mention Albert Claessen for overall system knowledge, Bruce Tuch for radio technology, and Willem Diepstraten for protocol development.

But good experienced engineers were not the only reason for success. Wireless LANs were embraced at the top ranks of NCR. Chuck Exley and Tom Tang were both believers of the potential power of the technology. Despite the politics that developed around the subject, and its breakthrough being postponed for years, they continued to function as corporate sponsors – an absolute requirement for success. Without it, persistence at the bottom of an organization will be to no avail. Being located outside of the US headquarters turned into a positive thing here as well. Usually the “out of sight, out of mind” expression is used in a negative context, but in our case, it was definitely a positive.

Still my experience with headquarters in general is that their involvement tends to create more problems than they solve. The general “help” from headquarters can be very distracting, is usually out-of-touch from customers, and has a tendency to include personal “hobby horses.” The best advice I got from headquarters was, “do not listen to us; listen to the customer.” (But try telling that to your own headquarters boss!)

4.9 SUMMARIZING THE FIRST PERIOD (1987–1991)

Looking back on the first years of wireless LANs, from 1987 to 1991, it was a true pioneering time, as well as somewhat of a business and marketing disaster. The issues we faced were manifold.

The product technically worked fine but was essentially just a single card. The product line needed serious review to become a true system supporting a true mobile indoor environment. But that was the long-term strategic problem. Short term, we had to deal with serious marketing issues – the technology had

real market acceptance problems; security, interference and health created tremendous FUD. On top of that, we had pricing problems. The premium for wireless was recognized but was not \$500, so the \$995 price was too high, but we needed that price in order to make some margin on the product. At the same time, our manufacturing was not sorted out, as for political reasons we were forced to use expensive internal NCR manufacturing. Our distribution channels were up in arms because of channel conflicts, or they did not really understand the product, or were just ignoring it because a wireless sale was too complicated. It goes without saying that we did not make any money in these years.

But there were some positive sides. We had proven the basics of a completely new technology, we had started the standardization (more about this later), we got first-customer experience, including an understanding of what applications we could support. Plus, we had a clear concept and architecture for the product, which afterwards looked like having successfully walked through a large minefield of wrong decisions, compared to the traps that Motorola and Olivetti fell into.

We also had become absolutely convinced that with wireless LAN, we were onto something that we were convinced would have applications that would outrun our imagination, as soon as we could find “the tornado.” We were convinced that if we got it right, then the applications would generate the volume that would reduce the price of the card. Via this lower price, new applications would be enabled, further growing the volume and further reducing the price. The big question these days was how to get this momentum started.

Last but not least, by the end of 1991, we had built a team that was very committed to making the product line work. Actually, everyone who would form the key technical team, the key marketing team and the management team of the later 1990s was working on and committed to the WaveLAN product line at this time. Maybe that was the biggest asset to come out of this first pioneering phase. We got the idea that we were onto something big. We had a tiger by the tail, but the tiger was sleeping. Martin Bradley, our General Manager in these days, phrased it this way: “Bringing new technologies to market takes time, and whatever time you estimate it will take, it will take longer.”

THE SYSTEM AND THE WORLD MARKET (1991 – 1994)

It was clear at the end of 1991 that wireless LAN was an interesting concept, but that it also had some significant flaws that needed work. Also, the economic climate had become harsher, and rationalization of all activities was required. At the end of 1991 it was apparent that the days of “the big idea” were over and that persistence needed to set in.

These were also the days that AT&T Computer Systems fell into considerable trouble, so much so that AT&T decided to solve the problem by buying NCR and merging their organization Computer Systems together with NCR in an effort to gracefully exit their own computer business and bring their customers into NCR. However, because of the economic downturn, NCR was not that healthy either, and NCR management used the merger as an opportunity for a complete internal clean-up. Wireless LANs did not really come into focus, because we were still that small at that time. Otherwise I am sure the story would have ended here.

5.1 AN INCOMPLETE PRODUCT

In the course of 1991, it became clear that the technology was very interesting to everybody, but that the WaveLAN product was incomplete compared to the notion that people had about it.

On the technical side, WaveLAN was a wireless LAN card that eliminates the need for a cable to the central computer or server in the computer room. However, the computer room in many cases was too far away, even on another floor. We developed many ideas how to overcome this distance problem: repeaters, leaky cable, passive antenna amplifiers. However, none of them proved to be feasible, with the exception of the concept of the base station, which we internally baptized “Access Point.” Later we changed the name of the access point to “WavePOINT” for our WaveLAN product line, a name that came from Tino Scholman, one of our Product Managers, and that we all really liked. It said what it did – it provided a point of wireless access to the network.

But we immediately realized that we needed to have multiple access points to cover larger buildings, plus the capability of roaming (also called hand-

off) between these access points. Walking around with a notebook through a building should keep you connected, getting out of range of one access point and coming within range of another should not interrupt the communication session. Actually, the user should not even realize when he reaches the end of a cell, the range of an access point, or even where the access points were located.

This concept was easily described, and had it parallels in the cellular phone world, so it was relatively easily adopted. However, the implementation turned out to be completely different compared to the cellular phone world, and for good reasons. At that moment however, we had no clue that the migration from a wireless LAN card to a wireless LAN system was an effort of the same size and magnitude as developing the original wireless LAN card. To be frank, nobody in the industry had an idea of the magnitude; it was completely new territory. But I remember we started the effort with enthusiasm, as if the whole concept of wireless LANs was just invented yesterday.

Coming back to my earlier description about the difference between the telecom world and the computer world, here we ran into another interesting difference between the two. A cellular phone is a very smart piece of equipment, but in the networking sense, it is pretty “stupid.” When having a telephone conversation, the decision to hand-off from one base station to another base station is taken by the base station. This base station is in close contact with its “neighbors” to make sure this happens seamlessly. The amount of software technology this requires is quite complex, in particular to maintain the right voice quality in such a system.

In the data world, this concept where access points (base stations) keep track of the terminals and the signal quality of the communication with these terminals, would be unnecessarily complex. Data requires integrity. Quality in the sense of timeliness is not of significant importance. Therefore, the decision-making to disassociate from one access point and reconnect through another access point takes place in the terminal. A PC or a notebook is powerful enough to keep track of the available access points, as well as being able to make the right decision of when to switch and to what access point. Actually, this is a relatively simple process – each access point regularly sends out a beacon, and the PC keeps track of the beacons. If the quality (integrity) of the current connection is not good enough, it looks in the table for a better one. Having found such an access

point, it then creates a connection with that access point, telling the previous access point to ignore whatever (weak) signals it still may receive.

The implementation of this concept was relatively easy. The real challenge, as is usually the case in these systems, is size. In principle, one PC roaming from one access point to another access point worked fine. But the real need was to have a robust protocol that would allow 10,000 students roaming over a campus with a 1,000 access points, as we would quickly learn when we closed an agreement with the Carnegie Mellon Institute in Pittsburgh, Pennsylvania. Testing such a system required us working with other companies, as it was virtually impossible for a small organization like ours to size-up a testing program to such an enormous scale.

5.2 THE MARKET OUTSIDE OF THE US

Beside the lack of roaming, more limitations had shown up in the WaveLAN product – it could not sell in Europe and Japan. It was clearly targeted at the wireless LAN market in the US, as it used the 915 MHz band, a band that was only approved in the United States for this type of unlicensed data communication usage. Europe had reserved this band space for one version of its cellular system (GSM), so after some quick inventory, it was clear that there was no way that we could even think about exporting this to many other countries outside the US. But for computer companies it was very clear that wireless networking products only have any level of viability when they can be sold and shipped to all parts of the world.

There was an alternative – a frequency band that was available both in Europe and in the US in the 2.4 GHz for ISM applications. But there were many complications to overcome. People may wonder what took us all so long. Well, in Europe the 2.4 GHz ISM band was excluded from usage for (data) communication. Actually, the most important application in the 2.4 GHz ISM band was to use it for microwave ovens. I didn't know, but eventually found out, that the common microwaves in our kitchens and restaurants run at 2.4 GHz. Despite this fact that would cause us serious grief later, we energetically started to pave the way for wireless LAN applications in the 2.4 GHz band.

My vague notion was that we needed a “lobbyist,” someone who could work the authorities and find the ways through which we could establish WaveLAN in

Europe on a similar concept basis as in the US. We found our lobbyist through a conference that he presided over, Doug Postlethwaite. A difficult name to spell, perhaps, but with the absolute right skills to get the job going. We hired him on a retainer basis and developed a plan to pursue the opportunity along two lines. One line was the direct approach, asking for approval. The second line was indirect by publishing a set of articles on the concept of wireless Ethernet, creating the image of how it had established itself in the United States, including positive quotes from companies that had started to use the WaveLAN product there. Then we used these articles to further support the direct approach that proved that this technology needed to find its way into Europe.

When this all started to work, NCR Japan picked up on these efforts and started to replicate them in Japan. Although these activities were somewhat later, the reality was that a mostly frontal worldwide attack took place in 1992 and 1993 on using the 2.4 GHz for data communication applications.

Although there were other companies in the market, in particular Proxim and Aironet (now Cisco), I am not aware of any serious efforts from them to develop the market outside the US, but this could have been the case. Asking them would probably lead to a positive answer; success has many fathers. And a success it became, first in the UK, Australia, Norway, Sweden, Japan, the Netherlands, and Germany. And then, slowly but surely, the worldwide acceptance started to roll in, even before any official product was available.

Not that the situation in Europe or the rest of the world was as trivial as it was in the US. What is called today conveniently the “2.4 GHz” turned out to be a patchwork of different rules and regulations in different countries. For instance, the allocated band in the US was from 2400 to 2473.5 MHz, in Australia it was from 2400 to 2450 MHz, in Japan it was from 2474 to 2500 MHz, and in Europe different countries had different rules. So, there was France that allowed the usage of the 2.4 GHz only in about 70 of its largest cities, as the French army was using the band for military purposes as well. Italy even brought the bureaucracy a step further – although the 2.4 GHz band was license free, when it was used for data communication a license fee was required to the government.

Today the bands are pretty much harmonized across the world, and so are the products. But in the original 2.4 GHz product definitions, we had several “flavors”

for the different countries and regions in the world, with different software tuning the radio to the required channels in the band.

Interestingly enough, in efforts to open up the 2.4 GHz band in Europe, we got a lot of hostile opposition, if not plain sabotage, from Olivetti. I still vividly remember the negative and defensive contributions to the discussion from Andrew Budd, Director at Olivetti. Olivetti's development was based on DECT and they clearly understood that ISM-based wireless LANs would mean the end of DECT-based products, their vested interest. In retrospect, the early license fee required in Italy could have had something to do with this opposition. However, when Olivetti withdrew from the market, the license fee requirement for the 2.4 GHz very quickly disappeared as well.

Our initial strategy, by the way, was to continue with the 915 MHz products in the US, and the 2.4 GHz would be marketed in the international markets. The reason was very simple – the 915 MHz product had a better range than the 2.4 GHz and also penetrated walls better. This is just based on the laws of nature: the higher the frequency the quicker the wave dampening, and therefore these limitations. The advantage of the 2.4 GHz was clearly the availability of more channels, three instead of only one channel in the 915 MHz.

Within the company, there were forces that stated that striving for one product worldwide (2.4 GHz) would be preferable, but as the percentage of travelling wireless LAN users was still very minimal, the business case for selling 2.4 GHz in the United States did not seem very compelling. On top of that, while the functionality is comparable, the cost of a 2.4 GHz product is higher than a 915 MHz product. So, it would be difficult, if not impossible, to ask a higher price, and therefore this would shrink our margins. But this is where all the logic ended.

Other companies that had no product in the 915 MHz band started to market products in the 2.4 GHz band, and with the story of more channels they gave the market the impression that 2.4 GHz was faster than 915 MHz, which was largely false and would only be visible with intensive usage. So instead of balking over a move from 915 MHz to 2.4 GHz, the market started to ask for a 2.4 GHz product. The range reduction was taken, and probably not really noticed because of the relatively low market penetration of 915 MHz products, and the illusion of higher speed overrode everything. The bottom line is we avoided an

expensive marketing campaign to create this migration from 915 MHz to 2.4 GHz, and in 1994 the shipments of 915 GHz had dwindled.

But in 1991 and 1992 the prospects for being able to reach any profitable business case were pretty dim. This was the time when the cost of a 2.4 GHz radio product was estimated to be higher than \$350, when the business case most likely needed a cost of product below \$250, or maybe even below \$200. This meant that we immediately started with product cost reduction activities via integration of many circuits into larger single chips. But this was still early, and for years to come, product cost was going to be a major issue.

5.3 INCOMPLETENESS ON THE MARKETING SIDE

Also, on the marketing side, the WaveLAN product was incomplete. The pricing was a major issue; our margins were horrible, usually below 10%, where at least 30% was required. And on top of that, the outlook was not positive, as going to 2.4 GHz, which was clearly required, would only drive up the cost. Higher frequency components for higher frequency products are just more expensive.

At the same time, a lot of the customer feedback indicated that price was a major obstacle. There were also clear signals about the price elasticity – a price reduction always immediately showed an increase in volume. These volumes initially were quite low. Usually the deals we closed were around 10 to 20 cards, say around \$10K orders. Sometimes opportunities came up for 1,000 WaveLAN cards. These deals got a lot of attention, but usually once they came through, the quantity had been lowered significantly to something like 200 cards.

The customers that we engaged with on these large deals were typically the larger US companies. I remember we tried to close a deal with American Airlines for a pre-configured travel agency computer system. The main reason for American was that about 40% of new travel agencies had stopped within one year and wiring a travel agency was usually a big write-off.

Based on NCR's market positioning we did a lot of work with retail customers like JCPenney, Littlewoods, Younkers, House of Frazier, Victoria's Secret, Stop 'n Shop, and Wal-Mart. The results were mixed at best. JCPenney was probably the most successful – we had three of their department stores with on average six Access Points and eighty point of sale terminals. The good news

was that the system not only was working, but also was failproof to support larger installations.

I remember being pretty proud of this result. I went to one of the stores and paid over a wireless LAN, taking some impertinent views behind the terminal, making the sales assistant somewhat unsure about my intentions, which I shyly decided not to share!

We were more successful initially with banks, another market segment that NCR was quite strongly represented in. During the early years we worked with many smaller banks like Chemical Bank and HighPoint Bank. The main interesting feature was called “replicated branch.” The central staging of a bank branch computer system could be done centrally and tested. Then overnight the total system was trucked to the branch and installed before the employees came in the next day. No disruption, and up-and-running instantaneously. Wherever we could sell this benefit it was a success. But in general banks turned out to be very concerned about security, despite the fact that we supported DES encryption.

Interestingly enough we also were quite successful in Mexico. I’m still not sure why that was. Probably the presence of the NCR salesforce played a positive role, and probably the publicity around NAFTA gave the Mexico sales positive support. It stopped immediately when Mexico fell into a financial crisis in 1995, but until then, we had been quite successful with larger sales to companies like Bancomext, Operadora Vips, and Bancaser.

We thought a lot about the successes in Mexico, and why the US was so much more difficult. One concrete thought was the fact that networking cabling in the US was an established and profitable business, so wireless was a threat to this business. In Mexico the networking, and therefore the cabling installation market, was significantly less developed, creating a much more open mind for newer technologies. Anyhow, after 1995 this all had become a moot point; the Mexican market was gone.

One of the markets where we seemed to have developed traction was healthcare, in particular reaching into hospitals. Despite the fact that there was serious concern about 2.4 GHz radios interfering with other medical equipment, the reality was that hospitals were quite behind in automation. Specifically, the

data collection part was rather primitive – hand-written notes at the bedside instead of collected by a central hospital database, which would then be retrievable from wherever in the hospital the patient would be. But even in the hospital world, good applications that would be conducive to wireless LANs were missing. Hospitals were very much based on mainframe machines. So, while we were trying to move hospitals from wired to wireless networking, they still had to make the first steps into networking. Interestingly enough, the successes that we achieved over time created quite some excitement, but not necessarily enough follow-up. We worked with large companies targeting PoC solutions like IBAX and Baxter, but progress was very slow.

Besides direct sales – NCR’s sales teams directly selling to the large retail and financials customers – we worked hard on developing a so-called indirect distribution channel, a network of distributors and Value-Added Resellers (VARs) who function as an intermediate between the product brand and the end customers. This was when all the large computer companies were looking for ways to reduce their expenses, especially in the sales area.

One particular tool that we used in the indirect distribution channel was the so-called sales kit. This was the time to pack two WaveLAN cards and an Access Point combined with a Novell network operating system in one package to sell this as a “Network in a Box.” This was supposed to make the installation very easy, but in reality, these were still the days that networking was not yet to any level of maturity to support these types of effort. Actually, these were still very much the days that Microsoft LanManager and Novell Netware, the market leader at that moment, were in a bitter battle for leadership in this market. No surprise that the simplification in the networking only came when this battle was decided.

5.4 REACHING INTO THE WIRED WORLD

In this respect, these were also the days that there was another big industry battle going on – the fight between Token Ring (IEEE 802.5) and Ethernet (IEEE 802.3). Token Ring was a deterministic protocol running at 4 Mb/s; Ethernet was a stochastic protocol running at 10 Mb/s. Initially one could wonder why 10 Mb/s was not an easy winner over 4 Mb/s. But the answer is simply that a stochastic 10 Mb/s may suffer throughput degradation quickly with many computers connected. In other words, in practice the 10 Mb/s was significantly

lower. But the other reason might simply be that Token Ring had the stamp of approval of IBM. There was very much a fight going on between IBM and the rest of the computer industry. The IEEE committee with their somewhat democratic procedures did not want to just adopt IBM's proposals for wired networking, while IBM did not want their proposals to go by without standardization. The compromise was that IEEE created a separate standardization committee for IBM outside of IEEE 802.3, and this became IEEE 802.5. This is all important, because in the wireless standardization (IEEE 802.11) something similar was going to happen.

Initially IBM could sell Token Ring well against the higher speed of Ethernet. The reason was that they could convince the customer base that IBM's Token Ring protocol was more efficient than Ethernet. And the choice for IBM was the safe choice, wasn't it?

This fight went on for almost a decade, mainly because IBM was so strong in the industry that many computer companies were forced to develop Token Ring solutions as well. They needed to be able to interface with other IBM Token Ring-based solutions. Where IBM had the volume, they could produce these cards with a profit, while most other companies lost money on this.

In the second part of the decade, this battle was decisively won by "the rest of the computer industry." The main reason was that the cabling of Ethernet significantly improved. Ethernet originally required coax cabling, but later it also support shielded twisted pair, and then later, unshielded twisted pair. At about the same time that Token Ring technology moved from 4 Mb/s to 16 Mb/s, Ethernet moved from 10 Mb/s to 100 Mb/s. That last fact really did it.

Ethernet had already been consistently less expensive than Token Ring, so when it was also easier to install and lower in cost, the war was over, and IBM slowly but surely had to accept Ethernet in their portfolio.

This battle in the wired world had impact on the development of the wireless market for two reasons. In the first place, the drive for higher speeds made it more difficult to get wireless accepted. When we started with wireless, and we found ways to get to 2 Mb/s through the air, the market was largely dominated by Omnet (1 Mb/s) and ARCNet (2 Mb/s). But when we came with

our wireless product WaveLAN to the market at 2 Mb/s, the market interest had already moved to Token Ring (4 Mb/s) and Ethernet (10 Mb/s), indicating that “networking marketing” was going to be a speed game. I say, “networking marketing,” because the applications that were using the network had to deal with the real speed, and that usually was significantly lower – sometimes even between 50 – 80% lower!

The most striking example was the Xircom Ethernet port adaptor. This was a very strong product, as it allowed computers to relatively easily get on the network by connecting this port adaptor to the Centronics port of the computer, creating the illusion of a 10 Mb/s network connection, where a Centronics port had a maximum throughput of 700 Kb/s. In my discussions with Durk Gates, the CEO of Xircom at that time, he carefully circumvented (“xircomvented” we called it) answering the question about speed. It was a great example for me about brilliant marketing (or shrewd customer deception), but it worked, and Xircom was a successful company until the day they were taken over by Intel.

So, the wired market moving up in speed was one thing that made it more difficult to get lower-speed wireless products accepted in the market. The other complication it created was the integration between wireless and wired networks via the described access points. To integrate in a wired environment, we had to bridge back to both Ethernet networks (IEEE 802.3) and Token Ring networks (IEEE 802.5). The technologies are very different, and therefore the bridge products required were very different, which increased the amount of development work and reduced the chance to turn to profitability quicker. The integration with Ethernet was the easier part. The main reason was that the protocol used for WaveLAN, our wireless product, was quite comparable with Ethernet. WavePOINT, our access point, was then also a wireless Ethernet bridge. For our Token Ring bridging, we created a relationship with a company called Persoft, based in Madison, Wisconsin. They provided our Token Ring bridge solution for as long as Token Ring was around.

But overall the integration into existing wired environments was painful and time consuming, and the fact that there was a split in the wired world had a complicating factor for the emerging wireless world.

5.5 OEM MARKETING

Another way to get a product sold is via so-called OEM programs. The idea of OEM-ing a product is to sell the product under the brand name of someone else, someone who did not make the investment to develop a technology. “OEM-ing” is very common in the PC industry, even to the point that most PC companies today are more “integrators” than technology developing companies. In this respect it is interesting to mention that Steve Jobs, the CEO of Apple Computers, had the view that there are only two computer companies: Apple and Wintel, Wintel being a combination of Windows of Microsoft and Intel. In his opinion all other PC vendors are just “distributors” of Microsoft’s and Intel’s technology, or OEMs, as the term that is often used here.

Under NCR and AT&T we had developed wireless LAN technology, but getting this technology into the market was not very easy under these brand names, as the real names in the networking industry these days were IBM, 3Com and DEC. So, we developed the strategy to try to sell our products under other brands. We unsuccessfully approached IBM and 3Com. IBM either had no clue about how to use this technology, or they were doing something themselves, deep down in some of their own divisions. Considering their size, both could actually be true. 3Com had an explicit strategy to only support standards and not to bother with anything new that was not standard. Wireless LANs at that time were far away from a standard, and therefore were ignored by 3Com.

Not so Digital Equipment Corporation. For one reason or another, wireless had caught the attention of Ken Olsen, the CEO of DEC, who was very interested in technology. Via his lieutenants, we were contacted for cooperation and this turned into a success that brought real products into the market. Actually, DEC became a serious competitor to AT&T, and we had some very serious channel conflicts over the years as a consequence of this OEM strategy.

But these companies were not the only ones we approached. Virtually every company in the IT industry was at least approached, from Europe to Japan and to the US. Nokia, Sony, Samsung, Toshiba, Dell, Siemens, etc. I still have the notes from most of those efforts. It was just amazing to see how little interest and response there was. Usually, if a company responded with sending someone to the IEEE 802.11 meeting, this was big progress.

I need to mention something special on Dell and Toshiba. After Apple Computer did a major launch with wireless in 1999, I received an email in my inbox that happened to have in its tail an original discussion-generating email from Michael Dell, who furiously expressed himself about the fact that Apple had beat them with wireless LANs. Going through my notes, I found out that I personally had called Michael Dell in 1992 to propose cooperation, but even more interesting, in 1993 Dell had tested some wireless LAN products internally. As they stated, they had been playing with it, but they were not convinced that there was a real market for this type of technology. Probably rightfully so, as the price in these days was a major stumbling block; however, also in that period, I think there might have been good premium applications that would have justified the investment.

We also were in very close contact over the years with Toshiba. Actually, Toshiba and NCR had initially been working closely together on developing some portions of the MAC protocol. Toshiba has been a firm believer in wireless connectivity since the early 1990s. However, they were also very aggressive in the price points that they wanted to achieve, maybe just too unrealistic. That probably drove them into the arms of Bluetooth in the later 1990s as well, but in this period, they were eyeing wireless LANs, trying to integrate this in both their computer and notebook portfolios.

I still pleasantly remember my first true “notebook” computer. It was the famous Toshiba one, with the half-screen. I used it for several years, and even though it was just the early 1990s, I personally got very committed to notebook computing. Actually, in 1996 I got rid of my desktop PC altogether and run on a wireless notebook only – but that should be no big deal, understanding the natural marriage between notebooks and wireless networking.

In retrospect, in these days we learned a lot about channel marketing and distribution channels via our OEM strategy. But with new technology, one is very much in the situation best described as, “damned if you do, damned if you don’t.” Executing an OEM-strategy created a lot of channel conflict, where the AT&T WaveLAN product line, and later the Lucent Technologies product line, competed with DEC or its successors Cabletron and Enterasys. But not doing it would reduce the number of distribution channels to market, as only the channels owned by AT&T were too limited.

DEC (now Enterasys) was one of the more successful brands under which we were able to have our technology reach into the market; others were Puredata, Solectek and NEC. Interestingly enough, these companies (with exception of Puredata) are still the companies today that are better positioned in this market than others. This leads to the conclusion that even in the case of not owning the technology, an OEM strategy might be useful for quick learning and building market share.

5.6 GETTING THE PRODUCT RIGHT

We did a tremendous amount of analyses to understand what we could do to become more successful. The key at that time was the ToC (Total Cost of Ownership) calculation, showing the total cost of a wired network system during its lifetime, including the cost of cabling and re-cabling, to show the cost/benefits of a wireless system. I have to admit, though, that these calculations were not always that convincing in the eyes of the customer. The weak point usually was that an investment in WaveLAN meant a higher investment upfront. And although companies in those years may not have been as cash conscious as many companies are today, this upfront investment was still was a difficult point in the total proposition. The real successful wireless LAN applications were the applications that were focusing on mobility, more than the ones that were about “getting rid of cable.” But the technology was only marginally supporting mobile applications. Usually the product was too bulky, and it still used quite a lot of power from the battery of the device that it was loaded on. Nevertheless, we were trying hard to get mobility features integrated in the development of the product.

We continued to look at Apple as a target customer and distribution channel. We worked this through a company called Digital Ocean, a startup group out of Lenexa, Kansas that was very enthusiastic about wireless LANs. Their focus was clearly the Apple market, and they had a lot of knowledge about the Apple technology to integrate this with wireless. Apple was pioneering these days with a new concept – the PDA (Personal Digital Assistant), the predecessor of the palmtop. It was clunky and heavy and overall not successful compared to the expectations, although I saw quite a lot of enthusiastic users who were really satisfied. Digital Ocean continued working on general wireless LAN connectivity between Apple computers, and they also developed a sleeve for the PDA that functioned as housing for a wireless LAN card. Unfortunately, when this product

became available, it did not create the excitement that we had hoped for, not even within Apple, that in those days was going through some very difficult times to just stay afloat. Our many visits and selling efforts were all clearly ahead of their time; it never broke through.

Actually, Apple's later cancellation of the PDA product program also killed Digital Ocean as a company. The company had been relying too much on Apple and did not survive their market share decline. In the meantime, Digital Ocean had been doing a lot of development work on the MAC protocol that would lead to interesting future standardization efforts.

These years we also worked together closely with Teledyne and Raytheon; it was the early days of GaAs (Gallium Arsenic). The idea was to build the first integrated radio chip. In these days, CMOS was not suitable yet to generate the high speed and high-frequency radio signals that are required. The program with Teledyne and Raytheon ran into severe complications. Initially the chip did not work, and revisions were required. But with each revision, we also saw the cost of the chip increasing dramatically, as there were serious manufacturing problems. For every good chip produced, there were at least two failures, so the cost of the chip, originally targeted below \$20, turned out to be over \$50 and could not compete with a so-called "discrete design." This turned out to be one of those cases where a design with traditional components (resistors, capacitors and inductors) was just significantly more cost-effective, even though it was larger in size.

But there were more technical complications that resulted in the slow progress. Developing networking products in these days was still a quite cumbersome effort for product developers. There were many networking "flavors," requiring many different software drivers. Novell and Microsoft were battling, together with a group of smaller network technology providers like Banyan and different flavors of Unix. Each provider had its own peculiarities, as well as different versions requiring startups like us who wanted to compete to try to interface with and through each of them. It seemed like every customer required a version that we just did not support. Software driver development was the key to success, and fortunately our team had a few smart guys to help us through these days. Unfortunately, though, the whole problem was exacerbated by the introduction of a new hardware interface by IBM and some of their fellow PC

companies. This interface was called the Microchannel bus, competing as a successor of the AT bus with EISA. So, all the software varieties existed for each hardware bus version, creating a serious challenge in determining what to develop first and how many resources needed to be spent on what. Getting the product right seemed to become an impossible challenge, purely by the number of permutations that had to be integrated with.

5.7 OUTDOOR PRODUCTS

I think it happens to a lot of product developers at least a few times. After having done all the marketing and product management homework, and launching the product, the customers who really start buying the product are using it in a way that you never thought of. Worse yet, they use the product in a way that it was not at all intended for. Especially hurtful for engineers is the situation where they could have developed the product much better, had they known about that usage. So was the case with WaveLAN. In the first years over 50% of the usage was “outdoors” – connecting the networks of two or more buildings together into a single network, by using the wireless LANs. These wireless LAN connections create “network bridges” between the buildings.

There were plenty of reasons why the use of wireless LAN cards to resolve these building interconnectivity problems. In the first place, it was simple and easy to install. The alternative was a T1-line (or E1-line in Europe) that needed to be leased from a telecom operator. The monthly rate for this lease was expensive, plus the multiplexing equipment that was required for such a connection was not simple. Amazingly, several VARs (Value Added Resellers) found that by adding special (and low-cost) directional antennas to the WaveLAN cards, they could (surprisingly) cross distances easily extending to a 3-5 mile range, or 5-10 kilometers. There were even situations that the distance would go up to 50 miles/75 kilometers. That was unheard of, as the product was developed to produce a robust link in-house, with its multiple indoor reflections, because of the internal walls and furniture or equipment in offices or in buildings in general. I believe that at a certain stage in 1992, more than 80% of our sales of wireless LAN products were used for outdoor connections that were unintended by the developers of the product.

The situation got out-of-hand when some VARs found ways to add “power amplifiers” to the product between the card and the antenna, boosting the output

power significantly above the maximum of 1 Watt that was allowed in the 2.4 GHz. We had a lot of internal debates about this. It was clearly illegal, but was it our responsibility? Should we bring this situation back under control, and how could we? What was the long-term consequence? Would it pollute the ether, so it would cause interference with legal products?

Fortunately, the FCC realized that this was getting out-of-hand and came to help by prescribing a special antenna connector in a special format, with a closed distribution so that selling was only allowed to those companies, who had signed for the legal use of this connector, in compliance with the regulations for the 2.4 GHz.

In the meantime, the outdoor market was flourishing. And despite the fact that the product was not developed for this application, the results were amazing. Since then, more specialized products have been developed for the outdoor market, not even compliant with the standards – but interestingly enough, that is not necessary. The Wi-Fi standard is a voluntary standard, compliant with the transmission rules of the 2.4 GHz. This means that as long as products are compliant with these rules, they will be allowed in the 2.4 GHz band.

In general, the need for standardizing the technology in an outdoor point-to-(multi)point connection is significantly less compared to the indoor situation, where different PC brands and different infrastructures need to talk together.

5.8 RELATED TECHNOLOGIES CREATING OPPORTUNITIES AND CONFUSION

At the same time wireless LANs were struggling to get into the market, wireless Wide Area Networks came up, also with struggling business models, but nevertheless trying hard. One was Ardis, a company that started as an internal network for Motorola and IBM support technicians. Ardis was privatized, and to reach economy of scale, was supposed to win a much larger customer base. One of the ways that Ardis tried to achieve this was to get more companies to develop modems, and so we were also invited to develop a modem for Ardis. Unfortunately, we could not accept the invitation, as we plainly did not have the resource base to support such a development. Our focus was very much “cordless” networking, whereas Ardis was pursuing more the concept of “cellular networking.”

Another company that really made waves was McCaw. Actually, McCaw had put together a wireless network and was about to be acquired by AT&T, when it launched its CDPD initiative. CDPD was going to be what cellular companies are trying today with GPRS. With McCaw, we were also invited to develop a wireless modem. Actually, at this point we were already part of AT&T, and our interest in this program was very high. We did some development on it, but we were unable to complete the business case with McCaw in a satisfactory matter, so the program was dropped.

Looking at today, the wide area wireless technology that pretty well survived is Mobitex. In its different formats, and with the help of some big operators, it managed to pull in enough subscribers to stay in business and has a presence across the globe. It remains a proprietary solution – Ericsson-based – and therefore is mainly used for closed applications rather than open applications.

This was also the period of “tablet hype.” There had been several breakthroughs in screen technology, ruggedness, and reduced scratch sensitivity. This had led to a new trend in the computer industry – keyboard-less. A computer became a tablet and writing on it with a pen was possible. Apple was pioneering this with their PDA (Personal Digital Assistant), along with Fujitsu, Grid, NCR Computers, new startups like EO and Momenta, etc. These companies were all immediately the target of the wireless LAN companies willing to “help,” to further improve the tablet business case as a mobile device. Unfortunately, the tablet technology overall was not really accepted in the market, probably because the handwriting technology was difficult to master and relatively slow in execution. The product opportunity fizzled and only re-emerged later as part of the palmtop technology – and then with great success.

In this period, another interesting event took place. In 1992, Bill Gates and Steve Balmer visited NCR Computers in Dayton, Ohio. They got a presentation about NCR’s product portfolio; NCR was working these days on a tablet computer, and they worked closely together with Microsoft on the Operating System software. They also received a presentation on wireless LANs, and I understood that they were quite intrigued by it. However, there was never any follow-up from Microsoft’s side. We tried hard, but probably not hard enough. It took until 1999 for Microsoft to think that the standard had developed well enough to embrace it for their own network, later touting that the Microsoft’s

Redmond campus was the largest wireless network in the world. I'm not sure whether this was true, but for Microsoft, this was all very new at that moment.

During this time, there was additional confusion because of new spectrum that became available in the US. The confusion between Direct Sequence (DSSS) and Frequency Hopping (FHSS) was exacerbated with a wireless LAN data communication proposal for the 1.9 GHz, the so-called PCS band that is used today for the US version of GSM, the cellular telephone network. For new vendors, new spectrum is always seen as a new opportunity. For companies like us, with vested interests in existing bands, new bands could easily mean a threat, annihilating or at best reducing returns on investments that we had done so far. The amount of energy sapped away by these proposals in industry committees, with consultants and with the press, is always very concerning. Looking back and realizing that nothing came out of this 1.9 GHz PCS band proposal is a reminder that a lot of time can be spent on something with nothing to show for it.

This was also the period that the first interest was raised for 5 GHz as the successor of the 2.4 GHz. This was as early as 1993, and it started in Europe, where in particular the 5.2 GHz was proposed for harmonization on a worldwide scale. This 5.2 GHz is now known as the "lower part" of the 5 GHz band that has been standardized via IEEE 802.11a.

More threats to the business were coming up during these years – and not only from outside. NCR was by now acquired by AT&T, and AT&T Consumer Electronics started eyeing the 2.4 GHz band for a new generation of digital cordless phones.

Research had shown that this band was appropriate, but that there was some wireless LAN activity in the band. Therefore, the AT&T design requirement was that the engineers must make sure to blast "WaveLAN" out of the band, when the ether is required for the cordless phone. It was clear that with this type of colleagues, you do not need any competitors.

We had a lot of back and forth discussion on this, but AT&T (I mean this in the larger sense, as we were part of AT&T ourselves...) took the position that if they did not take on this opportunity, then somebody else would. So, a discussion

leading to a win-win situation, and even the opportunity to co-market wireless voice and wireless data, was rejected.

Fortunately for us, this program ran into trouble, as the AT&T development team was not able to make a low-cost radio – at least not low-cost enough to effectively compete with lower cost alternatives – and the program was cancelled.

It was a good warning sign for us though, as still today, the 2.4 GHz band (the band that is used by Wi-Fi) is not “protected.” This is a major difference in comparison with, for instance, a cellular phone band. It means that still other applications can be developed that in essence interfere with a Wi-Fi radio. Fortunately, interference is usually a reciprocal activity, and in all the cases I have seen to date, new applications are interested in peaceful coexistence with Wi-Fi, not least because Wi-Fi nowadays has become pretty widespread.

5.9 SUMMARIZING THE SECOND PERIOD (1991–1994)

If the first product development phase (1987–1991) was somewhat of a disaster, at the end of the second phase, (1991–1994) our situation was not much better. In essence we had “doubled down,” but there were no real profits in sight. After the first period, we had a wireless LAN card. After the second period, we had a wireless LAN system. So, we could bridge into a wired environment, and we could roam through a building while staying connected. After the first period we had a US-only product; after the second period we had a worldwide product – or at least as worldwide as we could expect. This all was clear and encouraging progress in these areas.

But the cost was still too high, and the speed (2 Mb/s) was falling behind in comparison with the wired market. So, the overall market acceptance for wireless LANs was still very low. Also, there were serious market concerns about the lack of standardization. The major players in the market at this time were Proxim, Aironet and AT&T, plus a whole array of smaller companies like Breezecom, WaveAccess and Xircom (Netwave), and all these companies had different products and were using different technologies. Publicly, these companies were telling the customer base why they were the best and what was wrong with the technology of the competitor. But the consequence of this all was that customers did not trust anybody. “Data is precious, and waves are weird,” was essentially the wait-and-see attitude for most.

At this point, none of the companies had real marketing muscle or the supporting financials to launch a proper marketing campaign, although several companies did try a bit of advertising.

By the close of the second period, it was clear. We needed standards, higher speeds and lower cost. This meant more integration, so the size of the products would become smaller and therefore also suitable for mobile products like notebooks.

By this time, we all had read Regis McKenna's "Crossing the Chasm," a marketing book that described how in high tech markets there is a gap ("chasm") between the adoption of a new technology by the "early adopters" and the "mainstream users." We had by now won quite a few early adopters, but mainstream seemed to be far away. We were also reading Geoffrey Moore's "Inside the Tornado," but despite all the hard work, the market seemed to be extremely quiet. Enthusiasm all around, yes, but sales were rather limited. It was clear there was more work to do while keeping the faith. The market fundamentals were there – mobility and networking had increasing profiles but viewing wireless as the solution for both was still a ways off. Could we wait long enough?

THE INITIAL STANDARDIZATION (1995–1998)

Part of the strategy at NCR, and later at AT&T and Lucent Technologies, was that the process of standardization was wholeheartedly supported. The belief was that creating a standard would be critical to the acceptance of wireless LANs in the market. The whole process of coming to standardization, however, took quite a long time, because there were many different interests from many different companies. The result though, has been a very thorough, versatile and feature-rich standard that probably will be with us for the coming decades. I know – it is extremely dangerous to make future statements in the world of technology, but I am quite convinced that the IEEE 802.11 standard is also a standard that is a solid basis for future extensions and newer capabilities.

6.1 THE EARLY DAYS OF IEEE 802.11

Even in the very early days of our wireless LANs efforts, we were clear that this concept would only fly when we were able to get a standard established, like the Ethernet standard for the wired LAN world. Ethernet was established as a standard in IEEE 802.3, and by the mid 1990s it was well-entrenched. It had not gone uncontested, however, as initially there were two competing standards – IEEE 802.4, also called Token Bus, and IEEE 802.5, also called Token Ring.

Each networking standard had its advantages and supporters, which had initially created a lot of confusion in the industry. The original ideas for a wireless LAN standard found a home in the IEEE 802.4 standardization committee, in a working group called IEEE 8024L. But it became clear quickly that creating a wireless standard under IEEE 802.4 would not give an appropriate result.

In this respect it is important to understand what is standardized in an IEEE committee. A standard actually exists of two parts – a physical layer (also called PHY) and a media access control layer (also called MAC). Although this may sound somewhat technical, it is pretty straightforward. The physical layer describes how a computer connects physically to a LAN – what sort of cable is used, the number of wires in the cable, what level of power over the cable, the measures of the connector – in short, everything that has to do with the physics. With wireless, the physics also has to do with the band used, the frequencies in the band (the channels), and the type of transmission (how to put the data bits

on “radio”). The MAC describes how a computer connects logically to a LAN. You can think of the comparison of two ships on the ocean, within sight of each other, communicating before the time radios existed. The PHY defined the flags that were required for the communication, and the MAC defined how each flag had to be used and what the meaning of the flag was. For these reasons, the MAC is also called the “protocol.”

The MAC and the PHY are quite dependent on each other, and it became clear very quickly that the IEEE 802.4 Token Bus MAC was not suitable for wireless technology. The Token Bus protocol works with a “token” that is sent around from station to station, and only the station that has the “token” is allowed to “say something,” while other stations have to be quiet to avoid disturbing the communication and garbling the messages. The advantage of “wired” compared to “wireless” communication is that electrical signals over a wire are more robust, and so less susceptible to noise or interference compared to wireless radio signals. In both the Token Ring and the Token Bus protocol, the risk is always that the token gets lost. These protocols have capabilities to recover from such a loss, but these recovery procedures take time, which manifest itself as a slowdown in the communication. In a radio environment, the risk of losing a token is even bigger, because of the susceptibility of radio signals to noise. Therefore, the conclusion was that a token-passing protocol was not the right way to go.

Some serious thoughts have been given to the use of the Ethernet IEEE 802.3 protocol, also called CSMA/CD (Carrier Sense Multiple Access with Collision Detect). This protocol is quite ingenious and interesting, as it does not predetermine who is allowed to “speak” or “who has the token.” Actually, there is no token at all. The Ethernet protocol simply works like this – if a computer wants to send a data packet, it just “throws it on the cable.” At the same time, it is listening to whether another computer is doing the same at the same time. If that is the case, it is assumed that a collision has taken place and that both messages have been garbled. So, both computers will decide to send their packet again. To avoid the packets being sent at the exact same time, each computer waits a random amount of time, called the “back-off time.” These very short, random times are usually different enough that the retransmission are successfully separate. Quite a bit of thinking and work has gone into this, and although it may not seem that efficient, Ethernet has turned out to be a very

efficient and cheap protocol. It is one of the main communication protocols in place today in the data communication industry.

However, for the radio world, this “collision detect” mechanism of the Ethernet protocol had a serious drawback. Contrary to electrical lines, radios usually have trouble receiving while transmitting (i.e., listening while talking). In the early phases, an efficient variant of CSMA/CD had been thought of: CSMA/CA, where CA stands for “Collision Avoidance.” The variant here is that before starting to transmit, the radio listens for whether someone else is transmitting. If so, it waits a random time (like in Ethernet), and then the whole procedure repeats itself. Only if the channel is clear does the transmission start. Actually, this type of protocol is somewhat more efficient than Ethernet, as the chance for collision has been considerably reduced, and therefore so has the time lost on recoveries. However, the extra complexity had a cost penalty that made the initial products more expensive. Nowadays these features are largely fully integrated in chipsets.

The conclusion of all of this technical groundwork was that the standardization of wireless could not efficiently take place in one of the existing wireline standardization committees, and so a separate standardization committee was created. The number assigned was 802.11, and that was the beginning.

The interest in IEEE 802.11 has been quite high from the beginning, but it was largely driven by relatively small companies like NCR, Aironet, Intermec, Symbol, Xircom and Proxim. Other companies stepped in on a regular basis, including IBM, Apple and Motorola, but their interest seemed to be spotty and more about monitoring the progress. The regular changes in participants from these companies did not help the progress, as in the early days, many private agendas were pursued.

6.2 THE ORIGINAL IEEE 802.11 MAC STANDARD

Although today the IEEE 802.11 standards are known as IEEE 802.11b and IEEE 802.11a, the original standard was simply known as IEEE 802.11. This was a 1 and 2 Mb/s standard in the 2.4 GHz ISM band, and the basis for the IEEE 802.11b standard.

The MAC protocol largely came together from cooperation between NCR (AT&T at that time), Symbol Technologies and Xircom. NCR was the computer company that my organization was part of. Symbol Technologies was a company that was and still is largely known as a data collection terminal company with specific strengths in infrared bar-code reading. Symbol's contribution was specifically on the needs for power management. The need for "range" was also articulated – bar code scanners need to be held under pallets and seemingly out-of-reach of everything. Xircom was known in the market for their "pocket LAN" adaptors, which allowed computers to be connected to a network via a printer port. This was a nice innovation that gave them considerable insight on the intricacies of PC-networks. In 1999 Xircom was acquired by Intel. Working together, these three companies made a proposal that became the foundation of the MAC for what is known as Wi-Fi today. The amount of effort that went into this MAC was quite extensive, although the basis was straightforward and comparable to the IEEE 802.3 Ethernet standard. However, because of the nature of wireless, and with the need for fulfilling the promise of true mobility, there were quite specific extensions put into this MAC protocol.

In the first place there was the need for roaming, or as this is called in the telecom world, "hand-off." But "hand-off" is a term that clearly makes this action a base station responsibility. In the IEEE 802.11 MAC protocol, however, this is a computer (also called "terminal") initiated action. Although the implementations may differ with the different product providers, the essence is a process of association. A computer listens on all channels to find what access points it can connect with. Then it selects the best access point and channel. When one is walking around with a laptop, the quality of this connection may degrade. At that moment it checks again, and if there is a clear other base station and channel through which to connect to the network, the switch is made. In technical terms: the computer disassociates from the access point it was connected with, and it reconnects with a new access point. The access points, as real bridges between the wireline and wireless network, are told that the computer has moved to another access point, and they update their network topology tables to make sure that the reverse traffic (the data packets from the network to the computer) is going to the right place as well. To make this work more efficiently, there is also a protocol between the two access points, as they are both on the wired network, and they confirm with each other the rerouting of the traffic.

A second important feature that was thoroughly investigated and included in the protocol, was power management capability. This is where Symbol Technologies came in with their experience that the lifetime of data collection terminals, that is the time that the battery stayed up, is essential. With radio communications, we needed to avoid the stand-by time draining the battery unnecessarily quickly. In this respect it was important to recognize that a computer, or more specifically a handheld computer, may only really communicate at 11 Mb/s (Megabits per second). To put this in perspective, 11 Mb/s effectively translates into 500 Kilobytes per second, while the effective data transmission requirement probably is in the order of magnitude of 500 Kilobytes per hour. The Wi-Fi protocol contains a lot of special power management features, including requiring the access point to wait and store data packets, as long as the terminal unit is “asleep,” allowing the terminal to wake-up once in a while, give a signal to the access point and check for anything waiting, and then quickly go back to sleep. This allowed for creating a radio unit that effectively used very little battery power. This should not be confused with the fact that when the battery is used by the radio for transmitting or receiving, there is still quite some power required, usually in the area of 1 to 1.5 W. Still today, this is a problem for the batteries of palmtop computers. Palmtop usually cannot deliver this power and simultaneously keep its processor alive.

A last feature included in the standard was a feature that improved the reliability of the communication, but that turned out to be generally unnecessary. It was heavily promoted from Xircom’s side, and it is called RTS/CTS (Request to Send, Clear to Send). It makes sure that a computer or a terminal is not sending any information without first having received approval to do so. Although this feature initially had become quite redundant, it currently may find a revival in relation to recent work on a newer addition called Quality of Service, a feature that is necessary to enable solid voice communication over wireless LANs.

Other features that got special attention were the “ad-hoc” networks and the repeater function. The “ad-hoc” networks allowed two or more PCs to get together and spontaneously setup a network, without anyone controlling the network. This is like calling between two cell phones without the mediation of and/or a subscription with a telephone company. Think about it – two cell phones can probably “hear” each other within both phones’ reach of 2 to 3 Km (1 to 2 Miles). Technically, there is no reason to have a telephone company in the mix,

other than that the current generation of cell phones need the infrastructure to setup the conversation.

The repeater function allows each access point that is connected to the wireline network to have satellite access points that can function as a go-between for a PC and the wireline access point. This feature helps to extend the reach of a single access point without the need for extra wiring.

The MAC proposal that came together from AT&T, Symbol Technologies and Xircom was not the only one. Before the real voting took place, I believe there were 5 or 6 proposal going around.

One of the other proposals that is important to mention was from a group around IBM. This proposal was based on a quite different architecture and reminded many people of the old Ethernet versus Token Ring controversy. Just like with Token Ring, the wireless MAC proposal from IBM included a central mechanism that would control the network. It was called the PCF proposal, where PCF (Point Control Function) was the mechanism that controlled the wireless network – controlling who in the network was allowed to send and at what moment in time. One of the reasons that this proposal did not make it was the fact that it did not support “ad-hoc” networks, and the structure of the protocol was such that ad-hoc networks would not be easy to add.

Contrary to the earlier wired situation where IBM’s proposal in IEEE 802.3 was rejected and led to the start of a separate standardization committee, this time IBM took its loss. I can only guess at the reasons – uncertainty about the viability of the wireless LAN market, a reduced interest in networking in general as their market leadership had been significantly reduced, or the fact that IBM was going through a difficult period overall. Probably it was a mix of all of this.

6.3 THE STANDARDIZATION PROCESS IN IEEE 802.11

In this respect it may be interesting to describe some of the processes in IEEE, as agreeing on standards in a highly politicized body is usually a cumbersome process – particularly when engineers have to play politics, something that is not necessarily a natural strength. For true engineers something is right, or it is wrong; and if it is not clearly right or wrong, then it is a problem that needs to be resolved, so it can lead to a clear right or wrong determination. The real political

world contains a lot of nuances and ambiguities and managing engineers in a process that leads to something useful is not that easy to comprehend for true engineers. Well, not so in IEEE. Probably based on the rich IEEE experience that many of the IEEE members have, the politics in IEEE, especially around voting days, is incredible. Lobbying, coalitions – everything that one is used to in normal politics is happening here.

In this respect I need to mention Vic Hayes, who was part of our organization and who has been chairing the IEEE 802.11 committee for a decade. Although he was an employee of NCR, AT&T, Lucent and Agere over the years, he has been able to build a reputation that put him clearly above the maze of politics and led to a set of standards that turned out to be very robust and useful and also laid a solid foundation for future generations of higher speed wireless LANs.

The summary of the process of creating a standard is as follows. Once the higher order in IEEE has approved a PAR (Project Authorization Request), the subcommittee, in this case the IEEE 802.11, goes to work and holds week-long meetings, about every other six weeks. First there is a period when proposals can be issued, usually covering three to five months. These proposals are discussed in a few subsequent meetings and judged on their merits – completeness compared to the original PAR and feasibility for implementation into real products. When the time is ripe, that is, the proposals have been discussed enough and the “mix and match” of proposals has been considered, there is a voting round.

This voting is interesting, as the rule is “one person, one vote,” contrary to, for instance ETSI (European Telecommunications Standards Institute), where the rule is essentially “one company, one vote.” Actually, in the ETSI, voting is also dependent on the type of membership, and big companies that can afford high membership contributions can get “heavier” votes than lower contributing members. Not so in IEEE, which creates an advantage for small companies that have a chance to really contribute. (A disadvantage might be that these small companies can also stall progress...) The consequence is that large companies are pushed to send over large contingents of people to participate in the meeting, and they are also forced to be compromise-oriented. But just being a member and participating in the meeting is not enough. Participation in several meetings in a row is required to build up and solidify voting rights, as these voting rights quickly evaporate if too many meetings are missed.

The voting procedure itself can be voted on, but the bigger ones that I have been exposed to work like this. There is a set of voting rounds, and in each round, the proposal with the least votes drops off. So, when there are six proposals, there are five voting rounds. It should go without saying that while this may sound like a solid procedure, it leads to incredible politics. Who gets the votes that in the previous round went to the losing proposition? The reality is that there are big swings in the number of votes for proposals. Despite the rule that everyone needs to vote for the best proposal, quite often voting is done to make sure that the proposals of certain companies do not proceed to the next round.

The proposal that ultimately wins is the proposal that the IEEE committee continues to work with – it is the proposal that with collective effort will be worked on and improved upon, until 75% of the voters are in favor of it. Actually, the goal as a true engineering organization is to reach 100%, with the assumption that there is only one right solution to the problem. Fortunately, the reality has set-in that 100% is perhaps too idealistic, as there can be stubborn renegade members who are absolutely convinced that something in the proposal is absolutely unnecessary or redundant. Also, sometimes companies have an interest to make sure that no agreement is reached on a new standard – for instance, if they dominate the market with a closed de-facto standard, or when they want to stop a new standard that will make an older standard redundant.

Another rule in IEEE that also really helps to move the process of standardization forward, while making the proposed standard technically solid, is that one cannot vote against a proposal unless one clearly defines the reasons. From a technical perspective, an engineer will not quickly approve something unless he is reasonably convinced that it is sound. And when he disapproves something, he is challenged at the same time to come up with a proposal for a solution.

The process of standardization in IEEE has been heavily criticized, because it is relatively slow, and because it favors small companies. Or maybe it does not favor big companies – that is, at least, what these big companies say. But all in all, it is a solid and a fair process – as solid and fair as these processes can be.

6.4 THE ORIGINAL IEEE 802.11 PHY STANDARD

Well, the good news was the rapid common agreement on the MAC. But this was not the way it was going to be with the PHY. The definition of the physical layer,

in essence the radio, turned out to have two large groups that had very strong preferences either for DSSS (Direct Sequence Spread Spectrum) or for FHSS (Frequency Hopping Spread Spectrum). Neither of the two groups wanted to give in, and neither of the groups would be able to reach the 75%. Actually, the FCC had given the ISM band to the industry for helping to further grow the economy. But this gift contained a poison pill that created a discord in the development community. Each side defended its preference with fervor, but the arguments were never conclusive. As happens often in the technology world, the discussion quickly moved from an objective debate to a subjective positioning where people assume that what is best for them is the best for everyone.

The main difference between the two technologies was essentially the way to look at speed. FHSS was defined in the standard as 1, 1.5, 2, 2.5, 3, 3.5, 4 and 4.5 Mb/s, but it was only truly working at a speed of 1 Mb/s.

DSSS was defined at 1 Mb/s and 2 Mb/s, and also worked at both speeds – 2 Mb/s at somewhat closer range, 1 Mb/s at longer ranges. At first glance, this would favor DSSS over FHSS; however, DSSS was generally only available on 3 channels. For FHSS the number of channels for frequency hopping was higher, although because of the statistical nature of the technology, it is difficult to state how many effective channels are actually available. It is fair to say that for the end-user experience, there were at least 10 to 15 channels available. Therefore, the total capacity of a Frequency Hopping system is higher. So, this would lead quickly to customers, who got completely confused about raw data speed and system capacity.

The problem really was that the two systems, DSSS and FHSS, did not talk to each other. Even worse, they actually interfered with each other. In the MAC protocol, the safeguards were nicely built-in to reduce the number of lost data packets, as described earlier. However, this did not work across radio technology, that is from DSSS and FHSS and vice versa. So, with both DSSS and FHSS systems “in the air,” the data loss during transmission could become quite significant.

It goes without saying that furious debates in the IEEE have been held for one party to convince the other party. There were also some serious efforts to harmonize the two technologies – that is, to overcome having transmissions

from one technology be harmful for the other. But these efforts were not successful. The IEEE 802.11 committee ended up with a “let the market decide” type of cop-out, and the market decided – it waited a few years longer to accept this wireless technology and let the companies struggle with marginal revenues. This also pushed the telecom community into continuing to ignore this technology. If IEEE would have been stronger and more unanimous these years, wireless LAN would have been put more clearly and forcefully on the map. I wonder whether the UMTS (3G) could have really taken off, at least in the minds of the marketers and advertising people...The IEEE failure at this stage probably caused people to reconsider alternatives, and this was also probably when the thinking about Bluetooth as a low-cost LAN solution started to form.

6.5 THE INTELLECTUAL PROPERTY OWNERSHIP IN STANDARDS

Leading edge technology is always about IP (Intellectual Property), and getting into IP always gets very tricky, very contentious, and usually very complicated. Standardization is trying to balance making something open and beneficial for everybody, while at the same time satisfying companies that want to have “fair compensation for their investments in research and development.”

Different companies have different IP strategies, depending on the type of business they are in. Usually small companies can be very narrowminded and focused on one single patent. Their goal in life is then to make sure that large companies pay them – usually in relatively small amounts, a few \$Ms or less.

Larger companies are less focused on a single patent and usually build larger patent portfolios. These portfolios are used primarily in a defensive way to avoid being put out of markets, and secondarily, depending on the strength of the portfolio, to close large deals with other large companies – so-called “cross license agreements.” These deals usually involve the companies weighing each other’s total applicable patent portfolio and then striking a deal. If there is a significant weight difference between the two patent portfolios, the company with the lighter portfolio has to pay the other company an extra amount that can run into the \$100Ms. These large deals usually include a lot of negotiations and large groups of lawyers.

To make things even more complicated, some patents are only worth something

if a standard is adopted. If a certain technology is rejected as a standard, then underlying patents may be rendered useless at the same time.

The way the IEEE is trying to handle all of this is to assure that all companies that participate on the standardization of a technology would make their IP available in a “fair, reasonable and non-discriminatory” way. This is a legal phrase, and it means that the IEEE expects that all the members who are contributing to a standard also work to make sure that their IP policy supports this effort. Practically, it means that (the legal department of) each participating company writes a letter to the IEEE and makes this “fair, reasonable and non-discriminatory” statement. Interestingly though, this hasn’t been more specifically defined. It is fully left to the individual companies to sort out amongst themselves – or to pursue legal action for the courts to determine.

At Lucent Technologies, we got ourselves into an interesting dilemma. Lucent Technologies, the owners of Bell Labs, has one of the strongest patent portfolios in the world. Their strong and explicit preference was, and is, to close one-on-one cross license agreements with other companies and to never make a generic statement like “fair, reasonable and non-discriminatory,” because in the eye of the IP lawyers, that could only backfire in negotiations with other companies and/or jeopardize their position in other standardization bodies. In the way Lucent Technologies was and is organized, and in the way its legal structure is built, all the IP is automatically owned by the IP department. Although this department must be very careful not to interfere with the immediate business of the business entity that had developed the IP, the position was always clear. Just when the standard was about to finish, the Lucent Technologies IP department was starting to wake up and realize what was happening at that small entity that was doing “something on wireless data.” This situation created a serious contention that required a significant number of multi-level meetings.

Ultimately, we, the business unit, won the discussion after a fight and stand-off that took several weeks. I think the IP department realized that the patents may not have had much value without the standard, as they were very standard-related. They also assumed that the wireless LAN business would be small and insignificant compared to the total company business. So, the Lucent Technologies IP lawyers wrote and signed a “fair, reasonable and non-discriminatory” letter. But overall, I do not know whether Lucent as a company

realized the significance of what was happening. The rules in the computer and data world were clearly differently written compared to the rules in the voice phone and cellular world.

6.6 THE FORMING OF IEEE 802.11A AND IEEE 802.11B

As soon as the first IEEE 802.11 “double standard” was completed, many companies started to work on higher speeds. It quickly became clear that this first 2 Mb/s standard, despite all the good intentions and all the hard work, was going to be a failure in the market. To a large extent this was because this standard actually was a double standard, but also that the speed was perceived as too low. On top of that, the standard did not add much to what existed already in products on the market. The critical issues at the end of the previous period still existed – the needs for standardization, higher speeds and lower-cost products.

But there are things to mention on the positive side. Despite that the PHY was ambiguous at this stage, the MAC turned out to be a unanimous success. It proved very robust, and although over the following years small corrections and extensions have been made, it stood and stands for years to come. Probably it was over-engineered for what was required at this stage; probably it contained features that were only useful in small market segments. Nevertheless, it will go down in history as a monumental piece of work, and many people will have used it, whether they were aware of it or not.

Together the MAC and PHY definition that make the IEEE 802.11 are assembled into a large, neat document of 528 pages that is downloadable from the Internet.

By early 1998, it was clear to everybody in the IEEE 802.11 that higher speeds were the first thing that needed to be pursued. The MAC seemed to be robust enough to handle higher speeds, the challenge was the PHY. In the 2.4 GHz, the PHY was divided between DSSS and FHSS; and then there were the unused capabilities in the 5 GHz. But there was another pressure as well. Europe regulatory bodies had never been very much in favor of the 2.4 GHz, and the ETSI had developed a standard in the 5 GHz band called HIPERLAN/1. This standard was over 20 Mb/s but generally seen as too difficult to efficiently implement. The response of the IEEE was double-sided. The target was set to develop a high speed/10+ Mb/s standard in the 2.4 GHz and a very high

speed/50+ Mb/s standard in the 5 GHz. From an administrative perspective, the proposal for the 5 GHz actually preceded the high-speed proposal of the 2.4 GHz. That is how the more complex and later-implemented “.11a” designator came out ahead of the simpler “.11b” that conquered the market. But the IEEE had managed to do it again – after the confusion of a double standard in IEEE 802.11, there was now confusion between the IEEE 802.11a and IEEE 802.11b.

Actually, my organization was part of the problem, as we were quite influential in both standards. Being an American organization with strong historical ties into the IEEE, but at the same time being located in Europe and having started the wireless LAN efforts in ETSI, we were pushed equally hard to participate on both standardizations simultaneously.

So, we had two teams working in Lucent Technologies in the Netherlands, simultaneously and side-by-side – and sometimes competing. One team was working on the high speed 2.4 GHz in the IEEE, and the other team on 5 GHz in the IEEE and in ETSI.

The IEEE quickly realized and confirmed the decision to only change the PHY and keep the MAC the same for the 2.4 GHz and the 5 GHz. At the same time, the transmission methods in the 2.4 GHz and the 5 GHz had to be more-or-less different. The 2.4 GHz ISM band requirements put extra constraints around the way signals could be transmitted, and these constraints were not required in the 5 GHz. Therefore, dropping these constraints would allow for lower cost and simpler product in the 5 GHz compared to the 2.4 GHz.

The critical phases for IEEE 802.11a and 802.11b were both in the spring of 1998.

The least contentious was the IEEE 802.11a. There were two serious main proposals – one from Naftali Chayat from Breezecom (currently Alvarion), and one from Lucent Technologies and NTT, based on OFDM (Orthogonal Frequency Division Modulation) technology. The voting was won by the Lucent Technologies and NTT combination, so actually IEEE 802.11a, with its 54 Mb/s transmission technology in the 5 GHz is an older agreed standard than IEEE 802.11b.

The accepted proposal was very close to the proposal that Lucent Technologies had made in the ETSI. The acceptance by the IEEE of the OFDM technology

inspired the ETSI to follow the same route for HIPERLAN/2, as was proposed by Lucent Technology. This explains why the radio technologies of IEEE 802.11a and HIPERLAN/2 are almost similar.

The voting for the IEEE 802.11b radio (PHY) was very contentious, and the fighting was on the brink of tearing the IEEE 802.11 committee apart. The main contenders were Jim Zyren and his team from Harris (nowadays Intersil) and Lucent Technologies, and then there was an outsider proposal from John Cafarella of Microlor, a start-up company with very good radio knowledge. Actually, Microlor was largely supported by Clarion, who had serious plans at that time to go into wireless LAN. When the voting got down to these three companies, the Lucent Technologies proposal was voted out, and the final voting round between Harris and Microlor started. What happened next is hard to describe and challenged the democratic rules in IEEE. In the voting, Microlor came out with 52 votes, Harris came out with 51 votes, and there was one vote abstaining. According to one set of rules, Microlor had won the vote, but this was immediately contested, as Microlor did not have a majority – 52 of 104 votes is not “more than 50%.” A violent discussion unfolded about the interpretation of the outcome of the vote with many real and emotional arguments involved. Then Jeff Abramowitz, who was the 3Com Product Manager for wireless LANs at that time (he currently works for Broadcom, still on wireless LANs), stood up and made a motion that contested the whole voting procedure. His statement was that according to the rules of IEEE, an IEEE member engineer should vote for the best technical proposal, and according to his assessment, despite the fact that the voting was “closed,” the reality of the voting was that the individual members had voted along party lines, that is, along the lines of the companies they worked for. This was of course true to a large extent, but he phrased his motion in such a way that the Harris proposal should be declared the winner because the voting for Microlor was not based on technical reasons but instead came from the anti-Harris camp. It was very clear at this moment to everybody that 3Com was a Harris supporter.

So, in essence, 3Com asked the IEEE meeting presidency to reject the minimal majority vote for Microlor and to declare Harris the winner. The chaos this proposal created was incredible, and the whole meeting went down in flames. I think there must have passed a motion to adjourn the meeting, but in all the chaos, I do not really remember.

There was of course some truth to the statement of 3Com that most of the Lucent Technologies supporters had decided to side with Microlor, to avoid giving Harris and their supporters an unfair market advantage from having the IEEE sanction their proposal – the development of which was already pretty far underway. At the same time, I think that John Cafarella from Microlor had one of the more miserable days of his life. I do not think that he had expected to come that far with his proposal, but when he saw the outcome of the vote, he really thought that he had won. Actually, I think he did win, but the majority of people did not want him to win. If the voting would have been 52 against 51, with nobody abstaining, I am not sure that there would have been any room for discussion. I know that John Cafarella escalated it to the higher ranks in the IEEE committee, but when he started to get some traction there, he was already being overtaken by the facts.

In the same week that the IEEE meeting took place, Lucent and Harris sat together, realizing that they quickly needed to compromise and come up with something. To a large extent, it was Jim Zyren from Harris and Richard van Nee from Lucent who figured out a new radio transmission scheme, different from anything that had been proposed so far. It was called CCK (Complementary Code Keying). Actually, I think it mostly came from Richard, although I do not want to shortchange anybody else involved. The advantage of this proposal was that it did not give any real advantage to any player, in particular not to Harris and Lucent, and therefore was acceptable to both and to the larger part of the IEEE members. This meant that six weeks later, at the next IEEE meeting when this new proposal was brought forward jointly by Lucent and Harris, most of the membership had already decided they wanted to forget the outcome of the previous meeting's vote.

This is how the IEEE 802.11b was born – a difficult birth and initially not a particularly beautiful baby. But as happens often in history, this baby was destined to make a difference for wireless LANs, for 3G (UMTS), and for the whole telecommunication industry.

As far as John Cafarella goes, he came out smiling after all, as his company Microlor was acquired about two years later by Proxim. And when Proxim started suing half of the wireless LAN industry, this was partially based on the patents that Proxim had acquired via this acquisition.

6.7 EUROPEAN STANDARDIZATION: HIPERLAN/1 AND HIPERLAN/2

Europe had also started to look at wireless LANs, under the umbrella of the ETSI (European Telecommunications Standards Institute). The European regulatory agency was never really enthusiastic about the 2.4 GHz band for wireless LAN, as this was an unlicensed band for ISM (Industry, Scientific and Medical) applications. So, in Europe, the 5 GHz band was set aside specifically for wireless LAN communication, and the ETSI was requested to standardize the MAC and the PHY.

Originally my organization was quite enthusiastic about this opportunity, as we saw it as the possible next-generation, higher-speed product. Just like Vic Hayes in the IEEE, our organization delivered the chairman for the HIPERLAN/1 committee, Jan Kruys, and the work for this effort started quickly. Our goal was to keep IEEE and ETSI aligned – for simplification, and because we saw in the future that notebooks would be traveling all over the world. Putting the burden on the end-user to tell the computer in what country he or she is, did not seem very feasible to us. Unfortunately, the HIPERLAN/1 committee fell into the hands of some people and companies with interesting ideas that turned out to be very difficult to implement. The concept that was brought forward was the concept of an “access point-less” MAC protocol – a protocol that could work without direct intervention of access points. Messages (data packets) would just hop from one computer to another computer until finding the correct computer or an access point “to go on the wire.” In technical terms, it would mean that every computer would essentially be a bridge. The biggest supporter of this solution was Apple Computer, who at that time had a large office in Paris, where the research department had identified HIPERLAN/1 as a primary target. The idea is pretty neat, and probably could be made workable, but it would also run into major complications. Special attention would be required to make the ease-of-use good enough, but also the security concerns were very high. It would be necessary to avoid the interception of messages or the “injection” of messages as if they were coming from someone else. That would not be easy to implement. On top of this, there were concerns about performance. If your PC or notebook happened to be in a “central” spot, it could become quite busy with all the receiving and forwarding other PC’s messages, not only slowing down your PC, but also eating into the battery power. Features could probably be developed to reduce the “cooperation” to the wireless networking, but that would defeat the purpose.

I do not know of anyone who implemented HIPERLAN/1. I know that Proxim and Intel looked at it seriously for some time. In general, there was considerable skepticism about how well this would work, quite apart from the additional concerns. Would a completely open and fluid concept of radio terminals be able to maintain robust network connections? And what would be the impact on overall performance?

On top of this, the industry was very skeptical about the 5 GHz as a good frequency for wireless LANs. The higher speed was attractive, along with the somewhat cleaner spectrum. However, with the frequency going up from 2.4 GHz to 5 GHz, which is more than doubling, the range goes down by about a factor of four, so the coverage of an access point would go down by at least a factor of ten. This meant that it would require at least ten times as many access points to cover a comparable area in the 2.4 GHz, if everything else is kept the same. There were later methods developed to mitigate this problem, but with the state of the market in and around 1998, this was not yet in sight. Also, the scope of HIPERLAN/1 was clearly Europe, with maybe access in the US, which did not help to make this attractive either.

Actually, this was all about the MAC protocol. The radio of HIPERLAN/1 was quite straightforward. It ran in the 5 GHz, and this was not an ISM-band frequency like the 2.4 GHz that required special constraints around the allowed type of radio transmission. Practically speaking, it did not require spreading (i.e., the need for using DSSS or FHSS, as discussed previously), and it could be a high-speed, narrow-band radio. The actual result, then, was a raw data speed of over 20 Mb/s, actually 23.5 Mb/s under ideal circumstances. Let me put it this way – if nothing else, the work on the HIPERLAN/1 standard was clearly pointing to the direction that higher speeds in the future would be pursued. But during the completion phase of the HIPERLAN/1 standard, there was already an awareness building that this had been an interesting artificial exercise, and a remarkable trial, but essentially a total waste, as the IEEE was already starting to develop a 5 GHz standard with even higher speed.

The ETSI response to this was to start aligning themselves more with the IEEE. Jan Kruys of the Lucent organization headed the ETSI committee, and he was very close with Vic Hayes, also from the Lucent organization, who was heading IEEE. ETSI somewhat cleared the decks and renamed the working group for

these standards from RES-10 to BRAN (Broadband Radio Access Networks). Actually, they want to cover not only wireless LANs, but also the area of wireless local loop.

In the wireless LAN working group, Apple was gone, as things were not going very well for Apple in general at this moment. Initially our idea was to align ETSI as closely as (politically) possible to the IEEE. This was very successful for the radio part, as the 5 GHz ETSI-RES radio is very similar to the IEEE 802.11a radio – same transmission methodology (OFDM) and same channelization (the way that the band is split in multiple channels over which the transmission takes place). But for the MAC, there were clearly other powers-that-be in Europe. The flaws of the HIPERLAN/1 MAC were clearly accepted, but there was considerable criticism of the IEEE 802.11 MAC. For instance, the way Quality of Service was absent in the IEEE 802.11 MAC was considered a serious negative. Quality of Service is the feature that enables a protocol to transmit time-sensitive information like voice and video. This was a feature that was not originally part of the (wireless) LAN required feature list but was getting more attention as companies are starting to use the Internet, and therefore the LAN for telephone calling. Beside the lack of Quality of Service, there were also items like power control and frequency selection that the European regulators wanted to see implemented in a certain way.

As usual amongst engineers, the thinking was that the best way to fix this all was to develop a new MAC, again from the ground up. Considerable lobbying and voting took place, and after all, an IBM proposal was accepted as the basis. This proposal was in essence a newer version of the proposal that had lost the voting in the IEEE a few years earlier. Again, a lot of work was done to define this new MAC, and the result (again) was a nice standard. Some companies started to implement it, especially the larger European telecom product providers were initially very much in favor – Ericsson, Nokia and Philips in particular. At Lucent, we made the assessment about the acceptance of the IEEE 802.11a in Europe, and we concluded that most likely IEEE 802.11a would be enough in compliance with the European radio regulations to be acceptable. So, we decided not to waste any time on this. We turned out to be right. Slowly but surely, everyone turned away from HIPERLAN/2 again, just as with HIPERLAN/1, and as of today, I think no company is working on any HIPERLAN/2-based product development. In particular, the decision in IEEE

to add functionality to the MAC that addressed the needs of the European regulators, and also to start to work on Quality of Service, eliminated the need for anyone to seriously invest in HIPERLAN/2-based product development.

Despite the huge success of the European legislation with GSM to standardize voice communication, all the European efforts and investments to play a role in the standardization of the data communication turned out to be a big waste of money and time. Where top-down “guidance” in the wireless voice communications had worked, it was not working in the much more dynamic and versatile world of data communication, where IEEE and the computer industry turned out to be the winners.

6.8 WHAT ABOUT IEEE 802.11G

It is interesting to watch where a competitive world leads us. But for standardization, it does not necessarily mean that the best solution is found the quickest.

The industry had settled for two standards, one for today (IEEE 802.11b) and one for the future (IEEE 802.11a) – and as discussed, the future standard was there before the one for today. Most of the industry has run for implementing products in the 2.4 GHz following the IEEE 802.11b standard, while putting investments in the 5 GHz on the backburner. One exception was Atheros, a company that developed technology solely focusing on IEEE 802.11a. The future seemed to them the best bet. And in a way, Rich Redelfs, their CEO, has it right – IEEE 802.11a will arrive. The only question is when. The computer industry, however, is not planning to wait until all the regulatory issues have been sorted out, and they placed their bets on IEEE 802.11b, making it a real success.

This is causing a problem for the IEEE 802.11a standard, as customers want to have a migration. An enterprise that has installed IEEE 802.11b with 2.4 GHz base stations (access points) everywhere, does not want IEEE 802.11a, which essentially means a complete reinstall of the base stations. The extra complication is the footprint. Because the indoor radio characteristics of the 5 GHz are quite different compared to the 2.4 GHz, a switch from IEEE 802.11b to IEEE 802.11a is not very attractive. Some companies are trying to compensate for this by providing base stations that support both standards as an interim solution, but this is not more than a stopgap, as it does not resolve concerns about worldwide legislation for the PC industry or the difference in footprint.

The conclusion was that the industry needed so-called “a/b-combo cards.” These were radios that support both the IEEE 802.11a and IEEE 802.11b. They can be compared to dual band (or triple band) GSM phones. The user is ignorant about the infrastructure capabilities, but the radio card in the computer tries to find the best way to connect to the internet, whether it is via IEEE 802.11b in the 2.4 GHz, or via IEEE 802.11a in the 5 GHz. The problem is that such an a/b combo card uses a significant amount of power, and therefore drains the battery of the notebook very quickly. Beyond that, the computer industry would not be very interested in an a/b combo card for a notebook PC, since it would not be a worldwide-shippable product.

In the meantime, the pressure was mounting on the industry to come up with higher speed solutions than the 11 Mb/s of IEEE 802.11b. This led to a conclusion to try to get the 54 Mb/s OFDM technology approved in the 2.4 GHz. This was unheard of in the past, as the requirement for the 2.4 GHz had always been that only Spread Spectrum technologies would be allowed. But the FCC in the US had created a serious problem for itself – they had never narrowly described what Spread Spectrum was. They had never described clear rules to measure minimum spreading of a radio signal. The risk that the FCC started to run was being sued for inconsistent rulings, and their behavior became extremely cautious. In Europe, the situation was somewhat simpler. OFDM in Europe would be allowed, as long as the signal strength over time would stay below a certain value.

In 2001, a meeting has been arranged between the IEEE and the FCC to understand the necessary conditions under which the FCC would allow higher speeds in the 2.4 GHz ISM band (i.e., would allow dropping the spreading requirement without getting into trouble based on older rules.) The outcome of this meeting was a quiet statement by the FCC that they would approve a radio technology that was created by the IEEE. This was the way that the IEEE 802.11g standardization committee was born: it would create a higher speed wireless LAN standard for the 2.4 GHz, compliant with a more lenient interpretation of the FCC rules in the US and compliant with the rules in Europe. The choice for the PHY of the IEEE 802.11g was in favor of using the 5 GHz technology, called OFDM, and get it scaled down to the 2.4 GHz. This way there was an interesting migration path created between the two other standards, as IEEE 802.11g is using the same band as the IEEE 802.11b, but it uses

the transmission technology of IEEE 802.11a. The advantage is clear. While it still may take a few more years before IEEE 802.11a worldwide legislation is established, with the IEEE 802.11g standard, higher-speed products can be rolled out at a lower cost and with no legal restrictions anywhere in the world. The expectation, therefore, is that companies and the computer industry will start rolling out IEEE 802.11g products soon, and a few years later the so-called a/b/g combo cards will find their way to market.

By the way, this scenario is heavily contested by some companies whose interest in IEEE 802.11g is very low. In their scenario, they prefer a/b combo cards going into the market quickly, and then moving to a/b/g combo cards later. Although this is a potential scenario, its feasibility really depends on the market leaders of IEEE 802.11b products, Intersil and Agere Systems. If they have 802.11g chipsets, reference designs and products available in time, the route will definitely be from .11b to .11g, and to .11a/b/g later. If the .11g development programs are slowed down too much, and the interested companies find a way to make .11a/b combo cards attractive for a large enough market, the route will be from .11b to .11a/b, and to .11a/b/g probably later.

The market jury is out and will decide in the coming 12 months.

6.9 BLUETOOTH

While in the wireless LAN world we were all working very hard to get the technology of wireless LANs endorsed in the industry, in 1998 we all received a sudden surprise in the form of the Bluetooth announcement. The core technology for Bluetooth was developed by Ericsson in Sweden and the Netherlands, interesting as significant parts of the core technology of Wi-Fi (both IEEE 802.11a and IEEE 802.11b) were also developed in the Netherlands. But Ericsson rightfully realized that it would be very difficult to set a de-facto standard for this technology, so they started a standardization committee with several companies – a SIG (Special Interest Group) as it was called, to get this technology standardized. Besides Nokia and Motorola, they enticed Toshiba, Intel and IBM into this effort, in a serious attempt to cross the boundaries between the telecom world and the computer world. This crossing of the borders between telecom and computers was a real breakthrough, considering how both industries had moved separately from each other over the years.

The original idea of Bluetooth was a low-cost wireless connectivity device that would connect phones with computers, or get phones and computers connected with peripherals like headsets, printers, and PDAs. That was, and is, a very practical idea. Comparing it with Wi-Fi, the Bluetooth technology characterized itself with short range, up to 10m (30 feet) and low speed, under 1 Mb/s. For comparison, Wi-Fi covers up to 100m (300 feet) and 11 Mb/s. But the claim to fame for Bluetooth was that it was advertised as extremely low cost. The target of \$5 was mentioned in early Bluetooth press releases.

Besides the technology differences, there were differences on the business side. Looking at the wireless LAN world, IEEE 802.11 does the standardization, and the Wi-Fi Alliance (at that time called the WECA, the Wireless Ethernet Compatibility Alliance, puts the infrastructure in place to certify interoperability under the Wi-Fi logo and does the marketing/certification of Wi-Fi as the global wireless LAN standard. So, IEEE 802.11 does the wireless LAN technology, the Wi-Fi Alliance does the wireless LAN “business”.

This historical split does not exist in the Bluetooth world, where both technology and business are covered in the SIG.

This integration may look like an advantage, but in practice it has shown to be a big disadvantage. The marketing promotion of Bluetooth moved ahead much faster than the technology could keep up and created a promotion disaster. In the Spring of 2000 in CeBIT, the goal was to have more than 100 Bluetooth devices working together; but it failed miserably, forcing the standardization to retreat. And the \$5 price target never became clear – was it the price of the device for the end-customer, or was it the price of the chipset that goes into the device? When in 2001, the first products came into the market with prices over \$150 or beyond – significantly exceeding the price of Wi-Fi products – this was another blow to the credibility of Bluetooth.

In the meantime, a serious mistake was made in the Bluetooth positioning. Despite the fact that the Bluetooth technology tried to address low-cost, low-speed, short-range peripheral interconnects, the Bluetooth marketers positioned the technology as replacing LAN data connectivity, which is high-speed, longer-range connectivity between computers – the area that was covered by IEEE 802.11. Apparently, the cooperation in the Bluetooth SIG between the telecom world

(Ericsson, Nokia and Motorola) and the computer world (3Com, Intel and Toshiba) was not so well organized to avoid this. Bluetooth created an attitude of planning to “kill” IEEE 802.11 and Wi-Fi. This was probably fueled by another fact as well – both Bluetooth and Wi-Fi are using the same 2.4 GHz radio band. But both technologies are different in how they use the frequency band. Bluetooth is using FHSS (Frequency Hopping Spread Spectrum) technology, Wi-Fi is using DSSS (Direct Sequence Spread Spectrum) technology. These two technologies interfere with each other, to a certain level. This interference leads to a reduction of quality in the Bluetooth link, creating drops, noticed as “cracks” in the communication. It also leads to some reduction in the range of the Wi-Fi link and/or of a reduction in speed, because of the loss of packets, that then need to be retransmitted.

Between 1998 and 2001, serious initial efforts have been undertaken to “harmonize” Wi-Fi and Bluetooth, to avoid interference. Several tests have been done between Ericsson and Agere Systems to understand the level of interference and to build models for protocol adaptations that would reduce or eliminate this. Several other companies have also seriously addressed this with special features in software, but as of today, there is no fundamental solution for this interference. Initially the efforts to resolve this were effectively frustrated by the Bluetooth SIG and their mindset to “kill” Wi-Fi altogether. Probably, though, Bluetooth ended up having enough internal compatibility problems of their own to be able to pay too much attention to this.

In the 802.11/Wi-Fi world, there has always been a clear recognition of the fact that a low-cost, low-speed, short-range wireless solution would be feasible – and potentially an attractive market opportunity. Serious frustration about the fact that the Bluetooth community had no interest in working out the interference problems pushed IEEE 802.11 to create a new group called IEEE 802.15. This group’s charter was to define a standard for a low-cost, high-speed, short-range peripheral interconnect – actually a high-speed successor for Bluetooth. As part of the 802.15 charter, coexistence with IEEE 802.11 is explicitly defined. Because many companies nowadays have representation in both Bluetooth and in IEEE 802.11 and 802.15, the attitude towards harmonization has improved significantly.

Bluetooth itself is also recovering from its initial marketing problems. It has realized that there is a difference between peripheral interconnect and LAN connectivity,

and it has decided to focus on its core mission of a low-cost wireless peripheral interconnect, realizing that LAN connectivity requires so many additions to Bluetooth that it would end up with something like today's Wi-Fi.

Bluetooth's main problem was interoperability, and the levels of freedom in Bluetooth 1.0 were too high to force this interoperability to happen. So, part of the refocusing effort of Bluetooth was to clean-up and tighten-up the efforts and come out with the Bluetooth 1.1 – although general interoperability still has a long way to go. This has to do with the fact that Bluetooth not only tries to standardize the interface, but also tries to standardize the application used over the interface, in Bluetooth terms called “the profile”. One of the Bluetooth profiles, for instance, is the “wireless headset” application, connecting a headset (ear piece and microphone) wirelessly to a cellular phone. The ultimate goal is to have “any” headset of “any” vendor working with “any” cellular phone. To be frank, I am quite skeptical about this goal. Not so much because it would technically not be feasible, but more so because currently, the industry players are probably more interested to ensure that once you have bought a phone, you buy the same brand headset.

Nevertheless, Bluetooth is clearly on a path of recovery, and the expectation is that with the growing popularity of wireless headset applications, the volume will increase so much that the price will come down enough to trigger many other applications – especially the interconnect between phones and computers. Currently the computer industry is getting more seriously interested after the initial high price failures from Toshiba and IBM. The prices have come down, and the integration of a Bluetooth radio and a Wi-Fi radio into one device for the 2.4 GHz has come within reach. This is important, as it will create a seamless connection between the computer and the phone. This enables a computer to connect to the Internet via Wi-Fi directly in a so-called hotspot, as well as on other places via GPRS over the cellular phone.

6.10 SUMMARIZING THE THIRD PERIOD (1995–1998)

In this third period, the wireless LAN industry almost went under. They were nearly overwhelmed by the hype around the cellular phone industry. The cellular phone industry did not understand LANs and was starting to make the same mistake as their wired predecessors has made with ISDN – assuming that data was just a variation of voice. The companies that worked the wireless LAN field

were struggling to keep their heads above water. Although I think most of the companies were marginally profitable, the large expected growth was just not there. Every year, the forecast moved out another year.

Many executives were expecting that wireless data was going to be “something,” but the companies that could make a difference did not trust the unlicensed ISM band, and actually they did not trust an Ethernet-like protocol either.

The telecom industry represented by Lucent Technologies, Ericsson, Nortel, Nokia, Alcatel, Motorola and Siemens were riding their magnificent growth numbers and were already looking for the “next big thing.” Fiber was the area where the investments went, and 3G (UMTS in Europe) was the mantra for the next revenue spurt. Wireless LANs were too flaky, too uncertain, and too PC-centric, which means maybe interesting for the enterprise market, but not sexy enough for the consumer market.

The computer industry was too wrapped up in driving cost reductions. Compared to the telecom industry, the model in the computer industry is completely different. The computer industry is so cost-reduction biased that a lot of innovations have taken place in logistics and distribution. I think there are few industries, maybe with exception of the car industry, which have focused so much on cost reductions. Innovation has clearly suffered, and the PC of today hardly differs from the PC of five years ago. The only innovators in the computer industry maybe are Intel and Microsoft, and I think they are currently focusing more, probably rightly, on improving security, reliability and cleaning up “the bugs of the past” versus any real innovation. Still, the breakthrough for the wireless LAN industry came via the computer industry, from another true innovator in the field – Apple.

THE BREAKTHROUGH (1999–2000)

In retrospect, in the second part of 1998, upper management at Lucent Technologies was getting really fed up with wireless LANs. At this stage it was regarded as the eternal promise, never coming through. Slowly but surely, resources were being sucked into a promising new concept called Wireless Local Loop. Wireless LANs seemed to be destined for permanent abandonment. However, as happens so often, just when everything seemed to be lost, the resolution arrived. The real market launch of wireless LANs out of the vertical solution space took place this year and changed our world forever.

7.1 WIRELESS LOCAL LOOP

The name of the game in Wireless Local Loop is the price, or more specifically, the cost of the CPE (Customer Premise Equipment) – simply stated, the unit that goes on the house and wirelessly connects to the central base station of the neighborhood. The whole idea of this was quite simple. It is very expensive to dig up the street to get the access to every house to run a telephone wire; doing this wirelessly would save a lot of money. Unfortunately, the problem turned out to be that the CPE units needed to be installed outside of a house and carefully pointed to the central base station. The consequence was that the price per connection – the price per house – just shot up to a level too high to compete with wire. But we were not that far in the late '80s, when many companies invested highly in WLL, and so did Lucent Technologies.

For WLL systems, there were essentially two target markets. In the first place, the low-end market, where the immediate need was for plain telephone – the new telephone markets in the so-called third world. The second market was the high-end market, where the pressing need was for services, in particular DSL, that could not be run over traditional wire, because it was too old, for example, and noise sensitive.

Our target was to make a simple CPE unit in the 3.4 GHz that would cost less than \$120. Our core expertise was low-cost radio development and manufacturing, so we started the work in 1998 and made very significant progress. The challenge of such a product was not only to meet the cost targets, but also to comply with a very stringent temperature specification. The product should be usable in cold, below-freezing temperatures, as well as under hot, direct sun exposure, and in

extremely dry or extremely humid conditions. From a technical perspective, we were able to make such product – I think a major achievement. But when the product was ready, in late 1999, it had become clear that that the market was not buying the product. Most of the telecom companies had started to freeze their investments in WLL, as a precursor of the doom and gloom that was going to come over the telecom industry.

As of today, the WLL industry has largely come to a standstill. With the exception of a few courageous efforts, this market is dead. But it is important to realize that Wi-Fi may have the capabilities to revive this market. The main reason is that Wi-Fi could bring down the cost of the CPE unit significantly, or could combine the CPE unit with other functionality, like the DSL-box and/or the Wi-Fi residential gateway functionality. Wi-Fi itself, or its potential successor in IEEE 802.16, may also come to help here. But the telecom industry itself should first overcome its doom and gloom.

7.2 THE APPLE STORY

In 1998 and 1999, all the hard work to create understanding and acceptance for the technology suddenly came to fruition. Although I believe it is not necessarily the situation for all technology breakthroughs, for Wi-Fi wireless LAN, there was a clear set of events that really put this technology on the map – even when the name of Wi-Fi did not exist.

We tried hard with many computer industry critical companies. I remember personally calling AST Computers in their heydays. I also remember personally calling Michael Dell in the summer of 1994 and speaking with his secretary; and that I got a call back from a European sales guy who had no clue what I was talking about. I spoke with many people at IBM. I spoke with Toshiba, with Mr. Nishida. Actually, he had the right vision, but we were not ready or at the cost level that he considered necessary to have this technology breakthrough.

Last but not least, we spoke several times with Apple Computer. If there is one person who caused the wireless LAN world to happen, it was Steve Jobs, at that time interim CEO (he called himself the i-CEO) of Apple Computer Inc. He decided to select wireless LAN as a differentiating feature for the iBook launch in 1999. Credit to him and his organization for being at the center; and I think Apple has greatly benefited from this as well. Still today, Apple has a leadership

position in the PC industry on wireless LANs, although now other companies like Dell, Sony, IBM and Toshiba are catching up rapidly – in internal knowledge as well as in “attach rate,” the number that defines the percentage of (notebook) computers that have a wireless connection when shipping.

So, all praise to Apple, but not for having done anything to develop the technology. I think Apple’s quality was to “be there” when the technology and standardization had reached the level that economic viability was just around the corner. That was the moment when they struck, so the praise is about timing. Which is everything, after all.

In the years before, we had targeted Apple to sell wireless to, but with no success. Apple had their own thoughts about wireless and was more in favor of FHSS, the Frequency Hopping technology, as they kept erroneously thinking that this technology had better interference resistance than DS, the Direct Sequence. Apple had a whole history with wireless going back to the beginning of the decade, where they were secretly trying to develop this technology, inside and outside the standardization bodies. In the US they seemed to have had activities with Motorola, and in the UK, with Plessey. They also had several passes with a Bay-area startup company called Photonics, using infrared technology instead of radio. Interestingly Dick Allen, who led Photonics, spent many years in Apple and was the Frequency Hopping supporter who changed his mind in 1998 and helped Apple to create this wireless LAN leadership position.

Apple was the first PC vendor who launched wireless, causing their competitors, especially Dell, to take swift action. But it took all the PC vendors about a year to also include wireless LANs after Apple did it. Apple actually gave us quite a lot of heat for this at that time, as they thought we were not aggressive enough with other PC vendors to further popularize the technology. In reality, the implementation with Apple had taken more than one and a half years from the decision to implement to the launch, a fact that they tended to conveniently forget.

The whole interaction with Apple was an experience in and of itself, and, it seems, quite characteristic for any dealings with the company. After comparing notes with other vendors who tried to sell technology to Apple, there were a lot of similarities. First of all, they clearly have a mindset and schedule of their own, which leads to no movement at their side until the moment they make a positive decision; at that moment a supplier cannot move fast enough.

Early in 1998, Apple contacted us to express interest in wireless LANs. Actually, this was Steve Jobs directly, through one of his lieutenants. I have to admit, after unsuccessfully trying to sell to them for eight years, there was initially some level of skepticism about this call from Apple. In particular the way the interest was communicated – “Steve Jobs wants to have a meeting with Rich McGinn about wireless LANs.” Rich McGinn was the CEO of Lucent Technologies that our organization was part of that time, and Rich’s exposure to wireless LANs had been virtually zero. I believe Lucent Technologies at that time was about ten times as large as Apple (\$30B versus \$3B). Telecom operators were the key customers, large switching and cellular base station deals were the core business, and someone asking for access to Rich McGinn for a deal far below \$100M was somewhat odd. I remember that Rich McGinn’s staff was trying to keep Apple away from him, so we had to fluff up the wireless LAN potential, threw in some other technologies on the agenda (DSL), and managed to get Rich McGinn interested. Steve Jobs own persistence probably helped as well, as he kept his people calling on everyone in the Lucent Technologies chain of command. Personally, I believe that for Rich McGinn, it was only the DSL acceptance in the computer community that was an interesting subject, and not so much the wireless LAN sales expectations for the coming quarter (nothing) or maybe the year after (\$5M).

The meeting date was set for April 20, 1998, and a pre-meeting was scheduled in the Peppermill restaurant opposite the Apple headquarters in Cupertino. I still remember this meeting quite clearly. A small Divisional manager did not have frequently contact with Rich McGinn. He came in with John Dickson, the head of the semiconductor division; actually, if there was any division at Lucent doing business with Apple, it was this division. As he is known to do, Rich immediately took the initiative by firing off questions. “How much are we going to sell?” “Explain wireless LANs to me.” “Why is Apple interested?” We did a preview of the presentation we were planning to give to Steve Jobs, and got a “Let’s go.” Actually, Rich McGinn quite impressed me in the way he was able to use the right words at the right moment in the meeting, as if he had believed in the wireless LAN business for years and had been personally pushing the technology, which he had just learned about.

The meeting in the Apple board room in Cupertino was peculiar, but interesting. Lucent had drummed up probably ten of its brass, and Apple had done the same. Cordialities were exchanged, business cards, the usual pleasantries.

But as anyone who has done business with Apple knows, this only happens the first meeting. Subsequent meetings are a little more challenging, shall we say. The meeting started at 2:00 PM, with the companies sitting on opposite sides of the table. Lucent was sitting there with suits and ties; Apple was showing up California style. No Steve Jobs, which was a little awkward; Steve had been delayed.

Lucent wants to talk and present, but Apple is waiting for their king. Then the king comes in, Californian style too, walks over to the Lucent suits and shakes hands with everyone, without introducing himself. I was thinking, “who is this guy?” As a European, I guess I had not really been exposed to many photos of Steve Jobs, and in Europe, corporate people do not present themselves in the media as movie stars either. Then Steve Jobs sits down and starts talking, saying that wireless LANs are the greatest thing on earth, this is what Apple wants, for about ten minutes straight. I believe that Rich tried a few comments, but couldn’t get a word in. Then Steve Jobs asked, “Are there any questions?” I remember that I tried to prepare a few slides – key winners, market positioning, product offering, value creation, etc. Presenting with Steve Jobs is actually quite easy – you put up the slide, and he will do the talking, even if it’s not necessarily related to the slide. Then he asks for the next slide. Rich McGinn chiming in a few words, interested in DSL, he thinks 1999 will be the big year for DSL, “Will Apple be ready?” In other words, “Will Apple PCs have DSL?” Steve Jobs replied, “Probably not next year. Maybe the year after. Depends on whether there is one standard worldwide.” Turning the conversation back to wireless LANs, he says he needs the radio card for \$50, wants to sell at \$99.

Then Steve apologizes; he has to leave. Stands up, says “bye!” and goes. Room is quiet.

This was clearly the extreme of the telecom world meeting the PC world, although, yes, Apple is somewhat of an extreme sort of a company itself. For Steve Jobs the work was done. But for us, the work started in a way we had not seen before, and it was going to be very intense as our cost at that moment was probably above \$100. We had the chipset in development that probably would lead to a cost of a little above \$50, but it was not clear by how much.

In the subsequent months, we went through several rounds of product definition. Apple wanted a special interface. Actually, it took quite some time before the real negotiations started, as both Apple Computers and Agere Systems tried to figure out what the total system proposition should be, in particular what access points (base stations) needed to be added, which all confused the total picture significantly. Initially Apple wanted three different access points to be added – a high-end one, a medium one and a low-end one. But later in the project, they defined their own access point and dropped ours. Originally the plan was to launch in the Spring of 1999, but both Apple and we slipped into the later part of the summer of that year.

During this period, I learned a negotiation technique that may be interesting to share. The technique is that you ask your supplier to take a loss on one product and make it up with a higher margin on the other. Then after agreement, you decide that the product that makes the higher margin is giving the supplier an “unfair” windfall, so at an appropriate time the negotiations are opened again on that specific product, while the lower margin product is conveniently forgotten. In projects that run for a while, there always seems to be the opportunity “to muddy up the water,” by changing the requirements, for instance, or threatening to drop the whole business altogether. Or the most effective way – by threatening to take away other business in the future. I think Apple has the most sophisticated supply line process, bullying elevated to a form of art. Or we were just plain stupid. But key to the process is to have different discussions with different layers in the organization. And really, Lucent was not so sophisticated at that time – and also not that interested in wireless LANs. Getting DSL going was more important.

Fortunately, we had never agreed to take a loss on the radio card product. The negotiation agreement was cost plus 5%, where we agreed on doing joint negotiations with suppliers. When the product actually launched for \$99, the industry was shocked. We were accused of “buying” the market, and that we were losing money on every card. But we were not. The mechanism we used was to “forward” price the product, so the volume would go up substantially, the cost would go down quickly, and the market share gained would bring in the margin. That is the theory, and it worked. In fact, it worked very well, as we would see in the years after.

Another very confusing element was the fact that the agreement we closed with Apple was on the first IEEE 802.11 standard product, 2 Mb/s. Our thoughts were that we would lower price of the 2 Mb/s products anyway and be able to make better margins on the 11 Mb/s when the .11b would come out. During 1999, the 11 Mb/s standard solidified, and it became clear that the goal needed to be to break in the next IEEE 802.11b product, the next generation 11 Mb/s. When these negotiations were going on, the real question was, “When is the next generation standard going to hit, and will we be ready for this?” What ended up happening, after some stretching and delaying, is that Apple went straight for the next generation standard. We tried to negotiate a better price for this newer standard, as we felt that the market would bear that. This was clearly not in Apple’s interest. Actually, wireless LANs were not in Apple’s interest, as such. Their goal was selling more PCs. So, we were given the choice – continue with the program and deliver 11 Mb/s (IEEE 802.11b) for the price of 2 Mb/s (IEEE 802.11), or see the whole program cancelled. We decided to give in, something that looked like the wrong decision at that time but has contributed to the very quick commoditizing of the wireless LAN business.

There was another real negotiation trick that was played in what was going on. The first round of agreement was based on every Apple notebook having a wireless LAN card. This was called EUJ – a wireless LAN connection as Every Unit Item for all Apple notebook PCs. This would guarantee to us the required volumes to create a profit, and that would be attractive enough to take a low margin. We would “make it up in volume.” This was agreed in December 1998 and was the justification for taking the business at a cost plus 5% level.

Well, in March 1999, Apple’s understanding of the program suddenly changed. All the notebook PCs would “be able” to carry a wireless LAN card, that is to say that all notebook PCs would have buy for a radio card and an antenna integrated at the back of the screen. But the deal to sell one radio for every notebook sold was off, and we saw the volume dwindling to about one-twentieth of what the deal originally was based on. I remember that we seriously considered canceling the program at this stage. But we had tried so long to break in into the PC industry, and I believed that once this breakthrough was established, the wireless LAN market could be huge. Apparently, Apple understood this all too well, and saw an opportunity to exploit it.

This negotiation trick was not the last one. Another trick was around a very important feature of the access point, called roaming. It allows users to roam around buildings covered by multiple access points and be continuously connected while the connection switches from access point to access point. In Lucent Technologies, we were working on a family of access products differentiated in power to support different numbers of users, ranging from the home user access points without roaming to very large access points supporting many hundreds of users. The agreement was that the original Apple Airport would work only in a home environment and had no roaming, as the roaming capabilities were a higher-end feature for which Apple would buy and resell a Lucent Technologies product. This was all fine and agreed, even put in contract, until the moment that the rollout started. Apple had thought about the concept again and preferred to get the roaming feature in their Airport base station, so they could sell it to schools as well. But what about the higher-end product that would be bought and resold? Clearly, Apple had changed its mind. One might think of this as a great endorsement. However, in business there is such a thing as a required price to make a decent margin. Adding a critical feature to a product that sells for a rock bottom price in a home market to allow it to be sold in a higher-end market like a school or enterprise is asking for a market to erode more quickly than one can afford to build it up. The simple possibility is always to say “no” to a customer and explain that you cannot afford to give this feature away for free, as it erodes the price in other market segments. And then you learn again the hard way that bullying suppliers really exists, when your CEO gets a call that he may lose business in a different, unrelated business unit if he does not provide this feature for free. Apple won this battle as well, despite our earlier agreements.

From a theoretical business perspective, the conclusion might be that there are also disadvantages to being part of a larger organization, as it exposes you to potential attacks on your business through a colleague’s business. But that was just theory, because what happened was that the price in the access point market now started to erode much more quickly than necessary. Apple’s comment was, “If we had not forced you to take the price down, someone else would have.” My reply was, “But not yet!” But I do not think anyone heard. It was another good lesson in supply management tactics though!

7.3 OTHERS PC VENDORS FOLLOWING

Despite Dell's furious reaction to Apple's announcement, it still took about a year to change this into real sales of products. Dell had assigned a tiger team and was trying to find the quickest way to resolve the situation. The problem at Dell, though, is that they always try to balance quick and quality. And while quick was required here, the team really started to focus on quality – what is the best we can get in the industry, and how are we going to integrate this. Dell, like other PC manufacturers, has a major disadvantage compared to Apple: their Operating System is coming from Microsoft (Windows). So, where Apple and Lucent and Agere could resolve software and Operating System interface problems straightaway, Dell had to forward Agere to Microsoft. Microsoft themselves had gotten a little fed up with all the interface problems was all third-party peripheral manufacturers they were exposed to and had developed a new procedure pushing the testing and responsibility away from themselves by defining a certification process called WHQL (Windows Hardware Quality Labs). WHQL, pronounced “wickel,” has probably kept more software programmers awake than any other word. It is a stamp of certification that is required to avoid having Windows state at start-up that you have uncertified software running on your machine, your computer or notebook. Well, this really did well for the PC industry that wants to reduce the number of telephone support calls as much as possible. So, everyone delivering peripheral products to a PC vendor, like a wireless LAN NIC card in our case, needed to make sure that the WHQL stamp of approval has been obtained.

Unfortunately for us, there were some requirements in the Windows certification program that could not be met with wireless LANs, particularly in the area of immediate “being alive” messages. So, this required us to engage with Microsoft to explain the concept of wireless LANs and the needs that this technology puts on operating systems. To be fair, it was also for us to learn about making rock solid drivers that kept on working under all circumstances. Initially, we had to make some compromises with Microsoft to obtain waivers to expedite our entry into the market. It took a while before the WHQL tests and the wireless LAN drivers could meet the required quality levels. It sounds easy to describe now, but the reality was somewhat more gruesome, as Apple was still feasting in the market with their wireless LAN solution.

Almost in parallel with Dell, we were approached by Toshiba, Compaq, HP and

IBM to provide wireless LAN solutions for them. So, in the summer of 2000, there were many announcements from many of the PC vendors about their wireless LAN solutions – wireless Ethernet had arrived. Not that everybody was immediately convinced. Issues around security and reliability kept hovering. But the price had come in below \$100, so the investment risk had become relatively low.

As Agere Systems, we had almost a clean sweep of the wireless LAN market for PCs – an amazing result after the initial efforts with Apple. The interesting experience with the PC vendors, though, is that they really are trying to claim exclusivity with a supplier. Not that this is in writing, or necessarily even spoken, but PC companies are running on such small margins that they have essentially pushed out all the technical work into their supplier base. The technical part of a PC company has very much become a management shell, where suppliers are managed to deliver their technology, and where the technologists in the PC company are trained to ask questions and make the right judgment calls in their supply base. The technical support requirement on a supplier, therefore, has become incredibly large – so large that it becomes a strain on every organization to support multiple companies at the same time. I suspect this is part of the supply line management strategy – if we get them to spend all their time supporting Dell, they cannot put in time supporting IBM or Toshiba, so keep on asking! It was clear that with the success of Wi-Fi, many companies rushed into this market. Suddenly, Wi-Fi was hot.

7.4 NETWORKING COMPANIES CATCHING ON

For many years, there had been only one networking company really engaged with Wi-Fi. This company was DEC. Actually DEC (Digital Equipment Corporation, or later Digital), in the late days of Ken Olsen, their famous CEO, had been personally engaged with wireless LANs. DEC has never invested in radio technology, but they had focused on the access point side of the business. In 1994, they launched their access point, RoamAbout, that had been developed in cooperation with AT&T, and later with Lucent. This DEC access point still exists, but it has gone through several transitions. These transitions had not so much to do with wireless as much as the woes that DEC was going through as a computer company. When DEC Computers was sold to Compaq, DEC Networking was already sold to Cabletron; and when Cabletron reengineered itself, the enterprise networking, including wireless ended up

with a company called Enterasys, which is still selling the RoamAbout product line and its successors today. Enterasys therefore is one of the longer existing companies in the wireless field.

Another networking company that was interested in wireless, and that really looked at wireless long and hard, was 3Com. I remember contacts very early on, but 3Com was so convinced about Ethernet that their position initially was to wait until the standard is there, and then we will take over the market based on our name and reputation. Strangely enough, the outcome of everything that 3Com tried for a long time was not good enough to meet their own standards – or they were just not capable of grasping what wireless meant as part of their total portfolio. They had worked extensively with Lucent to integrate Lucent's wireless portfolio, even to the point of acquiring the wireless LAN division from Lucent. Then they worked with Intersil and with Symbol, initially without a lot of marketing success despite the coolness of their product with the click-out antenna, an idea that led to their click-out V.90 modem (the one from US Robotics). Their company woes probably played a role, along with the repercussions that had on their wireless product development. Today, 3Com is in the wireless LAN business with their OfficeConnect product portfolio.

So what was Cisco up to? They had been a powerhouse in enterprise networking, and they definitely could have led the way. Actually, they did, but in the typical Cisco style and without any vision. Although the details will never be public, Cisco was pushed into wireless by Microsoft. The sequence of events was that after the launch of AirPort, Apple's wireless LAN product, Microsoft got really interested. Their relationship with Lucent and later with Agere was very open and constructive, and Microsoft got excited about wireless LAN, even to the point that they wanted to get the campus in Redmond completely covered by wireless access points. The IT department at Microsoft is probably the most challenging place to work in the world, as most of their "customers" know more about IT than they know themselves. Anyway, they were told to go wireless and they approached, amongst others, Cisco, Aironet, 3Com and Lucent Technologies (the part that was on the verge of becoming Avaya Communication) to bid on the contract. I cannot think of a contract that we worked harder on or flew in more people to sell our capabilities. However, only success counts, and the success went to Aironet. This was probably one of the most painful moments in my career, as we had excellent relationships and capabilities in Microsoft, but apparently not good enough to win over their IT department.

There was more to it. During the final bidding phase, the rumors went around that Cisco was going to acquire Aironet, as they were appalled at not having anything to offer. Apparently, this linkage was already in place when the deal was given to Aironet. It seems that Microsoft more-or-less forced Cisco to go into the wireless LAN business, and as of today, this has been quite a success for them, as Aironet is one of the most successful wireless LAN offerings.

Another big Cisco success was the way they won Boeing as an account. Boeing was interested in unwiring their facilities, and they were out with a very large contract that almost went to Lucent. The IT department had gone out and evaluated all the products on the market and clearly favored the Lucent product. Then, according to rumor, Michael Chambers, the CEO of Cisco, went in to the top IT guy at Boeing and explained to the Boeing IT department that it probably made more sense to buy the Cisco wireless LAN product, as their whole network was already Cisco. These were not even rumors; this was the inside track of the participants of the evaluation committee, who had tried to make an objective evaluation of the functions and capabilities of each product. Chambers seemed to have promised that everything that was not provided by the Cisco Aironet product would be resolved in another six months, and there went another big deal for us. In my early days in business I had been taught to ask myself the soothing question, "What is fair about a 600-pound gorilla?" I have tried to avoid these gorillas since, sometimes with less success, as in this case.

There was another reason why this move of Cisco into wireless LANs was so important, and that reason was IBM. Although IBM had grown into the largest network provider in the industry by the late 1980s and early 1990s, their betting on and promoting Token Ring technology had slowly but surely pushed them out of the game, leaving the field to 3Com and Cisco. Looking back, this could create some interesting thinking about IBM's moves to contest their loss in IEEE 802.3 (Ethernet) and their move to IEEE 802.5 (Token Ring). IBM's strategy at that time came from a perspective of controlling the market through owning the standard. But the world had changed, also for IBM, and the free forces that promoted Ethernet and TCP/IP (Internet) pushed out IBM's Token Ring and SNA (Systems Network Architecture) altogether. (I should probably make a sub note about SNA, of which the targets were far beyond the goals of TCP/IP.)

For all practical and financial purposes, during the mid and later 1990s, IBM slowly resolved all its owned networking technologies capabilities in favor of Cisco. Therefore, their wireless activities have been minimal, although they have tried several times to carve out their own niche.

A few words here about Proxim are appropriate here. As mentioned before, Proxim has always tried to do something different. Their claim to fame in 1999 was HomeRF, and they had managed to line-up several larger companies behind the idea – Intel, Siemens and Motorola are probably the most important ones to mention. But they saw their plans go up in smoke when we closed the deal with Apple that really pushed down the cost of IEEE-based products so drastically that it took away the key benefit of HomeRF. Actually, only Siemens today is pioneering some HomeRF as successor of DECT. Both Motorola and Intel have abandoned the idea. Dave King, the CEO of Proxim, has admitted that the launch with Apple was the real blow to HomeRF, and he moved on with Proxim to start to embrace IEEE 802.11 – first with the acquisition of Farallon, and later (in 2002), with the acquisition of ORiNOCO out of Agere Systems.

7.5 WECA

The wireless LAN community had always looked with a certain level of envy to the Bluetooth world. Bluetooth was already a household name years before a first product hit the streets. The opposite was the case with IEEE 802.11. This was a very solid technical standard, but not recognized outside of the wireless LAN in-crowd. This was generally realized as an industry-wide problem that needed resolution. But there was another, unaddressed problem – the IEEE has formulated a set of standards (IEEE 802.11 and its flavors), but there was not a uniquely defined and accepted way of standard compliance. Practically speaking, this could lead to a situation where two companies could both claim IEEE 802.11 standard compliance, but products from these two companies would not work with each other, and the two companies could easily blame each other for this lack of interoperability.

This situation forced the leading wireless LAN companies to sit together, amongst them Harris Semiconductor (now Intersil), Lucent Technologies Semiconductor (now Agere Systems), 3Com, Aironet (now Cisco), Symbol Technologies and Nokia. That last name may be a surprise, but Nokia over the last few years had started to test the waters of IEEE 802.11. Partly driven by the idea that wireless

LAN technology may become a factor for their terminal (cellular telephone) business, Nokia had started to develop technology that would later lead to the first smart phone. A smart phone is loosely defined as a cellular telephone combined with a PDA that supported a personal management system – a contact list and agenda.

The constitution of this group was quite remarkable. It existed out of Intersil, with four of their customers, plus Agere Systems. Still Agere Systems had a real need to participate. They were already part of WLANA (the Wireless LAN Association and pronounced as “walana”). However, WLANA had become more-or-less a frontier promotion institute for Proxim and Proxim’s OEM customers. All efforts to make WLANA a neutral body in this respect had turned out to be in vain. So, for Agere Systems, partnering up with Intersil was the right thing at that time, and it has worked out positively. WECA (today it is called “The Wi-Fi Alliance”) organized itself quickly and made a smart move by not allowing a complete and per person democracy, as was and is the case with IEEE. Governed by a small board, WECA was quickly out to establish an interoperability testing procedure and seal of compliance via a tested logo – the Wi-Fi (Wireless Fidelity) logo. Mentioning people here will surely result in other people not mentioned, but I think full credit needs to go to Angela Champness (then Agere Systems, now Proxim), Jim Zyren (Intersil), Phil Belanger (then Cisco, now Vivato) and Greg Innes (Symbol Technologies). They knew what was required to get it done and how to get it done. And although IEEE-like politics were around the corner, they established the Wi-Fi logo with credibility – to the point that the technology today is widely known under that logo, and the IEEE 802.11b statement is disappearing.

In this respect, it is interesting to mention Jim Brewington, who was Sr. Vice President at Lucent Technologies, and my boss from 1995 through 1999. Before the technology had broken through, he had already complained about IEEE 802.11. In his mind GSM, TDMA, CDMA, GPRS, UMTS and 3G were much easier for the general public to grasp than IEEE 802.11, and he personally gave me an action item “to fix this.” Well, he got what he wished for and probably something with even more marketing appeal than that he could imagine. I have heard statements that he is now an enthusiastic Wi-Fi supporter, which is somewhat of a change compared to his position in 1999.

THE HAY DAYS AND BEYOND (2001–2002)

At the end of 2000, when it started to accelerate – the bubble burst in the IT industry. Actually, this happened in two ways. The Internet bubble burst with the dotcoms going bankrupt very quickly, and with that the telecom bubble burst as well – the need for capacity reached its saturation point. It became clear that the “new economy” had to reckon with the harsh laws of the old economy, and the simple fact that the goal of business is to make money, not just spend it. But while the crisis developed to its full extent in 2001, this was also the real thriving year for wireless LANs and the peak years of all the companies that had survived the pioneering times of the decade before.

8.1 MANAGING GROWTH

After the successes with Apple and the other PC vendors, as well as the overall progress of the wireless LAN industry, we had to deal with a completely new set of problems. The whole organization was designed to pioneer and establish new markets. Now that the broader market had accepted wireless LANs, the name of the game became flawless execution. This turned out to be a major challenge and an exciting opportunity at the same time. Suddenly volume became the key item in manufacturing, flexibility and lead-time reduction, inventory management, and optimizing test capabilities. Agere Systems had for about a decade worked with USI (Universal Scientific Inc.), a Taiwanese manufacturing organization. This organization helped us through the early manufacturing startup, invested in manufacturing like we did in R&D, and became at this time the largest Wi-Fi radio card manufacturer in the world. Development working closely with manufacturing engineering and with test engineering, as well as with the manufacturing organization itself, made very significant progress. Compared to the first radios in 1990, the progress was even more astounding. In those years each wireless LAN card “radio” had about 15 test and calibration components – little screws on some of the components of the design that helped the product to work properly. Nowadays wireless LAN cards are fully software-tested and have no hardware calibration points anymore. Also, the first products had over 300 components, which is down to under 30 components today. With the progress of further integration, the goal of less than 10 is in sight. This was all the result of growing from 100 wireless LAN cards a week in 1991, to 100,000 wireless LAN cards in 2001.

Managing this growth was a challenge, but managing what was called “customization” was even more of a challenge. Despite the fact that all PCs look the same, and despite the fact that they all run Windows, there is a desperate drive in each PC company for differentiation – for just looking and feeling a little different than everybody else, or for just another step further in software integration, or for complying with just one more test in the complicated and somewhat archaic PC architecture. PC companies have become masters in technology supply-line management. They are demanding not only to reduce their own cost, shifting all work from themselves to their supplier, but also to reduce the time that any supplier can work on resolving differentiation issues for any of their competitors. At the same time, the tremendous competition amongst PC vendors is replicated in the PC supplier base, putting even more pressure on margins.

This was also the time that the PC industry as a whole was trying to go through a major quality improvement. I think it was a very good initiative, and it forced the industry to do what they had to do anyway – make sure that products meet customer expectations.

One of the other mechanisms for this was the earlier-mentioned WHQL certification, an abbreviation that for some software engineers still brings up nightmares of long and lonely nights when they were debugging and testing many times without any visible progress. Or worse, ending the night with more test failures than when the night began. Every piece of hardware that gets added to a PC interfaces with this PC via a piece of software called a driver. So-called “WHQL-ing” the driver certified that the driver had gone through a set of tests and met the requirements set by WHQL. Although Microsoft seems to be pretty independent in setting the requirements for these tests, the reality is that the rest of the industry understands that contesting these requirements is not the right way to approach the general notion of quality improvements. Probably rightly so – when a PC or a notebook hangs, we tend to blame Microsoft for this, even though there is a serious likelihood that the problem is caused by the peripherals of the PC and their interaction with the operating system. At the same time, it is interesting to see that this is called Windows Hardware Qualification Labs. One might wonder, “what about a Window Software Qualification Lab?” Maybe this is in the works. Or would it cause too many problems because then Windows would have to go through this Qualification itself? However,

considering all the long hours that go into working on compliance, the hope is that the quality of PCs will go up significantly.

8.2 NEW MARKETS: RESIDENTIAL GATEWAYS

Our vision for wireless LANs had been very PC-focused – how to connect mobile PCs to the enterprise network. However, through Apple it had become clear that there was a home market as well, as the Apple Airport product had been focusing on the home users (and the education market) in the first place. This helped us to completely rethink our marketing strategy, and that led to redefining the product line that was still called WaveLAN. The new concept was around Internet connectivity everywhere, and we looked at the Internet as a river of information that flows through our society and that everybody can tap into at any moment. So we came up with a new name for the product line – “ORINOCO,” as this was an interesting sounding name for a river. It was a river in Venezuela that according to old stories had on its banks the city of El Dorado, where everything was made of gold. ORINOCO stood for wireless connectivity to the Internet, not only as a LAN in the office, but also as wireless DSL or wireless cable at home, or as a wireless “modem” in hotels or other hotspots. We started working these new markets, the home market and the hotspot market, with mixed successes.

The home market was very interesting, and we learned a few lessons quickly. In the first place, the home market requires a product that needs another class of ease-of-installation and ease-of-use than we were accustomed to in the enterprise market. As the wireless to the home provider, we seemed to get the first brunt of this. But actually, in the enterprise environment, the IT department knew how to install wireless to the existing LAN and was used to the tools that Microsoft offered. In the home environment, no LAN existed – and no IT department either – and as is generally known, the Microsoft environment is not as simple as the closed Apple environment. So, the first launch of the product in the home market led to a flurry of support calls and made us realize that there was still a lot of work to do, including closing the loop with Microsoft, who fortunately, now knew how to spell wireless as well. This was a great help, and Windows 2000 and XP include good support tools to make the wireless installation a lot easier. At the same time, this market has now completely commoditized, and the home market has become a major target for Taiwanese and Chinese low-cost product providers. Also, the companies that are used in delivering cable set top

boxes or DSL modems to the end-users are working now diligently to include wireless in their offerings. It will be just a matter of time until “cordless” Internet will be as common as cordless telephone has been for decades.

8.3 NEW OPPORTUNITIES: HOTSPOTS

An even more challenging new proposition was the hotspot market. The idea for this had been floating around since 1997. It did not as such require so much special wireless technology; it was more an interesting application of existing technology. Unfortunately, many trials have been done so far, all without success. The issue is not so much the technology or the products, but how to build a profitable business case.

Many existing service providers have made trials, and we have participated in several ourselves. This is another example of a chicken and egg problem – when there are enough computers with Wi-Fi cards, then it will hit a threshold where the business case of hotspot Internet access will become attractive.

In the meantime, there are many failures on the hotspot trial path so far. The most prominent and visible may be Mobilstar. Mobilstar got great press with its Starbucks announcement – Starbucks in the US would serve Internet combined with coffee. Unfortunately, the deal was struck in such a way that the investments for Mobilstar were not justified by the revenue it created, and as a consequence, the company went under. However, this deal and announcement sparked interest to a next level. The efforts of Mobilstar have not been totally wasted either, as Voicestream, a US cellular provider, acquired the assets out of bankruptcy, and then were acquired themselves by T-Mobile, so it looks like the hotspot market has taken off. T-Mobile, a German cellular provider is active in Europe, as is British Telecom.

The good news is that the number of PCs with a Wi-Fi card is growing steadily. At the end of 2002, the statistic is one million units per month and growing. So, it will be a matter of time. At this moment there are about 20 companies offering hotspot services, reminiscent of the early cellular days. Worldwide I estimate the number of hotspot service providers to be at least at 250, maybe even ranging to 500. According to market research, the actual number of hotspots has grown in 2002 from 2,000 to 10,000 – clearly a booming market. The big challenge, though, is still to make money in it.

Interestingly enough, the companies in the late 1990s that did not want to have anything to do with Wi-Fi, including Lucent Technologies and Ericsson, are now seriously exploring this business. Their biggest challenge is how to make this seamlessly work with their CDMA and GPRS offerings. From a business perspective, a combined customer offering that includes wireless WAN and wireless LAN seems to be the way to make the original idea of hotspots effective – that is, enabling hotspot service providers to make money.

8.4 SOCIAL CONSEQUENCES

One of the reasons we know that the Wi-Fi market has really arrived is the fact that news articles have started debating the social consequences of instant and immediate networking. The consequence of wireless networking is that the Internet gets pervasive – that is, it is present with almost all the activities we are doing. A few examples can be given here.

With the arrival of notebooks and wireless LAN connectivity in universities, a professor in a classroom now has competition for the attention span of the student. In a recent interview, a professor complained that as soon as he has to pause to look up something, a chorus of keyboard clicks tells him that his audience has moved on to other activities – chatting, email, or even watching a movie. Stories are going around of lecturers furiously sabotaging access points to make sure that no one could access the internet during the lecture. In some cases, students have been asked not to make notes on their computers, but on paper. A more realistic professor took it as a challenge to make his lectures more entertaining to captivate his audience. His conclusion was that this internet competition had forced him to improve his presentation style.

Similar things happen nowadays at conferences where boring presenters are immediately penalized by finding out that the audience has left for “cyberspace.”

In enterprises that are completely wireless notebook-enabled, something similar is happening. Meetings are running the risk that most of the attendants have drifted away into their stacks of emails, which is usually not a good sign for the quality of the meeting. In several companies, codes of conduct are in place that include no emailing during meetings. Interestingly, however, this has also resulted in meeting attendants demanding more to-the-point meetings with agenda items that are relevant for them. Still, this is clearly a downside of

continuous connectivity. In the past, being early for a meeting allowed people to network and get some of the latest and often relevant information. Now if you are early for a meeting, you find everyone hidden behind the screen, which clearly creates a barrier for communication. We clearly need to get more used to our new wireless notebook toy, and we need to accommodate the related strengths and weaknesses into a new working style.

At home, we are already past this point. The television has been a dominating force in many families for a few decades, and people have learned to live with it and found ways to preserve some non-television life. The complaints of the past that an ever-increasing amount of time is lost watching television are gone, maybe because people are watching less television these days because of the increased amount of advertising. Perhaps the Internet has become a serious competitor for the television, and people won't really notice if television time is replaced with Internet time.

Life is changing, though. Getting hotspot access to the Internet in the Chiang Kai-shek airport in Taipei, and seeing an instant message "Hi dad!" from my daughter in the Netherlands over the chat while she is in the middle of attending a lecture at the university, is a signal of the start of a new and different era – one where continuous connectivity is paramount. This is largely because of the nature of the Internet itself, but Wi-Fi has clearly become an essential enabler.

8.5 OVER THE TOP: TROUBLE LOOMING

By the beginning of 2001 at Agere Systems, we had reached our top as Wi-Fi provider of the world. The market had grown to about one billion US dollars, and it was probably the only segment in the high-tech market that showed active growth signs this year. In Agere and our distribution channel, we owned about 50% of the market. The technology had now become well accepted, with the computer industry leading the way. The entrance in the home market and in the hotspot market and then finding the next step into consumer electronics was within reach.

But with our success, we had created our own conflict at Agere Systems. We had become chip supplier (of wireless LAN chipsets), a product manufacturer (via our subcontracted manufacturing house USI in Taiwan), as well as product brand supplier (with the ORiNOCO product line). However, as a chip supplier

we were selling chips to companies that competed with our own brand, creating continuous priority conflicts – external customers always had the impression that they were served later than ORiNOCO themselves. Also, our relationship with a subcontractor manufacturer made it difficult for us to work with the group of Taiwanese ODMs (Off-shore Development and Manufacturers).

This last item may need some explanation. The way the PC industry has been set up is very particular. Despite the fact that the PC market appears to be dominated by companies like Toshiba, Dell, HP/Compaq, IBM, NEC, Sony and others, the reality is that almost all these PCs are developed and manufactured by a limited group of Taiwanese companies called ODMs. The PC companies themselves are essentially companies that manage the supply and distribution chain with, in a lot of cases, a deep knowledge of the technology. But contrary to what an outsider would think, just as they don't manufacture, they don't develop products either. That is what the ODMs are doing for them. Most of these ODM companies (Ambit, Compal, Alphatop, Gemtek, etc.) are quite unknown to the general public, as they have not invested in marketing beyond selling their expertise to the PC and networking companies. Working with USI, a manufacturing house, excluded us from general access to the Taiwanese ODMs.

8.6 THE SOLUTION

By the end of 2001, it became clear to us that sooner or later the heydays would be over, whether we liked it or not. Although 2001 was still an excellent year in terms of growth and profitability, the competitive pressure was starting to kick in. A trend started that every quarter(!), the prices came down by 10%. This had begun in 1999 but continued with a steady pace into 2000 and 2001. Up to a certain level, the margin decline could be compensated with cost reductions, but it became clear that the competitive pressure was going to require a reformulation of the business. There was also another reason – with Wi-Fi becoming fully accepted, the next step was to integrate Wi-Fi as a technology with other technologies, like Bluetooth and Gigabit Ethernet, and in the longer term to integrate Wi-Fi into the general computer I/O chip (input/output chip), also call the Southbridge. Something needed to happen with the business to avoid having the market move on and leave us behind.

The year 2002 became the year of the “unraveling,” something that needed to be managed very carefully.

In the first place, by the end of 2001, the organization was carefully split into a chip division and a product division. The chip division was going to focus on developing and selling chips, just like many of their competitors that had started to appear left and right – Intersil, Broadcom, Infineon, AMD and others. The product division was becoming a customer of the chip division and was focusing on selling the ORiNOCO brand and the infra-structure products – access points, residential gateways, and outdoor routers. As a third activity, the USI was going to transform themselves from a manufacturing subcontractor into an ODM, making the necessary investments to do so.

The first phase is complete. The chip division of what was the famous WaveLAN product has now been integrated with Agere Systems' microelectronics division for client products, further building the technology into new chipsets and selling the wireless LAN chipset technology to other chipset providers who lack the wireless technology but have the capability of combining the wireless technology with other I/O technology. At the same time, there is wide open space in the consumer area. The wireless technology in the home is largely limited to cordless phones, and many of us have to go through the trouble of wiring stereo equipment, computers and other home gear.

The second phase was completed by mid-2002. The ORiNOCO business unit in Agere Systems was acquired by Proxim, who had been their biggest enemy in the last decade. This was a friendly takeover for \$65M, although Angela Champness, the General Manager of the ORiNOCO remarked ironically, "If you cannot beat them, buy them!" to Dave King, who had been Proxim's CEO for a long time. Currently Angela Champness is Senior Vice President of Proxim's LAN Division.

The last phase with USI is going full-steam and will be completed quite successfully. USI is migrating from being a subcontractor manufacturer into becoming an ODM, with special expertise in wireless LANs. Although the market share of USI may have gone down somewhat, they are still the largest wireless LAN card producer in the world.

THE FUTURE OF NETWORKING

To be able to talk about the future of networking, or more specifically the future of wireless networking, it is important to understand the situation that the communication industry is in today. This future in general is solid and bright, as communication is a key human requirement. That basic need cannot be satisfied other than with more bandwidth and more gadgets, and despite the current softness of the market, there will be a steady need for more bandwidth. But as usual, to get there will take just take a little longer than everyone expects. Plus, where “there” is, will only be known and understood after arrival. Until that time, “there” is usually quite foggy.

But trying to look through the fog and understand this communication future a little better, it is important to understand the key drivers today.

9.1 THE TELECOM NETWORKING INDUSTRY TODAY

Currently the telecom world is in a state of turmoil. Actually, it is in a sort of death struggle, since it will never again be the way it was in the late 1990s. This is the way I see what has happened over the last decade. The telecom industry was a slow growing industry, a few percentage points per year, steady and predictable. They were rooted in a monopoly type culture, not only AT&T in the USA, but in most countries worldwide. Telecom companies were actual or pseudo state-driven companies, slow-moving and only somewhat exposed to the dynamics of the market and the competition. Then both in the United States (initially by the surge of the Internet) and in Europe (initially by the surge of the cellular phone industry), some unsettling events happened that created first a frenzy and then an aftermath that we still have to overcome. These events were different in the US and in Europe, but they started to feed on each other.

The frenzy in the United States was based on the very rapid growth of the usage of the Internet. It was a frenzy, because the Internet business model was totally flawed for the telecom companies. They saw their traffic increasing, but not in a way that paralleled their increase in revenue and cost. Quickly the Internet Service Providers had figured out how to enable local calling for providing nationwide access, and local calling in the US was virtually free of charge. The Internet Service Providers themselves also had a major flaw in their business model. They did not make their primary revenue stream come

from the service they provided (Internet access); customers got Internet access from a service provider almost for free. Instead, the Service Providers thrived on advertising money as the primary revenue source. And on the back of this, many product services on the Internet started to make money primary from advertising and not from the service they provided. This created the well-known “dotcom mania,” aka the “Internet bubble” or the “new economy illusion,” which in essence was living in the belief that advertising would pay for everything.

But this bubble burst, as over time people were looking for true service and not for advertising. And with the bursting of the bubble we were punished for our collective greed – also in the Internet economy, stock prices will grow only in line with value created.

In Europe, the Internet frenzy always stayed within reason compared to the United States. The main reason was the fact that local calls were always charged. This meant that there was a natural limitation on the usage of the Internet, as many parents found out the hard way when they got the telephone bills for the Internet usage by their kids! I had to personally put my children on “an Internet diet” after being hit by a few astronomical phone bills. And by the way, this was not because they spent their time collecting useful information for their study or education. Their time spent was on chat programs and other Internet games that ate connection time from early in the evening until the wee hours of the morning. Internet addiction showed to be real, although at the same time, as far as I can see it was a quite benign disease as at least all our children grew out of it.

However, despite the more measured growth of the Internet industry in Europe as compared to the US, the telecom industry in Europe had its own frenzy and hype and it became victim of a comparable form of collective greed. This greed, however, did come not from the capitalistic stock market, but, ironically, it was collectively “organized” by the mostly socialistic European governments. In the mid 1990s, the cellular phone market had become very successful – more successful than anyone had expected – and because of a worldwide cellular phone standard (GSM) that was successful almost everywhere in the world. Almost everywhere, as there were a few large exceptions, like Korea, Japan and the United States, who really struggled a few years longer to get to a reasonable uniform coverage.

This success made the European governments organize two activities. One was privatizing (parts of) their cellular phone companies, which actually was already planned for or started. The other one was selling to these phone companies more spectrum for future cellular phone service. Both activities targeted many multiples of billions of dollars to go into the European states' finances and helped to create further European integration with the final introduction of the Euro in 2002.

The strong belief in Europe was that UMTS (Universal Mobile Telephony System, now also called 3G, for 3rd Generation Mobile Phone System), probably best described as a combination of a cellular phone plus wireless Internet access (again, essentially a combination of voice and data), would be the next "gold-rush." The need for more spectrum was very clear to everybody. The phone companies had worked with their network infrastructure suppliers (like Lucent, Nortel, Ericsson, Nokia, Alcatel and Motorola), and they unfortunately had completely missed the point about the difference in voice networking and data networking, and how difficult the convergence of the two is and has been since the days that AT&T and IBM tried and failed miserably. UMTS (3G) was the name for the solid belief that in the future all, or at least significant portions of the Internet access, would take place wirelessly by using the same technology as was used for cellular phones.

I cannot describe how skeptical everyone in the data networking world, including myself, were about the viability of UMTS, and we have given multiple presentations at many conferences to show the flaws of this route. And it was not only the lack of success of the convergence of voice and data in the past that could have predicted the overestimation and overvaluing of UMTS. Another reason was clearly the speed of the network. UMTS was defined as 2 Mbit per second under ideal circumstances – while stationary and close to a base station. Wired data networking is already moving from 100 Mbit per second to 1 Gbit per second, DSL wired networking is over 5 Mbit per second, wireless networking via Wi-Fi technology is moving from 11 Mbit per second to 54 Mbit per second and beyond. What would be the basis for expecting a lot from a 2 Mbit per second UMTS connection in 2005 anyway?

But the most convincing reason for the lack of viability of UMTS (3G) was the fact that a cell phone is not a computer. A cell phone, with its roots in the "dumb"

terminal industry, was built for providing high-quality voice, including being able to have a low weight, small size and long battery life for long talk-times and standby times. Supporting certain data applications, like messaging, is fine, but “true data” requires a seriously more powerful processor, even beyond the level of the PDA’s today. And more powerful processors require heavier batteries. And while we’re at it, a bigger screen and a larger keypad would also come in handy. Before we know it, we have redesigned a laptop computer. But a laptop computer already has a wireless connection – significantly faster than UMTS (3G) is going to provide.

In summary, the technology needs brought in via data networking requirements are in serious conflict with the original requirements for a cell phone – weight, size and battery life. A full-fledged UMTS (3G) cell phone would be a compromise resembling the one that the industry tried to make in the past for voice and data networking, and that never succeeded. Most likely it will be too clunky for a phone and too primitive as a computer.

And the television commercial showing video on a cell phone is somewhere between a joke and seriously misleading. Unfortunately, the makers of it did not realize that this is not going to happen for a long time. The ones who paid for the advertisement probably could have known, so they did not have to write off the \$Bs of their UMTS (3G) spectrum investments.

Summarizing what has happened, the European governments “gave away” spectrum for free in the ISM bands (2.4 and 5 GHz) to be used for “license-free” data communications. At the same time, they sold spectrum for billions of dollars to the telephone companies for licensed data communication. Also, as shown many times, telephone companies did not understand data communication well enough, and had not learned from the mistakes in the past about trying to integrate voice and data, so they bought the spectrum, only to find out that it is pretty worthless with the pace the technology is developing.

The UMTS (3G) frenzy fed on the Internet industry and the belief that voice and data would converge quickly and easily in the wireless world. This drove the price up that the wireless Service Providers wanted to pay to the European governments for buying so-called UMTS (3G) spectrum. Large bidding exercises have been set-up in the UK, Germany, the Netherlands and many

other European countries. The total sum of money that has been spent on new UMTS (3G) spectrum by the European cellular phone companies exceeds the staggering amount of \$120B.

However, in the aftermath, the awareness is slowly setting in that this money is probably all going to be written off. Some telephone companies already started doing this, like Telefonica in Spain. Unfortunately, the financing of these transactions, and therefore the writing off of this, involves a lot of banks and financial institutions, and the consequences are not necessarily fully understood and absorbed by them.

The reality of today is that even a few years after the UMTS (3G) spectrum was sold, little or none of it is being used, and it most likely will not be used for anything else than where GSM and GPRS are being used today.

It may be good to make some further comments about GPRS (GSM Packet Radio Service), or better yet, the data networking capability of GSM. Actually, the technological capability is pretty good and theoretically capable of supporting around 100 Kbit per second. Comparing this with the 56 Kbit per second modem speed of a wired telephone line, it actually is so good that many of today's low-end mobile data communication applications, mainly sending and receiving of messages and/or email, can already be supported.

Another good thing of GPRS is the fact that it can be implemented in the base station of the cellular phone networks via relatively simple upgrades, making it relatively inexpensive. This in contrast with the UMTS (3G) network infrastructure that needs to be implemented from scratch – not only is the equipment totally different, it also scales differently. In other words, it is not the same cell size per radio station. This all means more “bubble” trouble for UMTS (3G), as the speed of GPRS is pretty good for the low-end data applications required, and the cost to bring up the infrastructure is relatively affordable for the service providers.

The real killer for UMTS (3G) is now becoming the combination of GPRS with Wi-Fi in a public environment. This was not the planned concept of the late 1990s, but if we have learned anything from history, this combination has all the ingredients to cost effectively deliver voice and “true” data at appropriate

speeds and quality. It keeps voice and the needs of its quality where it belongs, including cellular telephone weight and battery life. It also provides mobile computers what they need – high speed, “low quality” and cost-effective access to the Internet.

9.2 THE CONSUMER NETWORKING INDUSTRY TODAY

Somewhat similar to the battle between telecommunications and computing, there are two other industry battles going on, both targeting the consumer. One battle is between computers and consumer electronics, and the other is between telephone networking and cable networking. Actually these battles are being fought on two battlefields close to the consumer, one is who is going to provide access to our homes (is it going to be DSL or cable?), and the other battlefield is our homes themselves – our living rooms and our study rooms (is it going to be the television or radio, etc., or will it be the computer?). Let me explain this a little further, starting with the last one.

Between televisions (and let me include radios and consumer electronics in general for a more complete picture) and computers, there is already some older convergence history, as both industries made efforts to crawl into each other’s territory.

The computer industry tried to enter the consumer space with the so-called multimedia PC. The idea was that a desktop computer can easily be converted into a radio, television, CD-player, movie theater encyclopedia, all at the same time. And indeed, looking at a desktop PC nowadays with sound blasters (integrated stereo speakers), CD-RW (readable and writable CD players), high speed Internet connection, and a television card, we can listen to the radio and to music and watch television on a desktop machine, while roaming the Internet all at the same time. It’s to the point where you would wonder whether there is still space for consumer electronics.

At the same time, the television industry has been looking at data for a longer time. One of the services on (at least European) television is something that is called “teletext.” This allows the television viewer to find with the remote control “written” background information on the programs that are being watched, the important sport and news items, weather, stocks, etc. It also makes programming of the VCR easier, as recording a movie can be done by clicking on its teletext

announcement. In a way, it reminds me of an early “one-way” Internet. So, this was a first step of the consumer industry into the world of data in the early 1980s, at least from an application level, and a quite useful one. I still use teletext almost daily for the different applications mentioned above. Interestingly it is much quicker than via Internet.

To quote Steve Jobs, Apple’s CEO, “Watching television is switching off your brain, while with a computer what you want to do is to switch your brain on”. He was trying to point to the fundamental difference between computers and televisions, at least at that time – this quote is from 1998. Today I am not sure anymore. There are many areas where the difference between the televisions and PCs start to blur.

An example of this is providing news. A few years ago, I swapped my paper Wall Street Journal subscription for an electronic (Internet) subscription. The reason for this switch was pretty thin – I wanted to read the paper at 6 AM, and the delivery service had moved from 6 AM to 7 AM. Until that time I had always preferred the smell of freshly printed paper, and I still do.

However, I did not want to move reading the paper to later in the morning, so I cancelled the paper subscription and saved myself some money by going electronic. Frankly, now that I’m used to the electronic version, I do not think that I want to go back to a paper version at all. Electronic is stronger in many of the Internet news services today. For instance, the BBC Online News Service provides television clips alongside the written news articles.

Nowadays on modern television one can split the screen to be half television and half teletext. Would it be nice to be able to split it into half television and half browser, so one can watch the news and simultaneous check items being mentioned in the news, or watch a sporting event while simultaneously following the scores of some other games?

These are just a few examples. But there are many more, including obvious ones like downloading music or movies on the computer and playing them, even on the television screen, if required. Or the more recent Philips’ Internet-based radio and CD player. Actually, the claim is that there are already more than 1,000 (digital) Internet radio stations.

One may debate whether reading a paper or watching a movie is something that requires your brain switched on or off. But the reality is that the borders between the usage of a television (as representative for the consumer electronic industry) and the computer are blurring quickly. This has definite consequences for home networking as such. Currently the consumer electronics industry does not have a true home networking standard. Actually, the consumer electronics industry does not really “think wireless,” either. This is in contrast with the PC networking industry that already established a solid place for in-house networking (Ethernet) and in-house wireless networking (Wi-Fi), although the network implementation is mainly between a computer and so-called “Residential Gateway,” a box that connects to the wide area network.

There is a significant difference between the consumer electronics industry and the telephone equipment industry, however. The consumer electronics market has been over the years very competitive, and therefore I do not think that the PC industry will just walk in and take over. Still, the rate at which the consumer electronics industry is embracing the computer industry seems to be relatively slow, and the explanation could be that the consumer electronics is much more focused on design, cost reductions and large-scale innovations, like the recent DVD players or digital televisions. From the view of consumer electronics, the computer industry is probably too small and too erratically driven by innovations. At the same time in consumer electronics, a lot of thinking is going on about networking. This is particularly driven now by the arrival of plasma displays. They make it possible to have a television screen hanging flat on the wall and connected via a wire to the television tuner.

The residential gateway market has developed itself quickly out of the wireless LAN industry and is coming at it from two angles. The first angle is basing this gateway on an enterprise access point and adding a wide area connection (modem) to this. It can be just plain telephone, known as V.90, DSL or a cable modem connection. The other angle is from the set top box provider, where a Wi-Fi wireless LAN connection is integrated in the set top box (for cable) or the DSL box, in addition to or actually replacing the Ethernet connection.

In this light, it is probably good to spend a few words on the cable television industry and their position in the consumer space, as they have leveraged this to start successfully providing Internet services to the consumer. This clearly has

created a marketing battle with the telecom companies, where the consumer has two options – an Internet connection coming from a telephone company with a DSL (Digital Subscriber Line, from the DSL Forum) connection, or from a cable company with a cable connection (using the DOCSIS standard: Data Over Cable Service Interface Specification from CableLabs).

There was more to this war though, as both the telephone companies and the cable companies are trying to conquer the homes with a box – a DSL box or a cable modem box – and then use this box as a “beach head,” where over time more services can be added to it. In particular, the cable industry has never made it a big secret that their target is to deliver voice telephone services as well. And the opposing view clearly is to have the telecom DSL providers distributing television. This last one is not insight on my part, but how feasible the first one is going to be may be questionable. As indicated earlier, voice “quality” is something that is from a different magnitude than could be offered over a “cable” network, for the same reasons as voice over the Internet has clear quality limitations.

Both technologies (providing Internet via a DSL box or via a cable modem box) have turned out to be very expensive propositions for the service providers. The actual consequence of this is that the growth of the Internet by using these high-speed access technologies is seriously falling behind expectations. What is happening at this moment is that this industry is resetting itself to find sustainable price points for the boxes, as well as the Internet access services to create a sustainable business model. Another point to mention here is that the rollout of DSL infrastructure proved to be much more expensive than expected and has taken much more time, due to all kinds of quality and distance problems with the physical cabling.

One of the questions going around in the industry is, “who is going to win – DSL or cable?” Both technologies are a good step forward compared to the slow speed access of the past, but neither of them has a real advantage over the other. The cable industry has some advantage of an “instant” connection, but it is a shared connection and potentially suffering serious throughput degradation. The advantage of a DSL connection is a cleaner, high-performance switched connection that at least is made to look like an instant connection. Another element is that the telecom service providers in general have a better

relationship with the end-customer, but they have to spend some serious money to provide a DSL line. The cable provider has a less direct relationship with the end customer, but probably a somewhat lower investment to make in order to connect an end-user.

Still it is unclear where the balance will swing. It is not unlikely that a newer technology will start to threaten both – interestingly a technology coming from the computer world call ETTH (Ethernet to the Home.) This technology should reach into the 100 Mb/s to the home, and it carries a lot of promises. The IEEE already has already a standardization working group in place working on the definition of it.

At the same time, the IEEE is also working on the wireless equivalent for this ethernet to the home, called BWA (Broadband Wireless Access). This is currently under definition by IEEE 802.16.

In summary, the networking industry getting into the home is a big battle between the different sorts of technologies that could provide this. Cable, telephone DSL, and in the future ETTH or its wireless equivalent, are the real contenders here.

But also looking in the home we are starting to create a serious wiring mesh (or “mess”). There is a box from which the cables are run that distribute the television (and radio) signals through the house. These cables run parallel to the telephone cables that also go through the house and support the telephones in the different rooms. And sure, if you want to run a wired ethernet network through your house to connect all the different PCs of all the family members, then it requires another set of cables running from room to room. Each PC of each family member needs to be connected to the Internet as well. Not to mention if there is a burglar alarm system in the house, which then requires another set of cables to the different motion detectors in the different rooms.

It is clear that home networking is a major opportunity for wireless networking, but here is the question – what is this network going to look like, and will it be integrated, or will it stay separated?

9.3 POTS VERSUS VOIP

There are more industry battles going on that are important to understand for the total picture. The two that I want to mention are POTS (or ISDN) versus VoIP, and the cell phone versus the palmtop.

POTS stands for Plain Old Telephone System and is probably what most of us are still using today for landline phones at home. In the '70s POTS was upgraded from analogue to digital, or from pulse dialing to tone dialing. This upgrade to digital was necessary to implement ISDN. However, the adoption of this technology took so long that before it has already been replaced by ADSL. By the way, occasionally on vacation in some out of range area, I have had the pleasure to run into a real, old tone-dialing telephone system – a nostalgic pleasure but also a nuisance. It is impossible to check your voicemail with such a system. This migration will probably not see a real next generation. What is generally known as Internet Telephony, but what in technology terms is called VoIP (Voice over the Internet Protocol) is planning to disrupt this gradual flow of technology development. Although this promise was made quite some time ago, it is my expectation that slowly but surely this technology will break through, despite all the skepticism.

There are clear reasons why the technology development is taking longer, but there are also clear reasons why this technology unavoidably will break through. The reason for the delays includes the lack of user friendliness of IP telephoning, the quality of the connection, and the lack of migration capabilities from POTS to VoIP. These shortcomings limit the volumes of IP phones, actually making an IP phone too expensive and reducing the interest of companies to invest in this technology. That's the bad news. But in the meantime, the volumes of IP phones are slowly but surely growing to several hundred thousand per year. The reason for this growth is that there are very strong underlying economic factors that keep on pushing this technology forward to the point that the current shortcomings will have been designed out of the system, after which this technology will break through. These underlying economic factors are, first, the fact that an Internet call is essentially free once one has an Internet connection. Furthermore, the cabling for Internet telephony is much simpler than for POTS telephony. In an office today there are separate outlets for voice and data, in the future these outlets will be fully interchangeable. This will reduce the installation cost, it will make management easier, and on top of that, Internet calls and

voicemails can be more easily managed and routed with and via computers than today in their separate domains.

Slowly but surely the roadblocks are moving out of the way. One serious roadblock was the fact that a phone needs power and that this power is usually provided over the POTS (or ISDN) telephone line. Well, in the last year, the concept of PoE (Powered over Ethernet) is making its entry into the industry. The battle for the standardization is not yet fully resolved. Cisco is trying to force the industry into their “standard,” while IEEE 802.3af is trying to set a general public standard. But in the meantime, products are shipping that allow IP phones to plug into Powered over Ethernet ports, just like normal phones, not needing their own power.

Places where IP phones are in use today are in multinational companies that have their own Intranet domain within the Internet. Good high-speed connections make the quality of telephone connections quite acceptable, and a fraction of the cost of an external call. Internet telephony will slowly but surely take a position in enterprises, and from there make its way into the broader consumer market – at least for these consumers with an Internet connection.

Why is this POTS versus VoIP battle important for the wireless industry? The reason is that it is not unlikely that wireless technology will help VoIP gain ground in the enterprise and from there into the worldwide telephone market. Some more background is required here. The current standard for cordless phones, as mentioned earlier, is DECT (Digital Enhanced Cordless Telecommunications). Although the penetration of DECT technology in the enterprise is limited, the technology is mature and has come down in cost very significantly. It actually has even become more successful in the consumer and residential markets. But DECT as a technology is isolated, because it is not integrated in the IP world. The big advantage of wireless VoIP, or to be more precise, of Wi-Fi based VoIP, is that it can seamlessly use the existing Wi-Fi infrastructure, as all the roaming protocols provided by DECT are integrated in Wi-Fi, as well. But Wi-Fi based VoIP will bring DECT a step further. A Wi-Fi based DECT phone can also work while visiting other Wi-Fi environments, for instance Wi-Fi hotspots. Actually this creates an interesting proposition – a cordless phone that can be used in the office, that then also can be used as a phone in Wi-Fi hotspots, as well as at home. But interestingly this phone has another implicit feature that has been

touted by the telecom industry for almost a decade, but that has never found a way for easy implementation – this phone has automatic follow-me capabilities. It is a Wi-Fi cell phone that can be used in all the Wi-Fi hotspots.

The next question then will clearly be whether this Wi-Fi phone will be a serious competitor for the cell phone? Actually, I do not think so. But I expect the arrival of cell phones with Wi-Fi capabilities that can make calls on the Internet if in the hotspot, or at least in a Wi-Fi-enabled area. The next step is quite critical – being able to move in and out of a hotspot and have the GPRS network taking over the call from the Wi-Fi network, and vice versa, when moving into the network. This is actually how I see the future for the cell phone industry – it will have to further consolidate and accept Wi-Fi as a complementary opportunity for voice as well. It may be clear that for the cellular service industry, there is a future, but that future is embracing the Internet. As a standalone solution, it will over time eradicate itself.

Interestingly enough there is another trend that needs watching – the coming of the smart phone. Since the initial palm computers, there has been the notion that a palmtop computer and a cell phone are pretty close, actually overlapping. So, the cell phone industry has been looking for a while to see what it takes to make a cell phone into a palmtop. On the other hand, the palmtop industry has been looking at what it takes to add telephone capabilities to a palmtop. Of the first category, Nokia is the real pioneering company, although as with many new concepts, acceptance takes a while. The main roadblock for such a piece of equipment is usability. Shorter battery life for the cell phone, because it is used as a palmtop, is a serious drawback. Bigger batteries would be the solution, but they have as a drawback that the cell phone becomes larger and heavier, and not really a nice experience to keep at your ear. Here the wireless headset, using Bluetooth, comes into the picture, where the smart phone is a palmtop computer, but has a wireless connection to a headset. If only the price of the Bluetooth headset would come down, and the battery life of this headset would be reasonable, these solutions would appear “in sight” for broader acceptance.

9.4 THE FUTURE OF NETWORKING

Having explored the different territories or battlefields, let’s see whether we can make some future predictions about where the networking industry is going to

be and what the role of Wi-Fi is going to be. The interesting part of predicting the future is that one can be sure that the reality will be different, whatever the prediction was. Maybe it might be better for our chances of success to discuss some scenarios.

The enterprise is more-or-less the place where Wi-Fi started in the mid-1990s. Over the years, the technology has improved and reduced in cost significantly, allowing it to go from its original data collection niches into mainstream computing. Some PC notebook models already have Wi-Fi as a standard feature. Just as Ethernet is today commoditized, Wi-Fi will be tomorrow. This means further integration of Wi-Fi technology into the core processing of the PC, a subject that Intel has already working on for a few years, although not very successfully so far. From there on and with the further cost reductions, clear expansions can be expected into other types of computing – really cutting the cable in desktop computing, as promised in the original marketing campaign in 1990.

The more interesting part will be the future of Wi-Fi itself. Just as with Ethernet, the future in Wi-Fi will be in higher speeds. Ethernet followed a clean path of 10 Mbit per second to 100 Mbit per second, and from there to 1 Gbit per second. Unfortunately, Wi-Fi is using some more exotic stepping stones. The first IEEE 802.11 standard was 2 Mb per second, the next standard (IEEE 802.11b) was 11 Mbit per second, and the next standards (IEEE 802.11g and .11a) run 54 Mbit per second. Work is going on beyond the 100 Mbit per second; 108 Mb/s is one of the stepping stones to be looked at.

Important here is the backward compatibility of all these standards, something that is worked into the standard definitions. This is particularly important in an enterprise environment, where one will not be very willing to completely re-implement the wireless infrastructure (access points) when a newer and higher-speed standard becomes available.

The infrastructure side in the enterprise is still a major trajectory to be followed and completed. This has everything to do with the fact that both Ethernet infrastructure (switches and hubs) and wireless infrastructure have completely independent trajectories. An Ethernet switch does not specifically expect an access point on one of its connectors, and the Ethernet port of an access point does not expect a switch. More interestingly, an access point is a fully

functioning networking bridge that is running on full-blown PC hardware. In essence, an access point is a PC itself.

Take the example of a larger facility with, say, a configuration of twenty access points. All are hanging on a few Ethernet switches or even on just one Ethernet switch. One can think of considerable cost reductions, concentrating key functionality of the access points into the switches and reducing the functionality and the cost of the access point to a bare minimum. The size of a complete access point can be as small as a match box and should be directly pluggable onto the Ethernet connector or hub of the standard cabling system. At the same time, the price could be reduced to significantly below \$50 per access point. One step further would be to find a convenient way back to the hub, for instance through a wireless backhaul including solar cells for power.

The functionality can also be further improved. In particular in larger facilities, considerable radio planning work is required. Together with the further integration of the functionality, auto configuration capabilities are required. These configuration capabilities include the channel settings as well as the power output, managing the range of each access point. With the fact that access points can also hear each other, they can balance the traffic that they need to support, and with that further optimize the total configuration and response time of the total system.

This will require significant development work, not only on the access points themselves, but also in the interface definition between access points and Ethernet switches, where most likely the necessary standardization work needs to be done.

In the coming years, we will see that GSM migrates to GPRS and gets combined with Wi-Fi for short range data services. This is not so much news, as well as that in my opinion there is currently no role for UMTS (3G), as the UMTS (3G) business case does not make sense. Even after writing-off the investment for the license, the investment of the infrastructure will be too high.

My expectation is that the business person will have four “gadgets” in the near future – a computer notebook with Wi-Fi and Bluetooth, a palmtop with Bluetooth, a cell phone with GPRS and Bluetooth, and a wireless headset with Bluetooth.

My computer notebook will connect directly to the Internet via Wi-Fi over a high-speed connection, or indirectly via GPRS with Bluetooth to my cell phone, for a low-speed connection. My palmtop will connect via Bluetooth to my cellular phone or to my notebook computer for Internet access.

From a hardware perspective, this is all pretty straightforward. From a software perspective, however, it is much more of a challenge. How does my address list (or my agenda) on each of the devices stay synchronized, and how to avoid compromising that integrity? For whoever has thought about this synchronization problem, it is far from trivial – in particular if an address card is updated on both devices before synchronization has taken place. The better solution clearly is, to avoid duplication, and to make sure that there is only one card stored altogether, but that is not necessarily trivial either.

This synchronization problem will grow one dimension more complicated if there is a desktop PC at home, as well as a desktop PC in the office. Keeping track of what is where will become more complicated, and actually end up with the notebook computer becoming the dominant business tool, despite the fact that it is about twice as expensive as a desktop – a difference that is expected to continue in the future. But if there is a synchronization requirement between these PCs, the connection between these computers will be dominated by Wi-Fi wireless technology.

The immediate question that comes forward is whether the number of devices can be reduced. One combination could be to integrate the cellular phone and the palmtop. There have been already many efforts going in this direction, actually coming from both directions – adding palmtop functionality, in particular short email, agenda, contacts and calculator capabilities, to the phone or adding phone capability to a palmtop. One of the roadblocks has been the battery life and the weight of this device, but assuming that a Bluetooth headset will become more-or-less the standard, this will be a feasible solution from a weight perspective – one at least does not have to keep all this weight to the ear all the time. Still the battery life and the weight will be somewhat of a challenge, both for the Bluetooth headset and for the combined cell phone palmtop, but I expect that this will be solution that will be quite accepted in the marketplace. However, in the sheer number of units, I think the simple cell phone and the simple palmtop will beat this combined unit for quite some time.

Let's in the meantime not forget the watch. Is a watch jewelry, or is it a functional device that could and should be synchronized with your notebook computer, preferably with wireless? I think there is maybe room for both approaches, although to get something useful out of a watch will require a sophisticated interface to be user-friendly.

But what about combining the notebook computer, the cell phone and the palmtop into one complete device? I think that this would be a feasible solution, however, the weight and the battery life of the laptop to carry this around is not very attractive. Still, if the price of a GPRS radio has come down enough that the PC industry is interested in integrating this, it seems to be a logical solution. The Bluetooth headset talks directly via Bluetooth to the notebook in the briefcase. It would be interesting if the notebook could be even in sleep mode as well.

9.5 WIRELESS NETWORKING ALL OVER

This journey of Wi-Fi may have started in the late 80s, and it was off to a slow start, as many technical and market barriers had to be conquered. In reality it has so far been just a small taste of what is expected to come in future years. The coming years will show a significant growth in Internet usage, and much more equipment will be connected to the Internet. The expectation is that in the future more "non-computer" equipment, actually peripherals or devices, will be connected to the Internet. More about this later. Actually, the expectation is that the number of devices connected to the Internet will rapidly outnumber the number of computers.

To avoid cabling becoming a major limitation to this growth, the need for wireless technology will become even more critical. Wireless technology in itself will follow the lines of faster products that are becoming very small in size, that will use less power, and for which, the price will continue to drop to a bare minimum. What is really going to be interesting is the growth in the usage of the Internet, with more management and control type applications creating the pull-through for wireless technology as a key enabler for those applications. In this respect both GPRS and Wi-Fi will play a role, and these applications will have an awareness of a slow connection (GPRS) or a fast connection (Wi-Fi). This knowledge will take care that in case of the slow link, only the most necessary information will be transferred, and a "full synchronization" will take place when the fast link is available.

The consequence of wireless in a business environment will be a permanent connection of every employee to the network, both for voice and for data, independent of whether this employee is an inside or an outside worker. Paperless will become the norm, electronic forms the standard way of operating, and immediate access to key information necessary to execute the job will become vital – even more than it is today. It is going to take a few more years before the application space in the business environment has caught up with the technical capabilities of anytime/anywhere (wireless) access to the company's resources. In that respect it is interesting to mention that I completely disagree with anyone saying that last decade's spending on automation was completely wasted. Sure, there was waste, but comparing the ways we ourselves operated in the late 1980s compared to recent years is showing a key acceleration in accuracy, timeliness and in general being on-top.

I do not want to deny that with every two steps forward, there was also sometimes a step back. In this light I vividly remember the complaints about the hundred-plus emails that daily filled my inbox in the office. It is not unlike complaints from the 1980s about the (paper) inbox piling up. I think it is a "law" that work will always pile up to the level that someone is barely able to handle, regardless of the time period. But this does not take away the fact that today more work can be done more efficiently with less people, which is a form of creating wealth, while at the same time an opportunity to do more.

But more than in the business environment, I am excited about the progress that can be made in the consumer environment, which is an attractive environment for the Wi-Fi business development, because of its volumes. But here we will also have to wait until the applications are developed to support new Internet consumer applications, and maybe also the time it will take for new behaviors to establish themselves in consumer environments. In this respect it is interesting to see where consumer spending is going, where the major consumer "annoyances" are, and how efficient improvements can be made.

Energy management in a household will become more important, as energy for the decades to come will continue to be a scarce resource, and the reduction in energy usage will also reduce environmental pollution. Despite our sizable energy bills, energy management in our homes is limited, if not absent. Usually we have little idea what equipment is using what energy, and although probably

20% of our household equipment is using 80% of the energy, I wonder whether many people know what the 20% consists of.

But probably even more than usage, the forecasting of energy usage is key to controlling waste of energy. Inventory of available energy has the tendency to “age” if not used quickly. So major energy savings can be achieved by more awareness and better prediction. For both, wireless and Internet connectivity will be the key, which requires that every piece of relevant electronic gear have an energy control “chip” and a radio “chip,” and a computer with the necessary software will operate as the control board. This is coupled to an application that helps to predict energy usage, and one can imagine that a correctly predicting user gets his energy at a lower rate than when this mechanism is absent.

A directly related application can be the preventive maintenance of electronic equipment, although I must say that I would be somewhat more skeptical about such an application, as quality of products is improving so much that this may not be a justifiable implementation – although including a car or the house’s heating and air conditioning installation probably can be controlled and preventively managed more than is the case today.

Another application may be the inventory management of goods in the house. All goods today have a UPC (Universal Product Code), identified with a barcode. Despite the dropped prices for barcode scanner, I do not yet know about any private use of these codes for home inventory management and reordering. Keeping a manual shopping list is still the preferred way.

At the same time, ordering of the weekly home shopping and delivery over the Internet is growing in popularity. This application could be the basis of an extension into management, and then coupled with retail organizations, it could be connected to their forecasting systems. Both wireless and the Internet play a role here, as one of the prerequisites would be another computer in the kitchen. I know people who are using this already – also to get new recipes from the Internet.

Other applications in the home environment are starting to develop in the area of medical control and support. These applications today are usually not wireless, and do not run via the Internet, but directly via the computer into a hospital system. This requires more individual and expensive application

development than necessary compared to a more standard Internet-related application development. When homes are “wireless-enabled” the automatic positive consequence is also a standard infrastructure can be assumed that is not bound to a specific room in the house or require cabling. When the infrastructure is simplified this way, it will be easier to check and control patients without having to have them in a hospital – another way to reduce the cost of our healthcare systems.

I think that the number of applications in the household will grow quickly once the infrastructure is in place. Regulating room temperatures, home security, Internet radio and television – the sky is the limit once the basic conditions of a solid and standard infrastructure are established. But these examples are not even mentioning the role that wireless Internet, and the easy access to it, can play for home entertainment in general, or on the flip side, for home-based education. But the bottom line is the fact that home networking has just started and will create many new horizons.

9.6 AN ADDITIONAL BUT DIFFERENT VIEWPOINT: TELEMATICS

It is always fun to run into a different viewpoint and try to think where this would lead us. This different viewpoint is based on the notion that one is usually not further away from his car than say 1,000 feet. Living in downtown Utrecht in the Netherlands, where parking is a higher art, this may not be completely true, but in general it is amazing how true this is.

So, the different viewpoint comes from thinking about the car as the center of the data networking. This is already required, in that the car in the world of telematics is already a complete info and communication center. Understanding that my notebook computer is already a personal extension, likely my car is as well. What about a configuration in which my car has a GPRS connection as well as a Wi-Fi connection to the outside world, and my notebook has a Wi-Fi connection to my car?

EPILOGUE

Going through the history of Wi-Fi, it is interesting to think back about the critical moments and what could have gone differently. The “what-ifs.” Is history filled with lost opportunities, or is it a miracle that we made it at all? Sometimes there is that oblique answer – the future will tell. Very quickly Wi-Fi will be gone into the mainstream of communication technologies and protocols like Ethernet, DSL, USB, V.90 and the like.

A critical moment in the industry was clearly getting the IEEE 802.11b standard agreed on, as at that moment, the standardization committee was very close to going up in flames. If that had been the case, most likely the industry would have taken another route, probably via an extension of Bluetooth?

Another critical juncture was the acceptance of the technology by at least one PC vendor (Apple), creating a beachhead from where the market has been built. At the same time, the question is how much of a blessing in disguise this was. The price was pushed down so heavily that commoditizing happened fast, so relatively few companies have enjoyed sustainable growth. But this may not be something to look for in the world of Information Technology development anyway.

A critical juncture for us personally was when Lucent Technologies decided that wireless data was going to be GPRS and UMTS (3G), and that Wi-Fi would never find acceptance in the market. Many things could have happened to us – the most logical would have been that all the WaveLAN would have stopped right then and there. Probably the industry would have proceeded without us without a hitch, although if a well-known company that is market leader in a new technology decides to abandon ship, that would definitely give the market a serious blow.

The conclusion might very well be that one can compare the world of Information Technology with shooting stars. A new idea sublimated into a product looking for a market, suddenly starting to shine like a star and then finding itself absorbed by its environment before one even realizes.

Let's go for the next idea!

ACKNOWLEDGEMENTS

The problem with acknowledgements is twofold. With a project the size of Wi-Fi and its duration over so many years, and with contributions large and small from so many people, the list could be endless. And even if such a list could be made, I am sure it would still be incomplete and end-up disappointing those not mentioned. (After all, as the saying goes, “success has many fathers (and mothers), while failures are orphans.”) In light of these risks, I am going to intentionally mention only a few people who I think did special things.

From a technical perspective, thanks to Bruce Tuch and Albert Claessen for their initial work on the radio, and then Richard van Nee for follow-up generations working on a worldwide scale; Wim Diepstraten, for his work on the MAC – and how solid it was, straight out of the gate; Vic Hayes for navigating all the dangers towards an approved standard. From an organizational and commercial perspective, thanks to Wiek Schellings for believing in us as a team and getting us going; Tom MacTavish for his skills in communicating a vision and helping us to market the concept internally; Angela Champness and Frans Frielink, who persistently and tirelessly worked on the external sales and marketing; Andre van Hees in operations and testing; and Paul de Wit, who helped tie up our financials. It was great to be part of such a great team!

Without mentioning any further individuals, I do want to acknowledge our international customers (including Apple), partners and “coopetitors,” as well as our manufacturer – USI Inc. in Nantou, Taiwan, who supported us through the early years and is still a leading force in Wi-Fi manufacturing today. They all helped us to make the difference!



Cees Links

Cees Links is a Wi-Fi pioneer.

Under his leadership, when he worked for NCR, AT&T and Lucent Technologies, the first wireless LANs were developed—ultimately becoming household technology integrated into phones, tablets and notebooks around the world. He pioneered the development of access points, home networking routers and hotspots. He was involved in the establishing of the IEEE 802.11 standardization committee and the Wi-Fi Alliance. He was also instrumental in establishing the IEEE 802.15 standardization committee as the basis for the Zigbee® sense-and-control networking. After that Cees became the founder and CEO of GreenPeak Technologies, now part of Qorvo, and currently serves as the General Manager of the Wireless Connectivity Business Unit. In 2017, Cees received the Golden Mousetrap Lifetime Achievement Award, and he was inducted into the Wi-Fi NOW Hall of Fame in 2019.

ISBN 9789090324326



9 789090 324326



qorvo
all around you