January 23, 2006

VIA ECFS

Ms. Marlene H. Dortch
Secretary
Federal Communications Commission
445 12th St. SW
Washington, D.C. 20554

Re: Erratum to Comments of the Computer Companies
CS Docket No. 97-80

Dear Ms. Dortch:

The attached erratum makes minor, non-substantive changes to the above referenced comments filed January 20, 2006, by ATI Technologies, Inc., Dell Inc., Hewlett-Packard Company, and Intel Corporation (the "Computer Company Comments"). Due to a file conversion error, portions of the text contained in Figures 1 and 3 and on pages 9, 10, and 11 of Appendix A became corrupted and were not rendered correctly. The attached erratum consists of a complete, corrected version of the Computer Company Comments, including Appendix A. Please substitute this erratum for the original filed Comments. A copy of the Electronic Comment Filing Service confirmation page for the original filing also is attached for the Commission's reference.

Should you have any questions regarding this matter, please contact the undersigned.

Very truly yours,

/s/ David Strasser

David Strasser
System Architect
DTV Products Group, Consumer Division
ATI Technologies, Inc.
33 Commerce Valley Drive East
Markham, ON
CANADA
L3T 7N6

Before the
Federal Communications Commission
Washington, D.C. 20554

In the Matter of	)	
	)	
Implementation of Section 304 of the Telecommunications Act of 1996	)	CS Docket No. 97-80
	)	
Commercial Availability of Navigation Devices	)	
	)	
Cable Industry Report on Downloadable Security	)	DA 05-3237

COMMENTS OF THE COMPUTER COMPANIES

ATI TECHNOLOGIES, INC.

Paul Lypaczewski
Vice-President and General Manager
Multimedia Business Unit
PC Division
ATI Technologies Inc.
1 Commerce Valley Drive East,
Thornhill, Ontario L3T 7X6
Canada

DELL INC.

Neeraj Srivastava
Director, Client Architecture & Technology
Dell Inc.
One Dell Way
Round Rock, Texas 78682-8033

HEWLETT-PACKARD COMPANY

Adam Petruszka
Director, Strategic Initiatives
Office of Strategy & Technology
Hewlett-Packard Company
2055 State Highway 249
MS-110225
Houston, TX 77070

INTEL CORPORATION

Jeffrey T. Lawrence
Director, Content Policy and Architecture
Intel Corporation
JF3-147
2111 N.E. 25th Avenue
Hillsboro, OR 97124-5961

January 20, 2006

TABLE OF CONTENTS

SUMMARY ii

BACKGROUND AND INTRODUCTION 2

I. The Commission Should Require Changes to the DCAS Agreement To Recognize that the PCI Express Interconnect Is Not a User Accessible Bus. 6

A. At Least Two Provisions of the DCAS Agreement Would Unreasonably Preclude Use of the PCI Express Interconnect. 7

B. Due to Its Design and Topology, the PCI Express Interface Is Not a User Accessible Bus. 9

1. The PCI Express Interconnect Is a Dedicated Link Between Two Components Whereas a Traditional PCI Bus is a General Pathway Used by Multiple Components. 10

2. PCI Express Utilizes High Frequency Signaling, Data Encoding and Scrambling, and Reception Techniques That Make Acquisition Using Probes or Interposers Infeasible. 12

C. The Difficulty of Recovering Content from the PCI Express Interconnect Cleary Demonstrates That It Is Not a User Accessible Bus. 13

II. NCTA Has Not Provided Sufficient Information for the Computer Companies To Conduct a Complete Review of the DCAS Agreement. 14

III. Approval of NCTA's Proposal Its Current Form Would Not Be in the Public Interest. 15

CONCLUSION 16

SUMMARY

The Computer Companies appreciate the challenge of the cable industry's continuing efforts towards developing the technology necessary for the secure delivery of advanced digital television content to consumers. While there are many issues and challenges associated with the NCTA's most recent DCAS Agreement and associated filing, the Computer Companies have focused this particular discussion on two provisions of the DCAS Agreement that, in and of themselves, could preempt the personal computer industry from participation in the proposed DCAS system and its proposed role in the creation of a competitive market for digital television navigation devices. Specifically, the DCAS Agreement (1) defines the PCI Express interface, which forms a critical part of a personal computer's internal architecture, as a "user accessible bus" over which unencrypted Controlled Content may not travel, and (2) requires the interface between discrete decryption engines and discrete video decoders to be encrypted, regardless of the robustness of the interface. Both these features of the DCAS Agreement, if applied to devices that will implement DCAS or to devices that will connect to DCAS-compliant devices, would negatively impact consumers without providing any benefits to the cable industry or their content suppliers.

Compliance, if even possible in the foreseeable future, will require personal computer manufacturers to alter internal system architecture at great cost, and ultimately consumers will receive equipment that is more expensive and less flexible. In this context, while some or all of the Computer Companies have raised many other critical issues and provided detailed comments on the DCAS Agreement, the Computer Companies want to specifically draw the Commission's attention to the unnecessary escalation of robustness requirements being imposed on our industry, and highlight the fact that the DCAS Agreement must be amended to exclude PCI Express from the definition of a "user accessible bus" as one condition of any approval of the DCAS

Agreement. The Computer Companies further request that the Commission require NCTA to provide additional information described herein regarding important licensing terms and conditions that were not included in the NCTA Report. The Commission should permit interested parties an additional opportunity to comment on any supplemental information that NCTA provides.

Before the
Federal Communications Commission
Washington, D.C. 20554

In the Matter of	)	
	)	
Implementation of Section 304 of the Telecommunications Act of 1996	)	CS Docket No. 97-80
	)	
Commercial Availability of Navigation Devices	)	
	)	
Cable Industry Report on Downloadable Security	)	DA 05-3237

COMMENTS OF THE COMPUTER COMPANIES

ATI Technologies Inc. ("ATI"),[1] Dell Inc. ("Dell"),[2] Hewlett Packard Company ("HP"),[3] and Intel Corporation ("Intel"),[4] (collectively, the "Computer Companies") hereby submit the following comments regarding the Report of the National Cable &

---

[1] ATI is a leading supplier of digital television and visual image processing products for the personal computer and consumer electronic industries. ATI's Digital Television Business Unit has shipped more than 10 million chips destined for High Definition Integrated Digital Televisions. In addition to being one of the world's largest computer graphics chip suppliers, ATI develops and sells add-in boards for the personal computer that allow customers to watch and record analog and digital television on their computers.

[2] Dell Inc. is a trusted and diversified information-technology supplier and partner, and sells a comprehensive portfolio of products and services directly to customers worldwide. Dell, recognized by Fortune magazine as America's most admired company and No. 3 globally, designs, builds and delivers innovative, tailored systems that provide customers with exceptional value. Company revenue for the last four quarters was $54.2 billion. For more information about Dell and its products and services, visit www.dell.com.

[3] HP is a technology solutions provider to consumers, businesses and institutions globally. The company's offerings span IT infrastructure, global services, business and home computing, and imaging and printing. For the four fiscal quarters ended Oct. 31, 2005, HP revenue totaled $86.7 billion. More information about HP (NYSE, Nasdaq: HPQ) is available at http://www.hp.com.

[4] Intel, the world leader in silicon innovation, develops technologies, products and initiatives to continually advance how people work and live. Additional information about Intel is available at www.intel.com/pressroom.

Telecommunications Association on Downloadable Security ("NCTA Report") and the accompanying Downloadable Conditional Access System Host License Agreement (the "DCAS Agreement") filed by National Cable & Telecommunications Association ("NCTA") in the above-captioned proceeding.[5]

BACKGROUND AND INTRODUCTION

The Computer Companies are a group of personal computer designers and manufacturers that are developing the next generation of converged personal computer equipment, featuring digital television processing capability. We are designing the hardware and software necessary to give consumers what they want: multipurpose tools capable of providing video, voice, and data services through a single device. The Computer Companies have a strong interest in this proceeding because once consumers have these machines they are going to want to use them to receive cable and interactive television services. This converged universe cannot happen, however, until the content providers and cable operators are satisfied that their content will be protected when it flows through personal computers. The Computer Companies offer these comments on the NCTA Report and DCAS Agreement.

The Computer Companies generally support efforts to move toward downloadable security.[6] The Computer Companies, however, have great concern with the NCTA Report, which was completed without consulting major computer industry participants. The NCTA Report introduces troubling new Robustness Rules for Certified Host Devices, a term that is not clearly defined in the DCAS Agreement and that could be construed to include, for example, the

[5] *See* Media Bureau Announces Dates for Filing Comments and Reply Comments on Cable Industry Report on Downloadable Security, *Public Notice*, DA 05-3237 (released December 20, 2005); Implementation of Section 304 of the Telecommunications Act of 1996, *Order*, CS Docket No. 97-80, DA 05-3316 (released December 23, 2005).

[6] *See, e.g.* Comments of Dell Inc., Hewlett-Packard Company, Intel Corporation, and Sony Electronic Inc., CS Docket No. 97-80, at 6-8 (filed January 20, 2006).

four corners of a multi-function computing device and everything therein (rather than those portions of such a device that instantiate an actual DCAS implementation)[7] or even a device connected to a Certified Host Device in a particular CableLabs "profile." Indeed, the Computer Companies believe that implementation of the DCAS Agreement as submitted ultimately will effectively preclude television content originating from a DCAS-compliant cable system from being recorded, played, or otherwise processed on a personal computer.

The DCAS Agreement grants the licensee the right to "use, reproduce, and distribute the DCAS Specifications for the purpose of making Host Devices, including Prototypes, Licensed Components and Certified Host Devices."[8] The associated Robustness Rules are defined in Exhibit B of the DCAS Agreement. There are several provisions of the Robustness Rules that would make computer industry participation in the digital television navigation device market very difficult, if not impossible. This discussion does not address all of the challenges to our industry in the DCAS, but is narrowly focused on a few issues, including the DCAS Agreement's treatment of the Peripheral Component Interconnect Express ("PCI Express") interface, which is a critical piece of evolving personal computer internal architecture and represents the state of the art in component interfaces designed to handle video, graphics and

[7] The cable industry's exclusion of personal computer designers and manufacturers from the development of the DCAS Agreement follows the pattern established by NCTA in developing the unidirectional plug-and-play compatibility agreement that the Commission approved in 2003. The lack of consultation regarding plug-and-play cable compatibility led to confusion and standards that, although on their face and as stated in the proceedings were intended to apply to personal computers as well, were arguably prejudicial to the development of personal computer-based unidirectional plug-and-play devices. *See* Comments of Dell Inc., Hewlett-Packard Company, Intel Corporation, and Sony Electronic Inc., CS Docket No. 97-80, at 13-14 (filed January 20, 2006). Similarly, the lack of consultation on the DCAS Agreement has led to a license agreement containing provisions that threaten to make participation of the computer industry in the market for commercial navigation devices all but impossible.

[8] See DCAS Agreement, § 2.

other data traffic. Specifically, the new Robustness Rules introduced by the DCAS Agreement: (a) explicitly name PCI Express as a "user accessible bus" and (b) require encryption of the interface between discrete decryption engines and discrete video decoders, regardless of the robustness of the interface.

Although the requirements of the DCAS Agreement apply only to DCAS Certified Host Devices, that term is not clearly defined; and the new Robustness Rules may impact not only implementation of DCAS directly into personal computers and other multi-function devices that may seek certification as Host Devices, but also may impact devices connected to Certified Host Devices, including personal computers.[9] In addition, these new Robustness Rules may set an inappropriate precedent for other conditional access or content protection license agreements, effectively precluding personal computer makers from developing devices for other content and other markets, limiting consumer choice, and increasing the cost of devices that are produced. These new requirements represent a dramatic escalation in robustness requirements generally, and turn years of cross-industry negotiations and understandings on their head.

Because compliance with these new Robustness Rules would be extremely expensive if not almost impossible to implement, applying them to personal computer-based navigation devices (either as Host Devices or adjunct to a Host Device profile) would lead either to the exclusion of computer manufacturers from the market or to a prohibitive increase in the price of DCAS-compliant personal computers with no discernable increase in content security. If the Commission does not require the cable industry to revise its new Robustness Rules in a manner that avoids the foregoing, the personal computer will be effectively precluded from participating

[9] A personal computer could be impacted either as Host Device (where Host Device means that the DCAS functions are integrated directly into PC architecture) or as an extension to a 'dongle' or 'card' style Host Device if CableLabs requires consideration of the larger PC device when approving a dongle/card "profile."

in the navigation device market for the foreseeable future. Additionally, as DCAS-compliance will be cost-prohibitive for personal computer manufacturers, consumers will be deprived of the fully converged, video-enabled personal computers that they desire.[10] In addition to the obvious ill-effects for consumers, impeding personal computer manufacturers' entry into the market for digital cable devices could jeopardize existing high technology jobs, undermine the creation of new high technology jobs for American citizens, and severely limit the development of future innovative solutions for the display and use of digital video content.

None of these results is justified on the basis that Controlled Content (as defined in the DCAS Agreement) could theoretically be intercepted and copied from a PCI Express interface. That risk is simply all but nonexistent. Neither the transport of Controlled Content over the PCI Express interconnect nor the storage of content for buffering purposes in the personal computer internal architecture will render such Controlled Content readily susceptible to interception and copying. Any attempt to intercept content traveling over a PCI Express interface would require a level of technical sophistication and financial commitment that only the most dedicated, well-funded commercial pirate with substantial engineering and manufacturing resources would possess. Consequently, inclusion of PCI Express in the list of internal device interfaces over which unencrypted content is not permitted to travel will not improve the security of content in DCAS implementations or in content delivered to consumers via DCAS-compliant cable systems.

---

[10] As demonstrated at the recent Consumer Electronics Show in Las Vegas, consumer technology is moving towards permitting consumers to perform video, voice, and data functions on a single device. The competition in development of these converged devices is fierce and consumers will continue to benefit from reduced prices and increasingly flexible equipment so long as anti-competitive commercial and regulatory barriers do not hamper technological progress. The TV-enabled PC will be a cornerstone of tomorrow's consumer television experience. But consumers will never get there if restrictions like those imposed by the DCAS Agreement are permitted to make such innovations cost-prohibitive.

Likewise, computer industry participation in the navigation device market would be made cost prohibitive by a requirement that the connection between decryption engines and video decoders be encrypted, because such a requirement would require unnecessary and prohibitively expensive changes in the use of PCI Express (and the implementation of DCAS in personal computers generally). In short, NCTA has proposed Robustness Rules that will limit consumer choice without providing any discernable benefit to the cable industry or content providers. The Commission should direct the cable industry to engage with the appropriate personal computer stakeholders to resolve these issues. Moreover, at a minimum, the Commission should require NCTA to amend the DCAS Agreement to expressly exclude PCI Express from the definition of a "user accessible bus" in all Plug and Play agreements.

As detailed below, NCTA also has not provided the Commission with important information regarding the terms and requirements of the DCAS license. The DCAS Agreement references standards that are not publicly available and it provides no information on essential terms such as royalty rates. The Computer Companies request that the Commission seek additional information on these issues and provide interested parties an additional opportunity to comment on any supplemental filing NCTA submits.

I. The Commission Should Require Changes to the DCAS Agreement To Recognize that the PCI Express Interconnect Is Not a User Accessible Bus.

The Computer Companies accept and understand that the protection of content within the cable system and indeed, throughout a Certified Host Device, is of great importance to cable operators and their content suppliers. The compliance and robustness rules in licenses such as the DCAS Agreement long have served as a negotiated means to "commercially enforce security and content protection obligations."[11] The Computer Companies strongly disagree, however,

[11] NCTA Report at 5.

with NCTA's attempt to accomplish this legitimate end by including in the DCAS Agreement escalating Robustness Rules, which include among other things an overly restrictive definition of the term "user accessible bus" and a requirement that the connection between encryption engines and video decoders be itself encrypted even when inherently robust links like the PCI Express interconnect already are sufficient to protect content traversing them.[12] As explained below, and setting all other issues aside for the purpose of this narrow discussion, these excessive requirements threaten to make personal computer industry participation in the market for navigation devices cost-prohibitive without providing any additional protection to Controlled Content.

A. At Least Two Provisions of the DCAS Agreement Would Unreasonably Preclude Use of the PCI Express Interconnect.

The Computer Companies have (among many others) concerns with two provisions of the Robustness rules in the DCAS Agreement that would restrict the use of the PCI Express interconnect. First, unlike any other content protection agreement of which the Computer Companies are aware, Section 3 of the Robustness Rules expressly includes PCI Express in its definition of "user accessible bus." This is inappropriate. The term "user accessible bus" is a term that came into use almost a decade ago as a means to allow device manufacturers to determine when additional mechanisms must be utilized to protect content while in transit over data buses. The term is not meant to provide an exhaustive list but rather a general description and a short series of examples.[13] As in every other content protection agreement that we are

---

[12] *See* DCAS Agreement, Exhibit B, Section 3.

[13] The traditional definition of a user accessible bus recently was adopted by the Commission in the Broadcast Content Protection Proceeding:

> "'User Accessible Bus' means a data bus that is designed for end user upgrades or access, such as an implementation of a smartcard interface, PCMCIA, Cardbus,

aware of, the purpose of defining the "user accessible bus" is to ensure that content will not be readily susceptible to interception or be present in a location where it could be easily accessed. In this context, PCI Express simply is not "user accessible" in the way that term always has been understood or in any other reasonable construction. As explained in detail below, PCI Express is simply not a "user accessible bus".[14]

A second provision of the DCAS Agreement that threatens the computer industry's use of PCI Express (and the implementation of DCAS in personal computers generally) appears in the Section 3 of the Robustness Rules, which state (in part) that:

> "…If the video decoder of a Licensed Product is not located inside the same silicon device or ASIC as the video decryption engines, then the interface between the two chips must be encrypted…"

> or PCI that has standard sockets or otherwise readily facilitates end user access. A "User Accessible Bus" does not include memory buses, CPU buses, or **similar portions of a device's internal architecture that do not permit access to content in a form usable by end users**." *(emphasis added)*

Digital Broadcast Content Protection, *Report and Order and Further Notice of Proposed Rulemaking*, 18 FCC Rcd 23550, 23587 (2003) (emphasis supplied) (adopting 47 C.F.R. § 74.90000(r)), *rev'd on other grounds*, *American Library Association v. FCC*, 406 F.3d 689 (2005).

The reference to PCI in the FCC's definition is to the now obsolete PCI local bus, *not* the advanced PCI Express internal connection. Unlike the DCAS Agreement definition, the above definition is contained in more or less the same words in every content protection agreement. The operative portion of the traditional definition is that the particular connection does not permit users to access content in a form they can use. The PCI Express did not exist when the definition was first crafted during the DVD Content Scrambling System negotiations ("DVD CSS"). Part of the PCI Express interconnect design effort was to extend the protection afforded by other point-to-point internal connections. Thus, PCI Express is not, by definition, a user accessible bus in that it does not permit access to content in a form usable by end users.

[14] Technologists familiar with the PCI Express interconnect know that it is not a "user accessible bus" as that term has been long understood. To assist the Commission in its analysis of the issues presented in this submission, the Computer Companies have attached a white paper prepared by members of the personal computer industry that explains the features of the PCI Express interconnect and demonstrates that it is a robust means of transporting all forms of content.

The Computer Companies are concerned that the above language will be read as an overriding requirement that content transmitted across a PCI Express interconnect between video decoders and video encryption engines must be encrypted. This requirement would be unreasonable because it would apply regardless of the robustness of the interface. Because the PCI Express interconnect is a robust interface, requiring additional encryption is unnecessary and would result in prohibitive compliance costs for the computer industry. Moreover, the encryption requirement appears to be so broadly conceived that it could affect other, potentially more robust but unencrypted interconnects that may be developed in the future. There is no justification for restricting future personal computer innovation and development by requiring encrypted interconnects between components when equally robust protection of the content is already provided.

One step toward appropriately balancing the concerns of content owners with consumers' interest in flexible and inexpensive personal computing equipment, is that the Commission should require that the DCAS Agreement (and other Plug and Play agreements) exclude PCI Express from both its definition of "user accessible bus" and remove the rule requiring encrypted interfaces between video decryption engines and video decoders.

B. Due to Its Design and Topology, the PCI Express Interface Is Not a User Accessible Bus.

Appendix A provides a technical comparison of the standard PCI local bus (which actually is a "user accessible bus") and the PCI Express interconnect. While these two connection technologies are similar in name, the similarity ends there as they are vastly different technologies. The principal difference between the two technologies stems from the fact that the PCI local bus is a multi-point parallel interface and the PCI Express interconnect is a point-to-

point serial interface. As a result, these technologies encapsulate two totally different physical and logical topologies as well as two totally different physical connections.

1. The PCI Express Interconnect Is a Dedicated Link Between Two Components Whereas a Traditional PCI Bus is a General Pathway Used by Multiple Components.

As illustrated below, a PCI local bus links several components using one global wiring connection. This is a traditional bus design consisting of one sideband clock signal associated with the data and control signals. Since all of the components on the PCI local bus are connected to the same wires and the clock is an independent signal, any component attached to the PCI bus can easily access any 32 bits of information traversing the bus. This bus protocol and topology makes any data transmitted on the bus vulnerable to data interception.[15]

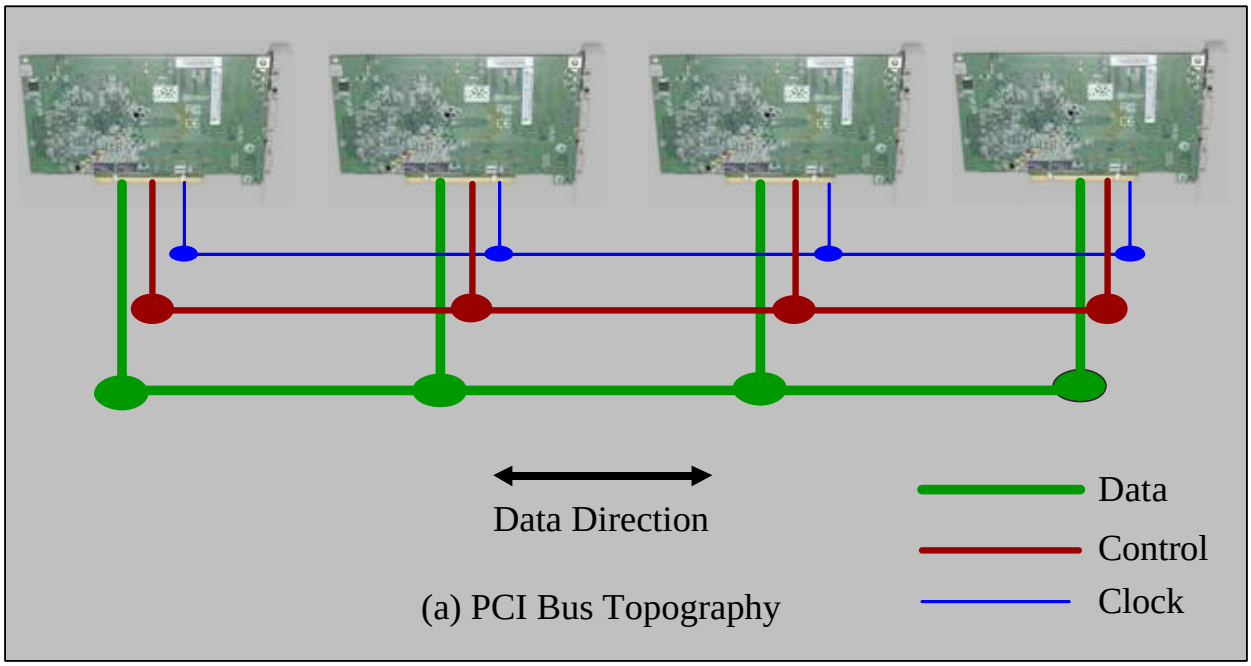

(a) PCI Bus Topography

[15] Another key element that makes data interception relatively easy in PCI is its moderately low clock speed. The vast majority of PCI local buses currently deployed in devices have a clock frequency of 33 MHz. A low frequency makes multi-point buses easy to monitor and debug using relatively inexpensive equipment.

In contrast to the PCI local bus topology, the PCI Express interconnect topology exclusively links two components over serial connections. The simplest type of PCI Express interconnect consists of a single serial data connection for transmission and another single serial data connection for reception. The combination of these two connections is called a "lane." Since different devices have different bandwidth requirements, the PCI Express interconnect can support up to 16 lanes for a particular device. The provision for multiple lanes makes the PCI Express interconnect ideal for components that require high volume data communication, such as video and graphics devices. The following drawing illustrates the topology of the PCI Express.

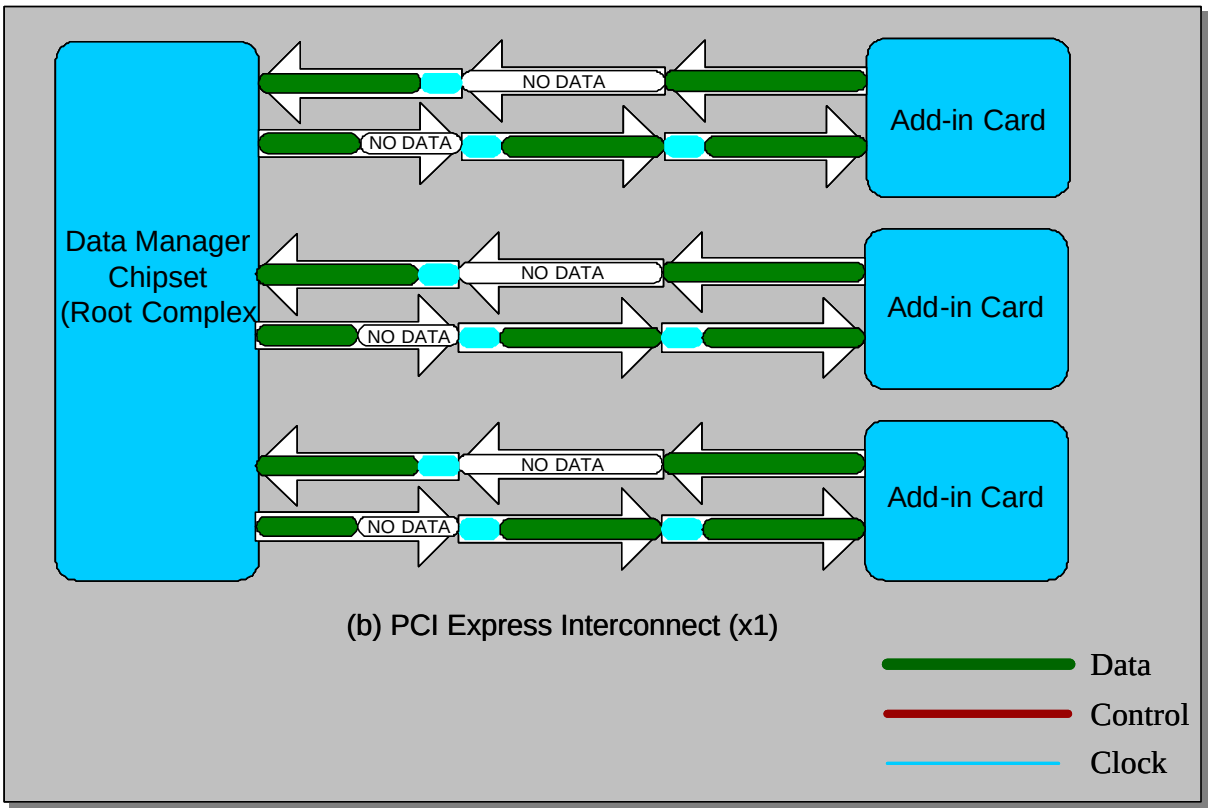

Since data is transmitted directly from one component to another, the only foreseeable way to intercept data traveling over the PCI Express interconnect is by inserting multiple acquisition probes or a single interposer between the two connected components. Acquisition probes are designed exclusively for use with logic analyzers or other similar professional

equipment. In addition to being extremely expensive, such devices are difficult to use, require significant expertise in handling, and are capable of capturing only small amounts of data after the expenditure of considerable effort by a skilled professional. Thus, an end user could not access in usable form even a second or two of a program traveling over the PCI Express interconnect, let alone an entire program.[16]

2. PCI Express Utilizes High Frequency Signaling, Data Encoding and Scrambling, and Reception Techniques That Make Acquisition Using Probes or Interposers Infeasible.

Even if a would-be commercial pirate has the time and financial backing necessary to attempt to steal content traversing a PCI Express interconnect, several features of the PCI Express design make it extremely unlikely that he will be successful. For example, to provide adequate data rates over its serial interface, the PCI Express interconnect must operate at a very high frequency (2.5 GHz). Due to practical technical limitations, the PCI Express interconnect transmitter component must embed the clock into both connections in each lane by 8-bit to 10-bit encoding of data into symbols. The PCI Express interconnect receiver component recovers the clock by "watching" symbol bits switch every 400 picoseconds, and aligning an oscillator's rising edges so that they occur at the mid-point of the symbol bit. Each connection in each lane must be independently "trained" to determine the mid-point of its symbol bits. It cannot be overstated that recovering the 2.5GHz clock on the PCI Express interconnect is an ***extremely*** difficult technical endeavor.

[16] The Commission also should note that there is no incentive for a commercial pirate to undertake the expensive proposition of attempting to intercept content traveling over a PCI Express interconnect as there are cheaper and easier ways to for commercial pirates to gain unauthorized access to protected content. The operation of stealing content traveling over a PCI Express interconnect is technically and financially daunting and commercial pirates would likely have no incentive whatsoever to take up the challenge.

There are other properties of the PCI Express interconnect that make probing impractical. To maintain the electrical signal integrity of the interface, data is randomized (scrambled) before being transmitted. Packets transmitted across the PCI Express interconnect contain both control information and data, thus complicating data capture. The data carried over the PCI Express interconnect is both proprietary and machine specific. The higher operating frequencies also make the PCI Express interconnect much more susceptible to signal reflections, which results in garbled data.[17] All of these factors reduce the possibility that acquisition probes can readily be used to capture transactions crossing the PCI Express interconnect.

Modern add-in graphics processing units ("GPU") and digital television chips are extremely complex components, with unique, proprietary, and non-public programming interfaces. Any mechanism which could capture transactions crossing the PCI Express interconnect to a GPU would also have to understand how the GPU works internally and how the operating system manages computer system resources – a tall order indeed.

C. The Difficulty of Recovering Content from the PCI Express Interconnect Cleary Demonstrates That It Is Not a User Accessible Bus.

As experts in the development of PCI Express interconnect components, the Computer Companies submit that an interposer made to capture content traversing the PCI Express interconnect would require a major development effort with significant financial backing. Interposers are discussed in greater detail in Appendix A to support this argument. Furthermore, to the Computer Companies' knowledge, there has never been a theft of content as it traversed even a "user accessible bus" – let alone anything approaching the technical complexity of the PCI Express interconnect. The Computer Companies contend that this is further evidence that such an added level of robustness is not required.

[17] Signal reflections make it almost impossible to recover the clock, and thus data.

Given the thorough analysis of the PCI Express interconnect and its protocols described herein and in Appendix A, the Computer Companies assert that PCI Express interconnect is not vulnerable to data interception and should be excluded from the definition of "user accessible bus" included in the DCAS Agreement.

II. NCTA Has Not Provided Sufficient Information for the Computer Companies To Conduct a Complete Review of the DCAS Agreement.

As potential licensees under the DCAS Agreement, the Computer Companies are being asked to make critical business decisions with incomplete information. NCTA has not provided sufficient information to permit interested parties (including the Computer Companies) with the opportunity to conduct a thorough review of the DCAS architecture pertinent to the development of navigation devices. In particular, the following documents are incorporated by reference in the DCAS Agreement, but, as of the date of this submission, have not yet been made available for review.

- Bootloader API Specification
- DCAS Security Specification
- DCAS Host Security Specification
- OpenCable Host 2.5 Core Functional Requirements

Furthermore, it is unclear when and how the keys for Licensed Products, Transport Chip Processors, and Security Processors are generated, distributed, and injected into DCAS-compliant components. In addition, the NCTA Report lacks information about licensing costs and key acquisition costs for such DCAS-compliant components. To the Computer Companies' knowledge, this information has not been made available to potential licensees of the DCAS Agreement. Without this information, it is impossible for the Computer Companies to determine

royalty-related costs in manufacturing DCAS-compliant products and therefore it is impossible for the Computer Companies to determine whether such costs would be reasonable.

The Computer Companies therefore request that the Commission require NCTA to provide this additional information regarding the DCAS Agreement to the Commission. The Computer Companies further request that interested parties be afforded an additional opportunity to comment on any supplemental filing that NCTA submits.

III. Approval of NCTA's Proposal Its Current Form Would Not Be in the Public Interest.

The Computer Companies cannot support FCC approval of the DCAS Agreement in its current form. As the foregoing demonstrates, the NCTA Report proposes (among other things) Robustness Rules that would severely impede computer industry participation in the market for digital television navigation devices with no corresponding gain in content security. Moreover, NCTA's failure to provide important information regarding the logistics of DCAS licensure makes it nearly impossible for interested industry participants to comment on other aspects of the DCAS plan. These shortcomings in the NCTA Report put consumers at risk of diminished choices without providing any corresponding benefit to content providers. Consumers increasingly expect their personal computers to have the capability of displaying video content received via cable and the Internet. They are not expecting – nor should they be forced to endure – dramatic increases in the price of computers that incorporate this functionality. Unfortunately, under the NCTA's current proposal, that is exactly what consumers will experience. The Computer Companies respectfully submit that under these circumstances, the Commission should not approve the DCAS Report as consistent with the public interest.

NCTA has provided the DCAS plan to the Commission to gain an additional postponement of the date by which it is required to separate the security function from cable set-top boxes. The current proposal from the NCTA does not justify that extension.

CONCLUSION

The Computer Companies ask that the Commission not further delay the integration ban, but invite the NCTA to address the issues raised in this and other filings provided by computer manufacturers.

Respectfully submitted,

THE COMPUTER COMPANIES

ATI TECHNOLGIES, INC.

/s/ Paul Lypaczewski
Paul Lypaczewski
Vice-President and General Manager
Multimedia Business Unit
PC Division
ATI Technologies Inc.
1 Commerce Valley Drive East,
Thornhill, Ontario L3T 7X6
Canada

DELL INC.

/s/ Neeraj Srivastava
Neeraj Srivastava
Director, Client Architecture & Technology
Dell Inc.
One Dell Way
Round Rock, Texas 78682-8033

HELETT-PACKARD COMPANY

/s/ Adam Petruszka
Adam Petruszka
Director, Strategic Initiatives
Office of Strategy & Technology
Hewlett-Packard Company
2055 State Highway 249
MS-110225
Houston, TX 77070

INTEL CORPORATION

/s/ Jefferey T. Lawrence
Jeffrey T. Lawrence
Director, Content Policy and Architecture
Intel Corporation
JF3-147
2111 N.E. 25th Avenue
Hillsboro, OR 97124-5961

January 20, 2006

APPENDIX

A

PCI Express Interconnect Robustness

Executive Summary

This paper explores the robustness of the PCI Express interconnect as it pertains to copy protection system implementations in an open architecture device such as a personal computer. Device robustness for an open architecture device requires one to determine whether an interface is vulnerable to data interception. Historically, a vulnerable interface was classified as a "user accessible bus", an example of which is the PCI local bus. In comparing the PCI Express interconnect to the PCI local bus this paper proves that the PCI Express interconnect not only is significantly more complex than the PCI local bus but is clearly in a different category as it relates to susceptibility to interception of data during transit. Using this analysis, it can only be concluded that the PCI Express interconnect should not be classified as a "*user accessible bus*".

Introduction to a Typical Desktop Computer Architecture

It is expected that the PCI Express interconnect will become as ubiquitous in open architecture devices, such as personal computers, as the PCI local bus is today, although the purpose and nature of the two buses is very different. The purpose of introducing the PCI Express interconnect into the personal computer was to improve the platform's performance by increasing the component to component data rate, while the PCI local was introduced to enable a variety of upgrades and add-on capabilities. This section will provide a brief overview of typical desktop computer architecture and how the PCI Express interconnect is integrated into its design.

The figure below shows the internal architecture of a typical desktop computer based on an Intel Central Processing Unit (CPU). It consists of a motherboard (not shown), an Intel CPU, a north bridge, a south bridge, system memory, a graphics processing unit (GPU), a hard disc drive and various add-in cards. As is shown in the figure, the PCI Express interconnect ports are located on the north bridge chip and the south bridge chip to connect to internal peripheral components. The north bridge chip is optimized for dedicated high performance interfaces and accommodates the CPU, system memory, the GPU and the south bridge. The south bridge chip connects to the north bridge chip and accommodates a wide variety of peripheral components including hard disc drives, Universal Serial Bus (USB) hubs, Ethernet controllers, etc. The north bridge and south bridge chip communicate via a PCI Express interconnect. The north bridge acts as the control point or "root complex" for the PCI Express interconnects in the computer.

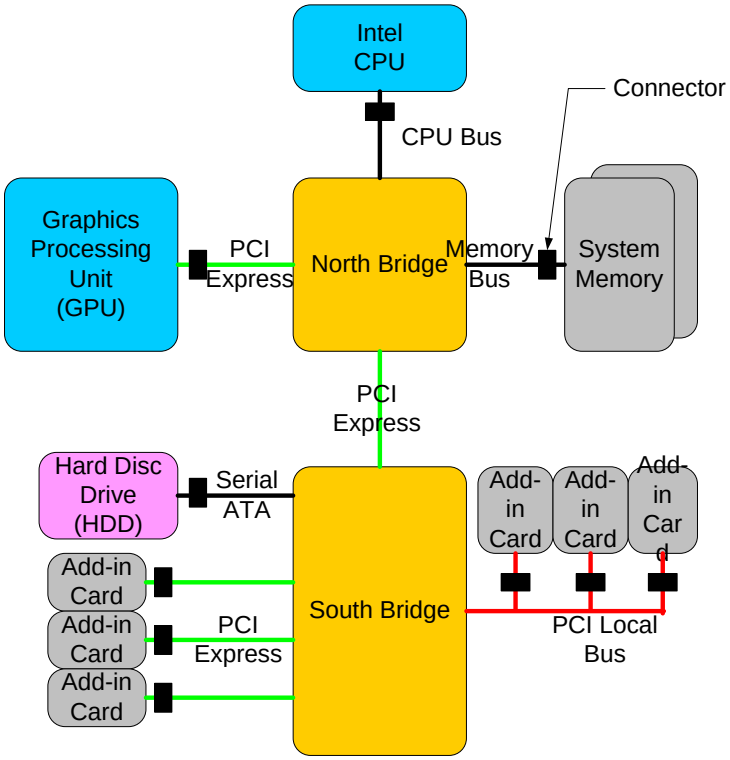

Figure 1: Typical Intel Desktop Computer Architecture

The AMD desktop computer architecture, shown below, is very similar to the Intel desktop computer architecture except that the CPU bus is an open standard interface called HyperTransport™. HyperTransport™ and the PCI Express interconnect share many of the same characteristics. This similarity is demonstrated by the possible application of either interface between north bridge and south bridge chips.

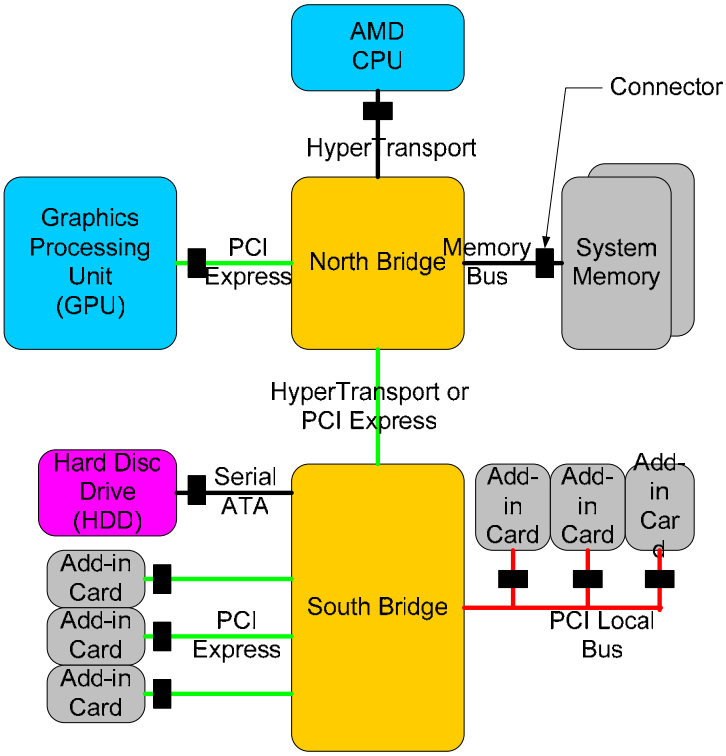

Figure 2: Typical AMD Desktop Computer Architecture

Defining Device Robustness

To understand how the PCI Express interconnect is in a different category than the PCI local bus as it relates to interception of data during transit, one must understand the method of measurement used to validate such a statement. The method of measurement is defined as device robustness and has a long history beginning with the Common Scrambling System (CSS) license agreement. This agreement is the license for the copy protection system used to protect content on DVD-Video discs.

The CSS license agreement was the outcome of years of inter-industry negotiations. Negotiations were started in 1996 and produced an interim license agreement in 1997 to enable the commercial sale of playback devices. The outcome of these negotiations was not only the CSS license agreement but the definition of device robustness, entitled Robustness Rules, used as a starting point for Robustness Rules in almost all other copy protection license agreements today including the Digital Transmission Content Protection (DTCP) license agreement and the High-Bandwidth Digital Content Protection (HDCP) license agreement. Respectively, these technologies are used today to protect the carriage of premium content across an IEEE 1394 interface, a Universal Serial Bus (USB), a Bluetooth interface or via the Internet Protocol (IP) transport; and protect the carriage of premium content across a Digital Video Interface (DVI) or a High Definition Multimedia Interface (HDMI). These rules were also the basis of the Robustness Rules defined in the DFAST license agreement.[1] Digital television receiver manufacturers must sign the DFAST license agreement in order to build Unidirectional Digital Cable Products also known as "Digital Cable Ready" receivers.

Creating the definition of the overall device robustness for DVD-Video playback systems was one of the most arduous and intense debates of the original CSS license negotiations. The group established rules describing the level of threats that devices must be designed to deter and some of the measures to be used in deterring them. The level of threat was described in terms of expertise and tools available to the attacker, with the general understanding that consumer products like DVD players would not be expected to thwart attacks by parties with expert training and professional equipment (note that such parties may be deterred through criminal litigation). They are the main reasons why, to the benefit of consumers, the cost of DVD players have reached a low of $25 US.

What is a "User Accessible Bus"*?*

To provide a balanced and competitive market for DVD play back devices, the group had to devise a scheme whereby open architecture devices (e.g. personal computers) were considered to have an equal amount of device robustness as closed architecture devices (e.g. standalone DVD players). This scheme was achieved in part by defining how to determine what busses in a device may or may not be vulnerable to data interception. Busses that may be vulnerable to data interception were prohibited from carrying unscrambled, compressed CSS content. The concern was that content could easily be

[1] The U.S. Federal Communications Commission (FCC) has assumed regulatory oversight of the Compliance and Robustness Rules set forth in the DFAST License Agreement.

intercepted by rogue software or external hardware while in transit over a bus which may be vulnerable to such data interception. The carriage of scrambled CSS content over such busses was permitted and interfaces that were not considered vulnerable to data interception were allowed to carry unscrambled, compressed CSS content.

The Procedural Specifications of the CSS license agreement established the scheme described above by introducing the term "*user accessible bus*" to classify interfaces. The definition of "*user accessible bus*" in the Procedural Specification of the CSS license agreement is as follows:

> *A "user accessible bus" means a data bus which is designed for end user upgrades or access such as PCI, PCMCIA, or Cardbus, but not memory buses, CPU buses, and similar portions of a device's internal architecture*

The spirit of this definition was to be interpretive rather than definitive. As times and technologies have changed since the original crafting of the definition, the only reasonable test that can be used to determine if an interface is vulnerable is through direct technical comparison. First, however, it is important to dispel some of the now ineffective tests used to determine the vulnerability of an interface.

One of the most common tests for determining if an interface was a "*user accessible bus*" was the presence of a physical connector on the interface. A convenient determination would be to presume that any interface with a physical connector is a vulnerable interface; however it is the accessibility of the *content* that crosses the interface that is of prime importance not the accessibility of the *physical component* that connects to the interface. As an example, personal computer users are twice as likely to upgrade system memory as to upgrade a graphics card[2] but since the data communications across the system memory bus are complex and fragmented it cannot be deemed a vulnerable interface.

Another test for determining if an interface was a "*user accessible bus*" was the existence of a private license agreement for the interface. The effectiveness of this test is diminished by its inconsistent application among interfaces within a device. For example, in a personal computer, a system memory bus may be publicly available whereas an HDMI interface is privately licensed. In contrast, a system memory bus is *not* considered a vulnerable interface (it is excluded by the definition of *user accessible bus*) whereas an HDMI bus is considered a vulnerable interface (it is considered a device output that requires protection such as HDCP). This contradiction invalidates the use of this test.

The final and most ineffective test for determining if an interface was vulnerable was a count of the number of pins or connection points on the interface. It is believed by some that: the fewer the pins, the more accessible the data on the interface. This test is not only arbitrary, but obsolete. Today's parallel buses, with large pin counts, have now reached their practical technical limits, and serial interconnects, with low pin counts, are now favored throughout the technology industry. To increase data rate over serial interconnects, designers have to markedly increase operating speed and packetize data transmitted over the interface. As will be discussed later in this paper, both of these

[2] See NPD Intellect Report 2004

compensations significantly increase the complexity of an interface and demand very sophisticated tools for monitoring and debugging. Thus, as the technology industry moves towards serial interconnects to meet ever-increasing data rate improvements, this test has become obsolete.

PCI Express Interconnect versus "User Accessible Bus"

There are three buses called out in the contemporary definition of "*user accessible bus*". The PCMCIA and CardBus buses are reasonably simple buses lacking the more advanced features of modern bus designs, such as the ability for any component on the bus to become a transmitter, therefore, the only "*user accessible bus*" that is worthy of comparison to the PCI Express interconnection is the PCI local bus.

The PCI local bus and PCI Express interconnect are similar in name but vastly different in technology. The vast difference between the two interfaces stems from the fact that the PCI local bus is a multi-point parallel interface and the PCI Express interconnect is a point-to-point serial interface. These attributes create two totally different physical topologies as well as two totally different physical connections. These differences are shown pictorially below. The top figure (a) shows the physical topology for the PCI local bus and the bottom figure (b) shows the physical topology of the PCI Express interconnect. Note that the PCI Express interconnect figure shows only a single point-to-point lane (x1), however, some components, such as GPUs, have sixteen point-to-point lanes (x16).

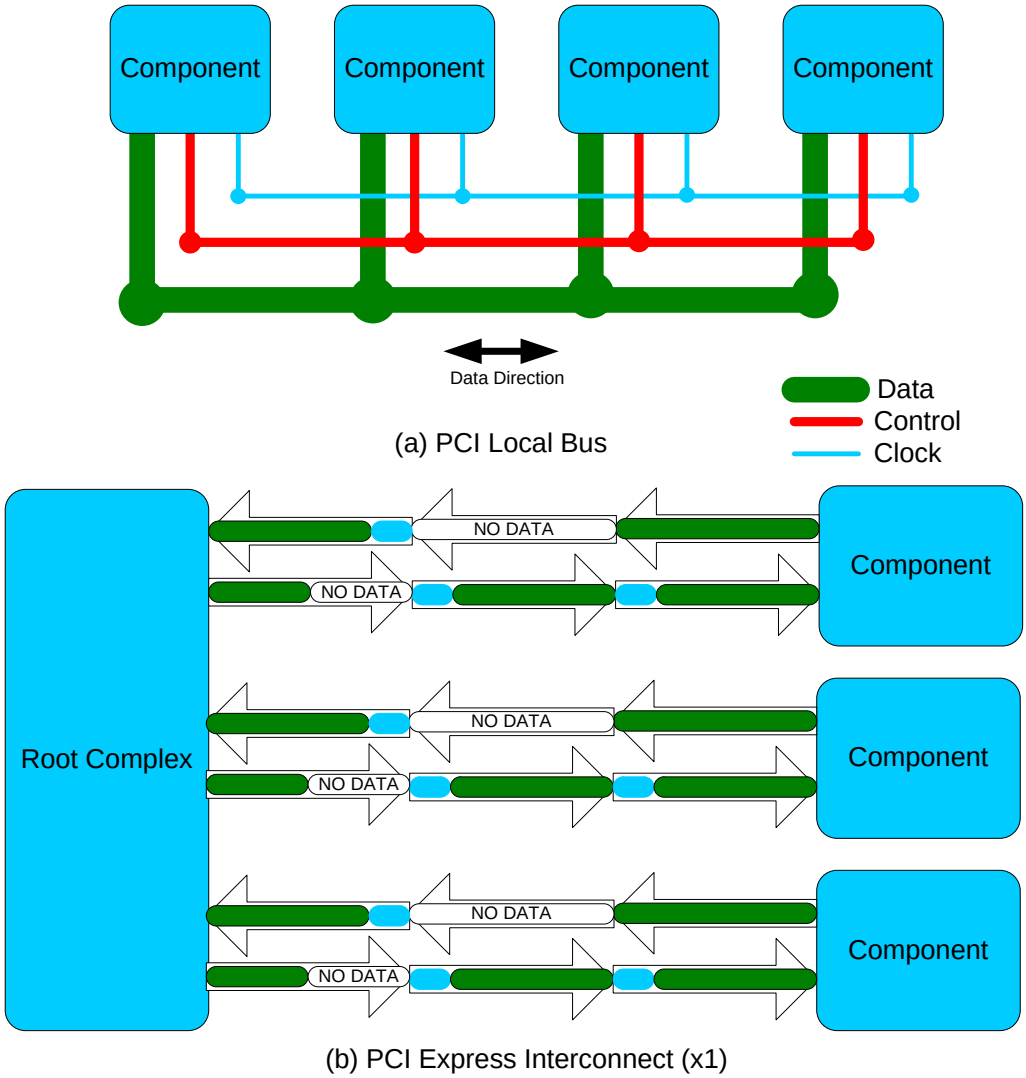

Figure 3: PCI Local Bus Topology vs. PCI Express Interconnect Topology

PCI Local Bus Description

In its physical topology, a PCI local bus links several components using one global wiring connection. This is a traditional bus design consisting of one sideband clock signal associated with the data and control signals. A master component transmits information over a PCI local bus and slave components receive the information. This topology is analogous to a cable television network whereby the cable system transmits a single signal and all televisions connected to the network receive the same signal.

Receiving data on the PCI local bus is a relatively easy technical operation. As part of the bus protocol, the master component readies 32 bits of information concurrently on the bus prior to a rising and/or falling edge transition of the clock signal. The master component ensures that the 32 bits of information are stable on

the bus before the clock edge transition so that slave components can use the transition as a trigger to sample the information. Since all of the components on the PCI local bus are connected to the same wires and the clock is an independent signal, any slave component can easily access any 32 bits of information traversing the bus. This bus protocol and topology makes any data transmitted on the bus vulnerable to data interception.

As a final point, due to practical technical limits of today's printed circuit board design, chip packaging design and chip interface design, all multi-point bus topologies with a sideband clock signal operate at a moderately low frequency. The PCI local bus clock frequency, for example, does not exceed 66 MHz although the vast majority of PCI local buses currently deployed in devices have a clock frequency of 33 MHz. A low frequency makes multi-point buses easy to monitor and debug using relatively inexpensive equipment. Monitoring and debugging can be made even easier by reducing the clock speed through a simple system board component change; however, clock speed reduction has implications on the overall stability of the device. In summary, external equipment may readily access data traveling over a multi-point bus.

PCI Express Interconnect Description

In contrast to the PCI local bus topology, the PCI Express interconnect topology links only two components over serial connections. The simplest type of PCI Express interconnect consists of a single serial data connection for transmission and another single serial data connection for reception. The combination of these two connections is called a "lane". The PCI Express interconnect can have up to 16 lanes. This point-to-point topology is similar to a phone call between two people over a scrambled line where the conversation is private and only the participants of the phone call receive each other's information or signals.

To provide adequate data rates over its serial interface, the PCI Express interconnect must operate at a high frequency. Since data is transmitted directly from one component to another, the only way to intercept data traveling over a PCI Express interconnect is by inserting multiple acquisition probes or a single interposer device between two components. The provision for multiple lanes makes the PCI Express interconnect ideal for components that require high volume data communication, like a GPU, but it compounds the data interception problem due to its high volume data traffic.

As discussed in the PCI local bus section above, due to the practical technical limits of routing a high frequency clock over printed boards, chip packages and chip interfaces, the PCI Express transmitter component must embed the clock into *both* connections in each lane by 8-bit to10-bit encoding of data into symbols. The PCI Express receiver component recovers the clock by "watching" symbol bits switch every 400 picoseconds, locking an oscillator to the minimum period of the symbol bit changes, and aligning an oscillator's rising edges so that they occur at the mid-point of the symbol bit. It cannot be overstated that recovering the 2.5 GHz clock on the PCI Express interconnect is an extremely difficult technical

endeavor. As an added effect, clock embedding also requires the interconnect to *always* run at its nominal frequency of 2.5 GHz—over 75 times faster than the PCI local bus—thus the interface cannot be slowed down for monitoring or debugging purposes. As a result of clock embedding, debugging and intercepting data transmitted over the interface requires specially designed printed circuit boards with multiple, dedicated acquisition probe connections and sophisticated capturing equipment.

Higher operating frequencies also make the PCI Express interconnect vulnerable to signal reflections resulting in garbled data[1]. To maintain the electrical signal integrity of the interface, data is required to be scrambled or randomized before being transmitted. This technique enhances the interface's resiliency to reflections by allowing non-periodic data to be carried over the interface. This technique reduces the possibility that multiple acquisition probes could be used to capture the transactions crossing a PCI Express interconnect.

Since electronics running at high frequencies consume significant amounts of power, the PCI Express interconnect specification allows components to enable and disable transmissions through a power management scheme. As shown in the figure below, the electrical signaling on the interface can appear and disappear dynamically based on the load. The power management scheme is performed by autonomous hardware directed communication and does not involve software interaction whatsoever. It eliminates the possibility that multiple acquisition probes could be used to extract data traveling over the PCI Express interconnect

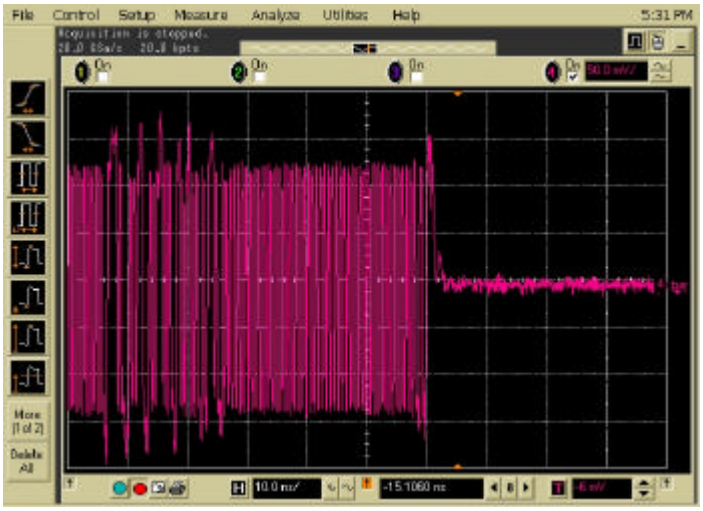

Figure 1: PCI Express Interconnect Power Management

.

Another consequence of serializing the PCI Express interconnect is that it requires data traversing over *each* lane connection of the interface to be individually packetized before transmission and to be individually de-packetized upon

[1] Signal reflections make it almost impossible to recover the clock, and thus data.

reception. The packetization process adds complexity to the transactions by adding control information to the data being transmitted. The control information is applied over three layers of abstraction. The layers and their added control information are shown in the figure below. The topmost layer is the Transaction Layer. The main purpose of this layer is to indicate the type of transaction that is requested. The next layer is the Data Link Layer. The main purpose of this layer is to ensure the integrity of the transactions between components. The bottommost layer is the Physical Layer. This layer is responsible for actual transmission and reception of transactions across the PCI Express interconnect.

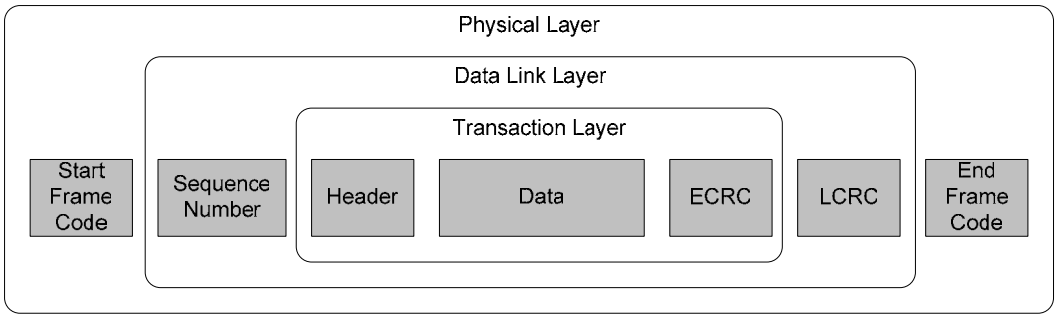

Figure 5: Layers of Packet Abstraction

Without describing the elements of each layer, one can plainly see in the figure above that, even if a device could capture all the transactions crossing the PCI Express interconnect, it would take a massive amount of processing to parse each packet and determine their type. It should be emphasized that capturing content in this manner is further complicated when PCI Express interconnect components are connected through multiple lanes as receivers must independently depacketize data from *each* lane and re-assemble data packets from all lanes into their original transmission order. The only way to process packets in this manner is via an interposer device. Interposer devices are described in the next section below.

Finally, although the PCI Express interconnect abstraction layers are standard; the data carried in the Transaction Layer is proprietary. Modern add-in GPUs for personal computers, for example, are extremely complicated chips and their programming interfaces are unique, proprietary and non-public. If a device could capture transactions crossing the PCI Express interconnect to a GPU it would have to understand how the GPU works internally and how the operating system manages the system resources—a very tall order. In addition, most GPUs do not process audio. The audio stream would have to be captured from the audio decompression system on the personal computer and re-packetized into the captured stream.

Interposer Device

So far the PCI Express interconnect discussion has focused on why multiple acquisition probes could not be developed to intercept data traversing across a PCI Express interconnect, but what about a single interposer device that could be placed in between a PCI Express root complex and a GPU? An interposer device is shown in light green in the figure below.

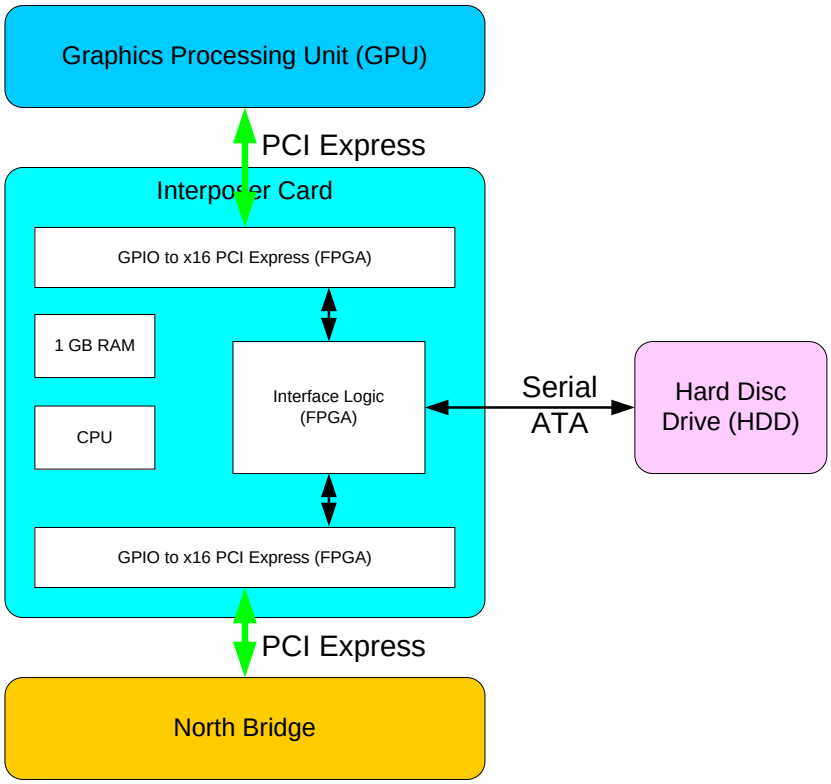

Figure 6: Intel Desktop Computer Architecture Showing an Interposer (Light Green)

An interposer device would have to have two PCI Express interconnects one to connect to the North Bridge and one to connect to the GPU. In addition to the PCI Express interconnects the interposer device would have to embed enough logic, processing power and RAM to process the data passing through the device to record content traversing the interface. The hard disc drive would have to be used to store the content. Since this is a theoretical exercise it is probable that other components would be required to control the interposer device.

Using the example system shown above, capturing PCI Express interconnect transactions for a 90 minute movie would require the interposer device to process several hundred billion proprietary GPU instruction/command packets interspersed with millions of control traffic packets, interrupt packets, PCI Express link management packets and other miscellaneous packets. Given the excessively high data rate it would only be practical to build a specialized chip or a board with several high-speed FPGAs. Developing a specialized chip is a serious undertaking as it is very expensive to produce such a chip in low volume.

Developing a board with several high-speed FPGAs would require extensive system development so that the FPGA could communicate with each other.

It is our assessment that, at the very least, developing a complete interposer would take the following number of engineers:

- *An expert engineer to design interposer architecture*
- *An expert engineer for board design*
- *Multiple expert engineers to reverse engineer the non-public, proprietary Transaction Layer format of a GPU*
- *An expert engineer for interposer control software*
    - *Needs to develop software for memory management, disk I/O control, file system and remote I/O control*
- *An expert engineer for PCI Express interconnection software*
    - *PCI Express interconnect flow control, parsing, store-and-forward mechanism*
- *Multiple expert engineers for video stream parsing software*
    - *Needs to develop a system to dynamically re-assemble and reconstruct video frames*

This is a significant undertaking especially given the fact that there is no guarantee that all of the content will traverse the PCI Express interconnect. There is significant probability that some of the video processing and all of the audio processing will be performed in another part of the personal computer. This fact will severely reduce the utility and success of an interposer device.

PCI Express Interconnect – A Vulnerable Interface?

The descriptions above plainly show that, technically, the PCI Express interconnect is significantly more complex than any defined "*user accessible bus*" and in particular the PCI local bus. Just given the topologies of these interfaces it should easily be understood why content traversing the PCI Express interconnect is extremely difficult to intercept. To presume that communication across any point-to-point interface is straightforward and, therefore, vulnerable to interception, is to ignore the facts. After a thorough analysis of the PCI Express interconnect and its protocol, it is clear that such presumptions are wholly unwarranted. The PCI Express interconnect, therefore, cannot be considered a vulnerable interface according to this analysis, and, thus, does not fit into the category of "*user accessible bus*".

Conclusion

This paper explored the genesis, necessity and definition of device robustness and specifically the vulnerability of data interception over an interface. The definition of "*user accessible bus*" is the measure with which all interfaces in an open architecture device, such as a personal computer, are compared to in order to determine the accessibility of content traversing over the interface. Since the definition is interpretive rather than definitive, one must perform technical comparisons of an interface in question with the buses identified as "*user accessible buses*". A summary of the comparison between interfaces described above and the PCI Express interconnect is shown in the table at the end of this paper.

This paper shows that the PCI Express interconnect is significantly more complex than the PCI local bus, and is clearly in a different category as it relates to susceptibility to data interception, since an acquisition probe to capture PCI Express interconnect transactions cannot be built and an interposer device that could capture PCI Express interconnect transactions for periods longer than a few seconds can only built using prohibitively expensive engineering resources and systems.

Summary Comparison of Buses versus the PCI Express Interconnect

Functions	User Accessible Bus		Internal Architecture		PCI Express	Implications for Data Interception
	PCI Local Bus	PCMCIA	Front-Side Bus (FSB)	System Memory		
Physical topology supports snooping	Yes	No	No	Yes	No	PCI local bus and system memory are multi-point busses, which route the interface cycle parameters to all connectors, thereby allowing a device plugged in another connector to "see" all of the cycle information of the targeted device. PCI Express interconnect does not, since it is a point-to-point connection
Data is encoded	No	No	No	No	Yes	PCI Express interconnect data is 8b/10b encoded
Data is randomized	No	No	No	No	Yes	PCI Express interconnect data is randomized or scrambled
Interface includes a clock to strobe the data	Yes	Yes	Yes	Yes	No	The clock for PCI Express interconnect must be generated by "watching" the data signal switching every 400 picoseconds, locking an oscillator to the minimum period of the data changes, and aligning the oscillators rising edges so that they occur at the mid-point of the symbol bit, all the while doing this at 2.5 GHz. Simply put this operation is technically difficult.
A single clock defines a transaction	Yes	Yes	Yes	Yes	No	Each lane of PCI Express interconnect has it's own clock, thus multiple lane interconnects must re-assemble packets, removing data skew between lanes
Transactions are packetized	No	No	Yes	No	Yes	Packetized transactions have to be parsed to determine packet contents. Some packets contain control information, some contain address information, and some contain data. Packetized transactions are inherently more difficult to "read".
Interface cycles associate read address with the read data	Yes	Yes	No	Yes	No	PCI Express interconnect sends the read request and the read data in two separate packets, spaced apart significantly in time, at somewhat random intervals. A PCI Express interconnect receiver must be able to re-associate the read data with the read request to determine what the address was, in order to determine what a linear portion of data image looks like.
Minimum operating speed	Depends on device stability	Depends on device stability	Depends on device stability	Depends on device stability	Fixed to 2.5 GHz	
Complexity of "snooping circuit"	Easy	Easy	Difficult	Medium	Very difficult	

Glossary

CardBus	A variant of the PCMCIA bus.
CPU	Central Processing Unit. The main component of an open architecture device that executes the operating system. Intel Pentium® and AMD Opteron™ are types of CPUs.
CSS	Common Scrambling System. The copy protection system for read-only DVD-Video discs. This technology is licensed from the DVD Copy Control Association (DVD CCA), The license is split into 2 parts CSS License Agreement and CSS Procedural Specifications.
DFAST License Agreement	A license agreement required to obtain DFAST (Dynamic Feedback Arrangement Scrambling Technique) technology licensed by CableLabs. Digital television receiver manufacturers must sign this license in order to produce Unidirectional Digital Cable Products also known as "Digital Cable Ready" receivers.
DTCP	Digital Transmission Content Protection. A link protection system used mainly to protect content from being copied over certain interfaces. This technology is licensed from the Digital Transmission Licensing Authority (DTLA).
DVI	Digital Video Interface. An interface that carries uncompressed digital video first developed by Intel and Silicon Image in the mid 1990's.
FCC	United States Federal Communications Commission
FPGA	Field-Programmable Gate Array. A type of logic chip that can be programmed. (*From Webopedia*). With an FPGA, a design engineer is able to program electrical connections on site for a specific application (for example a device for a sound/video card), without paying thousands of dollars to have the chip manufactured in mass quantities.
FSB	Front-Side Bus. A CPU bus for connection to a North Bridge chip. HyperTransport™ is a type of front-side bus.
GPU	Graphics Processing Unit. A processor that accelerates the rendering of graphics and decoding of video content. A GPU is equal to and, in some cases, exceeds the complexity of a CPU.
HDCP	High-Bandwidth Content Protection. A link protection system for uncompressed digital video interface such as DVI and HDMI interfaces. It is licensed by Digital-CP, LLC.
HDMI	High Definition Multimedia Interface. An uncompressed digital video interface developed by several companies and licensed by HDMI, LLC.
HyperTransport™	A front-side bus that connects a CPU to a North Bridge. It is licensed by the HyperTransport Consortium.
I/O	The term I/O is used to describe any program, operation or device that transfers data to or from a computer and to or from a peripheral device. (*From Webopedia*)
IP	Internet Protocol. The main transport protocol for the Internet and home networks.
IEEE 1394	Also known as FireWire™ and i.LINK™. A general purpose interface used to carry data or audiovisual content to hard disc drives. It is also used to carry audiovisual content to camcorders and digital storage devices (e.g. D-VHS)
PCI Express	A serial, point-to-point interconnect. Also written as PCIe or PCIE.
PCI local bus	Peripheral Component Interconnect local bus. A parallel, multi-drop point, parallel local bus developed by Intel in the early 1990s defined as a "user accessible bus". The PCI local bus is almost ubiquitous in personal computers today.
PCMCIA	Personal Computer Memory Card International Association. Also known as PC Card. A variant of this bus is called CardBus. A multi-point, parallel bus defined as a "user accessible bus". PCMCIA is almost ubiquitous in laptop personal computers today.

Federal Communications Commission

The FCC Acknowledges Receipt of Comments From ...

ATI Technologies, Dell, Inc., Hewlett-Packard Company, and Intel Corporation

...and Thank You for Your Comments

Your Confirmation Number is: '2006120619325'

Date Received: Jan 20 2006

Docket: 97-80

Number of Files Transmitted: 1

DISCLOSURE

This confirmation verifies that ECFS has received and accepted your filing. However, your filing will be rejected by ECFS if it contains macros, passwords, redlining, read-only formatting, a virus or automated links to source documents that is not included with your filing.

Filers are encouraged to retrieve and view their filing within 24 hours of receipt of this confirmation. For any problems contact the Help Desk at 202-418-0193.

Initiate a Submission | Search ECFS | Return to ECFS Home Page

FCC Home Page | Search | Commissioners | Bureaus Offices | Finding Info

updated 12/11/03