Characterizing Opportunities for Short Reach Optical Interconnect Adoption:
A Market Survey and Total Cost of Ownership Model Approach

by

Johnathan Jake Lindsey III

B.E. Computer Engineering
Vanderbilt University, 2006

SUBMITTED TO THE ENGINEERING SYSTEMS DIVISION IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE IN TECHNOLOGY AND POLICY
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2010

Signature of Author:

Engineering Systems Division
November 13, 2009

Certified by: ____________

Randolph Kirchain
Associate Professor of Materials Science and Engineering Systems
Thesis Supervisor

Accepted by: ____________

Dava J. Newman
Professor of Aeronautics and Astronautics and Engineering Systems
Director of Technology and Policy Program

Characterizing Opportunities for Short Reach Optical Interconnect Adoption:
A Market Survey and Total Cost of Ownership Model Approach

by

Johnathan Jake Lindsey III

Submitted to the Engineering Systems Division on November 13, 2009
in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Technology and Policy

Abstract

Over the past decade, the demand for digital information has increased dramatically with the rising use of the Internet and various types of multimedia data – text, audio, graphics, video, and voice. As a consequence, the technologies that connect and transport data have become critically important. Available interconnect technologies are broadly organized into two categories: electrical and optical. Although many digital systems use electrical interconnects, optical interconnects are becoming an attractive alternative as electrical connection has become increasingly difficult in terms of cost and performance. However, the transition from electrical to optical interconnects across multiple markets could still be hampered by its higher cost relative to interconnects in the mid-term. Thus, this work seeks to shed light on the following question: "What additional characteristics are useful to evaluate the attractiveness of optical interconnects in emerging markets?" This thesis seeks to explore and answer this question in three parts.

The first part of the thesis attempts to gauge the opportunities and barriers to optical interconnect adoption in emerging markets through an analysis of first phase interviews with professionals working in the datacom, automobile, consumer hand-held device industries. Initial review of the response set shows that of the five initial emerging markets for optical interconnect, datacom, specifically high-performance computing (HPC), has the greatest potential for increased optical interconnect adoption in the near future. To further explore the environment for optical interconnects in the HPC, a second, more detailed questionnaire was distributed to a limited number of interviewees. In response to this interview, some respondents noted that several metrics other than cost and performance, particularly power consumption, as being "very important" when deciding which technology to adopt.

The second part of the thesis is primarily concerned with investigating further the influence that power and performance concerns have on optical interconnect adoption in HPC data centers. Specifically, this part of the thesis seeks to explore whether power concerns in data centers could lead to increased adoption of optical interconnects. To that end, a cost model of an HPC data center has been developed to identify the possible economic impacts that the adoption of optical interconnect technologies would have in a power-driven scenario. The third part of this thesis presents a set of policy recommendations based on the results from the data center cost model.

Thesis Supervisor: Randolph Kirchain
Title: Associate Professor of Materials Science and Engineering Systems

Acknowledgments

First, I wish to acknowledge Dr. Randolph Kirchain for his role as thesis supervisor. I would also like to acknowledge Dr. Tremont Miao, Dr. Allen Benner, and Dr. Louay Eldada for their roles as Chairs of various technology-working groups in the Communication Technology Roadmap Project of the M.I.T. Microphotonics Consortium. Their advice and guidance was extremely helpful during the initial stages of this research project.

I would like to thank Professors Ken Oye and Nicholas Ashford. Their lectures and guidance have broadened my worldview and inspired me to pursue a different path.

I would like to especially thank my fellow TPP students and good friends: Rob, Junjay, Dai, Tim, and Kyle. I am incredibly grateful to these guys for their support and feedback. Without their help and friendship, these two years would have been much more difficult.

I would also like to extend my profound gratitude to Dean Steve Lerman and Dean Blanche Staton for their assistance in my quest to complete this thesis project.

Finally, I would like to thank my mother, Angela Lindsey, for her guidance over the years. In both good times and bad times, she has been a pillar of strength. Her words of encouragement have inspired me to persevere.

Above all, I would like to thank God for I know I can achieve nothing in this world without his everlasting love and grace.

Table of Contents

Table of Figures

Table of Tables

Chapter 1 – Introduction

1.1 Motivation

Over the past decade, the demand for digital information has increased dramatically. As a consequence, the technologies that connect and transport that information have become critically important. Available interconnect technologies can be organized into two broad categories: Electrical and Optical. Electrical interconnects are links that transmit digital information using electrical signals sent via copper wire, while optical interconnects use light to transmit signals through polymer or glass fiber. Because electrical interconnects are generally less expensive than optical ones to implement in short reach applications (figure1) copper interconnects have come to dominate this market space. However, as demand for bandwidth continues to rise rapidly, electrical connection has become increasingly difficult in terms of cost and performance.

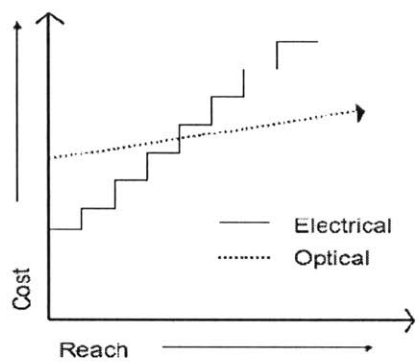

Figure 1: Cost vs. Reach [1]

Now, experts have begun to look at optical interconnects as a promising alternative to provide efficient, high-speed data transmission over shorter link lengths [2].

Historically, adoption trends for optical interconnects in various markets are driven by cost per performance and cost per performance x distance. Figure 2 shows the market adoption trends of optical and electronic interconnects in terms of 'bandwidth by reach distance.' While commercially available on the market, optical interconnects have historically been deployed in longer reach applications for which the cost to transmit digital information via copper cables becomes prohibitively expensive [3].

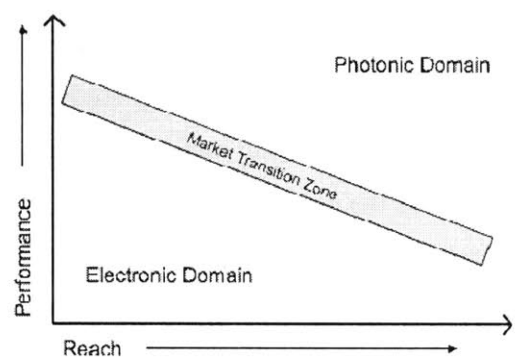

Figure 2: Market Transition [3]

For these applications, optical interconnects have, in fact, displaced copper solutions to become the incumbent technology. In contrast, optical interconnect adoption for shorter link distances have been hampered by their inability to compete with electrical interconnects in terms of cost. How technology will progress in such a way to lower costs of optical interconnects for shorter-reach applications is a question that greatly interest the photonics and electronics communities. In addition, predicting how and when various markets and applications will transition from electrical to optical solutions in the future has become a key research objective.

1.2 Current Research Efforts: Communications Technology Roadmap

To better understand the technology and market dynamics that impact future optical interconnect adoption, the Microphotonics Center Consortium at the Massachusetts Institute of Technology established the Communications Technology Roadmap (CTR) Project in 2000. The CTR is a collaborative research effort between research scientists, students, and industry leaders which works to create a roadmap that identifies and tracks key milestones in the evolution of optical components through a collaborative research effort. The CTR is in its second phase of research and analysis. In the initial phase (2000-2004), the CTR project comprised four technical working groups (TWGs): Next Generation Transceiver TWG, the III-V Materials TWG, the High-Performance Transceiver TWG, and the Silicon Materials TWG. In its report, *Microphotonics: Hardware for the Information Age*, the consortium found that "the future of components technology will be driven by electronic-photonic convergence and short-reach interconnection" [3].

In its current phase, the CTR seeks to comment further on the future direction of opto-electronic components by studying the economic viability of a silicon CMOS platform, the evolving integration and packaging methods for interconnect technologies, as well as the cross-market opportunities for optical interconnects. This thesis is most aligned with the Cross-Market TWG; the goal of this working group is to develop a ten-year qualitative roadmap that identifies opportunities for convergence in future architectural requirements for optical interconnects across multiple applications. The resulting roadmap could provide a significant contribution to the standardization effort for optical components, and in turn help to improve their ability to compete with other incumbent technologies such as copper and wireless.

1.3 Short Range Optical Interconnects as Disruptive Technologies

As discussed above, the cost per performance relative to incumbent copper technologies has been a significant barrier to the adoption of optical interconnects in a variety of short reach applications. However, for a technology that has the potential to be used in a diverse number of market segments, using cost or cost per performance as the deciding factor can sometimes lead engineers to make adoption decisions based only on short-term interests without fully examining the possible long-term benefits. In this thesis research, broadening the study of short reach optical interconnects to include its potential to disrupt current technology trends could yield some useful insight regarding the future integration in new applications. Recognizing and accounting for an innovation's disruptive potential in any discussion regarding the adoption of new technologies can benefit both producers and consumers within a given market. For consumers, accounting for a technology's disruptive potential when considering whether or not to adopt a disruptive technology can enable buyers to fully consider future technology requirements (other than cost) and better communicate those needs to their component suppliers. Farther up the supply chain, producers benefit from an increased awareness to those needs. Producers can use the additional information garnered from consumers and their own projections are to help balance against or even outweigh the cost barrier for adoption. In some cases, producers who neglect to consider the disruptive potential of emerging technologies risk loss in market leadership or failure [4]. Therefore, in an attempt to better understand the opportunities for the adoption of short-reach optical interconnects in a context beyond just relative cost per performance, this thesis will frame its examination of emerging optical interconnects as a disruptive technology.

Disruptive technologies are innovations, which use unconventional strategies to overturn an incumbent technology in a market. In his book, *The Innovator's Dilemma,* Clayton Christensen identifies three primary characteristics that are common among most disruptive technologies. First, disruptive technologies are usually lower in cost than the incumbent technology. Second, disruptive technologies almost always have lower traditional performance compared to existing solutions. And third, disruptive technologies display superior performance in ancillary characteristics. Superior performance with regards to these secondary metrics allows a disruptive technology to create a niche market for itself in which it can mature, further reduce unit production cost, or continue to build production volumes. Progress in a secondary market will allow a disruptive technology to mature (technology advances or reductions in production costs) and later overtake incumbent technologies in the original market.

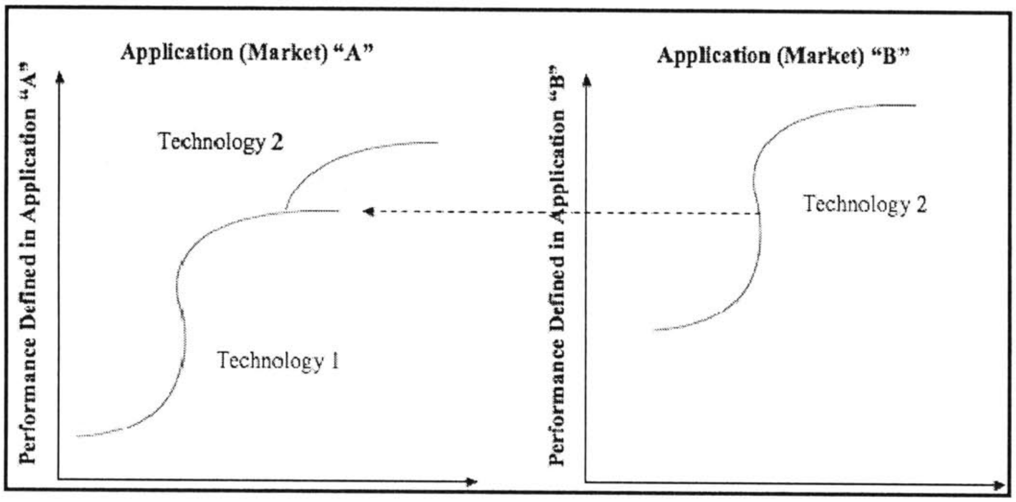

Figure 3: Disruptive Technology S-Curve Model [5]

This dynamic is illustrated in Figure 3 using generic s-curve plots; these plots illustrate how potentially disruptive innovations that are not ready for the initial market can evolve and mature in a secondary market before displacing an incumbent in a primary market.

In his study of disruptive technologies, Christensen greatly emphasizes the lower cost, lower performance scenario described above. However, this construction can be somewhat narrow and exclude those disruptive innovations that may not display one or more of the above characteristics. While exhibiting lower performance (in terms cost per bandwidth) and potentially greater ancillary performance, optical interconnects have an overall higher production cost than electrical interconnects which, by Christensen's construction, disqualifies them as a disruptive technology.

However, others contend that Christensen' characterization of a disruptive technology is perhaps too narrow. In a later study of discontinuous innovations, Utterback and Acee search for a broader view of disruptive technologies: "[B]y emphasizing only "attack from below" [scenarios] Christensen ignores other discontinuous patterns of change, which may be of equal or greater importance" [6]. In their study, the authors present several instances where discontinuous innovations do not conform to the "disruptive" pattern [6]. By demonstrating that disruptive technologies can displace the incumbent even with higher production cost, they offer an alternative view: an innovation is a disruptive technology if has the powerful ability to broaden markets and provide new functionality. Under this broader construction, even technologies that fail to meet one or more of the requirements defined by Christensen can still be categorized as disruptive if one could demonstrate its potential to expand markets or functionality.[1] Besides being more inclusive, a major consequence of this alternative construction is that performance in areas not traditionally valued by an initial market or customer base becomes a more important consideration when evaluating potential success in other secondary markets.

Using Utterback's and Acee's broader framework, identifying novel applications and evaluating the ancillary performance of short-range optical interconnects could provide valuable insight when researching their potential for adoption in a variety of markets. In particular, this thesis will explore the potential for optical interconnect adoption in the automotive, consumer handheld, and high performance computing markets.[2] The central questions that have guided this research effort are outlined in the following section.

1.4 Central Questions

This thesis will explore four objectives that have been distilled from the overall motivation and goals of the CTR project. First, this thesis seeks to identify general adoption trends and barriers to optical components in each of the three aforementioned markets. Using the insights from an initial survey of all three markets, the second objective of this thesis is to identify a single market or application that displays the

[1] James Utterback and Happy Acee examined seven cases (including compact discs, digital cameras, and fuel injection engines) that did not fit Christensen's disruptive technology pattern to support their conclusions.

[2] The Cross-Market TWG originally endeavored to examine five markets in order to create a 10 year roadmap: Automotive, Consumer Handheld, High-Performance Computing, Defense and Home Video; however, the Defense and Home Video markets were later excluded from this discussion due to low response to interview requests and limited availability of literature.

greatest potential for optical interconnect adoption within the near future. The third objective is to find and quantify ancillary characteristics that could prove to be useful when evaluating the attractiveness of optical components for a target application. The fourth objective of this research to present recommendations to guide the adoption of optical interconnects in short reach applications.

This thesis approaches the above objectives by addressing each of the four central questions that have been outlined below:

1. What are the opportunities for optical component adoption in several markets? And which market shows the greatest potential for near-term adoption?
2. What quantifiable ancillary characteristics are useful to evaluate the attractiveness of optical components in an emerging market? What market influences contribute to their relative importance?
3. With respect to the ancillary characteristics, what cost benefit could the implementation of short reach optical interconnects have in a target application?
4. How can the results of this research be used to create coherent policy recommendations regarding the implementation of short-reach optical interconnect technology in a target application?

What are the opportunities for optical component adoption in several markets? And which market shows the greatest potential for near-term adoption?

As discussed above, the primary purpose of the Microphotonics Consortium's Cross Market TWG is to identify opportunities for convergence in future architectural requirements for optical interconnects across multiple applications. Realizing this goal, however, would require an expansive interview effort of both suppliers and consumers of optical interconnects in each industry as well as comprehensive analysis of the response set. The research effort required to achieve that goal is far too broad to pursue given time and resource constraints of this thesis. Instead, this thesis will begin by identifying opportunities for the increased adoption of optical interconnects in short reach applications in several industries using qualitative and quantitative phone interviews.

Chapter 2 describes the surveys used to identify opportunities and challenges to optical component adoption across the automotive, high-performance computing, and consumer handheld industries. Furthermore, this chapter provides a brief overview of the results from the initial qualitative surveys. Using the insights from the survey responses and a literature survey, the scope of the thesis research is narrowed to a closer examination of a single application that shows the most potential for increased optical interconnect adoption in the future.

What quantifiable ancillary characteristics are useful to evaluate the attractiveness of optical components in an emerging market? What market influences contribute to their relative importance?

For many working with optical components, a major barrier that has prevented widespread adoption is the prohibitively high cost per performance required to apply short reach optical interconnects. However, making decisions as to whether or not to adopt a technology based solely on the initial cost can be narrow, and in some cases, could result in long-lasting negative repercussions in the future. One way to broaden this perspective is to view optical interconnects as a disruptive technology. Reframing the question to consider ancillary characteristics beyond the initial startup costs could change the way decisions by balancing short-term interests with long-term concerns.

Once the project scope has been narrowed to single target application, chapter 2 provides results from a second, more quantitative, survey used to identify trends for key performance benchmarks as well as ancillary characteristics that could possibly impact optical component adoption within the target market. Chapter 3 explores the technological and market influences that contribute to the importance of those ancillary characteristics in the target application. Furthermore, chapter 3 outlines in what ways implementing optical interconnects could impact the target application with regard to the ancillary characteristics identified in the preceding chapter.

With respect to the ancillary characteristics, what cost benefit could the implementation of short reach optical interconnects have in a target application?

Simply identifying ancillary characteristics is insufficient to promote the increased adoption of short reach optical interconnects. It is also necessary to demonstrate what, if any, additional value can be gained from using optical interconnects instead of copper ones with respect to potentially increased ancillary performance. This thesis research pursues this objective by demonstrating a financial impact of optics adoption in the form of cost savings relative to incumbent copper technology. Chapter 4 introduces a simple cost model structure that has been tailored to the previously identified target application and ancillary characteristics while Chapter 5 gives an analysis of the model's results.

How can the results of this research be used to create coherent policy recommendations regarding the implementation of short-reach optical interconnect technology in a target application?

Chapter 6 of this thesis integrates the major insights from the model used in the preceding chapters to discuss how the current policy framework might affect optical interconnect adoption and outline potential policy recommendations regarding future adoption of short-reach optical interconnects within the selected target application. At the end of this chapter, the thesis concludes with brief summation of findings and suggestions for future work.

Chapter 2 – Market Insights[3]

As outlined in the preceding chapter, one of the primary objectives of this research is to identify additional metrics other than cost or cost per performance that might be useful when evaluating the attractiveness of optical interconnect adoption. However, this question is largely dependent on the context in which it is posed. In this case, the choice of metrics to be examined and the market factors that make them important is largely dependent on the application chosen for further study. Therefore, the goal of this chapter is to identify a target application based on information collected during qualitative and quantitative surveys.

In this chapter, the results from an initial set of qualitative interviews are presented. For the purposes of this research, the interview results are used to identify possible opportunities (in the form of novel applications) for optical component adoption across various industries. To complement the interview results, a broad literature review was done to provide additional information regarding technology attributes, potential market volumes, and challenges to optical component adoption. Together, the insights from the qualitative interviews and literature are used to narrow the project's scope to a single, target application (High-Performance Computing); the study of optical component adoption within this market will be the focal point of discussion for the remaining chapters of this thesis. This chapter also presents results from a second set of quantitative surveys administered to industry leaders developing technologies in the target market. From the quantitative surveys, ancillary characteristics that could prove to be useful for evaluating the attractiveness of optical interconnects are identified for further analysis in later chapters of this thesis.

[3] This chapter was previously published in the second Communications Technology Roadmap (CTR II) Report in May 2009, *available* at http://mpc-web.mit.edu/index.php?option=com_content&view =article&id=63&Itemid=78.

2.1 Qualitative Survey and Insights

As described above, qualitative phone interviews were conducted with a broad range of participants working in each of the three markets (automotive, HPC, and consumer handhelds) in order to quickly identify opportunities for the adoption of short reach optical interconnects. The survey, found in Appendix A, is roughly broken into four parts. The goal of the first part of the survey is to identify the current challenges a firm may face within the market. The purpose of the second and third parts of the questionnaire is to ascertain how optical components are currently used and how they are perceived within a firm. The survey also asks interview participants about the pervasiveness of positive or negative perceptions of optical interconnect technologies within a firm. In the final part of the survey, the interviewee is asked to predict how many optical interconnects will be deployed in future applications. Together, answers to the survey's four parts create a preliminary view of each market that will serve as a foundation for further analysis.

Table 1 displays the companies and research organizations that were targeted to participate in the qualitative survey. Representatives from nine companies participated in phone interviews.[4] Based on the insights received from these interviews and from

Table 1: Target Companies

Automotive	Consumer Handhelds	HPC
Ford	Samsung	IBM
Toyota	Motorola	SUN
BMW	Nokia	HP
SMC Europe	Research in Motion	Cray
Relnetyx	LG	Cisco
	Benefon	
	HTC	
	Sony Ericsson	

[4] Of the nine qualitative interviews, two were conducted with a home video and a defense company as part of an original effort to get information from the Defense, Home Video, Automotive, Consumer Handheld, and HPC markets. Since these two markets have been excluded (per footnote 2), these interviews have also been excluded from Appendix A.

available literature, the following sections identify emerging applications for optical interconnects in the automobile, consumer handhelds, and high-performance computing (HPC) industries.

2.1.1 Automotive Market

In a 2006 social trends report on the travel behavior of people in the US, the Pew Research Center observed that people are spending more time in their automobiles each year; the number of trips and vehicle miles traveled per person as well as the average time spent in traffic delays per person have all increased between 1991 and 2003 [7]. Over the past decade, the amount of time that the average person spent waiting in traffic increased 56 percent from 16 hours in 1991 to 25 hours in 2003 [7]. The increase in the time spent in cars has helped promoted an increase in consumer demand for additional "electronic functions that benefit drivers directly: safety, entertainment, information and comfort" [8]. The response of automobile manufacturers to this shift in consumer expectations could provide opportunities to incorporate optical interconnects in automobiles. To gain a preliminary understanding of the automobile market, qualitative interviews with the leading automakers and industry associations were sought. Ultimately, representatives from two companies participated in qualitative interviews. A broad literature search was also executed to further investigate the technology attributes and market volumes of the emerging applications identified in the interviews.

2.1.1.1 Opportunities for Optical Interconnect Adoption

The insights gained from the two completed interviews suggest that most optical interconnect adoption opportunities in the automobile market are driven by a growing number of wired in-vehicle networks installed in automobiles. Available literature regarding wired in-vehicle networks indicates that twelve different network protocols are currently offered in the automobile market [9]. Depending on their intended use within a vehicle, the different types of networks are organized into three broad categories: Body Control, Advanced Driver Safety, and Infotainment [10].[5] In any given vehicle, one or

---

[5] Alternative in-vehicle network categories have been offered. SAE classifies in-vehicle networks based mostly on speed (Class A – low speed, Class B – medium, Class C – high speed). Frost and Sullivan has

more of wired networks "can co-exist to deliver the right combination of data rates, robustness, and cost" as shown in Figure 4 [10].

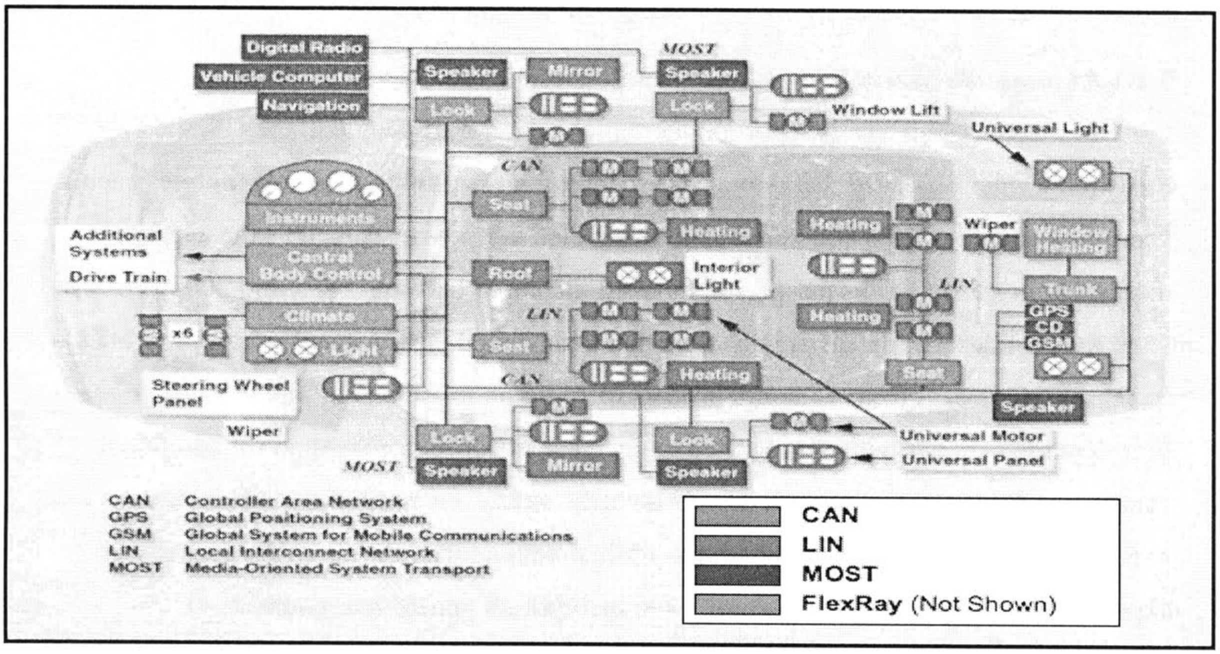

Figure 4: In Car Complementary Networks [10]

Body Control Networks

Most "Body Control" networks installed in passenger vehicles today control safety and comfort applications that span from engine management to anti-lock brakes to power locks and windows. In this market space, the Controller Area Network (CAN) and Local Interconnect Network (LIN) protocols are the most used.

CAN is a two-wire serial bus system that supports high-speed communication among microcontrollers, sensors and actuators through out a vehicle using a multi-master architecture where devices can broadcast messages asynchronously. CAN operates under two internationally recognized standards: ISO 1159-2 standard for low speeds up to 125 Kb/s and ISO 11898 standard for high speeds up to 1 Mb/s [9]. Known for its strong reliability and robustness, CAN is utilized mostly for power train management and safety applications [10].

---

group network protocols into four categories (General Purpose, High-Speed Safety, High-Speed Infotainment, and Low-Speed Smart Sensor) [8].

LIN is a low cost serial communication system that connects intelligent sensors and actuators throughout the vehicle via a single wire. LIN operates up to 20 Kb/s over a maximum cable length of 40 meters. The devices supported by LIN have very specific functions and do not require the higher bandwidth made available by CAN; as a consequence, body control functions such as power locks, power windows and mirrors, is the primary application area for the LIN protocol [9]. These networks are generally implemented as subsystems of a larger CAN network structure in an automobile.

Advanced Driver Safety Networks

Advanced Driver Safety networks in vehicles allow sensors from multiple subsystems (braking, steering, and suspension) to communicate with each other in order to achieve a safer, more comfortable driving experience. In order to operate effectively, these systems must transfer data at high speeds as well as include multiple layers of redundancy. One example of an advanced driver safety network is Flex Ray, which includes an optical bus to control serial communication in a vehicle's adaptive drive chassis.[6] In 2007, BMW became the first to use this technology in standard production vehicles when it began installing Flex Ray in its X5 series.[7]

Infotainment Networks

Infotainment networks allow multiple media devices such as phones, MP3 players, video, and navigation systems to connect and communicate with each other in an automobile. The dominant technology standard that has emerged in this market space is the Media Oriented Systems Transport (MOST) Network. MOST is a networking standard that was designed to connect and integrate various consumer electronics to an automobile seamlessly via a single optical fiber bus [9]. The most recent specifications require possible network bandwidth to reach 25 Mbps over optical fiber. Current roadmaps forecast network bandwidth to increase 150 Mbps over optical fiber in 2013 with the goal of supporting 64 total consumer electronic devices in an automobile [11].

---

[6] In vehicles with the Adaptive Drive Chassis System, each axle has its own independent motor that can adjust the dampers on each individual wheel. When the road surface under one side of the vehicle differs in character from that on the other side, the Flex Ray instantly adjusts the suspension settings of the wheels on one side to practically eliminate any perceptible unevenness in the road.

[7] BMW. http://www.bmw.com/com/en/insights/technology/technology_guide/articles/flex_ray.html

There are two major advantages that are contributing to the success of MOST in this market space. First, MOST networks are designed to scale easily to accommodate new components added to the system. Second, these networks are designed to minimize added costs to total cost of the automobile; its use of plaster optical fibers instead of copper to control communication between a device and the on-board system decreases the overall weight of the vehicle [9].[8]

Of the three types of networks, infotainment systems represent the greatest opportunity for optical component adoption in the automotive market. Although CAN and LIN network components enjoy greater market penetration among passenger vehicles in North America (shown in Figure 5), the low data transfer rates achievable using these protocols are insufficient for future applications. As CAN and LIN continue to reach maturity in the market, the overall number of applications supported by each of these protocols could decrease over time with the continued development of higher-speed, higher-bandwidth protocols (such as Flex-Ray and MOST) that have the multiplexing capabilities to support several high-end safety and infotainment applications on a single bus [9]. However, the deployment of such networks in the light vehicle market has been limited mostly to the luxury vehicle segment of the automotive market over the next five years. While many of the advanced driver safety networks that require optical components and optical fiber have still yet to be introduced, the MOST network remains the only major viable opportunity to incorporate optical interconnects in the automobile industry.

[8] Copper is much heavier than plastic optical fiber.

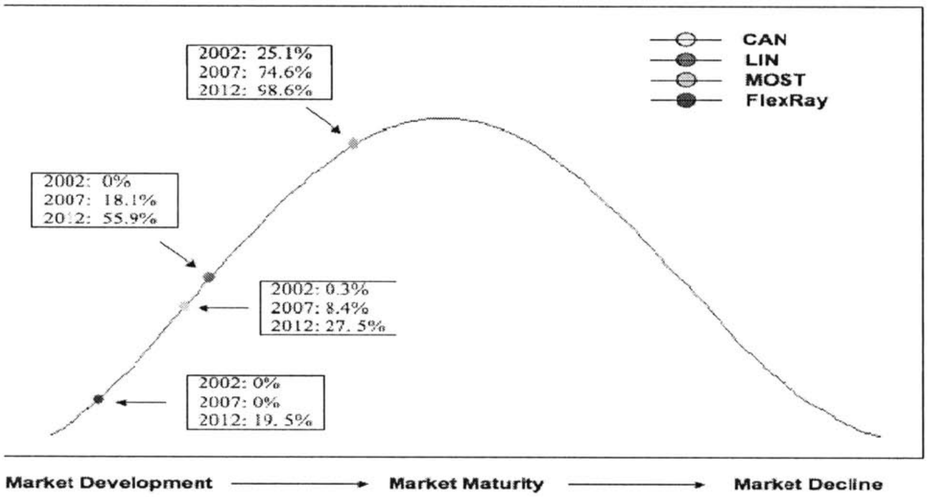

Figure 5: Market Analysis of Networking Technologies in US Automobiles [9, 11]

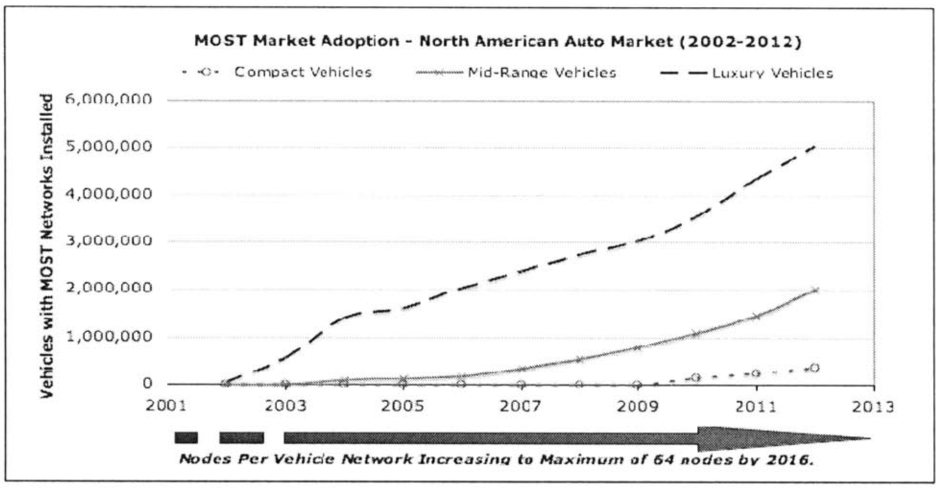

Figure 6: MOST Components Market Adoption, 2000-2010 [9, 11]

Figure 6 shows the acceptance of MOST components in the automotive market. Since the introduction of MOST components in the market in 2000, the number of shipped nodes has steadily grown. In 2004, MOST components began to reach mass production levels as its target customer base started to expand from high-end to mid-range vehicle [11]. Furthermore, extended forecasts from 2005 to 2010 suggest that the number of MOST components installed in vehicles worldwide will continue to rise as infotainment networks become more common in mid-range and possibly low-end, compact vehicles. The possible growth of this market space in the near term represents an opportunity for increased optical component adoption in the automotive industry.

2.1.1.2 Challenges to Optical Interconnect Adoption

As described above, the automobile industry has already begun to utilize optical interconnects in MOST and Flex-Ray networks; however, cost remains the major limiting factor to the increased installation of these wired networks across a broad range of light vehicles in the North American market. The average cost per node for these emerging networks remain much higher than cost levels desired by automakers. At the end of 2005, both MOST and Flex-Ray node costs averaged around six dollars per node while the node costs for more mature network protocols such as LIN averaged around $1 per node [9]. High node costs have thus far limited the adoption of MOST and Flex-Ray networks to the luxury vehicle market segment.

2.1.2 High-Performance Computing

High-Performance Computing (HPC) refers to computer systems that provide close to best currently achievable sustained performance on demanding computational problems. As defined by the *Dictionary of Science and Technology*, an "HPC system can be an extremely powerful, large capacity mainframe computer that is capable of manipulating massive amounts of data in an extremely short time or a [single] computer within a larger set of networked computers" [12]. Originally limited to scientific research only, use of HPC systems have been extended to a broad range of applications including business, defense, and even media. Over the last decade, there has been continuing rapid improvement in the capability of HPC systems; a cursory review of the Top 500 fastest systems in the world show that mean performance has improved by roughly 80 percent annually since 1993 [13]. As computing processor performance for systems in this market space continues to improve, the performance of interconnect network has become a critically important factor that must be considered when evaluating the overall performance of an HPC system. In order to gain some insight regarding the opportunities for optical component adoption within the HPC market, five qualitative interviews were completed.

2.1.2.1 Opportunities for Optical Interconnect Adoption

Servers in a modern data center environment can employ a hierarchy of interconnects that span a wide range of link distances, costs, and bandwidth requirements. Based on the results from the qualitative questionnaires, interview participants identified three areas in the HPC market where optical interconnects could be implemented: Box-to-Box, Board-to-Board, and Chip-to-Chip.

Box-to-Box interconnects are used to connect multiple server backplanes together within a single rack or across several adjacent racks in order to transfer information between each other. Depending on size, there maybe tens to thousands of box-to-box interconnects between individual servers depending on the size of a given system. Link lengths for box-to-box interconnects can span wide range. For shorter connections, link lengths can range from one to ten meters. For much longer connections, link lengths can range from ten meters to as much 100 meters. Optical box-to-box interconnects can be implemented using optical transceivers natively installed in the server backplane or by using Active Cable Assembly (ACA) connectors to perform the electrical-to-optical (or vice versa) conversions externally to the server.[9]

Board-to-Board interconnects are used to connect subsystems on different boards. The normal link length range of these interconnects is 0.3 meters to 1 meter. For larger systems, two to 16 boards are often plugged into a central backplane that contain the interconnect links between the boards. Because of the high cost required to implement optically, board-to-board connections and backplanes are generally electrical. For the highest performance systems, "the pluggable interconnect between a board and the backplane are pushing the limits of incumbent copper interconnect technology; [in some systems], the total bandwidth through the board edge can be bottleneck" [14]. In order to enable the photonic transfer of digital information at the board-level would require the integration of optical components on or at the board's edge.

Several interview participants also expressed interest in the possibility of creating optical interconnects between multiple chips on a board. Current chip-to-chip

[9] Active Cable Assembly (ACA) Connections are optical fiber connectors that are delimited by optical transceivers. These connectors plug directly into servers, network switches, and storage devices that use electrical interconnects; ACA's perform the electrical-to-optical conversion externally to the device, and transmits data using light.

interconnects are usually implemented with electrical pathways between two or more processor chips on a single board. The link lengths for this category of interconnects range from 1 cm to 50 cm [15]; and depending on its size, a given system could employ thousands of chip-to-chip interconnects [14]. Replacing the incumbent electrical technology with optical components would require a solution that tightly integrates the printed wire board and optics; during the interview process, participants largely agreed that the creation of chip-to-chip interconnects would require a new and novel design, but most did not know specifically what that design would entail.

Figure 7 shows the Optoelectronics Industry Association's (OIDA) estimates of production volumes for optical interconnects at the different levels within the HPC market. As one moves closer to the board and realize decreasing link distances, the potential market for short-range optical interconnects within the HPC market increases dramatically.

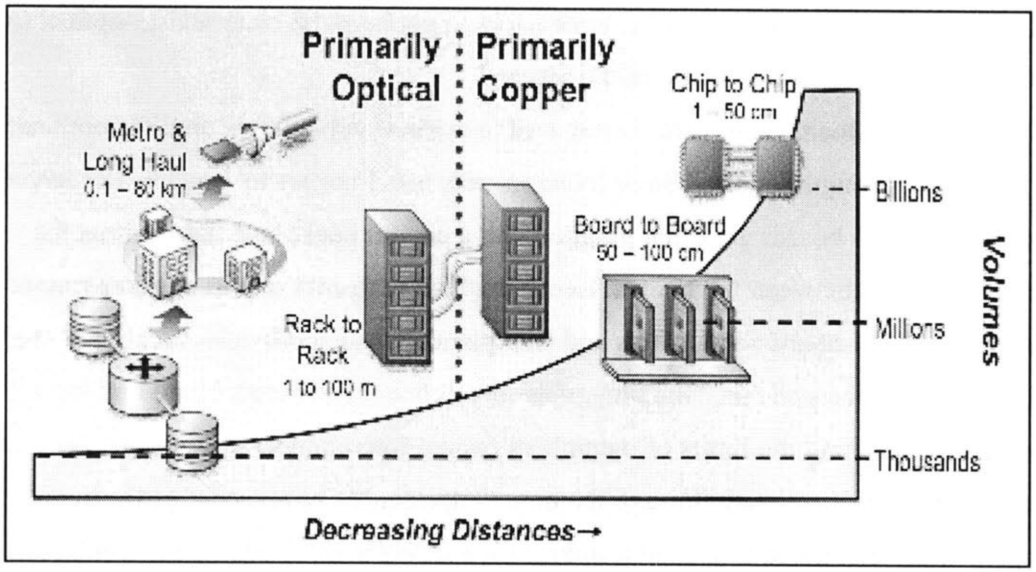

Figure 7: Potential Optical Interconnect Volumes in HPC Market [15]

2.1.2.2 Challenges to Optical Interconnect Adoption

Challenges to optical interconnect adoption exist at all three levels of the HPC market. At the box-to-box level, optical transceivers and optical active cable assemblies are currently mass-produced; however, the relatively high cost premium persists as a major factor limiting increased adoption of these components in servers, networking, and storage equipment employed in data center facilities. The interviews and literature review suggest that biggest technical hurdles to further optical interconnect adoption lie at the horizon of board-to-board and chip-to-chip interconnects. For those studying both board-to-board and chip-to-chip interconnects, the biggest challenge is to design printed circuit boards that integrate optical components as well as create suitable manufacturing processes.[10] While the potential number of electrical pathways that can be replaced by photonic technology increases dramatically at the board and chip levels (shown in Figure 7), the degree of integration between photonic and electronic technologies on printed circuit board (PCB) have only been demonstrated in laboratory research [16, 17]; however, many of those solutions have yet to be economically feasible for mass production.

2.1.3 Consumer Handhelds

The consumer handheld industry encompasses a broad range of devices such as cameras, mp3 players, mobile phones, and personal digital assistants. In order to gain some insight regarding the opportunities for optical component adoption within the consumer handheld market, a different questionnaire (shown in Appendix A.2) was created by Dr. Louay Eldada, co-chair of the IPI TWG in CTR II.[11] Although numerous requests for phone interviews with key market players were submitted, no live interviews were completed. However, Dr. Eldada compiled the previous public statements of several leading scientists working in the industry. Even though no live interviews were done for this market, the compiled set of public statements can still be used to draw certain conclusions about this market.

2.1.3.1 Opportunities for Optical Interconnect Adoption

In the consumer handheld industry, mobile phones have become a particularly interesting market segment as the current trend to consolidate multiple communications

[10] In interview process, some participants noted that the PCB manufacturing is a dirty process that is often unsuitable to the strict requirements for a clean manufacturing space necessary for optical components.
[11] The IPI TWG is the technical working group investigating the Integration, Packaging and Interconnection issues facing optical interconnect technologies.

and web applications into a single device continues to strengthen. Real-life examples of this trend include camera phones and smart phones[12] that can access digital information through wireless Internet and proprietary cellular networks. As a major consequence of that trend, higher bandwidth requirements for high-definition handset displays and higher download speeds to transfer information from a personal computer to the device have become major drivers for optical interconnect adoption in the mobile phone market.

In looking at the table of compiled public statements made by research scientists at the three largest mobile phone manufacturers (refer to appendix A.2) all agreed that future bandwidth requirements for handset displays were a driving motivator to integrate optical interconnects in their devices. Figure 8, below, shows the handset display bandwidths

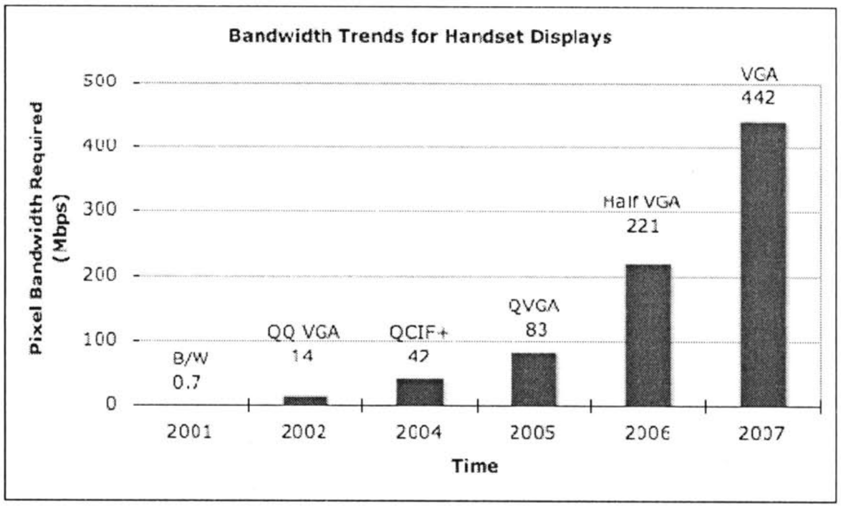

Figure 8: Bandwidth Trends for Mobile Phone Displays [19]

required by high-end phones between 2001 and 2007. Since 2002, pixel bandwidth in phone displays has roughly doubled annually. A similar trend has also been observed in the pixel-resolution of phone cameras; image quality has increased from 110 K pixels in 2000 to three million pixels in 2004 [19]. Over the next five-year period, several analysts and reports predict that the growing bandwidth required by imaging applications will continue; by 2012, high-definition displays could be a common feature in high-end mobile phones.

[12] No industry standard definition exists for smart phones, but many smart phones resemble miniature personal computers as they can include operating systems, wireless internet/email, digital organizers [18].

Though not cited as often, some have also noted that a higher speed connection to external devices as a driving force for utilizing optical interconnects. Converged mobile devices and smart phones are quickly becoming portable gateways that enable constant wireless connectivity to a wealth of online information and communication. Though there are multiple factors pushing this trend, trends in memory technology for consumer devices is a disproportionately significant factor. As one of the top three cost components in camera phones and smart phones, the mobile electronics industry has benefited enormously from "dramatic increases in storage densities of flash memory" [20]. Larger storage capacities allow smart phones and other mobile devices to take advantage of more complex software and media; thus, the files that need to be transferred from device to another are growing and the time needed to transport those files is taking longer. Optical interconnects could be useful in accommodating the larger transfer loads at higher data rates capable with current copper connectors.

A cursory examination of worldwide mobile phone sales is useful to illustrate the shift in technology trends that later create the aforementioned opportunities for the adoption of optical interconnects in the mobile phone market. Figure 9 shows the mobile phones sold between 2004 and 2007; the estimated sales forecasts up to 2011 were derived from the historical data using linear regression. By 2004, sales of camera phones more than doubled that of basic phones. Between 2004 and 2007, this segment of the market continued to see strong growth in popularity; sales of these phones jumped from 442 million units to 742 million units annually. By the end of 2008 and beyond, sales of camera phones are predict to grow but at a more moderate pace of 4 percent CAGR compared to 14 percent CAGR between 2004 and 2007. During the same period,

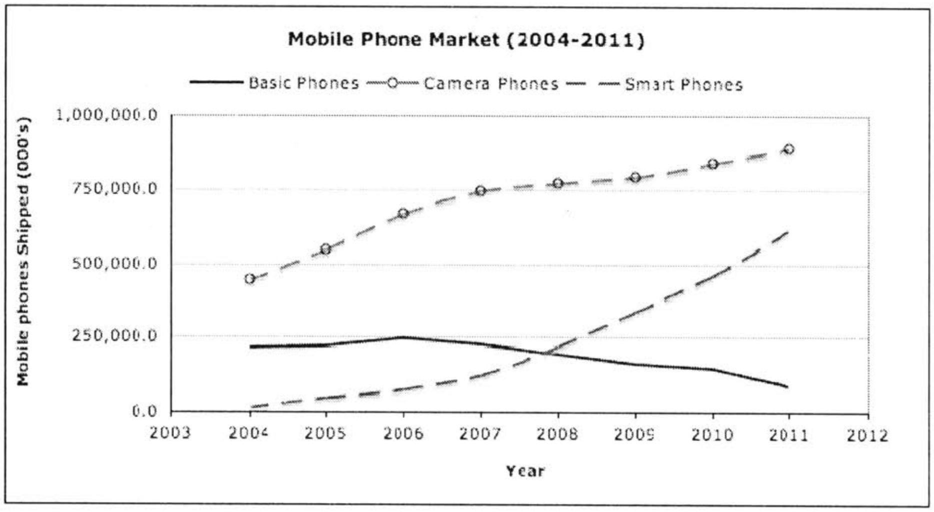

Figure 9: Worldwide Mobile Phone Sales, 2004-2011 [21]

smart phone sales rose rapidly from 16 million in 2004 to over 121 million in 2007. And, between 2008 and 2011, smart phones sales are expected to grow at 21 percent annually. In contrast, demand for basic phones flattened in 2004 and 2005; however, sales in this market segment have begun to fall and will continue to do so in the foreseeable future.

2.1.3.2 Challenges to Optical Interconnect Adoption

Like the HPC market, significant technological barriers limit optical interconnect adoption in mobile phones and consumer handhelds; the successful integration of optical interconnects in future mobile devices will require substantial architectural and manufacturing challenges to be solved. However, incorporating optical interconnects in mobile electronics has its own unique set of challenges. First, manufacturers will have to find a way to incorporate optical interconnects without increasing the overall footprint of the device. Since the introduction of the first cell phone to the market, manufacturers have consistently found ways to decrease the physical dimensions; between 1983 and 2006, major mobile phone producers have reduced the average size and thickness from 50 mm to currently 10 mm [19]. Second, efficient power consumption is critical to component adoption. Currently available mobile phones are expected to achieve a certain amount of talk time and idle time using a single charged battery. Phone manufacturers must design an optical interconnect that, when used, will not significantly reduce the level of battery performance currently expected by consumers. To achieve this goal, principal scientists at the Nokia Research Center estimate that power consumption by future interconnects would need to undergo "an order of magnitude reduction" from current levels to approach 1mW/Gbps for each connector [22].

Furthermore, mobile phone manufacturers will also have to address the above technical concerns under tight unit cost constraints caused in part by intense competition for market share among producers of mobile phones and other handheld electronics. In the compiled set of public statements, all subjects expressed the necessity for future optical interconnect solutions to approach a cost comparable to current electrical interconnect solutions. Displaying a breakdown of the 14 most expensive components in two different phone models, Figure 10 gives a rough estimate for the target costs for potential optical interconnect solutions. In each case, the printed circuit board and connectors (the two more important components in an integrated opto-electronic solution[13]) comprised approximately six percent of the total unit production cost. For future optical interconnects, the research scientists estimate the unit cost to approach $1 per Gbps in three to five years for internal serial connections reaching 5 Gbps in their public statements. Assuming the average smart phone or camera phone will require 5 Gbps in bandwidth in the next three years, mobile phone manufacturers will need to decrease the cost of an integrated solution by at least one dollar [23].

[13] In an April 2007 presentation, *Opportunities in Optics for Mobile Devices*, Leo Karkkainen noted: "Main possibility of taking into use higher level of optical interconnections in mobile multimedia computers is seen with the development of direct optical interfaces integrated in CMOS chips. These would provide high bit-rate serial connections, mainly for imaging applications." [22]

Nokia 6151 | 12/1/2006	Data		
Component Type	% of Total Component Cost $USD	Sum of Total Component Cost $USD	Sum of Qty
Memory	13.45%	$11.80	2
Analog	11.77%	$10.33	12
Display	11.08%	$9.72	1
Camera	9.00%	$7.90	1
Plastics & Elastomers	8.72%	$7.65	32
Logic	8.66%	$7.60	3
RF	6.10%	$5.35	1
PCB	4.79%	$4.20	3
Rechargeable	4.26%	$3.74	3
Accessories	3.93%	$3.45	2
Integrated Active/Passive	2.55%	$2.24	14
Connector	2.40%	$2.11	26
Oscillator	1.77%	$1.55	2
Literature & Packaging	1.60%	$1.40	6

Motorola E770V | 12/1/2006	Data		
Component Type	% of Total Component Cost $USD	Sum of Total Component Cost $USD	Sum of Qty
Logic	20.00%	$21.89	14
Memory	15.99%	$17.50	3
Analog	14.72%	$16.11	11
Display	9.57%	$10.48	1
Plastics & Elastomers	5.65%	$6.18	28
Camera	5.63%	$6.17	2
Rechargeable	4.55%	$4.99	3
PCB	4.18%	$4.58	4
Accessories	3.33%	$3.65	4
LED	1.89%	$2.07	13
Acoustics	1.40%	$1.53	3
Transistor	1.31%	$1.44	9
Filter	1.31%	$1.43	4
Connector	1.29%	$1.42	9

Nokia 6151
Total Cost: $87.72
Total Comps: 491

Motorola E770V
Total Cost: $109.47
Total Comps: 595

Figure 10: Camera Phone Teardown Report Excerpt [23]

2.2 Targeted Application

In the preceding section, the qualitative interviews with key players in the automotive, consumer handheld, and HPC industries provided valuable insights regarding the opportunities and barriers to optical component adoption in each market. Table 2 summarizes the barriers, potential market volumes, and estimated time to adoption of optical components for each market. Considering these three factors, HPC Box-to-Box

Table 2: Primary Market Insights

Industry	Barriers	Potential Market Volumes	Target Time to Adoption
Automobiles	Cost	Tens of Millions	Now
Mobile Phones	Cost, Power, Size	Hundreds of Millions	5 Years
HPC Box-to-Box	Cost	Millions	0-5 Years
HPC Board-to-Board	Architectural	Tens of Millions	5+ Years
HPC Chip-to-Chip	Architectural	Billions	10+ Years

and the automobile markets are perhaps the most interesting areas to study. Although they may have smaller potential volumes relative to other market segments, the barriers to adoption do not seem to be as challenging as those for mobile phones, HPC boards and computer chips. In each of those industries, the adoption of optical interconnects would require a level of integration between electronics and photonics that is yet economically feasible to produce; furthermore, both the literature and interview responses collected thus far suggest that the technology and manufacturing processes necessary to create the integrated opto-electronic PCB will not be available in the mid-term (less than 5 years).

Of the two remaining applications, HPC box-to-box is the more interesting market segment for further study with respect to the project's four central questions outlined in the previous chapter. Unlike the automotive market, an extended study of HPC box-to-box interconnects in this thesis has the potential to uncover observations that could be beneficial for future research in the other HPC and consumer handheld market segments. This potential for the transfer of knowledge across markets arise from similarities in two key areas that HPC box-to-box interconnects share with the other market segments that the automotive applications do not. First, automotive networks use plastic optical fiber to transfer information rather while other emerging solutions for HPC and consumer electronics devices are expected to utilize glass optical fiber. Second, the data rates required for current and future HPC box-to-box interconnects are comparable bandwidths required by the board-to-board, chip-to-chip, and consumer handheld applications while data rates available with wired in-vehicle networks are orders of magnitude lower. Current data rates for automobile infotainment networks transmit data at 25 Mbps; and by 2016, bandwidth for these networks is expected to reach 150 Mbps.[14]

To reach higher data rates in vehicles, the automotive community is looking to wireless or satellite networks [24]. Using wireless internet and satellite infrastructure to better integrate the automobile with the outside world has become a particularly attractive option in this market segment; by incorporating these technologies, manufacturers can open automobiles to a broad array of new applications as shown in Figure 11.

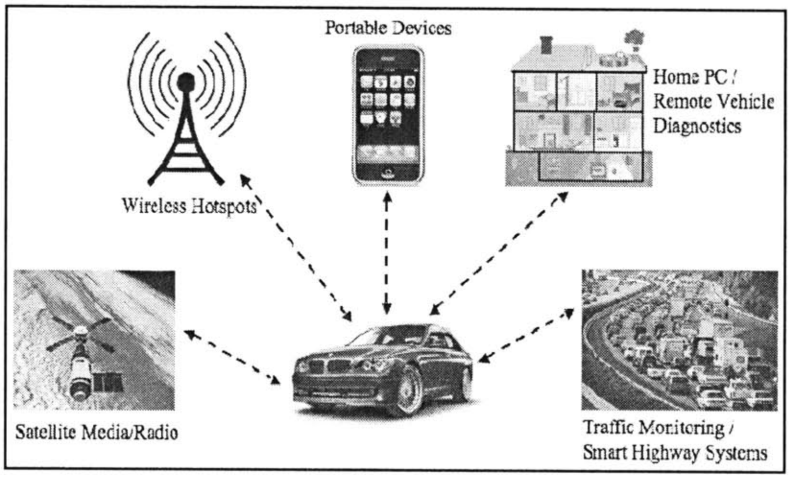

Figure 11: Wireless Infrastructure for Automobiles [24]

While satellite media and remote vehicle diagnostics systems have been introduced in some automobile models, researchers are exploring the technical and economic feasibility of implementing wireless peer-to-peer networks in vehicles and advanced highway monitoring systems [24, 25].

[14] MOST reports net data rates and board data rates. The board data rates account for information overhead which is usually twice the amount of the capable net data rate.

2.3 Quantitative Survey and HPC Market Insights

After the qualitative interviews were completed and a target application selected for further study, a quantitative questionnaire was used to identify general technology trends of HPC box-to-box interconnects. The quantitative questionnaire is also used to identify attributes that engineers deem important when evaluating the merits and disadvantages of various interconnect technologies. Throughout this exercise, multiple versions (both versions of the questionnaire shown in Appendix B) of the survey were administered to key players in the industry.

2.3.1 General Trends for HPC Box-to-Box Interconnects

While both quantitative questionnaires attempt to measure the trends in traditional metrics (bandwidth demand, cost, architectural challenges) that will drive optical component adoption within the HPC box-to-box, board-to-board, and chip-to-chip market segments, the two questionnaires differ slightly; the first version seeks to ascertain detailed specifics regarding the future design of HPC interconnects while the second version seeks to identify general trends only. Ultimately, fifteen participants provided comments for one or both of the quantitative surveys. Of those, six participants provided specific insights regarding the technology trajectory, particularly system link density, bandwidth, and cost, for HPC box-to-box interconnects. Those insights have been summarized in the following four figures; the raw data for those figures are contained in Appendix B.

Figure 12 displays how phone interviewees predict link density per system will evolve over the next ten years. In this survey effort, the term, "link," was defined as a single optical fiber with a dedicated transmitter and receiver attached to it. As seen in the figure, most respondents believed that the number of links per system could grow gradually from 100 links per system in 2007 to a projected average of 400 links per system by 2013. Beyond 2013, however, the interview participants are almost equally divided in their predictions for future growth; while half of the interview participants were conservative in their extended forecasts, the rest of the participants believe that the number of links per system will increase dramatically to 10,000 links per system by 2016.

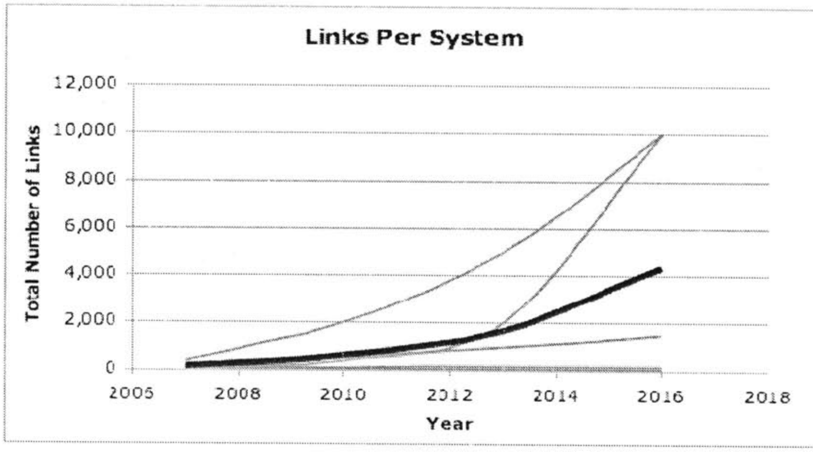

Figure 12: Links per HPC System, 2007-2016

Figure 13 shows the participants' predictions of bandwidth increases per link over the next ten years. Again, a high degree of variance between the responses was observed. Although most respondents agree that current data rates available is approximately 10 Gbps per link; but, looking forward three years into the future, the responses begin to

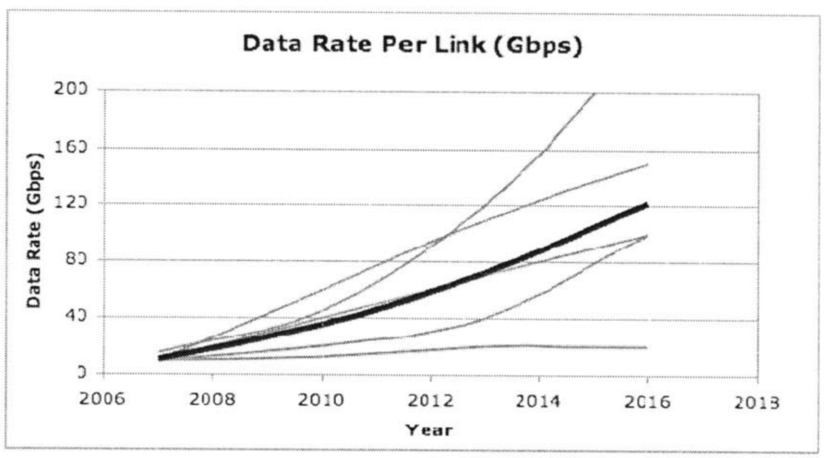

Figure 13: Bandwidth per HPC Link, 2007-2016

diverge as some participants believe link data rates will be 12.5 Gbps while others believe data rates will approach 60 Gbps per link. At 2016, the range of responses widens tremendously – 20 Gbps to 240 Gbps per link. However, the average trend of the response set shows that bandwidth per link could grow fairly linearly over the next ten years to approach 120 Gbps per link by 2016. Figure 14 displays the predicted aggregate bandwidth per HPC system, which is the result of multiplying the response sets of Figure 12 with that of Figure 13.

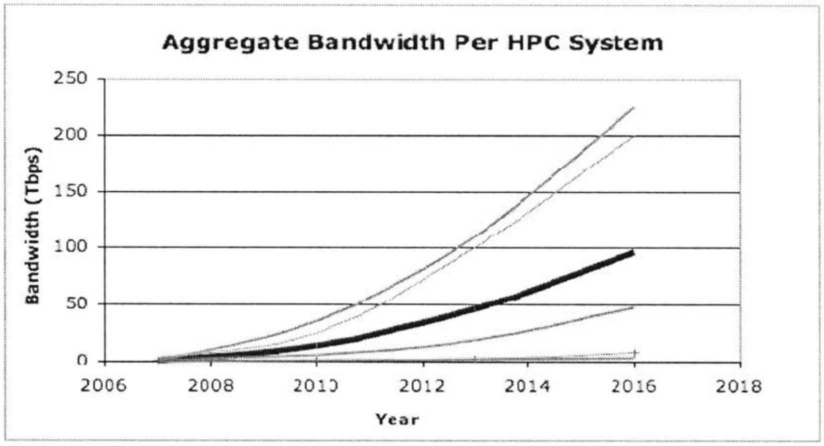

Figure 14: Aggregate Bandwidth per HPC System, 2007-2016

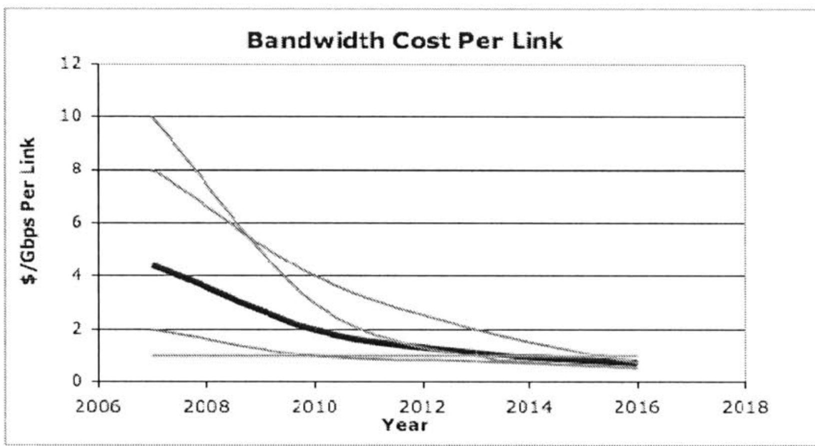

Figure 15: Bandwidth Costs for Adoption, 2007-2016

Figure 15 shows how the participants perceive how cost for HPC box-to-box interconnects will evolve over time. During the quantitative interviews, participants used bandwidth cost ($/Gbps) per link as a proxy to gauge how the market's expectations for component cost will change. In answering this question, most participants believe that bandwidth cost will generally decrease by a factor of two or three every three years regardless of current perception of cost. The average of responses shows that bandwidth cost will decrease from slightly over four dollars per Gbps to approach $0.75 per Gbps.

Despite the initial large divergence in their answers, there appears to be a strong consensus among the participants that the component prices for optical interconnects need to drop to one dollar per Gbps or less in order to encourage broad adoption.

While the above survey results are helpful to describe how basic attributes of the box-to-box interconnect (system size, bandwidth, and cost) will change over time, the large spread in interviewee responses also indicate a need for further survey efforts to produce a larger response set with more consistent results. In each of the above four figures, the wide divergence can be attributed to the inconsistencies between how each interviewee interprets key terms and the assumptions based on their unique professional experiences used in answering the question. Furthermore, as demonstrated by the great amount of divergence in response set, the uncertainty regarding the future technology landscape only amplifies the effect that the aforementioned inconsistencies can have on outlining general trends for HPC box-to-box interconnects.

2.3.2 Identifying Key Ancillary Attributes

Given the general trends outlined above, the interview participants were asked: "How important are the following metrics for trade-offs made during interconnect technology selection?" (Appendix B.1, Question 9). Despite the large divergence displayed in the interview responses regarding technology trends in the industry, the participants largely were largely in agreement when asked to identify metrics considered useful in their evaluation of HPC interconnects. Table 3 shows the average of responses

Table 3: Important Metrics for Component Adoption in HPC Markets

Metric	Possibly the Most Important	Very Important	Important
Faceplate Density (Gb/s/inch)		2	
Board Density (Gb/s/in^2)		2	
Energy Density (Gb/s/W)	1		
Cost Effectiveness (Gb/s/$)	1		

for this question. While nearly all respondents ranked cost effectiveness as possibly the most important metric, most viewed that energy density as a very important ancillary metrics to consider when evaluating the trade-offs for an optical interconnect technology.

Energy is a major concern that is also echoed at higher levels in the HPC market as data center managers and professionals have expressed similar concerns at larger, system-wide levels. In a 2007 market study conducted by Gartner, power consumption was identified to be a major concern from by more 100 data center professionals interviewed for the study [26]. From those interviews, the study concluded that energy has become a major area of concern for businesses. Highlighted in Table 4, the growing pressures on power distribution systems and cooling infrastructure top the concerns that

Table 4: Data Center Concerns [26]

What is the greatest facility problem with your current data center?	2007	2006
Insufficient Cooling	27%	35%
Insufficient Raised Floor Space	10%	12%
Insufficient Power	47%	33%
Poor Location	3%	10%
Excessive Facility Cost	5%	6%
None of the Above	8%	4%
Total Number of Responses	119	125

data center managers have when assessing the capacity of their facilities. Furthermore, interview participants also expressed concern that their businesses were being put at risk as the tools and processes needed to manage growing infrastructure demands have not evolved even as data centers have grown in complexity. For the purpose of this thesis, the results from the quantitative surveys and the Aperture Research Institute study have identified ancillary attributes that could perhaps illuminate key insights regarding the attractiveness of optical interconnect adoption in HPC/data center environments.

2.4 Key Chapter Findings

Qualitative survey results and literature suggests that HPC box-to-box interconnects is an emerging market with the greatest potential for the future adoption of optical components. Automotive infotainment networks is another area that also shows great potential for increased adoption of optical interconnects; but unlike HPC box-to-box and the other applications reviewed in this chapter, automotive infotainment networks use plastic optical fiber to transmit digital information at much lower data rates. For this thesis research, HPC box-to-box interconnects was chosen for further study.

Results from additional quantitative interviews conducted with professionals specializing in this market segment indicate that energy is a primary concern that is considered when evaluating the merits and disadvantages of new interconnect technologies. With the project scope sufficiently narrowed, the following chapters will explore three questions:

1. What market and technology influences contribute to the relative importance of power consumption as a major concern in HPC data center installations?
2. What cost benefit could the implementation of short reach optical interconnects have in HPC data center installations? Specifically, would potential per-server savings in data center cost components result in a positive return on investment?
3. Based on the results of this research, what policy recommendations be made to promote the adoption of short-reach optical interconnects in HPC data center installations?

Chapter 3 – Power Concerns for Data Centers

In order to model the potential cost impact of optical component adoption in a HPC data center, it is important to gain a better understanding of what constitutes a data center as well as the emerging power concerns that face IT managers and professionals. In this chapter, the following sections will review basic data center characteristics, the current power trends occurring in data centers as well as identify primary factors driving those trends. This chapter will conclude with a description of how replacing the incumbent electrical interconnects with optical ones could affect data center power usage.

3.1 What are Data Centers?

Data centers are facilities used to house computer systems that either handle the primary business and operational data for an organization or provide off-site backup services. These facilities range in size; they can be housed in small rooms within larger, conventional buildings or they can occupy large industrial spaces. Regardless of the size, all data centers house rows of IT equipment racks that contain compute servers, storage devices, and network switches. Other basic components of a data center include power distribution and cooling subsystems. The IT, power distribution, and cooling subsystems together comprise the vast majority of capital and operating expenditures as well as power load, and space requirements.

3.1.1 IT Equipment: Servers, Storage, and Network Switches

Again, the primary IT equipment used in data centers are servers, network switches, and tertiary storage. Most IT equipment is stacked on top of one another in racks that are 19 inches wide and 73.5 inches (or 45 U) tall [27].[15]. Servers (computers that can run certain applications under heavy workloads for extended periods of time without direct human supervision) are the most abundant equipment and can occupy 75 percent or more of the total rack space. To make servers more suitable for an unattended

[15] Rack height is often defined in terms of U's. A "U" is approximately 1.75 inches tall; therefore, 45U rack is 78 inches tall. This terminology is also applied to size of equipment that takes up space. For example, a 1U server occupies 1U of rack space [27].

industrial environment, servers lack many of the user-friendly features utilized by the average personal computer specialized features not found in the average personal computer; instead, servers are built with one or more processors, redundant hard drives, and redundant power supplies to ensure continual operation year-round.

As shown in Figure 16, server systems are often classified into three broad categories (volume, mid-range, and high-end) depending on the number of processors, physical footprint, and cost. Volume servers typically consist of one to two microprocessors chips mounted on a single board that, along with power, cooling and

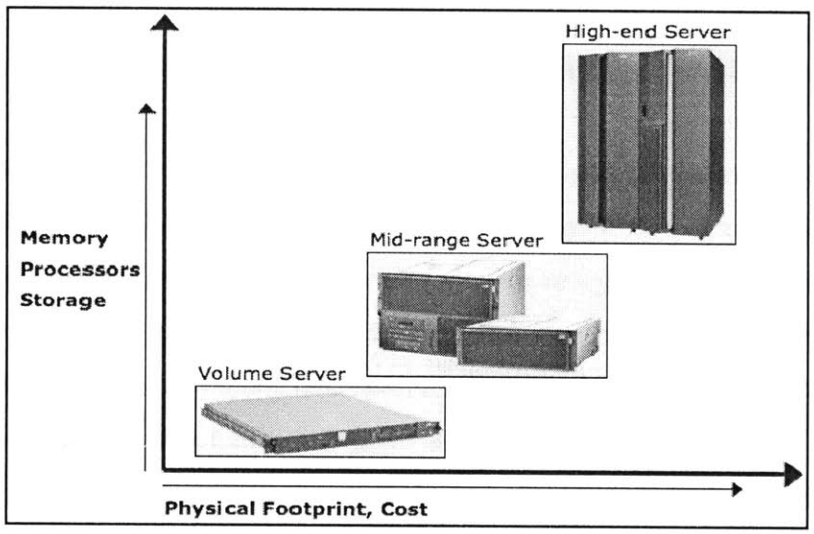

Figure 16: Classification of Server Systems [14]

cabling infrastructure, is contained in one or two U's of a rack-mount shelf [14]. Generally, these servers are designed with cost-optimized components in order to achieve good price-performance. Based on a small sample of volume servers currently available on the market, these servers have an average price range of $2000-3000 per server[16]. Mid-range servers contain one to four boards that each has its own memory, I/O ports, and up to four microprocessors [14]. Mid-range servers usually take more than 2U's worth of rack space, and have a broad price range between $25,000 and $500,000. High-end servers, generally recognized as large mainframe computers, support a greater

[16] A brief survey of the prices for HP Proliant DL 360, Dell Power Edge 1950 and 2950 series of servers was taken to determine average price range to demonstrate price range for volume servers currently available on the market.

number of microprocessors, storage, and extensive internal I/O networks. These servers typically contain four to 16 boards that each has four tightly integrated microprocessors mounted on them [14]. Priced above $500,000 per individual server system, high-end servers can be networked in an array of other high-end machines or installed alone.

Occupying the remaining rack space are network switches and storage equipment. These devices enable servers to complete their assigned tasks; network switches allow servers connected to a local area network to communicate with other servers in the network while tertiary storage allows important, non-urgent data to be saved without taking up finite space on the server's internal hard drive.

3.1.2 Power Delivery Subsystem

For most data centers, the power delivery subsystem utilizes three primary components: a backup generator, an uninterruptible power supply (UPS), and power distribution units (PDU). The size of the power delivery subsystem and the amount of redundancy depends on the data center's size as well as the minimum amount of downtime allowed.

As shown in Figure 17, the critical power load required by a data center's servers, network switches, and storage devices passes through the switchgear, UPS and PDUs.

Before reaching the IT equipment rack, electricity from the switchgear is first supplied to an UPS unit, which is a battery backup that allows the IT equipment and cooling subsystems to continue operating for a finite time period during unexpected power disruptions. Having a battery backup as part of the power delivery subsystem prevents the possible loss of revenue from disruptions in normal business operations or data loss. In order to keep the batteries fully charged at all times, electricity following into the UPS is converted from alternating current (AC) to direct current (DC). Prior to leaving the UPS, power is reconverted back to AC. Afterwards, the PDU takes the outgoing power from the UPS and sends it directly to the IT equipment in the racks. Furthermore, most servers are shipped with a power supply unit that further conditions incoming electricity

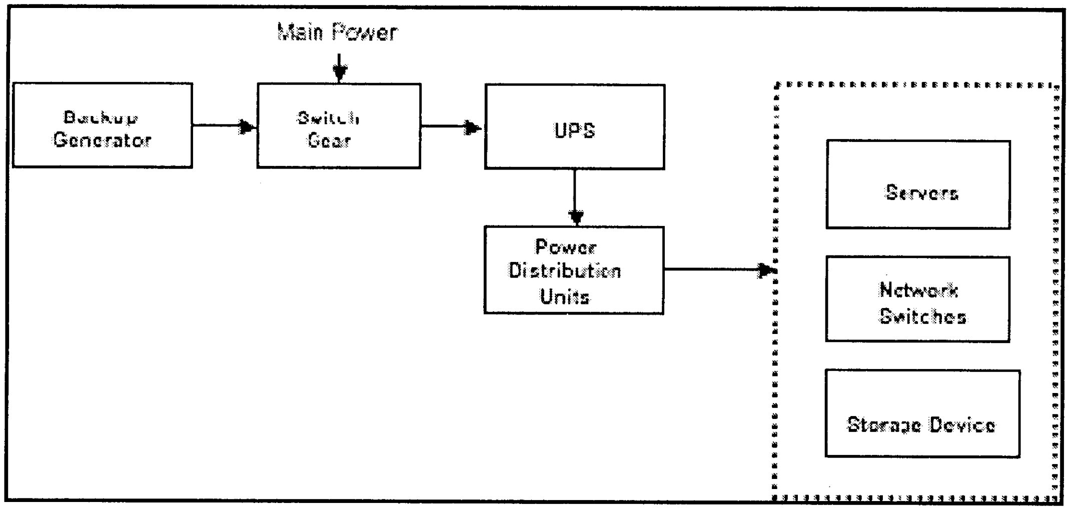

Figure 17: Power Delivery Subsystems Components [28]

by converting the incoming AC to a low-voltage DC power to charge the batteries. The low-voltage DC power is used by a server's internal components, such as the central processing unit, memory, disk drives, chipset, and fans.

3.1.3 Cooling Subsystem

The continuous operation of IT equipment and power delivery systems can generate a significant amount of heat that must be removed from the data center in order for the equipment to operate properly. For smaller facilities and server closets, adequate cooling can be achieved by using fans alone; but for larger ones, cooling is often provided by computer room air conditioning (CRAC) units and an air handling unit (AHU), both of which are situated on the electrically active space within the data center floor. In the case of larger data centers, IT equipment racks are organized to form "hot aisles" and "cold aisles". In "hot aisles" the back of racks are placed facing each other, which allows dissipated heat from servers to vent, while in "cold aisles" the front of racks face each other (shown in Figure 18). This layout helps the AHU (which contains fans, filters, and cooling coils) to draw warm air from the hot aisles towards the ceiling and to the CRAC unit. As warm air rises, it is sucked into the top of the CRAC unit where it is

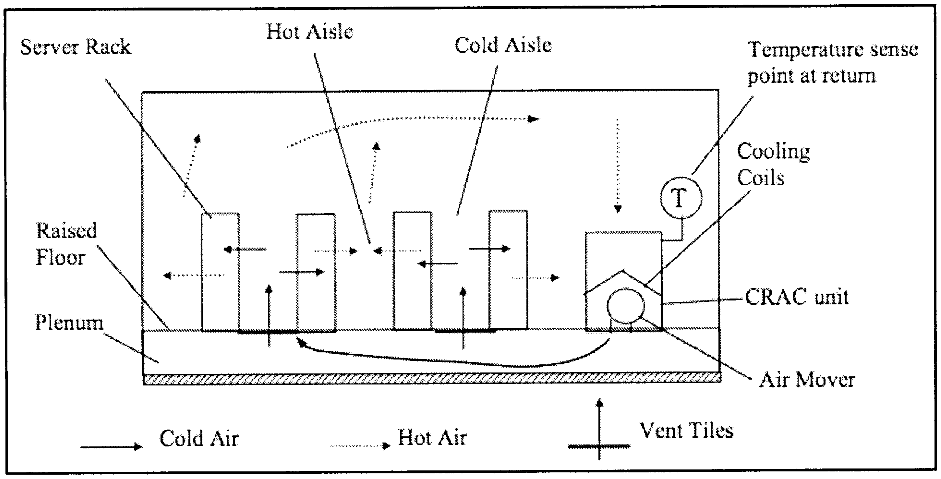

Figure 18: Typical Raised-Floor Configuration [30]

conditioned and cooled as it passes across coils containing chilled water pumped from a chiller located outside of the data center room. The conditioned air is then supplied to the IT equipment (usually through air ducts under the raised floor). As the cooled air is pushed up through the raised floor's perforated tiles in the cold aisles, server fans pull the cooler air across server components and push the warm air back into the hot aisle.

For large facilities, cooling subsystems can be relatively power-intensive and often difficult to calculate accurately without using computational fluid dynamic models. To simplify that problem, some industry professionals have estimate that the total amount of power consumed by cooling subsystems is a fraction of the total power consumed by the IT equipment. Other professionals choose to use a more indirect method by first determining the total amount of heat that the IT equipment produces which would allow managers to then determine the appropriate size of the cooling infrastructure. The rated power of each piece of cooling equipment could then be used to get the total amount of power consumed by cooling subsystems [31].

3.2 Rising Electricity Demands

Recently, energy consumption has become a source of major concern for both IT managers and policy makers. In 2006, Congress voted overwhelmingly for the passage of bill HR 5646, which mandated the Environmental Protection Agency (EPA) to "report to Congress on the growth and energy consumption of computer data centers by the federal government and private enterprise." In its final report, the EPA estimated that electricity consumption by the nation's high-performance computing data centers has more than doubled since 2000 to constitute more than 1.5 percent of the nation's total electricity consumption [28]. Not counting electricity consumed by the additional equipment in the power delivery and cooling subsystems, servers alone accounted for 0.61 percent of total US electricity consumption in 2005.

Figure 19 shows US and worldwide data center electricity consumption relative to total power consumed for all applications. With historical data for both total electricity

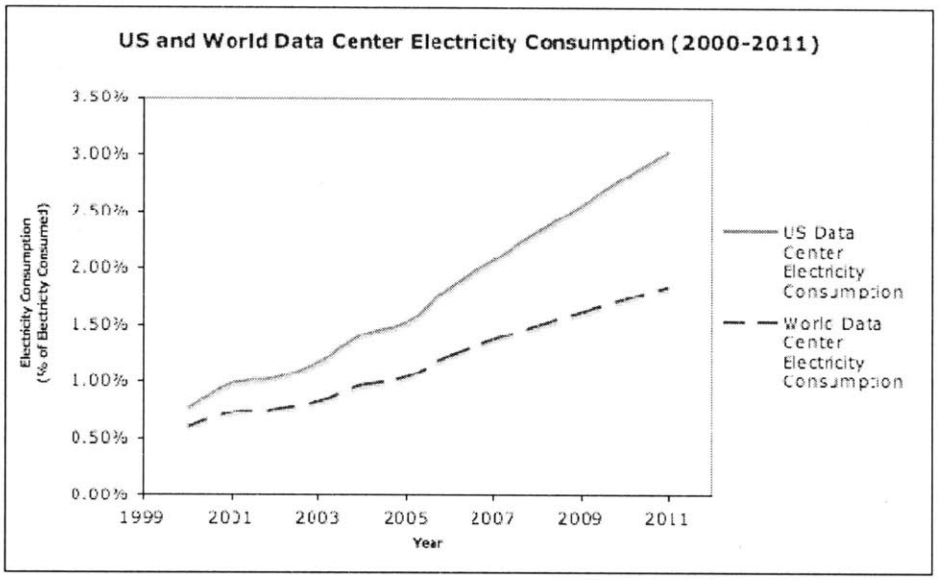

Figure 19: US Data Center Electricity Consumption, 2000-2011 [32, 33, 34]

consumed and power consumed by the installed server base between 2000 and 2005, linear regression analysis was used to predict future power consumption to 2011. The results show that the percentage of electricity consumed by the US and worldwide data centers could double by 2011 to reach 3 percent and 1.85 percent respectively. This trend appears to be driven not only by the 16 percent average annual growth in power consumption but also by the relatively stagnant growth in overall electricity production.

As displayed in Figure 19, worldwide and US electricity use could increase by 2.4 and 0.8 percent between 2006 and 2001. Some analysts studying energy use in this market segment have begun to speculate if the current trend is even sustainable in the future. In its 2006 survey of IT managers and executives, the AFCOM Data Center Institute predicted that within the next five years data center operations at more than ninety percent of all companies could be interrupted due to limited power availability and power failures [35]. The growing demand for electricity by data centers in the US and abroad is driven by two primary forces: a rapidly growing worldwide demand for online information and changes in data center architecture.

3.2.1 Digital Information Demand

The demand for digital information is increasing at a rapid rate. Since 1995, internet usage has ballooned to more than 1.3 billion users worldwide; and at the current rate of 10 percent rise in new users annually, the number of internet users around the world could reach two billion by 2011 as estimated by the Computer Industry Almanac. Much of the recent growth in internet users can be linked to a world economy that is becoming increasingly global and interdependent. Rising internet use is more pronounced in those countries that have observed especially rapid economic growth. China and India, for example, are two of the fastest growing economies in the world[17]; in each of these countries, the number of internet users has increased annually at 28 and 62 percent respectively between 2002 and 2006. As shown in Figure 20, annual increases in new Chinese and Indian internet users have, at times, outpaced growth in the US internet population which seems to have flattened in recent years. This surge in internet use comes at a time when demand for high-bandwidth applications, such as music downloads, video-on-demand, and voice-over-IP communications, has spiked in over the past few years [28].

[17] Dominic Wilson and Roopa Purushothaman. China, India, Brazil and Russia have GDPs that are growing at a faster rate than the G6. Dreaming with BRICs: The Path to 2050. 2003. [37]

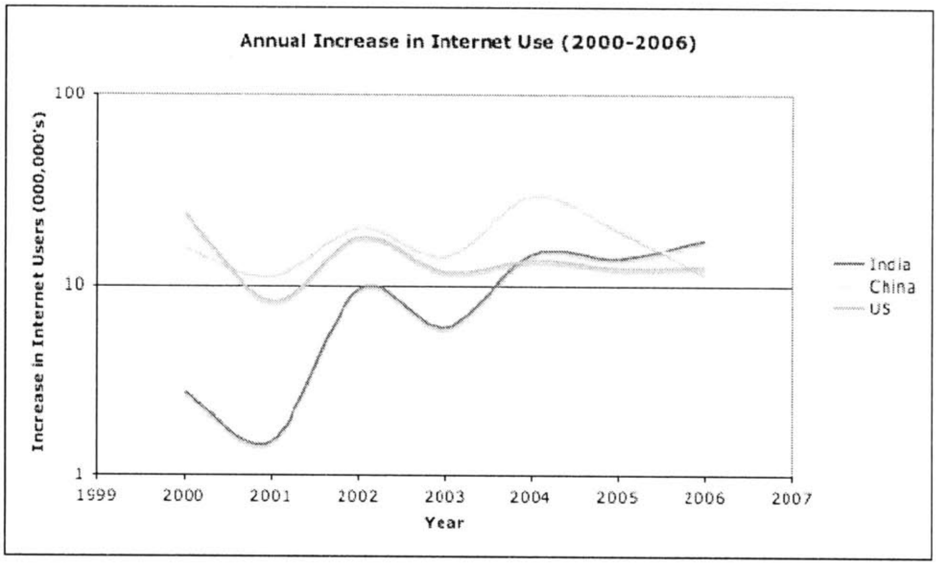

Figure 20: Annual Increase in Internet Use – US, China, India, 2000-2006 [36, 37]

The migration of businesses from paper-based to digital applications is also driving increased electricity consumption. Over the past ten years, information technology has become an important factor for the growth and operation of businesses. This shift from paper to electronic data has been observed in multiple market segments such as health care, insurance, and inventory transportation [28]. Finance and banking institutions, especially, have made high-performance computing data centers a critical component in their operations and competitiveness. For these institutions, HPC data centers allow analysts to execute sophisticated modeling of financial instruments, analyze complex risk portfolios, and detect currency shifts [38].

These trends require organizations to establish highly reliable data centers with enough compute capacity to handle current and future maximum processing loads. Furthermore, as organizations become more reliant on digital information, the need to collect and store that information becomes even more important. For some industries, disaster recovery measures beyond primary storage capabilities in the form of duplicate data and redundant off-site data center facilities are crucial to the continued operation in times of unexpected contingencies.

Together, the increase in demand for online content and internet use coupled with the migration to electronic data has lead to continued strong growth in server shipments over time and an inflation of the installed server base. Between 1999 and 2006, worldwide server shipments experienced 13 percent compound annual growth.

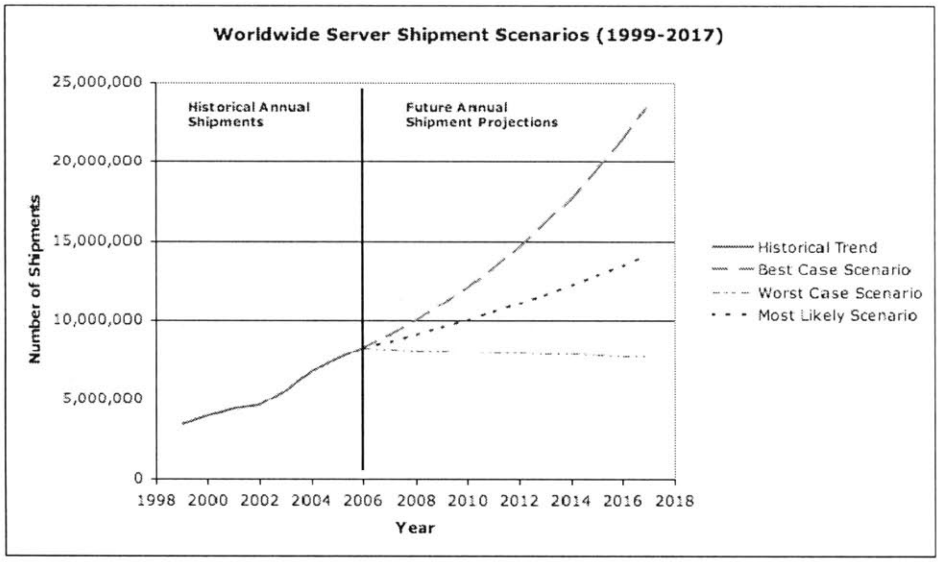

Figure 21: Projected Worldwide Server Shipments, 1999 – 2017 [34, 39]

In its quarterly server forecasts, Gartner DataQuest (a technology consulting firm) predicts total worldwide server shipments will continue to grow at a more moderate pace of five to ten percent in order to accommodate the continued rise in demand for online information and electronic data as displayed in Figure 21.[18] The substantial year-over-year increase in the number of servers shipped worldwide has contributed significantly to the large increase in electricity consumption by data centers, from 78.86 billion kWh in 2000 to 163.16 billion kWh in 2005. In the most-likely case of five percent annual growth in server shipments over the next five years, worldwide power consumption by datacenters could surpass 335 billion kWh in 2011.

3.2.2 Data Center Architectural Changes

The rising power demands can also be partly traced to the various architectural changes that have occurred in data centers over the past twenty years. During this period,

[18] In its quarterly forecasts, Gartner DataQuest outlines three possible scenarios for server shipments: Best Case (10 percent annual growth), Worst Case (-0.6 percent annual growth) and Most-Likely Case (5 percent annual growth) [39].

high-performance computing topologies have evolved from integrated mainframe systems to cluster-based systems. A cluster is a group of individual computers that, when coupled together in a network, functions as a single computer. Clusters are usually deployed to improve performance over that provided by a single computer when executing parallel application. Figure 22 shows the structure of a compute cluster using the generic IT equipment (described above in previous section). HPC clusters

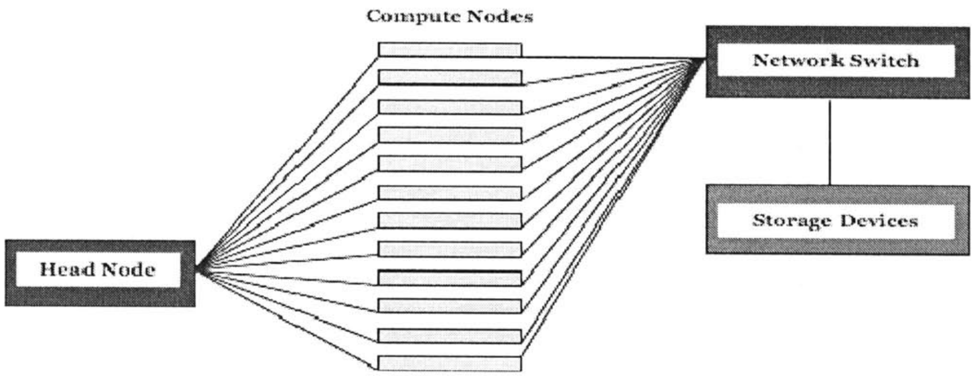

Figure 22: Simple Cluster Topology [40]

comprise servers, network switches, and storage devices in which servers are classified as either head nodes or compute nodes. Under this arrangement, a head node (or several head nodes) breaks up a large task into many smaller tasks and gives each subtask to a compute node. Each one of those compute nodes performs the required subtask and either returns the data to the head node or sends the resulting data to storage devices. Unlike larger integrated systems, clusters can be built with readily available off-the-shelf components such as volume servers to achieve superior cost per performance [40].

The growing popularity of cluster computing is reflected in the statistics for the Top 500 Supercomputers, a group of the world's fastest computers and a lead user for the wider high-performance computing market. As Figure 23 shows, the number of supercomputers in the Top 500 using cluster architecture has ballooned from seven in 1999 to well over 300 in 2007.

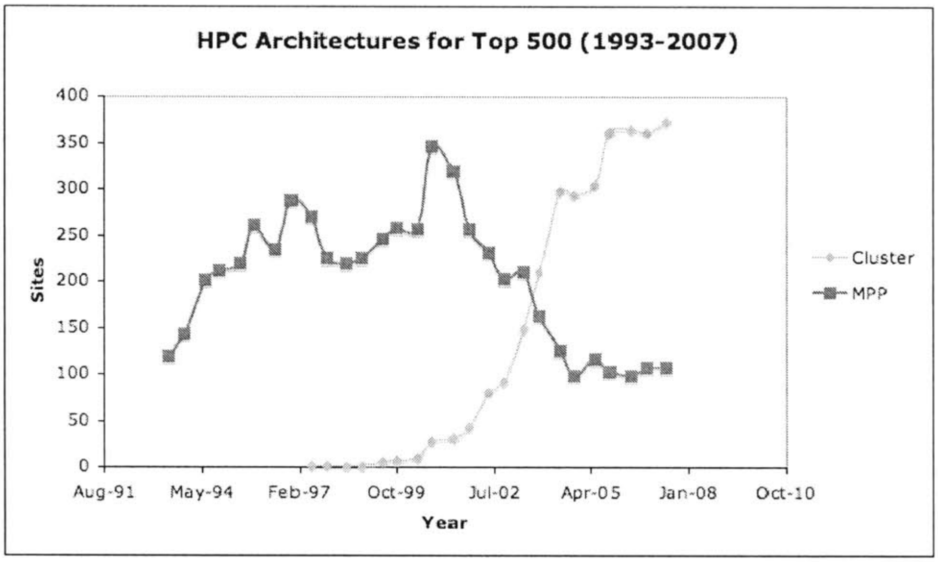

Figure 23: HPC Architectures for Top 500 Supercomputers, 1993 – 2007 [41]

One major consequence of this architecture's dominance in the high-performance computing market is the rapid increase of volume servers in the installed base. Figure 24 shows the growth of the volume servers relative to high-end and mid-range servers in the installed base. Historical data from Gartner's DataQuest worldwide quarterly statistics for servers show that volume servers have accounted for much of the growth in the worldwide installed base while the use of mid-range and high-end/high cost servers

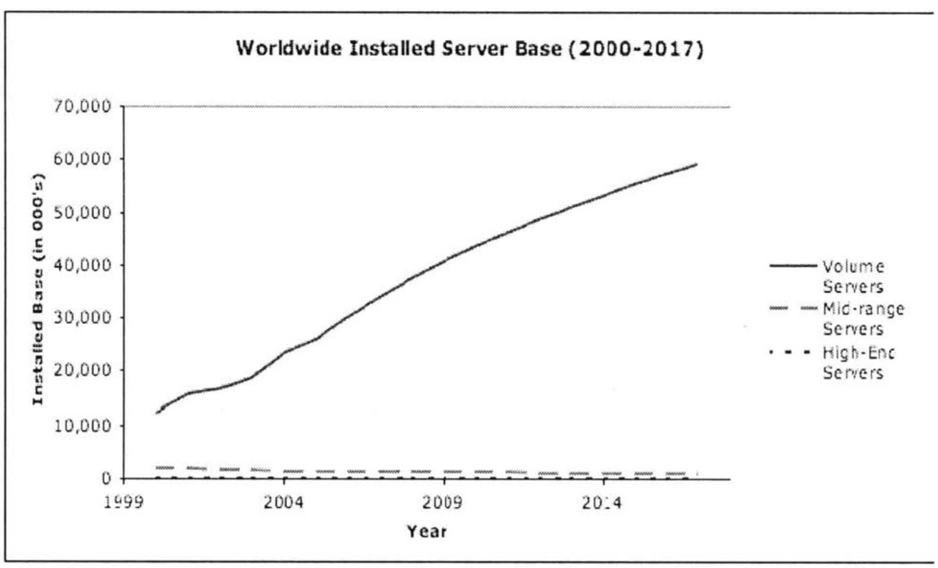

Figure 24: Projected Worldwide Installed Server Base, 1999 – 2017 [32, 34]

remains relatively flat. Furthermore, the results of linear regression analysis on the data set show volume servers will continue to be the dominant source of growth in the worldwide installed server base. As cluster computers using volume servers continue to be an attractive option for a broad array of parallel computing applications, the increased use of volume servers has contributed substantially to the rising electricity demands of data centers. Although power requirements in individual 1U volume servers (as represented by the top six selling servers) have increased by a modest 3.3 percent CAGR relative to the 18.57 percent CAGR rise in power- at-the-plug requirements for high-end servers as shown in Table 5, rising electricity consumption by servers is largely driven by volume servers. In 2005, for example, volume servers accounted for 82.5 percent of the total electricity use by servers while high-end servers accounted for 6.5 percent.

Table 5: Average Power Requirements - Top Selling Servers, 2000-2005 [32]

Server Type	2000 (W)	2003 (W)	2005 (W)	CAGR (%)
Volume	186	207	217	3.33
Mid-Range	424	524	641	10.24
High-End	5534	6428	10673	18.57

3.3 Optical Interconnects as a Solution?

In the preceding sections, the demand for digital information and a shift from monolithic, massively-parallel processing machines to a cluster platform based on smaller, cheaper servers have contributed to current worries regarding energy in data center facilities. To achieve improved power performance usage, data center managers are examining a variety of solutions such as implementing virtualization[19] and blade servers[20]; however, this thesis propose the increased adoption of optical interconnects as another possible solution.

In terms of energy usage in data centers, copper interconnects requires fairly high power consumption. For example, the extended operating frequency range of a 10 gigabit Base-T copper cable and interfaces can require as much as 10 to 15 watts of the

[19] Virtualization is a software method that allows multiple applications to be run on a single server by creating multiple virtual servers; this method may allow a reduction of physical servers, power, and space
[20] Blade Server is a "a server chassis housing multiple thin, modular electronic circuit boards, known as server blades. Each blade is a server in its own right, often dedicated to a single application. The blades are literally servers on a card, containing processors, memory, integrated network controllers, I/O Ports" [44]

total 25 watts capable per PCI slot in order to overcome increased insertion loss and cross talk [42]. On the other hand, emerging optical interconnects could require much less power relative to copper interconnects when used to transmit information:

Table 6: 10 Gigabit Transceiver Power Comparison [42]

10 Gigabit Transceiver Type	Power (W)
10-Gbit Base-T Copper Transceiver	10 W
10-Gbit X2 Optical Transceivers	4 W
10-Gbit XFP Optical Transceivers	2.5 W
10-Gbit SFP+ Optical Transceivers	1W

Table 7: Power Budget for a Typical Volume Server [25]

Component	Peak Power	Count	Total
CPU	40 W	2	80 W
PCI Slots	25 W	2	50 W
Memory	9 W	4	36 W
Motherboard	25 W	1	25 W
Disk	12 W	1	12 W
Fan	10W	1	10 W
System Total			213 W

Table 7 shows the power budget for a typical volume server. Each server may have one or two PCI slots that provides up to 25 watts of power for the transceiver that is connected to it [25]. Although the power savings gained from just exchanging electrical box-to-box interconnects is small relative to the larger power budget of an individual server, the cumulative savings in annual power and other indirect costs can be substantial for a facility that employs tens, hundreds, or even thousands of servers.[21]

3.4 Key Chapter Findings

Limited energy capacity has become a major issue that most concern data center managers. Electricity consumed by data centers in the US and the world has doubled

[21] Industry experts estimate that for every dollar spent on powering servers and equipment in a data center, firms could spend an additional $0.65 dollar or more to cool that equipment [30].

over the past five years; furthermore, various government and research studies show that energy use in data centers will again double over the next five years. Driving this steep rise in energy use is the rising demand for digital information and a shift in the HPC market from large, monolithic machines to clusters built with many smaller, cheaper servers. Implementing a box-to-box interconnect fabric with optical components instead of copper can reduce per-link power consumption. The following chapter will introduce a method to quantify the possible savings in power and cost that can be achieved using optical interconnects instead of copper ones.

Chapter 4 – TCO Model Description

4.1 Model Approach

IT equipment in data centers are organized into clusters of servers, network switches, and storage devices. However, the number of clusters that a facility can support is constrained by one or more of its characteristics. For analytical purposes, this model provisions IT equipment in a data center using power as a primary constraint. The model begins this process by first calculating the total amount of power available for use by IT equipment in a data center:

Available IT Power = (Max IT Load)*(Utilized IT Load) (Eq. 1)

Max IT Load is an input variable that specifies the maximum amount of power capacity (in terms of kilowatts) provisioned for a data center, while *Utilized IT Load* is an input variable that accounts for percentage of the power capacity that can be used for the IT equipment.

Once the *Available IT Power* has been calculated, the model then proceeds to determine the number of clusters that could possibly be installed. While a data center architect may arbitrarily determine the number of servers in a cluster, the number of network switches and storage devices is calculated using the cluster's network configuration and the number of servers attached to it.

4.1.1 Network Configuration

Clusters traditionally use two types of network configurations: Mesh or Tree. The mesh network is the type of network topology in which each switch is connected to multiple switches via point-to-point links [50]. A mesh network is a flat topology and usually results in the implementation of fewer network switches. The number of switches in a mesh-configured network can be found:

Mesh Network Switches = (1/Switch Utilization[22])*(Ports[23])*(Servers) (Eq. 2)

On the other hand, tree networks are not flat topologies. Depending on the number of servers in a cluster, tree networks may have multiple levels. At the highest

[22] Switch Utilization is the percentage of the ports actually used per network switch.
[23] Ports = The number of ports per switch.

level of the hierarchy, there is a root switch that is connected to one or more other network switches in a second level, which in turn, are connected to one or more other network switches in a third level [51]. This general pattern will continue until the base level of network switches has been reached. It is at the base level where network switches are actually connected to server nodes instead of other network switches. In the model, the number of network switches in a tree-configured network is calculated by first finding the number of network switches in the base level necessary to accommodate all server nodes in a cluster. From the base level, the function works its way up the hierarchy by calculating the number network switches necessary in each higher level until it reaches the root network switch. See Appendix for function code.

Once the number of network switches for a cluster has been calculated, the total number of storage devices utilized is found by dividing the number of ports per storage device into the total number of network switches. While intuition may suggest that the number of storage devices is controlled by storage capacity needs per cluster, the primary constraint on the number of storage devices is frequently the number of ports available per storage device that can be connected to network switches [51, 52]. The model assumes that a storage device has a fixed number of ports. Thus, the number of storage devices per data center is found using the following equation:

Storage Devices = (# of Network Switches) / (Ports Per Storage Device) **(Eq. 3)**

4.1.2 IT Power

IT power is the direct power that a device outputs once it is connected to a power source. In the model, IT power is given using the power requirements listed in the technical manuals of each device. In tracking how the IT power requirement per each device type adjusts with changes in interconnect type, the model uses the following general equation:

IT Power Output[24]= (Base IT Power) + ((# Interconnects)*(Δ in Watts $_{\text{Interconnect}}$)) (Eq. 4)

Base IT Power = Device Power Output (Copper)
Δ in Watts $_{\text{Interconnect}}$ = 0, if copper
Δ in Watts $_{\text{Interconnect}}$ = Interconnect (Copper) - Interconnect (Optical), if optical

[24] Base IT Power is the power requirement listed in each device's technical manuals. Base IT Power includes the power consumed by copper interconnects that are initially installed. "# Interconnects" notes the total number of interconnects in a device.

Here, it is important to note that for servers, Base IT power output is a function of how much its CPU is utilized. The power output by a server is linearly proportional to how much its CPU is utilized. The utilization of a server's CPU has no effect on interconnect power consumption, which is constant. Refer to appendix for the CPU Utilization data. The model also reports Cumulative IT Power Output Per Server. Unlike the equation 4, Cumulative IT Power Output Per Server is a metric that combines an individual server's IT Power Output with it's fractional share of the total IT power output produced by network switches and storage devices:

Cumulative IT Power Output Per Server = IT Power Output$_{Servers}$ + (1 / (# Servers Per Datacenter)) * [(IT Power Output$_{Switches}$) * (# Switches Per Datacenter) + (IT Power Output$_{Storage\ Devices}$) * (#Storage Devices Per Datacenter)] **(Eq. 5)**

4.1.3 Cooling Power

It is also necessary to calculate the amount of power necessary to cool all IT equipment in a given cluster. Because insufficient cooling can cause equipment to fail, the power required to cool must also be accounted for when initially installing servers, network switches, and storage devices. Calculating the power necessary to cool each type of device is accomplished through three major steps: (1) Calculate the Heat Output (in BTUs) for each device type, (2) Calculate the heat transfer rate of each device, and (3) Calculate the sensible cooling capacity and devices per computer room air conditioning (CRAC) unit. In the following approach, the model assumes airflow remains constant while the temperature of exhaust air from each device changes.

Determining the amount of heat produced by each device while operating is the first step to calculating the cooling power requirement. This calculation is found by first deriving the ratio between a device's maximum heat output to its maximum IT power output as shown in Table 8. For each device in the table, the heat output and IT power

Table 8: Heat Output to Power (HTP) Ratios for IT Equipment [53, 54, 55]

Device	Max. IT Power Output (W)	Max. Heat Output (W)	HTP Ratio
HP DL 360 G5 Server	290.70 (W)	843.48 (W)	2.9015[25]
Nortel 5510 Net. Switch	135 (W)	134.5 (W)	0.9963
Sun Storage J4500 Array	1100 (W)	1020 (W)	0.9273

output values are based on copper interconnects that are installed by default. The heat-to-power (HTP) ratios are useful to calculate how heat dissipation adjusts with changes in IT power output resulting from changes in interconnect technology. As shown in the equation below, the heat output per device can be found by:

Heat Output = (HTP)*(IT Power Output) **(Eq. 6)**

The second step to calculating the cooling power requirement is to find the appropriate heat transfer rate. For analytical purposes here, the physical model used to calculate the heat transfer rate is a rectangular tube as shown in the following diagram:

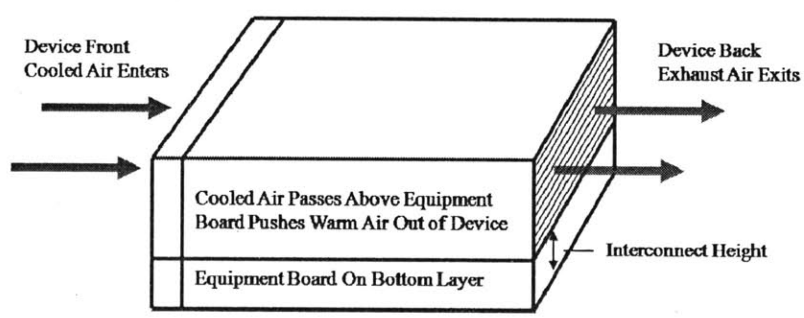

Figure 25: Physical Model of Air Flow

[25] Unlike network switches and storage devices, the HTP ratio for servers is greater than one. The 2.9015 value is consistent with prior conversations with IT managers who estimate that the power needed for cooling is three times the amount of IT power.

The heat transfer rate can be using Newton's equation for cooling:

Heat Transfer Rate (Q, in watts)[26] = h*A*(T-surface – T-source_air) **(Eq. 7)**

In the above equation, the area (A) is the partial cross-sectional area of a rack, which is defined to be the product of a rack's width and depth. A rack's depth accounts for changes in cable bend radii that occur from implementing optical rather than copper

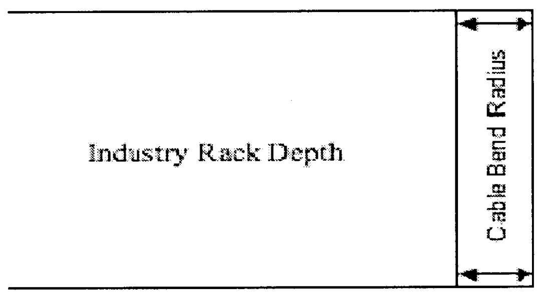

Figure 26: Rack Cross-Sectional Area

interconnects. From the above description and diagram, the equation for a rack's cross-sectional area:

Rack Cross-Sectional Area = (Rack Width)*(Rack Depth + Cable Bend Radii) **(Eq. 8)**

T_surface is the maximum operating temperature of the device while T-source_air is the temperature of the air that CRAC units provide to cool the racks. The heat transfer coefficient (h), which is dimensionless, is calculated using the following equation:

Heat Transfer Coefficient (h) = (Kw / Dh)*Nu **(Eq. 9)**

The heat transfer coefficient depends on the properties of the air coming from the CRAC. In order to solve the equation for the heat transfer coefficient, other equations for the thermal conductivity (Kw), the hydraulic diameter (Dh), and the nusselt number (Nu) must be solved first. The equation for the thermal conductivity of air is [58]:

Thermal Conductivity (Kw) = f(Temperature) = 0.0071*(T) + 0.025 **(Eq. 10)**

The hydraulic diameter (Dh) is a geometric measure of the open channel in the device through which cooled air flows [56]:

Hydraulic Diameter (Dh) = (2*Height*Width)/(Height + Width) **(Eq. 11)**

[26] Newton's Law of Cooling results in a heat transfer rate that is expressed in terms of watts. In order for the heat transfer rate to be used in this context, that value must be converted to BTUs. 1 Watt = 3.41 BTU.

The equation for Hydraulic Diameter also accounts for changes in interconnect-type through its preliminary calculation of Height. In the model, Height = Device Height – Interconnect Height (as shown in diagram below).

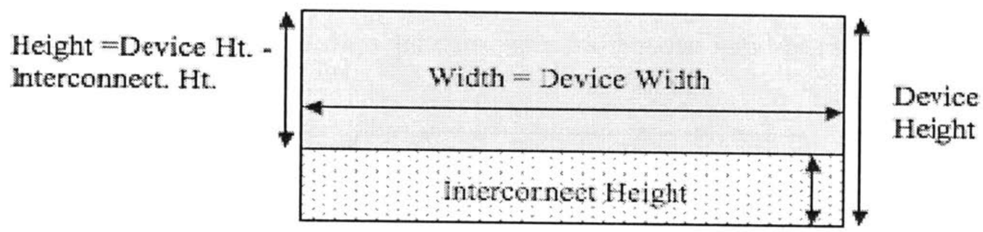

Figure 27: Hydraulic Diameter (Rear View of Device)

The nusselt number, a dimensionless number, is the ratio of convective to conductive heat transfer across the boundary of a device [57, 58]. In practice, though, the nusselt number is normally expressed as a function of the reynolds number (Re) and the prandtl number (Pr):

Nusselt Number (Nu) = 0.102*(Re^0.675)*(Pr^(1/3)) **(Eq. 12)**

Reynolds Number (Re) = ((Density*Air Velocity)*Dh)/(Viscosity) **(Eq. 13)**

Prandtl Number (Pr) = (Heat Capacity of Air*Viscosity)/(Kw) **(Eq. 14)**

In the above equations for the reynolds number and the prandtl number, density, viscosity, and thermal conductivity are functions of the temperature of air provided by CRAC units [59].[27] The Heat Capacity of the CRAC-provided air is constant at 1005 for all possible air temperatures.

Once the heat transfer coefficient is found, one can calculate the temperature difference between a device's maximum operating temperature and the CRAC-provided air (Delta-T) as shown in the following equation [59]:

(Q) = (h*A*Delta-T) = (Heat Output $_{\text{Device}}$) **(Eq. 15)**

(Delta-T) = (Heat Output $_{\text{Device}}$) / (h*A) **(Eq. 16)**

[27] These relationships between temperature and the above variables are linear [59]:
Density of Air (Temperature) = -0.00347*(T)+2.2184;
Thermal Conductivity (Temperature) = 0.0071*(T)+0.005;
Kinematic Viscosity (Temperature) = 0.0976*(T)-(1*10^-5)

After this change in temperature has been determined, that value can then added to the CRAC-provided temperature in order to calculate the temperature of the exhaust air flowing out of the device and back to the CRAC unit.

As part of the final step toward calculating the cooling power requirement, the exhaust temperature values are used to determine the sensible cooling capacity of CRAC units. The sensibility cooling capacity is the effective amount of cooled air that a CRAC can deliver. The relationship between the exhaust temperature and a CRAC unit's sensible cooling capacity is based on empirical evidence provided by the manufacturer's technical manuals and an interview with a technical support staff person. Generally, the sensible cooling capacity rises as the temperature of the exhaust air coming into the CRAC rises. Refer to appendix for CRAC sensible cooling data. Once the sensible cooling capacity for the CRAC is determined, the model proceeds to calculate the number of devices per CRAC unit by dividing the heat dissipation per device into the sensible cooling capacity of the CRAC:

Devices Per CRAC Unit = (Sensible Cooling Capacity)/(Heat Output $_{Device}$) **(Eq. 17)**

After the number of devices supported by a CRAC unit has been calculated, the cooling power (in terms of watts) can be found using the following equation:

Cooling Power Per Device = (CRAC Power)/(Devices Per CRAC Unit) **(Eq.18)**

The model also reports Cumulative Cooling Power Requirement Per Server. Unlike the equation 18, Cumulative Cooling Power Requirement Per Server is a metric that combines an individual server's cooling power requirement with it's fractional share of the total cooling power required by network switches and storage devices:

Cumulative Cooling Power Requirement Per Server = Cooling Power Requirement$_{Servers}$ + (1 / (# Servers Per Datacenter)) * [(Cooling Power Requirement$_{Switches}$) * (# Switches Per Datacenter) + (Cooling Power Requirement $_{Storage\ Devices}$) * (#Storage Devices Per Datacenter)]

(Eq. 19)

4.2 Data Center Cost Components

Once the number of servers, network switches, and storage devices have been determined by the power-based provision flow of IT equipment as described above, the other cost components can be calculated. In this model, IT equipment costs, infrastructure costs, and personnel costs are the primary cost components of a data center.

4.2.1 IT Equipment Costs

In order to calculate the purchase costs of servers, network switches, and storage devices the model is using the following general equation:

Device Cost = Retail Cost + (# Interconnects)*(Δ in Interconnect Cost) (Eq. 20)

In cases where copper interconnects are used, "Δ in Interconnect Cost" = 0. In cases where optical interconnects are used, "Δ in Interconnect Cost" = 100. Since most IT equipment are shipped standard with copper interconnects, the additional $100 is the required cost premium in employing optical interconnects rather than copper interconnects. [60, 61, 62].

4.2.2 Infrastructure Costs

The model contemplates a facility in which the vast majority of resources and space is dedicated to the constant operation of IT equipment. In order to measure as accurately as possible how infrastructure cost components change with changes in IT equipment, this model does not contemplate the costs of implementing a server room within a much larger facility built for purposes other than the operation of IT equipment. The total data center infrastructure costs include the cost to install racks and raised floors, the cost of an uninterruptible power supply, the cost of cooling equipment, and the cost to purchase and build the overall facility that houses the data center.

Rack Costs

One must also determine the costs to install a rack infrastructure that houses the IT equipment in stacks. In doing so, the number of racks required by the total amount of provisioned IT equipment is determined by first calculating how much space in terms of rack location units (U's) are required. In the model, industry standard 45 U racks are used to house equipment. Afterwards, the model calculates how many devices will fit in a rack. The model assumes that each type of IT equipment will be housed with equipment of the same type. Under this assumption, the number of devices per rack is found by dividing the rack space required by an individual device into available amount of space per a given rack. Once the maximum number of devices (either servers, switches, or storage device) per rack is found, the number of racks can be determined by using the general equation:

 # Racks = Roundup((# of Devices) / (Devices Per Rack),0) **(Eq. 21)**

Where Devices Per Rack = (Available Space Per Rack) / (Space Per Device)

Using this equation to find the number of racks necessary to house different groups of IT equipment will result in a calculation of the total number of racks in the data center. Once the total number of racks needed to outfit a data center has been found, the model can then calculate total rack costs.[28]

Raised Floor Costs

In this part of the model, the first step to calculating the raised floor costs is to first identify the total amount of area occupied by IT racks fully equipped with servers, network switches, and servers. By itself, an industry standard rack occupies 7.46 square feet of space [51]. With equipment and cables, however, the effective floor area for a utilized rack can be much greater. Calculating the effective floor area for a given rack is achieved by accounting for the additional horizontal space necessary to accommodate the bend radii of cables added to the rack's depth[29] as well as the required vertical space to channel cables along the height of the rack.[30] Calculation for the effective floor area of a rack is best illustrated in the following diagram:

---

[28] The cost to purchase and install a single rack include the cost of the physical structure ($3000), the cost of external hardwired connections ($5000), and the cost of rack management hardware ($3000). [62]

[29] Effective Rack Depth = Initial Rack Depth + (Bend Radius of Optical or Copper Cable).

[30] Effective Rack Width = Initial Rack Width + 2*((∏*(Optical or Copper Cable Radius^2)*(1+Cable Management Utilization)). Cable Management Utilization is a factor that tracks how much space in the data center cabling system used.

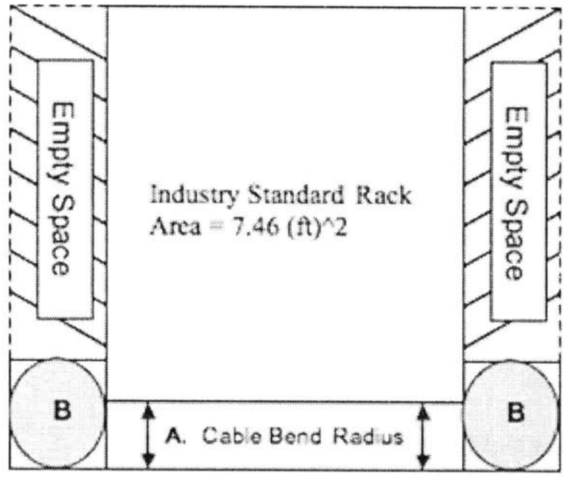

Figure 28: Effective Floor Area Of Utilized Racks

After finding the effective area for racks containing servers, network switches, and storage devices, the model calculates the amount of electrically active space (space covered by raised floors) within a data center. In most data centers, utilized racks occupy approximately 40 to 60 percent of the total amount of electrically active space [2]. The model uses a variable, *Floor Space Utilization*, which accounts for the amount of electrically active space that all utilized racks consume. Simply dividing this factor into the total area of utilized racks result in the total electrically active space in a data center. The industry standard costs for raised floors are constant at $220 per square foot [63].

Land and Construction Costs

In most data centers, the electrically active space comprises half (or more) of the entire facility's area. The model uses a variable, *EA Space Utilization*, which denotes a ratio between the total facility space and the electrically active space; multiplying this factor to the area of the electrically active space results in the facility's total area. With the total facility area determined, the cost of land and construction can easily be calculated. The cost of land in the model is assumed to be the national average of $45.20 per sq. foot [63]. The average cost to build on that land is $1100 per sq. foot [64].

Cooling Equipment Cost

The most critical component of any cooling power infrastructure is the CRAC unit. The model uses a 10 ton, 50 Hz, air-cool CRAC unit. The minimum number of CRAC units needed to support the cooling power requirements of a data center. On average, the cost per individual CRAC unit is $25,000.

Uninterruptible Power Supply Costs

As described in the preceding chapter, power infrastructure equipment includes an uninterruptible power supply (UPS). Of the elements of power infrastructure equipment, the UPS is the only critical component in the power infrastructure described above that does not depend on the number of IT devices provisioned. The minimum amount of UPS equipment needed to support the IT equipment is based on the initial power budget allocated to the data center. The number of UPS equipment provisioned is not a function how many servers are installed. On average, UPS costs are $20,000 per every kilowatt initially provisioned in a data center [64].

4.2.3 Personnel Costs

Personnel costs are determined using the number of technicians supporting fully utilized racks. In this model, there are three types of personnel staff supporting IT racks in a data center: IT Site Management, Facilities Site Management, and Maintenance. In each group, the model assumes that each staff member monitors 40 racks. In other words, at least one IT site manager, one facilities site manager, and a maintenance technician monitor 40 racks. For very small data centers, a minimal three person staff is allocated. In the model, the average salary allocated for personnel staff is $48,000/year. It is assumed that this figure for the average salary includes benefits (health insurance and 401k contributions) and taxes that the company must pay for the employee [63].
Thus, the total personnel cost is the product of the total number of staff and the average salary:

$$\textbf{Personnel Costs} = (\textbf{Site Management}+\textbf{Facilities}+\textbf{Maintenance})*\textbf{Salary}_{\textbf{Avg}} \qquad \textbf{(Eq. 22)}$$

4.3 Annualizing Infrastructure Costs

Once the actual cost of the infrastructure is determined, data center managers can account for the infrastructure cost as a yearly expense by first finding and applying the weighted average cost of capital (WACC) to the total value. The WACC is a weighted average of equity and debt that when multiplied by the market value of the capital results in the opportunity of owning the capital. The WACC can be found using the following variables and equations [65, 66]:

WACC[31] = (DE/(DE+1))*(cost of debt)+(1/((DE+1))*(cost of equity) **(Eq. 23)**

Debt-Equity Ratio (DE) = (Total Debt)/(Total Equity) **(Eq. 24)**

Cost of Debt = The effective rate a company pays on its debt **(Eq. 25)**

Cost of Equity[32] = The expected return on equity = Rf + β*(Rm - Rf) **(Eq. 26)**

After the WACC has been found, the annual cost of infrastructure can then be found using the following general equation:

Yearly Infrastructure Cost = (Total Cost)*(CFR) **(Eq. 27)**

CFR = (WACC*(1+WACC)^n) / ((1+WACC)^n-1) **(Eq. 28)**

Once the yearly infrastructure cost is calculated, the total annual cost is the sum of personnel, power, and infrastructure costs. The total non-annual cost is the sum of the costs to purchase servers, network switches, and storage devices.

4.4 Return on Investment (ROI)

Return on Investment (ROI) is one of several financial metrics that can be used to decide whether or not to proceed with a particular investment. ROI is the ratio between an investment's potential gains and its costs:

ROI = (Investment Gains – Investment Costs) / (Investment Costs) **(Eq. 29)**

---

[31] For the WACC, Akamai's financial information was used. Refer to its 2008 Annual Reports.

[32] The equation for the cost of equity is the Capital Asset Pricing Model: (A) The risk-free rate of return (Rf) is the expected return on a one-year treasury bill. This model uses the average return for August, 2008, which was 2.18%; (B) The average rate of return of the market is defined by Brealy and Myer's Principles of Corporate Finance to be 8.4%; (C) β, the covariance, can be found in Merrill Lynch's Security Risk Evaluation. Again, using Akamai as the representative company for our model, this value has been set to 2.65.\

In the above equation, "Investment Costs" is the additional costs to install optical interconnects rather than copper interconnects. "Investment Gains" will come from potential savings in annual personnel, power and infrastructure costs as well as additional revenue[33] realized from implementing optical interconnects instead of copper interconnects. In deciding between whether to implement optical interconnects, a data center manager will consider the results of an ROI analysis – a positive result indicates profitable investment whereas a negative result indicates losing investment. The ROI framework is implemented in the model in a user-defined function. Refer to appendix for function code.

[33] Additional revenue is based on the extra number of servers deployed from the implementation of optical interconnects rather copper interconnects. The model uses a constant for earned revenue per server.

Chapter 5 – Model Results

As discussed in the introduction, optical interconnects have generally failed to gain widespread adoption because they are perceived to be too expensive when evaluated on cost or cost per performance. This thesis explores whether other metrics could be used to evaluate the attractiveness of optical interconnect adoption for a particular market. Through the use of currently available literature and phone interviews, power consumption was identified as an additional metric useful in an evaluation of the profitability of installing optical interconnects in large HPC data centers. The preceding chapter described a total cost of ownership cost model, which would allow one to evaluate the attractiveness of optical interconnects in a much broader context than simply looking at cost per interconnect or cost per performance alone.

Using the model structure described in the preceding chapter, the following sections will examine the profitability (via ROI) of optical interconnect adoption in data centers. As stated in the previous chapter, ROI is a good metric that allows one to combine investment costs and gains over the projected period of time into one intuitive metric. For this reason, ROI is often used by IT and data center managers to decide whether to invest in new technology. This chapter will first describe a base scenario and present the results of the model using the inputs from the base scenario. Second, this chapter will present the results of sensitivity analysis performed by altering key input values.

5.1 Base Scenario

Because this model seeks to examine changes in the profitability of optical interconnect adoption based on potential savings in per-server power consumption, the chief constraint in the model is power consumption. This model begins with the assumption that each data center is planned with a maximum power consumption requirement of 1000 kW. Another critical input is the number of servers per cluster. By default, the model begins with 256 servers per cluster, which is an average value for a mid-size to large data center. With a power requirement and cluster size determined, the model begins with first calculating how many clusters of servers, network switches, and storage devices can be installed. The following table contains other default inputs that the model uses to assemble a data center.

Table 9: Default Model Inputs

Number of kW per Data Center	**1000**
Number of Servers Per Cluster	**256**
Server CPU Utilization	**50**
Network Switch Utilization	**100**
Ports Per Server	**1**
Ports Per Switch	**48**
Ports Per Storage Device	**2**
Rack Dimensions	**Depth = 3.27 ft.** **Width = 2.28 ft.**
Copper Cable Dimensions	**Cable Radius = 0.015 ft.** **Bend Radius = 0.417 ft.**
Optical Cable Dimensions	**Cable Radius = 0.006 ft.** **Bend Radius = 0.230 ft.**
Copper Interconnect Power Consumption	**1.1**
Optical Interconnect Power Consumption	**0.66**
Electrically Active Space Utilization	**0.5**
Rack Space Utilization	**0.75**
Floor Space Utilization	**2**

In the above table, Network Switch and CPU Utilization are initially set at 100 percent and 50 percent respectively. Network switch utilization is essentially the number of ports used in a network switch; as discussed in the preceding chapter, this value helps to determine how many network switches are installed in a cluster. On the other hand, Server CPU utilization tracks how much the CPU is used; this value affects power consumption per server – increasing CPU utilization increases power demand per server. While data center managers normally attempt to use as many of the ports in installed network switches, servers usually are not fully utilized; server CPU utilization can vary widely depending on workload traffic. In the model, server CPU utilization is set an average of 50 percent.

Other variables of note are rack space utilization, electrically active space utilization, and floor space utilization. Rack space utilization is the amount of space in a rack used by servers, network switches, and storage devices. Electrically active space utilization tracks how much of the total raised floor space accommodates the total number of racks. Floor space utilization is the ratio of total data center area to total raised floor area. The default values for the three variables are considered good data center design [31]. Component manufacturers provide default values for the remaining variables in the above table.

5.1.1 IT Equipment Provisioning

Based on the number of servers per clusters and on whether the cluster has a tree or mesh network configuration, the model then determines the number of network switches and storage devices as described in chapter 4.1.1. Once the number of servers, network switches, and storage devices per cluster has been determined, the model calculates the cluster's IT power and cooling power requirements. General methods to calculate the IT power and cooling power requirements for each type of device are given in chapter 4.1.2 and 4.1.3. Once the IT power and cooling power requirements per each type of device is found, the total power requirement per cluster can be found using the following general equation:

Cluster – Total Power = (# of Servers)*(Server-Total Power) + (# of Switches)*(Switch-Total Power) + (# of Storage Devices)*(Storage Devices-Total Power) **(Eq. 30)**

The total number of full clusters provisioned in a data center is determined by the following equation:

# of Clusters = (1000 kW) / (Cluster – Total Power) **(Eq. 31)**

If there is balance of available power following the initial installation of full cluster networks, the model will work to install additional servers if possible. In order to do this, the model will first determine if there is enough available power to support network switches and storages devices that could enable an entire cluster. If there is enough power to install additional network switches and storage devices capable of enabling one additional cluster, then the model will install as many servers as possible. If there is not enough power to install additional network switches and storage devices, then no additional equipment will be installed.

5.1.2 Calculating Racks, Raised Floor Space, Building Area

Once the number of servers, network switches, and storage devices is determined, the model calculates the number of racks housing each type of IT equipment according to equation 19 in chapter 4. After the number of racks per each type of IT equipment is found, the model determines how much raised floor space needed. In order to determine the total raised floor area, the model first determines the full effective area of racks holding each type of equipment. Refer to Figure 26 in chapter 4. The general equation for the full effective area of a rack is given by:

Full Rack Area = (Rack Depth + Cable Bend Radius) * (Rack Width + (# Cables) * (2*pi*Cable Radius)) (Eq. 32)

Once the full area for racks holding each type of equipment is calculated, then the total raised floor space can be found using the following set of equations:

Total Rack Area - Servers = (# of Racks - Servers) * (Full Rack Area - Servers) (Eq. 33)

Total Rack Area – Switches = (# of Racks - Switches) * (Full Rack Area - Switches) (Eq. 34)

Total Rack Area – Storage = (# of Racks - Storage) * (Full Rack Area - Storage) (Eq. 35)

Total Raised Floor Space = ((TRA – Servers)+(TRA – Switches)+(TRA – Storage))* (1/E.A. Space Utilization)) (Eq. 36)

After the total raised floor area is found, the total physical footprint of the data center is calculated by:

Total Data Center Area = (Total Raised Floor Space) * (Floor Space Utilization) (Eq. 37)

5.2 Base Scenario Results

Table 10 displays the results from the model when the default values (refer to Table 9) are entered. For a 1000-kilowatt data center facility, the results below show that installing servers, network switches, and storage devices with optical interconnects rather than copper interconnects may not be a profitable investment for data centers using mesh networks. On the other hand, deploying optical interconnect enabled equipment may be a profitable investment for data centers using tree-configured networks.

Table 10: Base Scenario Results

Model Output Variables	Copper Mesh	Optical Mesh	Copper Tree	Optical Tree
Number of Full Clusters	2	2	2	2
Number of Servers	608	609	601	603
Number of Switches	18	18	21	21
Number of Storage Devices	9	9	12	12
Total Power Req. Per Server (W)	1337.138	1334.694	1353.945	1351.018
Power Per Ft.^2 (w/ft^2)	582.35	622.06	576.22	618.24
Per-Server Power Costs ($/Yr)	$925.66	$923.97	$937.30	$935.30
Per-Server IT Equip. Costs ($/Yr)	$1,358.38	$1,469.92	$1,448.78	$1,571.58
Per-Server Infrastructure Costs ($/Yr)	$7,842.46	$7,766.89	$7,959.99	$7,864.39
Per-Server Personnel Costs ($/Yr)	$236.84	$236.45	$239.60	$238.81
ROI – Replacing Copper with Optical Interconnects	**Copper to Optical Mesh**		**Copper to Optical Tree**	
Return On Investment	-0.18		0.02	
Added Investment Costs ($/Yr)	$111.54		$122.79	
Gains From Investment ($/Yr)	$91.01		$124.68	

Table 11 shows the components of expected annual gains realized (as discussed in equation 27) from implementing optical interconnects rather than copper ones in mesh and tree networks. From Table 11, it is apparent that the bulk of the savings stems from reductions in infrastructure and construction costs as well as increased revenue per

Table 11: Break-Down of Gains From Optical Interconnect Deployment

Cost Component	Mesh Networks	Tree Networks
Power Costs	1.85%	1.61%
Raised Floor Costs	0.14%	4.88%
UPS Costs	12.07%	18.00%
CRAC Cost	0.22%	0.32%
Construction Costs	64.69%	53.25%
Personnel Costs	0.43%	0.64%
Additional Profit	14.68%	21.03

server. However, the steep decrease in per-server infrastructure and construction costs and the increase in per-server profit[34] can only be achieved from the rise of additional servers deployed that a result from a decrease in an individual server's overall power

[34] Per-server profit has been included in the results in order to give a full financial picture that data center managers must consider. The model was also run for same sensitivity analysis in the following sections without accounting for increases in per-server profit observed after implementing optical interconnects. However, those results differed in magnitude only; the decision to include per-server profit had no effect on whether ROI was positive or negative.

requirement. The additional servers that may be provisioned in a data center helps to lower per-server fixed costs such as infrastructure and construction.

5.3 Sensitivity Analysis

This section examines how return of investment is affected by variations in input variables. To see the effect, one variable is changed while other variables remain set to default values. The variables examined here are CPU utilization, network utilization, transceiver height, cable bend radius of optical fiber, transceiver power consumption, and cost premium of optical transceiver.

5.3.1 Transceiver Power Consumption

In the base case scenario, copper interconnects require 1.1 watts while optical interconnects require 0.66 watts. For mesh networks, saving 0.44 watts per interconnect result in a -0.18 ROI. In order to achieve a slightly positive ROI, optical interconnects can use no more than 0.64 watts. On the other hand, optical interconnects using 0.66 watts per interconnect will result in a slightly positive ROI of 0.02. Optical interconnect adoption in tree networks will not become unprofitable until transceiver power requirements exceed 0.68 watts.

Figure 29 shows how ROI changes as the power requirement per optical transceiver increases. Not surprisingly, this figure indicates that the larger the power savings in per-interconnect power consumption, the greater the return on investment. The figure shows that the return on investment for both mesh and tree networks decrease in a step-wise fashion indicating that there are ranges of potential per-interconnect power consumption that would result in the same number of servers deployed, and by extension the same ROI.

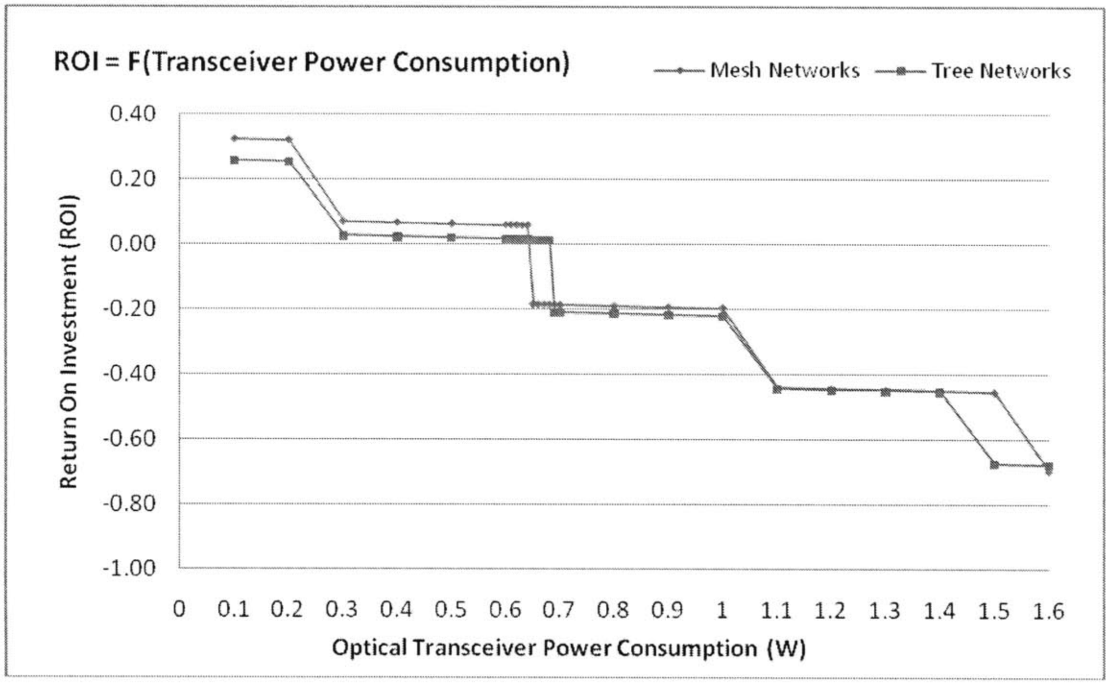

Figure 29: Return on Investment as a Function of Transceiver Power Consumption

Figure 30 and Figure 31 show how the components of per-server power requirement change as per-interconnect power consumption change between 1.1 watts and 1.5 watts. Even in cases where power consumption per interconnect exceeds 1.1 watts, mesh and tree networks deploying optical interconnects may not show a net loss in the number of deployed servers. For data centers with mesh-configured networks, the number of servers will not decrease until the per-server power requirement exceeds 1.76 watts. This threshold is met when optical transceivers use 1.5 watts. In tree-configured networks, a decrease in the number of deployed servers will not be observed until the cumulative power requirement per server exceeds 1.54 watts. This threshold is met when optical transceivers use 1.4 watts. While the changes in the cumulative power requirement per server is driven equally by increases in IT and cooling power demands, the different power thresholds between the two networks is caused by differences in the amount of enabling equipment required to provision a cluster. Because tree networks requirement many more network switches and storage devices than mesh networks to enable a cluster with the same number of servers, a lower threshold for increased cumulative power requirement per server is observed.

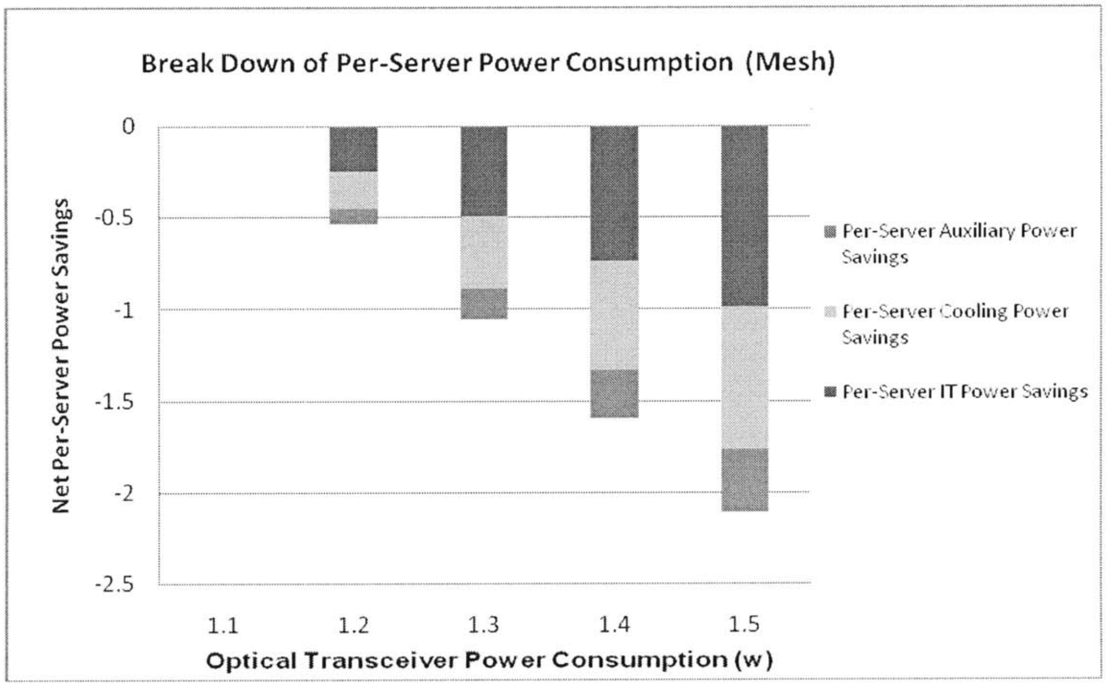

Figure 30: Break Down of Per-Server Power Consumption (Mesh)

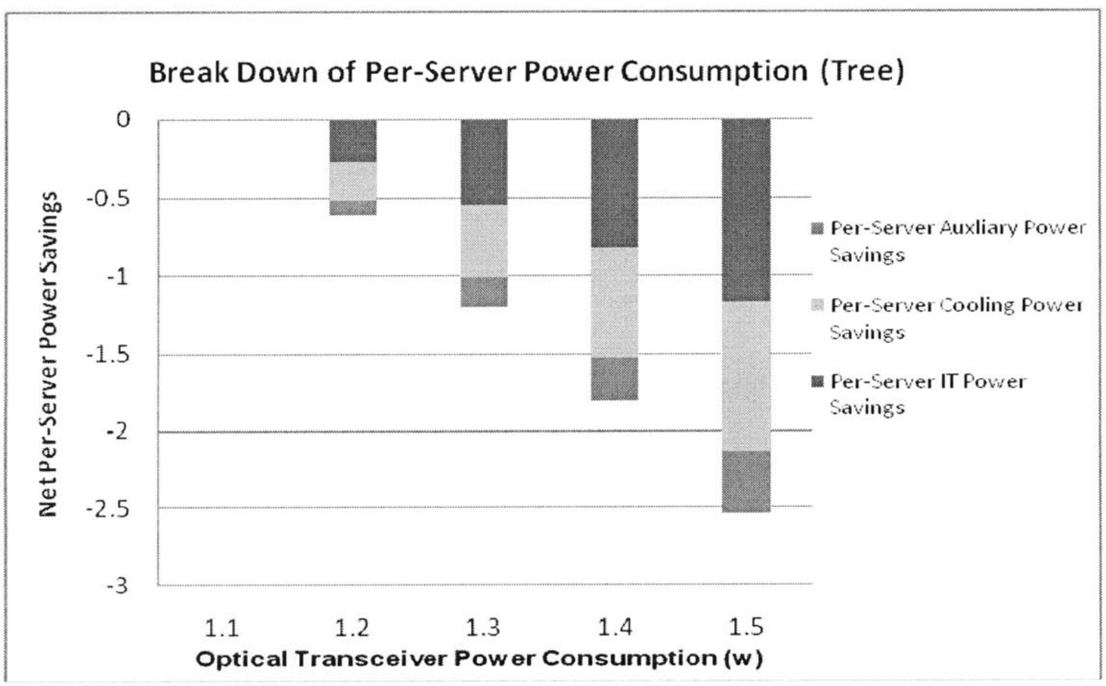

Figure 31: Break Down of Per-Server Power Consumption (Tree)

5.3.2 Transceiver Height

Figure 32 shows how the return of investment (ROI) as the transceiver's height increases. As evident in the figure, transceiver height does not change ROI. Instead, the resulting ROI from optical interconnect adoption remains constant at -0.18 and 0.02 for mesh and tree networks respectively. These results are the same as the base scenario results described above.

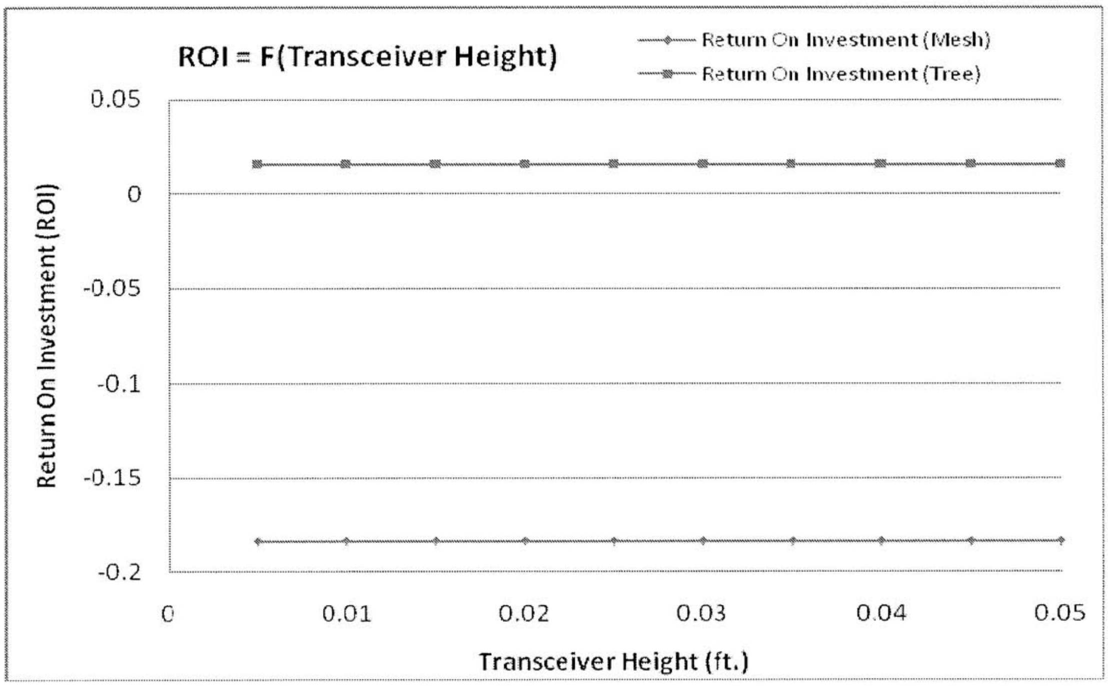

Figure 32: Return On Investment as a Function of Transceiver Height

In the model, transceiver height is important only because it could potentially change the cooling power requirement per server (refer to equations 9-16). In short, increasing transceiver height lowers CRAC capacity (shown in Figure 33). Consequently, lowering CRAC capacity could reduce the number of servers that could be provision and thus potentially increase power requirement to cool individual servers.

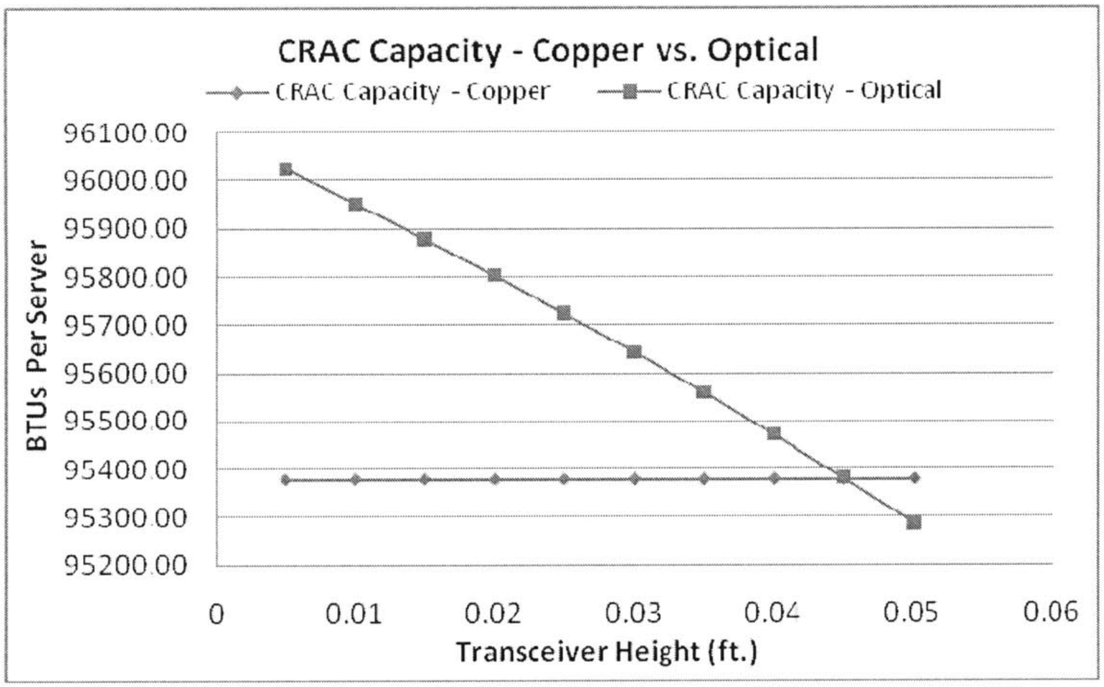

Figure 33: CRAC Capacity as function of Transceiver Height (Copper vs. Optical)

In this case, the number of BTUs that could be saved should optical interconnects be installed rather than optical interconnects do not remotely reach the number of BTUs necessary to provision an additional server per CRAC. Within the range of 0.005 ft. to 0.05 ft, the maximum heat output saved per server is 646.91 BTUs where as the number of BTUs that must be conserved in order to provision an additional server 2318.3 (shown in Figure 34). The threshold minimum of BTUs that must be met in order for a CRAC to accommodate one more server is the number of BTUs that an optical-enabled server outputs (refer to equation 6). Consequently, changes in the transceiver height do not change the number of servers per CRAC, and, by extension, do not change the cooling power requirement per server. Because the cooling power requirement does not change, no change in the number of provisioned servers per CRAC is observed. Thus, there is no change in the per-server costs of major data center components.

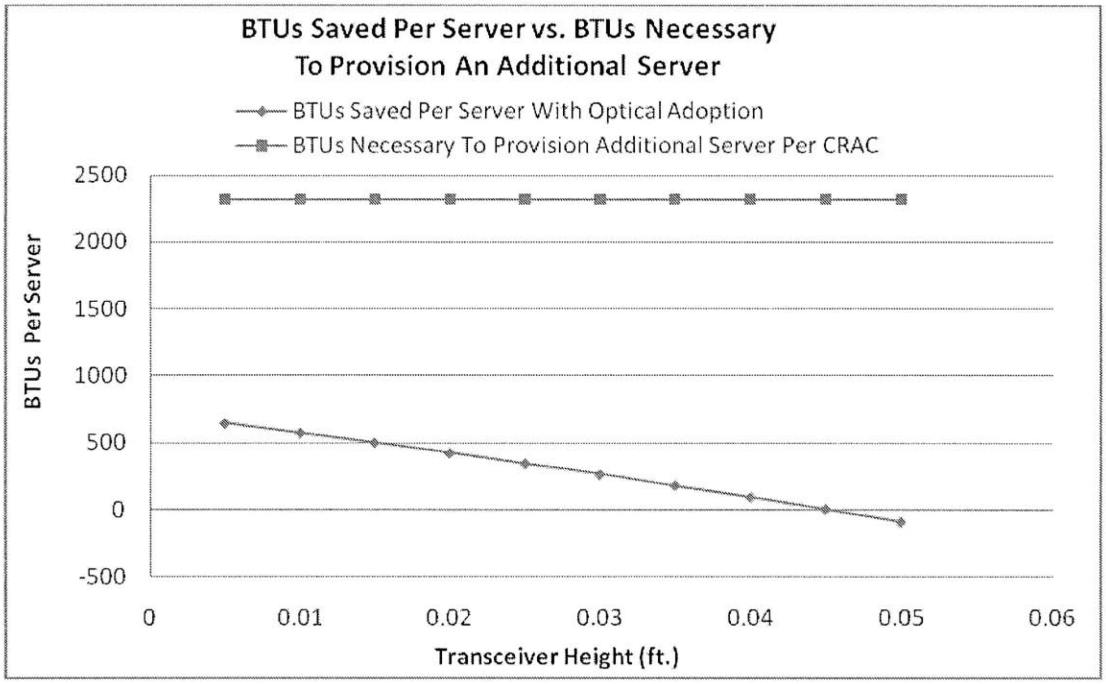

Figure 34: BTUs Saved Per Server as a Function of Transceiver Height

5.3.3 Cable Bend Radius and Cable Radius

Figure 35 shows how ROI changes as cable bend radius increases from 0.05 feet to 0.45 feet. Fiber cables attached to optical transceivers have a bend radius of 0.230 feet while copper cables have a bend radius of 0.417 feet. Figure 36 shows how ROI changes as cable radius increases from 0.005 feet to 0.015 feet. In the base case scenario, fiber cables attached to optical transceivers have a radius of 0.006 feet, while copper cables have a radius of 0.015 feet. In the base case scenario, ROI is -0.18 for mesh networks and 0.02 for tree networks.

In both cases, ROI decreases as either cable bend radius or cable radius increases. The pattern observed in both cases is best explained by the preceding chapter's discussion of raised floor costs (particularly footnotes 29 and 30). As either cable radius or cable bend radius increase, the costs of raised floor installation, land, and construction increase. Unlike transceiver power requirement or transceiver height, neither cable bend radius or cable radius affect the per-server power requirement or how IT equipment is provisioned in the data center. Consequently, any change in ROI of optical interconnect deployment is driven by changes in raised floor installation, land, and construction costs.

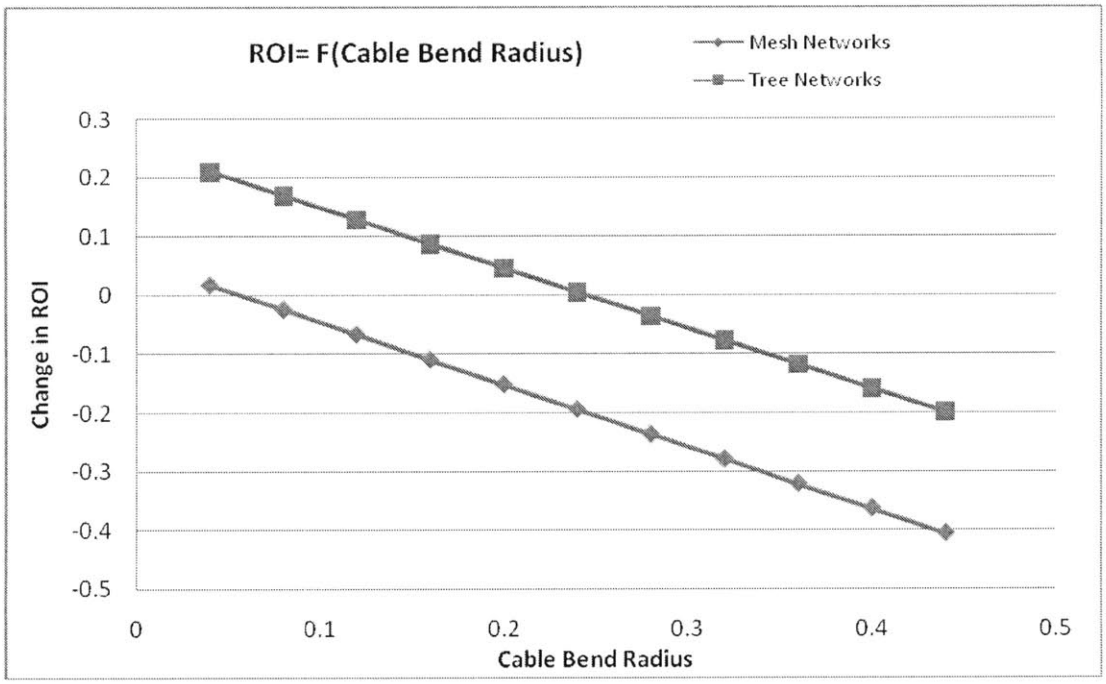

Figure 35: ROI as a Function of Cable Bend Radius

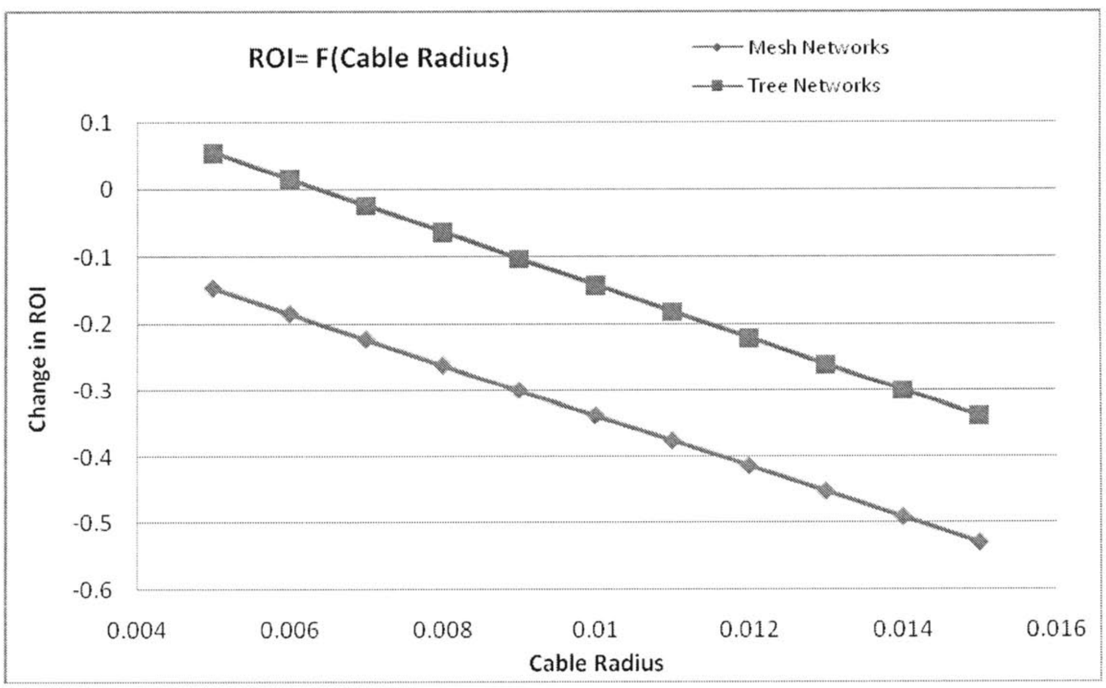

Figure 36: ROI as a Function of Cable Radius

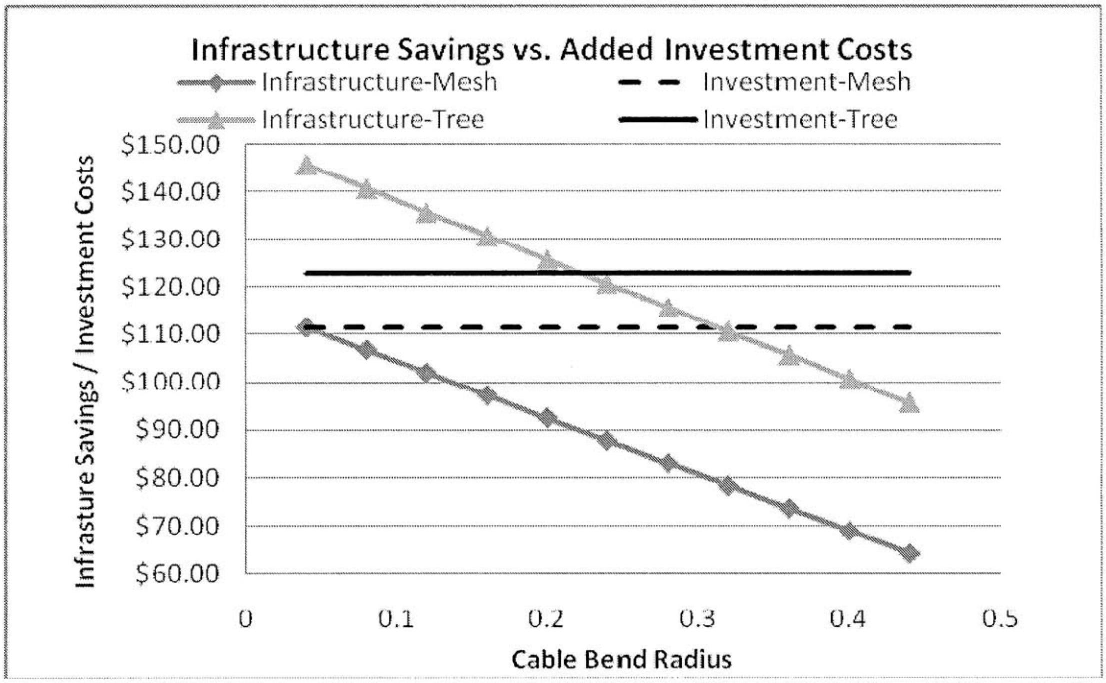

Figure 37: Cable Bend Radius: Infrastructure Savings vs. Added Investment Costs

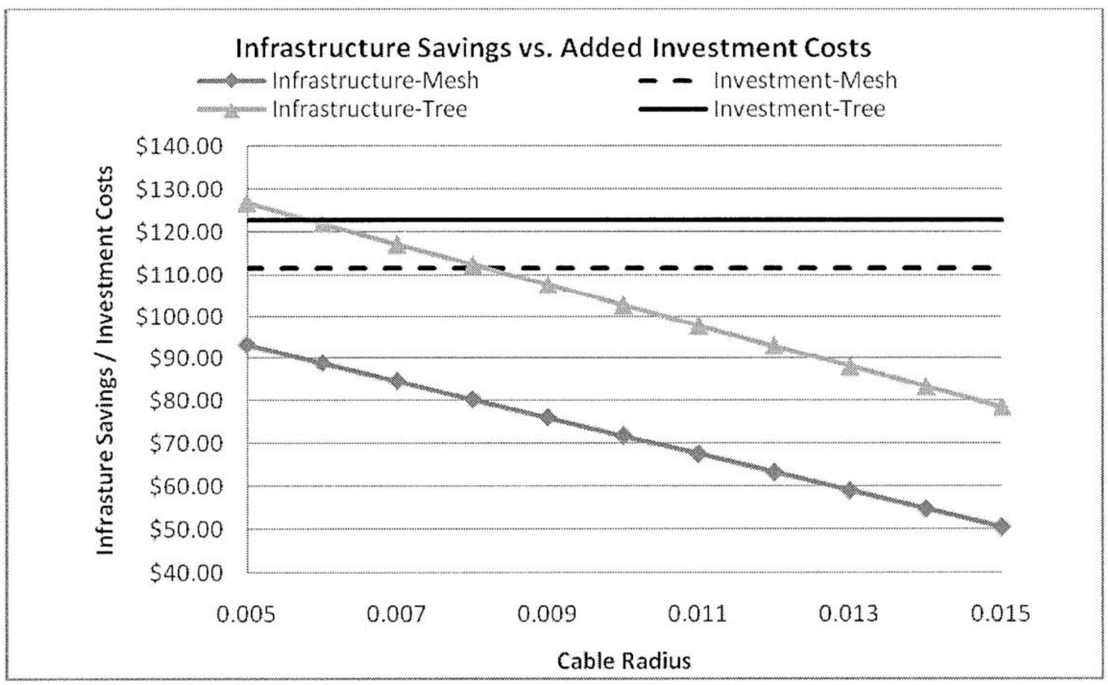

Figure 38: Cable Radius: Infrastructure Savings vs. Added Investment Costs

For mesh networks, decreasing the cable bend radius or cable radius of optical interconnects fail to produce enough savings in raised floor, land, and construction costs in order to recoup the initial investment costs as shown in Figure 37 and Figure 38. On the other hand, savings in infrastructure costs realized by deploying optical interconnects in tree networks are enough to cover the additional investment costs. Beyond the base case settings for cable radius and cable bend radius, the savings in infrastructure costs fall below investment costs.

5.3.4 CPU Utilization

Until now, the preceding discussion has focused on observing how altering key variables that describe transceiver characteristics (transceiver power consumption, height, cable radius, cable bend radius) change ROI. The following discussion seeks to explain how the prospects of optical component adoption are affected by changes in key network and computing variables. In other words, the following discussion focuses on ascertaining how data center conditions (server CPU utilization and network utilization) affect ROI of optical interconnect adoption.

Figure 39 displays how the ROI of optical interconnect adoption responds to

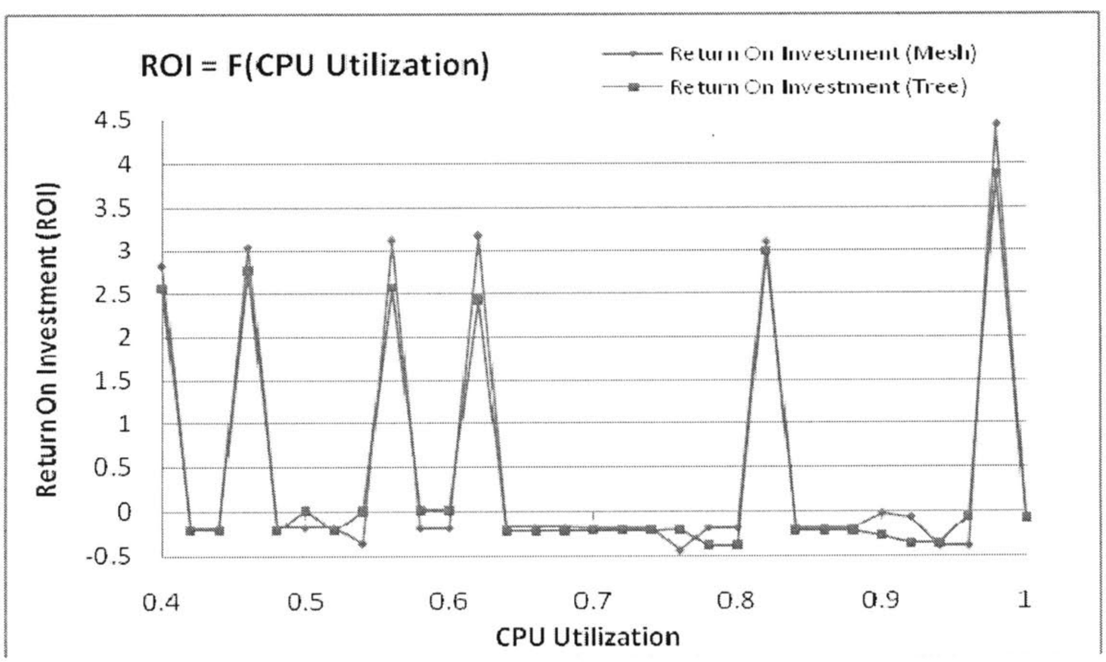

Figure 39: ROI as a Function of CPU Utilization

changes in CPU utilization in mesh and tree networks. CPU utilization affects model output by changing the cumulative IT power requirement per server (see equations 5-6) and the per-server cooling power requirement (see equations 18-19). In the base case scenario, CPU utilization in servers is set at 50 percent. As shown in Figure 39, the ROI of optical interconnect adoption for mesh and tree networks spike at certain points, while remaining relatively unchanged at most other points. At these particular points, the rise in ROI is caused by an increase in the number of additional servers provisioned in a data center (approximately 13 new servers for mesh and tree networks). As shown in Figure 40 and Figure 41, savings in cumulative cooling power requirement per server drives the increase of provisioned servers. The average threshold of combined IT and cooling power savings needed to deploy at least one additional server is 4.39 watts and 4.49 watts for mesh and tree networks respectively. At data points where ROI is positive (shown in the figures below), combined IT and cooling power savings per server surpass the average threshold for mesh and tree networks when cooling power savings spike.

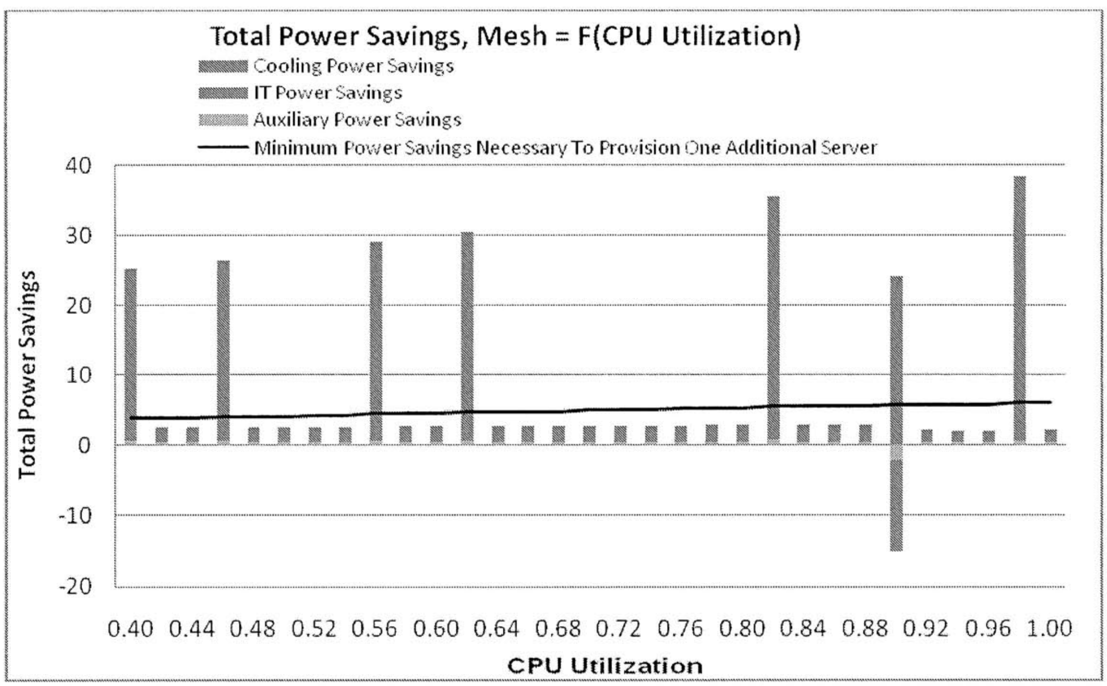

Figure 40: Total Power Savings as Function of CPU Utilization (Mesh)

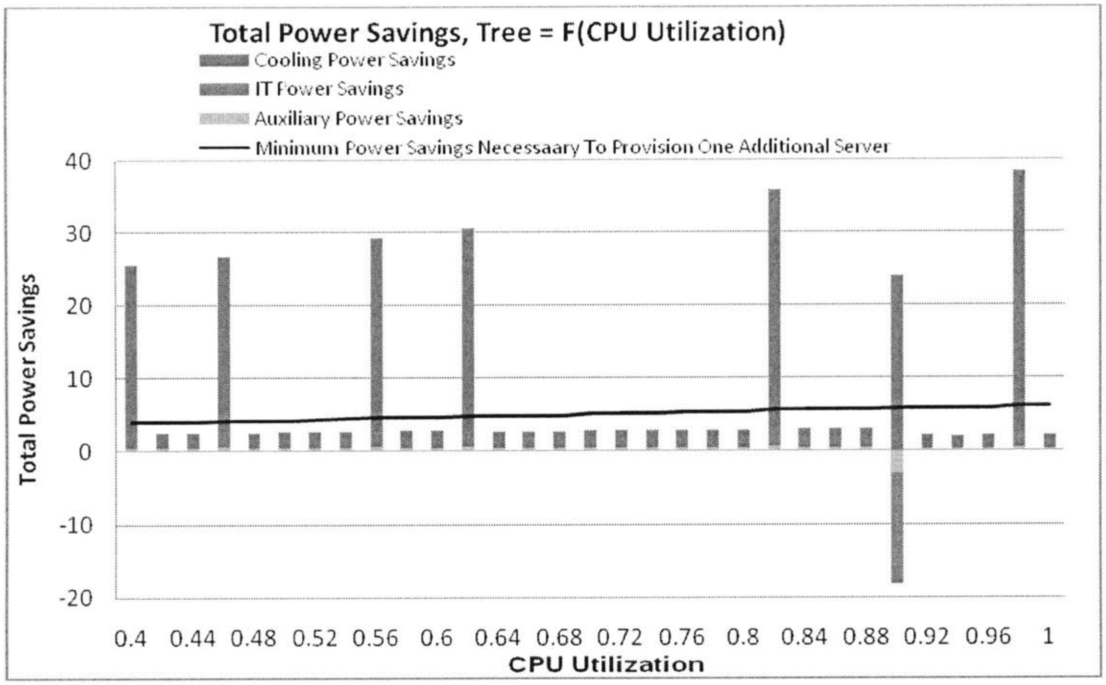

Figure 41: Total Power Savings as Function of CPU Utilization (Tree)

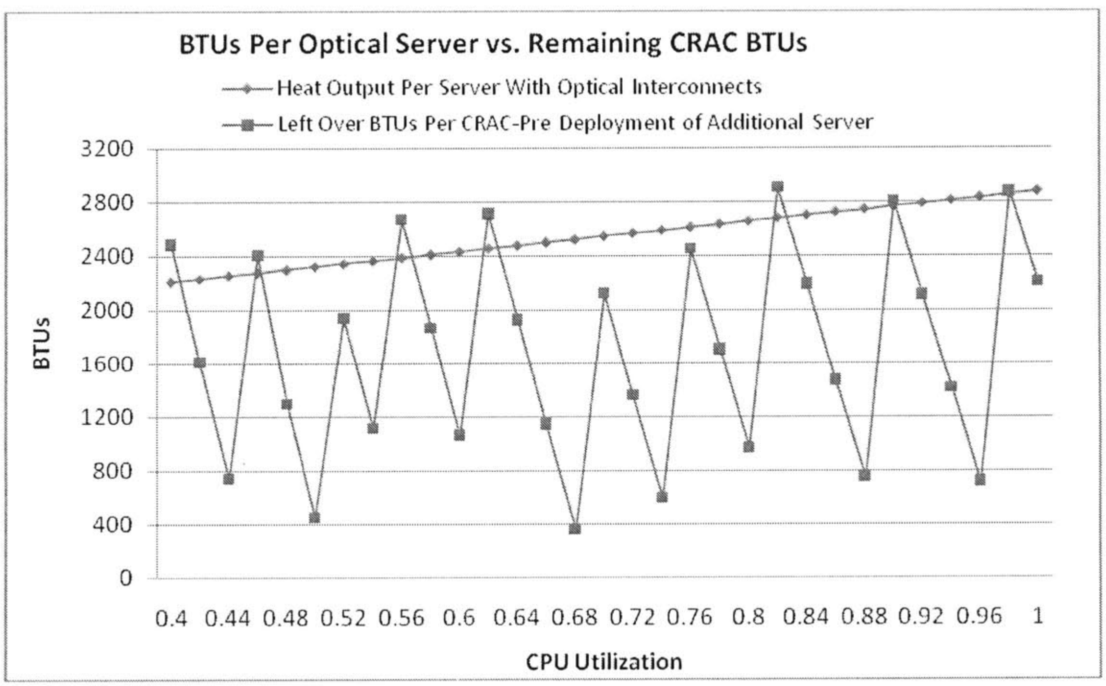

Figure 42: BTUs Per Server (Optical) vs. Remaining CRAC BTUs

Table 12: CPU Utilization at [0.40, 0.46, 0.56, 0.62, 0.82, 0.90, 0.98]

CPU Utilization	Servers Per CRAC	CRAC Capacity – Copper	CRAC Capacity – Optical
40%	43	275.03	2482.11
46%	42	120.87	2404.67
56%	40	284.82	2669.86
62%	39	256.02	2707.78
82%	36	228.26	2902.46
90%	35	36.15	2799.32
98%	34	21.31	2873.46

As shown in Figure 42, the spikes in cooling power savings at these particular data points are attributed to the additional capacity of CRACs found in data centers with optical interconnects. Although the heat output of servers increase linearly as CPU utilization rise, CRACs deployed in data centers with optical interconnects are able to accommodate one additional server per CRAC beyond what could be provisioned in data centers outfitted with copper interconnects. Table 12 provides further clarification by comparing the remaining CRAC capacity of data centers with copper interconnects with those of optical interconnects where the number of servers per CRAC is the same in both cases.

The one exception to the general pattern observed in figures 39 through 41 is the lack of a spike in ROI at 90 percent CPU utilization where ROI remains below zero for mesh and tree networks. Although CRAC units in data centers using optical interconnects are able to accommodate one additional server per CRAC unit beyond what could be provisioned in data centers with copper interconnects, the resulting rise in cumulative cooling power savings per server are offset by increases in the cumulative IT power requirement per server. This rise in the cumulative IT power requirement per server is due to an increase in the number of network switches and storage devices.

At 90 percent CPU utilization, data centers with copper interconnects deploy 512 servers for mesh networks and 508 servers for tree networks. In both cases, the number of servers provisioned is at or near the maximum number of servers (in this case, 512) that can be used without requiring the deployment of additional network switches and storage devices. By adopting optical interconnects, data centers deploy additional servers which causes them to surpass that 512-server threshold. To accommodate those additional servers, the model first installs enough network switches and storage devices to provision one additional cluster (refer to 5.1.1 of this chapter). There are two important consequences of this methodology. First, the additional network switches and storage devices use some of the potential power available that could be used to provision more additional servers. So, instead of observing an increase of 12 or 13 additional servers as seen with other the other data points listed in Table 12, only 5 to 6 additional servers are provisioned at 90 percent CPU utilization. Second, because the cumulative IT power budget per server (as reported in the model) consists of a server's initial IT power requirement (refer to equation 4) and its share of the amount of IT power used by switches and storage devices in a data center, deploying optical interconnects in servers operating at 90 percent utilization result in a marked increase in the cumulative IT power budget per server.

5.3.5 Network Utilization

Figure 43 displays how the ROI of optical interconnect adoption responds as the utilization of each network switch increases from 70 percent to 100 percent. As stated in

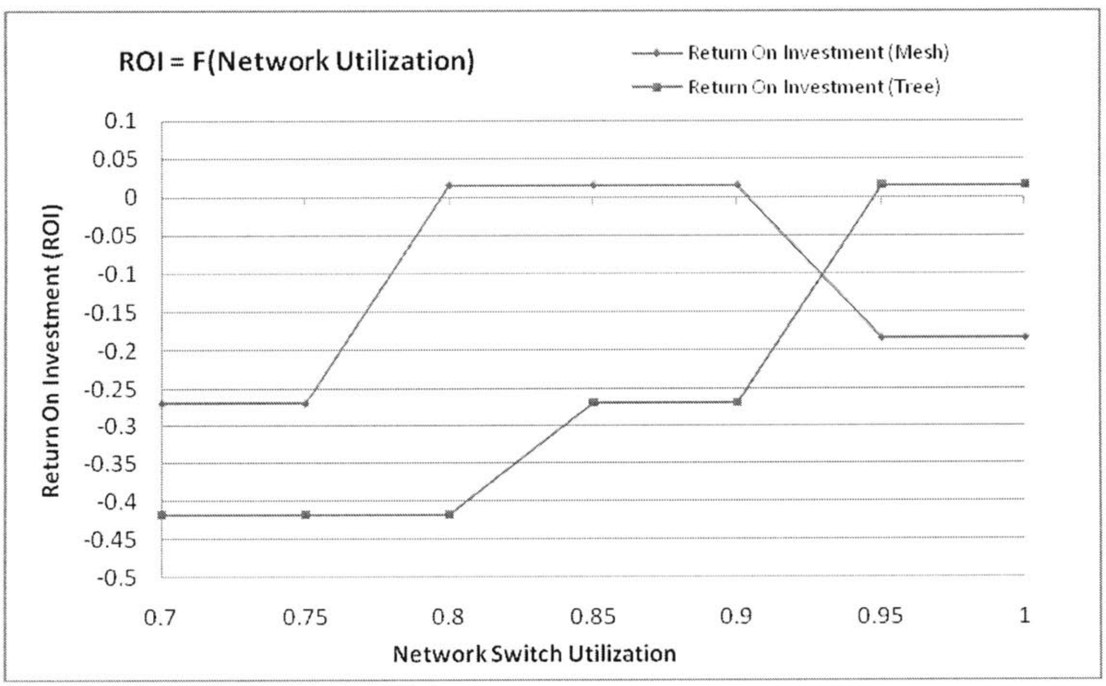

Figure 43: ROI as a function of Network Switch Utilization [0.7, 1.0]

the above discussion of the base case scenario, network switches are usually installed with most, if not all, of their ports fully utilized. Thus, in the base case scenario, network switch utilization is set at 100 percent.

As shown in Figure 43, ROI increases as network switch utilization rises. For mesh networks, ROI rises as network switch utilization increase; ROI then plateaus at 0.02 between 80 percent and 95 percent. Beyond 95 percent, ROI falls to -0.18. For tree networks, ROI increases as network switch utilization rises in a step-wise fashion. The ROI of optical interconnect deployment in tree networks reaches a maximum at 0.02.

The general pattern described in Figure 43 is driven by the decrease in investment costs and the rise of UPS cost savings per server. The decrease in per-server investment costs is expected; as the number of ports per switch increase, the number of network switches and storage devices required to deploy a cluster drops. Furthermore, the reduction in the amount of enabling equipment required per cluster increases the amount of power available to provision the total number of servers in a data center grows. As the number of enabling equipment decreases and the total number of installed servers grow, the costs to install individual servers decrease.

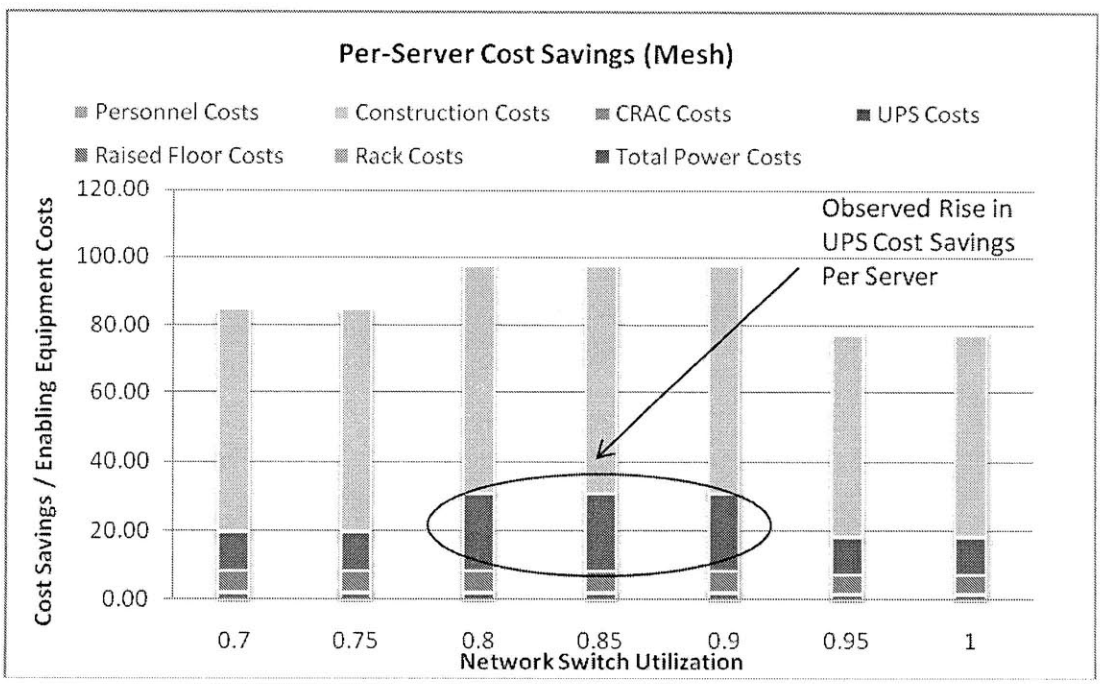

Figure 44: Break Down of Per-Server Cost Savings (Mesh)

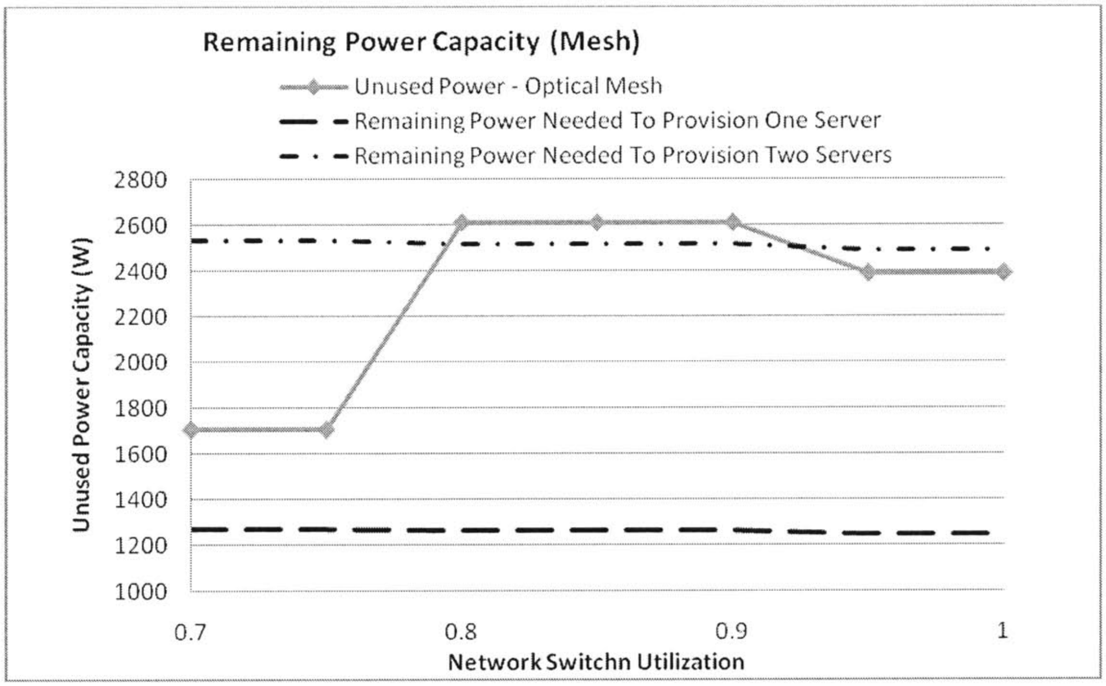

Figure 45: Remaining Power Capacity (Mesh)

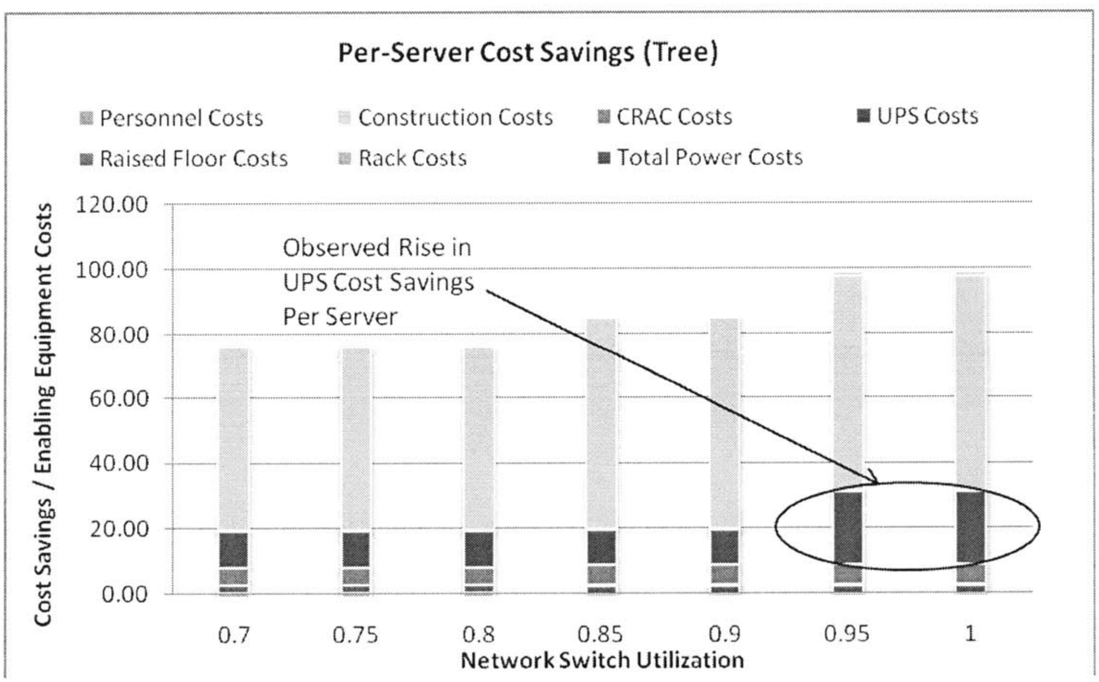

Figure 46: Break Down of Per-Server Cost Savings (Tree)

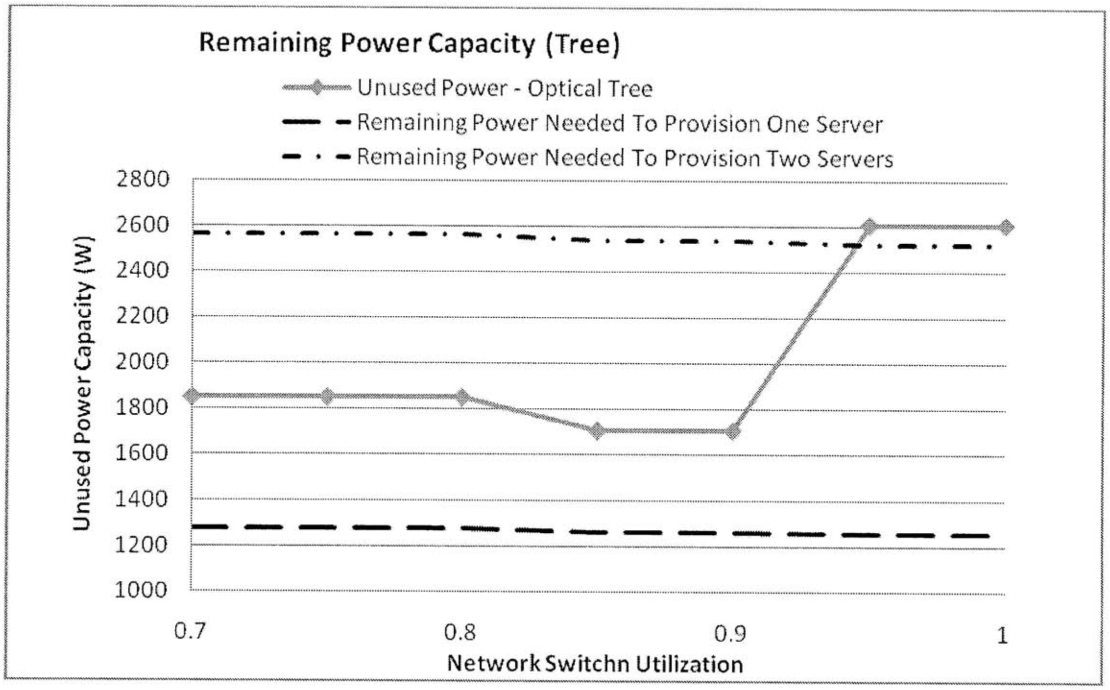

Figure 47: Remaining Power Capacity (Tree)

With respect to the rise in per-server savings in UPS costs observed in Figure 44 and Figure 46, the increase is due to a slight spike in the additional number of servers installed with optical interconnects rather than with copper interconnects. In each case where a rise in per-server UPS cost savings is observed, two additional servers are installed when optical interconnects are used; in all other cases, using optical interconnects result in only one server being installed. As shown in Figure 45 and Figure 47, the extra two servers realized between 80 and 90 percent CPU utilization in mesh networks and between 95 and 100 percent CPU utilization in tree networks is caused by the availability of additional power capacity to support two servers rather than just one.

5.4 Key Chapter Findings

This chapter has examined the output of the cost model, presented the results of the model when the default settings are entered for the input variables, and the sensitivity analysis of the input variables. The above analysis has shown that the deployment of optical interconnects could be slightly profitable in datacenters with tree networks, but not in mesh networks. As the results suggest, optical interconnect deployment is one way to optimize a data center's fixed resources and install several more servers than would otherwise be supported using copper interconnects. Because mesh networks are efficient by design, the degree to which optical interconnect adoption would further optimize a data center with mesh networks is limited as demonstrated in the above analysis. On the other hand, data centers with tree networks are less efficient because they require more network switches and storage devices to provision one cluster. Adopting optical interconnects in data centers with tree networks may result in a positive ROI.

From the above sensitivity analysis, transceiver power is the most influential characteristic in ROI analysis; thus confirming the initial conclusions from the qualitative surveys discussed in chapter 2. Unlike transceiver height, cable radius, or cable bend radius, changing the power required per interconnect has an effect that is seen throughout the model. Transceiver height has no effect on ROI; lowering transceiver height does not increase enough additional capacity in CRACs to accommodate more servers, which is crucial to the lowering of the cumulative cooling power requirement per server. Cable bend radius and cable radius both have a limited effect on ROI. Because both only affect the physical footprint of racks installed in data centers, savings are limited to raised floor, land, and construction costs. Changes in cable bend radius and cable radius do not impact the decision of how many servers should be installed; and thus leaving other per-server costs unchanged.

However, lowering the power required for interconnects (via the choice to install optical interconnects rather than copper interconnects) only impacts ROI indirectly. Reducing power cost, alone, will not produce enough funds to recoup initial investment. Instead, lowering power per interconnect influences optical interconnect adoption by reducing the cumulative IT and cooling power requirements per server. Particularly, reductions in cooling power requirements per server are critical to increasing the total number of servers in a data center, and by extension, achieving a positive ROI. As the sensitivity analysis results (specifically the discussion in 5.2.2 and 5.2.4), no additional servers can be installed without substantial savings in the cumulative cooling power requirement per server.

The combined savings in cumulative IT and cooling power requirement per server reduces a cluster's overall power footprint. In turn, the pool of remaining power capacity available to provision additional equipment following the initial deployment of one or more full clusters increases. In turn, the rise in the number of installed servers realized with the installation of optical interconnects further cause a reduction in the per-server costs of large, fixed capital expenditures such as UPS, CRAC, and construction costs. Together with the additional profit realized with the installation of extra servers, the savings in these per-server costs may help to yield a positive ROI for optical interconnect adoption in some cases.

Chapter 6 – Policy Discussion and Conclusions

Although the preceding chapter has shown that installation of optical interconnects could be profitable in some data centers without additional financial incentives, a total cost of ownership analysis would be incomplete without some discussion of current government and private energy saving financial incentive programs. This chapter briefly examines what government or power-utility-backed programs are relevant and available to help finance projects to install optical interconnects in data centers. Afterwards, this chapter presents key research conclusions and suggestions for future work.

6.1 Government Incentives

In its report on server and data center energy efficiency delivered to Congress in April 2007, the EPA conducted a comprehensive review of Federal and state laws outlining commercial tax incentives to promote energy efficiency. Only Federal statutes appear to provide for commercial tax credits that may be applicable to data center. Of the state programs examined in the report [35], nearly all referenced the U.S. Green Building Council's Leadership in Energy and Environmental Design (LEED) standards. Because the LEED promulgates standards requiring the use of energy-efficient building materials and do not include a standard for data center energy efficiency, it appears unlikely that a project to deploy optical interconnect enabled servers, storage devices, and network switches would qualify state tax credits or exemptions [68].

6.1.1 Energy Policy Act of 2005

Under current federal law, the Energy Policy Act of 2005[36] allows for a number of several "tax incentives for energy-efficient buildings, equipment and vehicles."[68] Section 1331 of the statute, which outlines tax incentives for commercial buildings, has

[35] The EPA report looked at state tax incentives programs offered in Maryland, New York, Nevada, and Oregon [28, pg. 90].
[36] President Obama signed the America Recovery and Reinvestment Act (otherwise known as the stimulus bill) in 2009; however, this bill did not alter or add to the incentive structure for commercial buildings established by the Energy Policy Act of 2005.

the most potential to be applied to data center projects. This section of the Energy Policy Act provides for tax credits of $1.80 per improved square foot to owners of new or existing buildings who construct or reconstruct their buildings to reduce the building's total heating, cooling, ventilation, water heating and interior lighting energy density by 50 percent or more compared to the ASHRAE Standard 90.1-2001 reference building [68]. For reductions in a building's total power density that fail to meet the "50 percent or more" threshold but manage to reduce total power density by 16.7 percent compared to the appropriate reference building, section 1331 allows for partial deductions of $0.60 per square foot [68].

Table 13: ASHRAE Standard 90.1-2001 [69]

Reference Building	2001	2004
Automotive Facility	1.5 W/ft^2	0.9 W/ft^2
Convention Center	1.4 W/ft^2	1.2 W/ft^2
Gymnasium	1.7 W/ft^2	1.1 W/ft^2
Hotel	1.7 W/ft^2	1.0 W/ft^2
Manufacturing Facility	2.2 W/ft^2	1.3 W/ft^2
Religious Building	2.2 W/ft^2	1.3 W/ft^2
Retail	1.9 W/ft^2	1.2 W/ft^2

Table 14: Tax Credit as Dollars Per Gained Energy [70, 71]

Measure Qualifying for Tax Credits	Tax Credit ($ Per Device)	Energy Saved (kWh / Yr)	$ Per kWh
Energy-efficient Exterior Doors and Windows	$ 200.00	1389	$0.14
Energy-efficient Electric Heat Pumps	$ 300.00	1400	$0.21
Energy-efficient Insulation Materials	$ 500.00	1980	$0.25
Energy-efficient Hot-Water Heaters	$ 150.00	556	$0.27
Electricity-generating Solar Panels	$ 2000.00	3400	$0.59
Metal Roofs Coated With Energy Star Hear-Reduction Pigments	$ 500.00	615	$0.81
Optical Interconnects (at $0.60 per ft^2)	$ 2.17	262.05	$0.01
Optical Interconnects (at $1.80 per ft^2)	$ 6.50	262.05	$0.02

Table 13 displays power density values (in terms of watts per square foot) for key reference building types defined in ASHRAE standard 90.1-2001 used by section 1331 of the Energy Policy Act of 2005.[37] Table 14 shows how the section 1331 tax incentives that may potentially be applied to optical interconnects compare with other types of projects under the Energy Policy Act of 2005. On a tax-dollars-per-kWh basis, tax

[37] Although ASHRAE has updated the values for each reference building listed in its 90.1 standard in 2004 (shown also in Table 13), the Federal statute continues to use the earlier 2001 values.

savings resulting from projects installing only optical interconnects lag far behind those tax credits that other types of energy efficiency projects may earn.

6.1.2 Challenges to Applying Section 1331

Applying for tax credits under section 1331 for projects to install new IT equipment with optical interconnects may be difficult. First, while the 90.1 standard defines the ideal amount of power consumed per area for several types of reference buildings, the standard does not provide corresponding values for data center facilities.[38] The exclusion of data centers from the list of reference building types leaves IT managers without guidance on whether a facility's improved power density resulting from multiple energy-efficiency projects will qualify for federal tax credits. Second, and perhaps more important, the threshold to qualify for the tax credits under section 1331 may be too high. As shown in Table 15, installing optical interconnects in a data center will only result in 2.16 percent and 2.25 percent reductions in facility power density for mesh and tree

Table 15: Installing Optical Interconnects – Change in Power Density

Model Output Variables	Copper Mesh	Optical Mesh	Copper Tree	Optical Tree
Power Density (Base Scenario)	355.32	373.76	347.3	373.66
Percentage (Base Scenario)	NA	+5.19	NA	+7.6
Power Density (Not accounting for base scenario conditions)	355.32	347.65	347.3	339.50
Percentage (Not accounting for base scenario conditions)	NA	-2.16	NA	-2.25

networks respectively. It is important to note that those reductions will only be observed in cases where: (1) deployment is occurring in an existing building so that the area is fixed, and (2) power capacity saved from installing optical interconnects is not being used to deploy additional servers. However, if the base scenario conditions – facility size depends on the number of servers installed and the number of additional servers deployed depends on freed up power capacity – are present, then optical interconnects will result in increased power density values for data centers using either mesh or tree networks. In

[38] Recent phone interviews with ASHRAE representatives confirm that the ASHRAE 90.1-2001 standard does not include data centers as a reference building nor does it provide power per square foot values.

either case, it does not appear that a project to solely install optical interconnects in a data center would qualify for the section 1331 tax incentives.

6.2 Utility Company Incentives

Another source of financial incentives that data center managers would normally look to help defer the costs of upgrades increasing energy-efficiency are power utility companies. For power utility companies, it is more cost-effective to utilities to "free up electricity through demand-side management (DSM) activities rather than to increase generating capacity." [68]. For the 2006 fiscal year, "power utilities companies spent more than $2.6 billion for energy-efficiency and demand-response programs." [68]. During the course of this study, several of the largest utility companies were surveyed to determine what kinds of energy-efficiency programs have been implemented[39]. Unfortunately, programs to deploy optical interconnects in data centers may not qualify for these incentive programs. Because most energy saving incentive programs surveyed function under relatively small annual budgets, financial rebates are provided mostly for projects that yield large savings in a building's overall power consumption such as retrofitting a facility with energy-efficient HVAC systems. While the use of optical interconnects may encourage the deployment of additional servers, it may not provide reductions in building-wide power consumption; in fact, the deployment of additional servers caused by per-server power savings may lead to increases in building-wide power consumption (refer to Table 10). In this environment, it is unlikely that optical interconnects will qualify for financial incentives provided by utility companies.

6.3 Policy Recommendations

From the above discussion, it is clear that current financial incentive plans offered by the Federal government and utility companies do not aid in the promotion of optical interconnects. Many financial incentive programs are structured to provide funding

[39] Refer to Energy Rebate Policies: Pacific Gas and Electric [72], Nevada Energy [73], CPS Energy [74], Texas-New Mexico Power Company [75], Austin Energy [76].

based on the amount of energy saved at the building level. While incentive structures aimed at rewarding building-wide improvements may be effective for most other types of commercial buildings, they may not be effective at promoting the adoption of energy-efficient computing equipment in data centers. As seen in the preceding sections of this chapter, extrapolating incentives at the device level from programs aimed at realizing large building-wide energy improvements is not effective. Because the extrapolated incentives are too small or the threshold to obtain the incentives are too high, installing high-efficiency computing equipment as a potential solution to driving energy usage is overlooked.

In order to bring greater awareness to high-efficiency computing equipment (and by extension to optical interconnects) as a potential solution to rising data center power consumption, government and private organizations should create effective financial incentives at the device-level. Providing financial incentives to help defer the purchase and installation costs of high-efficiency computing equipment may be particularly effective because it would work to address a problem that data center managers commonly try to answer: How to maximize current power capacity in an environment of constantly increasing demand for bandwidth or computing resources. Unlike tax credits for building-wide energy improvements, financial incentives for high-efficiency servers, network switches, and storage devices do not force IT managers to forego achieving maximum utilization in order to achieve a required power density level; instead, IT managers could use tax credits to increase a data center's overall computing efficiency without operating under additional power restraints. As shown in the base scenario and sensitivity results described in the preceding chapter, using IT equipment with optical interconnects (as one example of high-efficiency computing equipment) could save enough power to deploy additional servers without increasing provisioned power capacity in some circumstances.

Successful deployment of an incentive program for energy efficient computing equipment requires the development of energy measuring and labeling protocols. The EPA and the Standard Performance Evaluation Corporation (SPEC) are leading this effort; both organizations have developed initial energy measurement protocols that will demonstrate the power consumed by servers at different computing loads. However, the

EPA and SPEC protocols remain in their nascent stages and specify energy performance for servers only. To further development of these protocols and eventually create comprehensive energy measurement standards, equipment manufacturers need to cooperate and share equipment performance data with the EPA and SPEC.

6.4 Research Summary and Thesis Contribution

Although optical interconnects generally offer improved performance compared to copper interconnects, they have to achieve greater rates of adoption for several potential applications. Of the several applications examined in the second chapter of this thesis, HPC box-to-box interconnects is a market space that represents the greatest potential for adoption. As discussed in chapter three, data centers – the primary consumer market of HPC box-to-box interconnects – are under constant pressure to increase efficiency from the increased demand for bandwidth and the use of more distributed architectures while operating under fixed power capacity that can often be very expensive to expand. Despite this dynamic in the data center market space, optical interconnects are generally not viewed as potential solution; many analysts disqualify optical interconnects from consideration upon comparing the initial device costs for optical and copper interconnects.

However, the use of first-cost analysis may not always produce enough insight to allow IT managers to make well-informed investment decisions. For complex organizations such as data centers, it is important for IT managers and executives to know fully how investments in emerging technologies will affect their fixed and operational costs after the completion of installation. In chapter four, this thesis proposed the use of a Total Cost of Ownership framework to calculate ROI, as a more suitable method to discern whether optical interconnects should be adopted in a data center. Specifically, this thesis has sought to determine if the deployment of optical interconnects in data centers could be a profitable venture.

Base scenario and sensitivity results of the framework shows that optical interconnects could produce a profitable return on investment despite having higher initial costs per device than copper interconnects. Although the direct IT energy saved per interconnect maybe small when replacing copper interconnects with optical interconnects, per-server savings in cooling power realized after installation are great enough to lead to the additional deployment of additional servers in data centers deploying either mesh or tree networks. Model results also show that the deployment of optical interconnects could pay for itself within three years (the average lifetime of servers, network switches, and storage devices). And although model results show that optical interconnect deployment could be profitable without government-funded subsidies in some cases, programs to install optical interconnects would not qualify for incentive programs offered by the federal government and power utility agencies should additional financial incentives become necessary.

6.5 Suggestions for Future Work

Although this thesis research has explored some questions concerning the potential area for optical interconnect deployment, several interesting question arise as a result. First, further academic study may include how the server workload may affect a data center manager's decision to deploy optical interconnects. Since not all data centers perform the same functions or use the same types of information, workloads and CPU utilizations may vary wildly across different data centers. While the model provides a general framework capable of producing general ROI trends based on range of potential inputs for CPU utilization, the use of actual, real-time data for CPU utilization from multiple types of data centers will provide additional clarity in the decision to deploy optical interconnects. Depending on the nature and type of workloads, it maybe more profitable for some data centers to deploy optical interconnects than others.

Second, one could further study how optical interconnects would affect gains in revenue. The model uses the number of deployed server equipment to track changes in overall computing efficiency; the model uses power savings realized from the adoption of optical interconnects to deploy additional servers. However, it may be helpful to study how the deployment of optical interconnects affect the computing efficiency of individual servers. Coupled with the additional workload information identified above, one can calculate how the number of deployed servers change with the optical interconnects given a constant value for data center workload instead of an initial building-wide power constraint as done in the model.

A third area for future study would be to explore how changes in the energy required to generate the air flow necessary to achieve a required heat transfer rate affects the profitability of optical interconnect adoption. Given the sensible cooling capacity tables (refer to appendix), it may be likely that a system would be designed to work at the most efficient point (85C for that table). This would mean that the inflow and outflow temperatures could be the same for a device that use either copper or optical interconnects. For that case, improved airflow achieved using devices with optical interconnects could reduce energy use by requiring a smaller airflow velocity. Unlike the current model discussed in the preceding chapters, one would need to assume a given Delta-T and solve for air velocity. In order to solve for air velocity, one would need to gather information (via additional interviews with data center managers and CRAC manufacturers) regarding the power required to generate various amounts of airflow.

Appendix A – Qualitative Interview Questionnaires

A.1 Phase I Cross Market TWG Questionnaire

Goal: To map out the possible market space and driving product characteristics for optoelectronic technology.

Intermediate objectives:
(1) to develop an understanding of the current challenges facing firms in video, automotive, computer, and defense applications, and the role optoelectronics may be able to play in addressing those challenges;
(2) to develop an understanding of existing perceptions (true, false, or non-existent) of optoelectronics and of barriers to optoelectronic adoption within those firms; and
(3) to develop an understanding of the potential market demand for optoelectronics if it can meet these challenges and overcome existing barriers to adoption.

Understanding Current Challenges

1) Can you describe the major technical challenges facing your firm in the immediate future? (2 yrs out, 5 yrs out, 10 yrs out)
   a) Are any of these in...
      i) Communications?
      ii) Data transfer?
      iii) Connectivity?
   b) Can you quantify these challenges (in terms of speeds, feeds, density, packaging, reliability, and/or costs)? Alternatively, are these challenges in a different area (e.g. density, power)?

2) How do you think about cost in your product space? (metrics, scope)
   a) Is there a commonly understood cost-performance (market demand) or cost quality (manufacturing capability) trade-off?
   b) What to do you see as the major cost drivers, whether in manufacturing, test, life cycle, or otherwise?

Understanding Current Applications

3) Does your firm currently make use of optoelectronic technology (OE) in any major products?
   a) If so, can you list the applications in which you currently use optoelectronics?
   b) What function does OE provide?
   c) How many OE components are used per product?
   d) How many OE components do you purchase annually?
   e) What characteristics of the OE component make it attractive for your application?

4) Are there any optoelectronic technologies you are in the process of implementing, or any optoelectronic technologies you have tried and failed to implement? What is/was your experience? What challenges have you faced?

Understanding the Perception of Optoelectronics within Firms

5) What is your perception of optoelectronic technology?
   a) Is this view shared broadly within your firm?
   b) What do you see as the major advantages and disadvantages of optoelectronic technology?

6) What do you see as the pros and cons of applying optoelectronic components in your firm?
   a) Is this view shared broadly within your firm?
   b) What do you see as the major advantages and disadvantages of optoelectronic components?

Projecting Future Applications

7) Do you foresee the adoption of OE technologies in your firm's products?
   a) If no, why not?
   b) If yes, in what applications would you foresee considering the adoption of OE?
      i) What barriers (technical, economic or otherwise) would have to be overcome for optoelectronic technology to be adopted?
      ii) Is there a price point at which would OE be preferred over incumbent technologies? What other requirements would have to be met?
      iii) If optoelectronic technology could meet these requirements, how many OE components would be needed per product? What would you estimate to be the annual demand?
   c) Do you see your answer differing for 2, 5, and 10 years out?

A.1.1 Automotive Response 1 (October 2006)

Introduction

The company is major producer of American cars and trucks and automobile financing. The company's business involves three types of transactions: sell cars/trucks, sell spare parts, and renting money (?). The company has been around for approximately 100 years and officials there have noticed that externalities have suddenly changed and continue to change. The company has a product research/development organization that is comprised of 12 separate divisions. The interviewee currently runs the division overseeing software, electronics and communications in cars and trucks. The interviewee was originally recruited to understand how to enable the stand-alone vehicle to communicate with the outside world.

Challenges

-How does one blend the two experiences: Provide a stable platform of comfort, dependability, and durability while accommodating change and personalization. The company is attempting to blend these two experiences and install the right interfaces.

- Electronic: Life cycle management issues
- Materials: How does one inherently design the material that when it reaches the end of life cycle it can be recycled and used again?

- Transportation: Large Scale Atmospheric Sciences:
    - Depletion of fossil fuels
    - Hydrogen Based Economy / Alternative Fuels
    - Reduction of CO_2 emissions
- Improving information and control coordination with a larger crowd/system

Cost

No comment on Cost, Cost metrics, or Cost Drivers.

Possible Applications for Optics

-Lane Detection

-Beam Steering (possible disruptive change from filament based optics to solid state optics)

-Accommodate more features (fiber optics is a possible solution for the transfer of big amounts of high speed data across multiple nodes within the car).

-Larger Intelligent Transportation System

-Light Weighting (?): In 1950s, there was about 50 m of cable in an average high-end vehicle. Now cars have the 2 km of copper wire that weighs about 34-40 kg. He is not sure how much this will continue to grow but considers that alternatives such as wireless/optoelectronic technologies develop could slow the growth of cable in the car.

Perception / Outlook

No comment on perception or outlook.

A.1.2 Automotive Response 2 (November 2006)

Background
Interviewee currently works for Company A2. Company A2 manufactures auto systems that interconnect media components (MOST). MOST is currently implemented in 40 different vehicle models and is produced at a rate of about 10 million nodes/year.
Company A2 bought Company X (which founded MOST Consortium with BMW, Diamlet Chrysler, et. al to develop a fiber wire physical layer and accompanying software). He became involved with Company X in 2000 to work on the company's business development strategies. MOST implementation in cars began in 2001. Currently, he is the representative and technical liaison to the MOST Cooperation. Mostly its member companies run the MOST Cooperation.

MOST was originally developed to be fiber-optic network
-Current Speed – 25 mbps
-Research done to achieve 150 mbps by fiber optics
-Toyota wanted the 50 mbps with copper

Challenges
-Finding low cost methods to doing O/E to E/O conversion. A big part of the cost is O/E connectors. Relative to other parts of the system, these connectors can be quite expensive. The interviewee noted that these can be as much as $5.00.
-Fault Tolerance: EMI standards are a lot higher b/c of noise and close proximity with other devices. EMI is a major reason for the big push towards fiber optic networks.
LED Speeds and Light Source Issues: Currently experimenting with VCSEL and RCLED, but they are expensive and have experienced some problems concerning the robustness over a wide range of temperature changes.
-Other Generic Challenges: Faster, Cheaper, Smaller.

Cost Drivers
-MOST uses a lot of Flash Memory and the fabrication process can be quite expensive.
-Cost of Silicon, specifically the cost of cost of verification and qualification can be high.

Perception
-Fiber Optics is seen as a useful technology but remains largely a cost-driven issue.

Outlook
-Currently, 3 million cars have this technology, but interviewee believes that this will grow significantly in 5 years. He references Frost/Sullivan report that indicates that 27% of all vehicles will have this technology by 2012. He estimates that if cost can be cut in half then volume production will double. Other signs of growth includes reports of Toyota's decision to implement MOST in 2008 starting with the Lexus and Hyundai's decision to implement as well (Hansen Report, Oct. 2006).

A.1.3 Telecomm/Computer Response 1 (October 2006)

Major Issues:
- High Speed Data Communications

  As the speed of the processors and the size of the system increase, more high Speed interfaces have to be incorporated. He's mostly focused inside the Data center problems (supercomputing / main frame / clusters / server systems).

- How to cost effectively balance the speed of interfacing with the cost of the system.
  - The cost of the laser is ¼ the cost of the links
  - There's big difficultly in testing
    - Increased cost from yield and testing
    - Testing can cost b/t 20-25 % of delivering a link

- Lasers tend to fail; providing error correction for the DRAM is a major issue.

- Plugging optics can be disruptive to currently used architecture:
  - With optics, you get a different system design (Packaging, for example, is much different).
  - One solution: Provide pluggable modules for the optoelectronic technologies so that optoelectronics and electronics remain separate. This incurs big cost.
  - OR, if you decide to integrate optoelectronics, how do you do that in a way that does not incur much costs for customers that don't want it.
- Once one commits to the using optical interconnects, is there a chicken switch (i.e., can the customer easily revert back to electronic if he wants to)?

Interviewee expects to be working on optical interconnectivity in decades. The firm-wide perspective depends mostly on time; arguments are over "when" not "if." He predicts that the time range of implementation is 15 years out at the far end. One important note that interviewee makes is that there's a lot of flexibility for the industry to map to cheapest computing technology available. If the cost of optics stays higher than what we have now, then the industry will continue to re-architect to match the cheapest computer hardware.

A.1.4 Telecomm/Computer Response 2 (November 2006)

Background

Company 2 builds computer systems and specializes in hardware and software. It is an open systems company and depends heavily on industry standard components and interfaces. As a result, the company has very strong relationships with industry partners and suppliers that provide industry standard components. Much of the value added in its products lies in the strong software integration with its hardware systems. The interviewee is the director of the systems group. He became involved with this group when he began work on a DARPA HPCS program where he had written proposals that brought in phase-1 and phase-2 of the project. His team also drives a lot of the HPC deals in which they partner with a customer and design specific products to suit their needs.

Challenges

-I/O requires too much power
-In order to meet future requirements, systems need to become denser
-Bandwidth requirements continue to increase
These three are areas in which the company continues to hit head on as they with multiyear long-term system deployment in terms of what they need to accomplish.

Cost

-The interviewee identifies **Dollar per floating point operation (FLOP)** as an important cost metric.
-The interviewee identifies **cost of memory** as a cost driver. It used to be the case that memory was a small fraction of the cost of the system but it continues to grow. Why is that? A possible answer maybe that as more of the systems are based on high volume processors and industry standard components, the cost of systems is dropping dramatically (if you consider the cost of processor in terms of dollars per flop). The use of the open system and industry standard components can drive the cost down for many elements or aspects of the computer system. In doing so, memory continues to increase in terms of fractional value of computer systems.
-Validation and testing of components continues to be expensive.

Perception

-Optoelectronic components have the promise providing 1 or all of the following: more bandwidth, lower power, and greater density
-However, the company does not see devices with these components at level of mass manufacturing; it just does not beat copper at this stage.
-The company is a "taking a wait and see" approach.

Outlook

-HPC is the place where earliest adoption can be seen in the computer systems market segment. In this segment, the interviewee sees many customers that are willing to pay a mall premium (as much as 10%) for an optical solution that will provide some tangible performance benefit. He sees this happening within 2-5 years.

A.1.5 Telecomm/Computer Response 3 (October 2006)

Understanding Current Challenges

1. Can you describe the major technical challenges facing your firm in the immediate future? (2 yrs out, 5 yrs out, 10 yrs out)
    a. Are any of these in...
        i. Communications?
        ii. Data transfer?

Company 4 has a very broad product portfolio, ranging from printing and imaging equipment, personal computing products and networking products to a full line of Enterprise computing and data storage products. We are constantly looking for opportunities to improve data transfer across the product portfolio. Sometimes improvement comes in the form of absolute performance, more often it comes in the form of price/performance.

        iii. Connectivity?
    b. Can you quantify these challenges?

For some products, the major challenge is to provide data transfer using industry-standard protocols and bit rates (PCIe, Infiniband, FC, etc) at lower cost, higher density, and across longer distances than those afforded by electrical signaling. For other products, the major challenge is to provide performance enhancements for proprietary interconnects. In this area, cost-effective 10-20 Gbps/channel links are interesting. For still other products, 100Gb Ethernet appears promising.

In general, all product groups will agree that reliability is absolutely essential, but requirements and metrics vary across product lines as a function of the requirements of the industries those products address.

Power consumption on high-speed electrical links runs from over 1W/bit to under 100mW/bit. Optical data transfer mechanisms that can enable lower total power consumption at the system and device levels while enabling comparable or higher bit rates will be favored.

2. How do you think about cost in your product space? (metrics, scope...)
    a. Is there a commonly understood cost-performance (market demand) or cost-quality (manufacturing capability) trade-off?

Cost is a consideration for all products, with $/Gb/S being the common metric for comparing costs of short optical paths to electrical paths, and $/Gb/S/M the metric for comparisons on longer paths.

    b. What to do you see as the major cost drivers, whether in manufacturing, test, life cycle, or otherwise?

The major cost driver is per-unit cost. With some divisions shipping a server every 14 seconds, cost/volume is a tremendous driver. While some product line are somewhat less cost sensitive than others, in general the industry is very cost conscious.

Understanding Current Applications

3. Does your firm currently make use of optoelectronic technology (OE) in any major products?
    a. If so, can you list the applications in which you currently use optoelectronics?

Industry-standard optoelectronics are used extensively in storage networks.

    b. What function does OE provide?

Data transfer.

c. How many OE components are used per product?

This varies by configuration, with some configurations requiring no OE components, while others require many. Configuration size and distance between elements dictates whether OE or electrical data transfer solutions are applied.

d. How many OE components do you purchase annually?
e. What characteristics of the OE component make it attractive for your application?

The ability of an OE link to traverse tens of meters without boosting signals through repeaters or patch panels makes it cost-effective for certain paths.

4. Are there any optoelectronic technologies you are in the process of implementing, or any optoelectronic technologies you have tried and failed to implement? What is/was your experience? What challenges have you faced?

Very short haul data transfer optical links under 10M, in both inter-system and intra-system applications. Links of this type have not enjoyed sufficient cost reduction to compete effectively against SerDes electrical solutions. Industry-standard (MSA) solutions have tended toward adaptations of telecom technology rather than application-specific instances. Costs need to come down, reliability needs to increase.

Understanding the Perception of Optoelectronics within Firms

5. What is your perception of optoelectronic technology?

On one hand, optical data transfer within computer systems has been the next great thing for the last twenty years. Always just out of reach due to cost and continued cost/performance improvements in electrical signaling. On the other hand, the relentless drive toward higher frequencies has reduced the cost crossover point from electrical to optical signaling from hundreds of meters to tens of meters. Appropriately designed optical links could meet the need, both in terms of cost and capability.

a. Is this view shared broadly within your firm?

In general, the firm is taking a wait and see attitude. "Show me the economic crossover and I will design in optics."

b. What do you see as the major advantages and disadvantages of optoelectronic technology?

Advantages: Higher frequency data transfer, longer distances at high frequency before needing repeaters, possibilities for direct optical chip-chip communications, reduced system weight, reduced cable volume, improved system airflow, possibilities of unique system architectures built upon the unique capabilities of optical networks.

Disadvantages: Cost, cost, cost. There are also concerns with reliability, particularly as optics are considered for deployment in systems having tens of thousands of interconnections, and the need for common form factors.

6. What do you see as the pros and cons of applying optoelectronic components in your firm?

If the cost problem is conquered and reliability and performance meet appropriate targets, there will be many pros and few cons. It will then become a matter of product-specific requirements.

a. Is this view shared broadly within your firm?

Yes

b. What do you see as the major advantages and disadvantages of optoelectronic components?

The vast majority of optoelectronic components available today are variants of devices designed for telecom applications "before the bust." They are not and will never be cost-competitive. Although they offer data transfer over longer distances at higher data rates, there are electrical solutions that meet the needs in a more cost-effective manner. When optical components are designed specifically for cost-effective inter-system and intra-system connections, they will be widely adopted.

Projecting Future Applications

7. Do you foresee the adoption of OE technologies in your firm's products?
   Yes
   a. If no, why not?
   b. If yes, in what applications would you foresee considering the adoption of OE?
      i. What barriers (technical, economic or otherwise) would have to be overcome for optoelectronic technology to be adopted?

Cost.

ii. Is there a price point at which would OE be preferred over incumbent technologies? What other requirements would have to be met?

The cost point is price-parity. Other requirements would be comparable power and space. For some applications, such as chip-chip transfer, this may be difficult, for other applications, such as system-system and interconnected subsystems, it would be relatively easy.

iii. If optoelectronic technology could meet these requirements, how many OE components would be needed per product? What would you estimate to be the annual demand?

We ship a server every 14 seconds...

c. Do you see your answer differing for 2, 5, and 10 years out?

Yes

A.1.6 Telecomm/Computer Response 4 (November 2006)

Background
Company 5 makes supercomputers. One of their focuses is the interconnect channel speed where they try to maximize channel speed across the entire system. At the time of the interview, they had recently acquired a large contract from DARPA to deliver to market within 2008-2010 systems that will require channel speeds in excess of what is currently available on the market. This project, Project X, will produce the most advanced system on the market. Neither Interviewee 1 nor Interviewee 2 preferred to say what channel speed (or range) other than that is would be much greater than what is available in either optical or electrical markets. The value of this project to the customer would be higher processor to memory bandwidth and higher network bandwidth. Interviewee 1 is the technical project leader for Project X while Interviewee 2 is the optical engineer.

Challenges
-How far can we reach electrically?
-Where's the transition from the electrical to optical channel?
- How solutions to those two questions develop leads to getting aggressive channel rates. It in is Cray's business interest to continually develop channel rates and therefore, they have a very strong channel's group and developing expertise in optics.

Cost (Metrics/Tradeoffs in the Product Space)
-Always looking to maximize price performance (performance/price)
-Constrantly looking at ways to minimze cost in both optical and electrical channels
-Currently, optics is not very cost-effective; ***Company 5*** does use optics but not in their high channel systems. The cost of using optics is significantly higher per unit of data transfer compared to electrical by a factor of 5. ***Company 5*** will employ optics more when it drops to a cost competitive with copper.
-A major way to reduce cost is through large volume production. Many of the optics vendor are addressing larger volume markets; for example, some are addressing telecomm applications where the cost per unit of data transfer can be higher. As a consequence, these firms are competing in the data transfer market as well making it difficult for ***Company 5***. A further consequence is that optical vendors have not had strong motivation to work on low cost, effective, shorter transmission lengths and solutions.

Perception
-Optics is costly. Cost is a major barrier as management does not want to touch it except for special circumstances.
-In discussions with optics suppliers and users, they note that a lot of companies have adopted a conservative hedge position since the burst of an optics bubble
-Improving modulation rates and modular integration are areas in which they can see the cost per data transfer can be reduced.

Outlook
-At some point, there will be a crossover where optics might be used. The market is almost at that point but not yet.
-Within 5-10 years, they see optics going from cabinet-cabinet to within the cabinet.

A.1.7 Telecomm/Computer Response 5 (October 2006)

Background
In the telecommunications market, there is not much of a crossover point between transport and backplanes. In fact, copper remains the standard for backplane; they tried hybrid between copper and optics, but rejected it. For transport (line side), optics is used more often at 10G and beyond. For transport, the production volumes and prices depend very much on the product. At 1-2.5 G, the production volume is at 5-10 million. At 10 G, the production volume is at a couple hundred thousand while at 40 G, the production is low (<1000).

Challenges
-Disruption in Board Design caused by optoelectronics
-Newer technologies are unproven and have reliability issues
-As chip density increases, heat and temperature becomes a big concern
-Power Management; Operational costs are increasing for consumers
-Cost Development; B/c of increasing operational costs, they must lower price points.

Cost Drivers
-Optical Tech Supply Base is highly fragmented and unhealthy
  -The industry is unable to consolidate
  -For example, there are multiple manufacturers w/ multiple niche areas with the 10G space
  -Not many suppliers are in good conditioned
-Qualifying the devices from suppliers is costly.
  -Qualifications for testing include Telecordia Standards, temperature, vibrations, cycling.
  - Suppliers do their own reliability testing, but this company in telecomm does their own parametric testing

Perception
Overall, the decision to go optic is cost-driven and depends highly on enterprise, telecomm, and storage network sectors. Telecomm will drive toward optics for more bandwidth first. Enterprise is the biggest sector; here, it is believed that the 10 G will dominate. For the 40G space, one will continue to niche suppliers. In order to help push optics forward, one needs to really start looking at optics like semiconductors in terms of: automation, manufacturing, attaining better control, obtaining smaller packaging and lowering heat, lowering power, and increasing yield and quality.

Outlook
Within 2-3 years, the production volume for 10G will expand; no major expansion is predicted for the 40 G space because one could get the same performance by multiplexing (4) 10 G devices. The interviewee does not foresee an optical backplane within the next 5 years as the current technology works fine. Anything higher than 10G transport or optical backplane would be produced in low volumes. Interviewee suggests that a roadmap would focus more on the transceivers

A.2 Phase I Consumer Handheld Questionnaire

A.2.1 Consumer Handheld Interview

Mapping the Requirements of Current and Future Optical Interconnects:

Mobile Phone Interview

The Microphotonics Consortium at M.I.T. is conducting research on the adoption of optoelectronic components in consumer handhelds and mobile phones. As an element of this research, this interview attempts to ascertain the current and future requirements for optical and electrical interconnects used in mobile phones (basically what these components should do) and whether standards play a role in realizing those interconnects. With inputs from both the market side and supply side, we hope to create a 10-year strategic roadmap for optical interconnects in the Mobile Phone market segment.

Identifying Technological Needs & Barriers

1. What are the bandwidth needs in high-end mobile phones today? What are the applications driving these needs?

2. At what data rate within handhelds will optics (or other high-bandwidth interconnects) be needed? Which future applications will drive this need?

3. When do you expect this data rate to be needed? In what percentage of cell phones will copper no longer meet interconnect needs in 2008? 2010? 2015?

4. What performance or compliance requirements other than data rate might drive the adoption of optical interconnects over other potential solutions?

5. What barriers exist today for the use of optical interconnects in cell phones? flip phones?

Characterizing the Cost Competitiveness of Copper

6. At what cost ($ or %) of optical vs. electrical interconnects do you expect optics to start getting used in high-end mobiles?

7. Regarding multilayer air-gap copper-polyimide flex electrical connectors used for through-hinge connections in cell phones:
   a. What is the cost today ($ or % of cell cost)?
   b. How does the cost scale with bandwidth in this technology?
   c. When do you expect this solution to no longer be able to meet performance and/or compliance requirements?

Identifying Industry / Government Standards

8. Is there agreement in the industry that optical interconnects will be the approach of choice in cell phones? Flip Phones? Why or why not?
9. Do you expect convergence in the mobile industry on standard optical interconnect solutions? Are any MSAs planned?
10. What government regulations affect the selection of interconnects

A.2.2 Consumer Handheld Interview Responses

	Responses to Consumer Handheld Interview (Appendix A.1)	***Respondent 1 - Public Statement***	***Respondent 2 - Public Statement***	***Respondent 3 - Public Statement***
1A	What are the bandwidth needs in high-end mobile phones today?	700 Mbps	1 Gbps	
1B	What are the applications driving these needs?	Imaging applications	Imaging applications, connection to external devices	High resolution cameras (Samsung has today cells with 10 Megapixel cameras), high resolution displays, 3D sound
2A	What are the advantages of having optical interconnects in handhelds?	High bandwidth, EMI immunity, multiple media	High bandwidth, EMI immunity, distributed architectures	High bandwidth, EMI immunity
2A	At what data rate within handhelds will optics (or other high-bandwidth interconnects) be needed?	Today's data rates if solution ready and affordable	Today's data rates if solution ready and affordable	Today's data rates if solution ready and affordable
2B	Which future applications will drive this need?	Higher resolution cameras & displays	16 Megapixel images, HD video	High resolution cameras
3A	When do you expect this data rate to be needed?	1.25 Gbps in 2008	5 Gbps in 2012	
3B	In what percentage of cell phones will copper no longer meet interconnect needs in 2008? 2010? 2015?			
4	What barriers exist today for the use of optical interconnects in cell phones? flip phones?	System architectural compatibility, reliability, cost	Power consumption (need 1 mW/Gbps), cost (need $1/Gbps)	Power consumption, price, size
5	What performance or compliance requirements other than data rate might drive the adoption of optical interconnects over other potential solutions?	EMI immunity, multiple media, Specific Absorption Rate (SAR), Hearing Aid Compatibility (HAC)	EMI immunity, distributed architectures	EMI immunity
6	At what cost ($ or %) of optical vs. electrical interconnects do you expect optics to start getting used in high-end mobiles?	Goal for broad adoption is same price	Goal for broad adoption is same price	Goal for broad adoption is same price
7A	If optical interconnects are the solution of choice to resolve EMI, SAR, and HAC issues, when would the mobile industry start using this technology?	When cost and reliability are acceptable	When cost and power consumption are acceptable	When power consumption, price and size are acceptable
7B	What cost premium is acceptable over copper interconnects?	None (for broad adoption)	None (for broad adoption)	None (for broad adoption)
8A	Regarding multilayer air-gap copper-polyimide flex electrical connectors used for through-hinge connections in cell phones, what is the cost today ($ or % of cell cost)?			

	Responses to Consumer Handheld Interview (Appendix A.1)	*Respondent 1 - Public Statement*	*Respondent 2 - Public Statement*	*Respondent 3 - Public Statement*
8B	How does the cost scale with bandwidth in this technology?		Optics provides better price per bandwith than copper, it multiplies bitrates with only small increases in cost, but mobile applications are not seen to go soon to 10 Gbps or higher	
8C	When do you expect this solution to no longer be able to meet performance and/or compliance requirements?			
9	Is there agreement in the industry that optical interconnects will be the approach of choice in cell phones? Why or why not?	Yes, when cost and reliability targets are met. Because optics is best approach for solving EMI and SAR issues, providing high mechanical reliability, and enabling novel designs.	Yes, when cost and power consumption targets are met. Because optics is best approach for solving EMI issues and enabling distributed architectures.	Yes, when power consumption and price and size targets are met. Because optics is best approach for solving EMI issues.
10	Do you expect convergence in the mobile industry on standard optical interconnect solutions? Are any MSAs planned?		Standard processes need to be reached in optoelectronics industry.	
11	What government regulations affect the selection of interconnects?	Specific Absorption Rate (SAR), Hearing Aid Compatibility (HAC)		
	SOURCE: Louay Eldada, IPI TWG Chair			

Appendix B – Quantitative Interview Questionnaires

B.1 HPC Questionnaire Version 1

Mapping the Requirements of Current and Future Optical Interconnects:

High Performance Computing and Network Interview

The Microphotonics Consortium at M.I.T. is conducting research on the adoption of optoelectronic components in high performance computing (HPC) and Network systems. As an element of this research, this interview attempts to ascertain the current and future requirements for optical interconnects used in HPC and Network systems (basically what these components should do) and whether standards play a role in realizing those interconnects. With inputs from both the market side and supply side, we hope to create a 10-year strategic roadmap for optical interconnects in the HPC and Network market segments.

Identifying Possible Market Size for Interconnects

1. What are the characteristics of your High End and Mainstream products; and how do you differentiate between the two?
2. What will be the market size for high end/mainstream products in 2, 5, and 10 Years?

Product	2 Years	5 Years	10 Years
High-end Products			
Mainstream Products			

Identifying Technological Needs & Preferences for Interconnects

3. What are the characteristics of interconnects in your ***CURRENT*** High End/Mainstream products?

   a. What is the data rate (Gb/s/channel) for connectors used in each of the following scenarios:

      i. Rack-to-Rack? ii. Card-to-Card?

      iii. Module-to-Module? iv. Chip-to-Chip?

   b. For which of the following, do you use optical interconnects:

      i. Rack-to-Rack? ii. Card-to-Card?

      iii. Card-to-Card? iv. Module-to-Module?

      v. If Rack-to-Rack, what is the average and maximum cable length? Example: 3m/100m

   c. How would you describe interconnects your system architectures utilize in each of the following scenarios?

      Are they **Point-to-Point** OR **Massively Parallel** OR **OTHER?**

      i. Rack-to-Rack? ii. Card-to-Card? iii. Module-to-Module? iv. Chip-to-Chip?

d. Are the optical interconnects used in your current products **Fiber** or **Wave Guide (WG)**? If the interconnects are fiber based, do you use **Multimode Ribbon** or **Single Mode Ribbon?**

4. What will your needs be for Rack-to-Rack interconnects in 2,5,10 years with respect to the following:

Interconnect Requirements	**2 Years**	**5 Years**	**10 Years**
Bandwidth (Gb/s/Channel)			
Distance (m)			
Trigger Price (Dollars/Gb/s)			
Total rack-level interconnects deployed per product			
Do you envision your interconnects to be (**ALL/Mostly/Some/None**) Optical?			

5. What will your needs be for Card-to-Card interconnects in 2,5,10 years with respect to the following:

Interconnect Requirements	**2 Years**	**5 Years**	**10 Years**
Bandwidth (Gb/s/Channel)			
Distance (m)			
Trigger Price (Dollars/Gb/s)			
Total card-level interconnects deployed per product			
Do you envision your interconnects to be (**ALL/Mostly/Some/None**) Optical?			

6. What will your needs be for Module-to-Module interconnects in 2,5,10 years with respect to the following:

Interconnect Requirements	**2 Years**	**5 Years**	**10 Years**
Bandwidth (Gb/s/Channel)			
Distance (m)			
Trigger Price (Dollars/Gb/s)			
Total module-level interconnects deployed per product			
Do you envision your interconnects to be (**ALL/Mostly/Some/None**) Optical?			

7. What will your needs be for Chip-to-Chip interconnects in 2,5,10 years with respect to the following:

Interconnect Requirements	2 Years	5 Years	10 Years
Bandwidth (Gb/s/Channel)			
Distance (m)			
Trigger Price (Dollars/Gb/s)			
Total chip-level interconnects deployed per product			
Do you envision your interconnects to be **(ALL/Mostly/Some/None)** Optical?			

8. What are your most pressing problems with current interconnects? Examples: Density of interconnects (routing/packaging), management of increasing power consumption?

9. How important are the following metrics for trade-offs made for interconnect technology selection? Please use the scale (1 is the most important consideration while 5 is the least important).

Metric	Possibly the most important	Very important	Important	Somewhat important	Not very important
Faceplate Density: (Gb/s) per inch	1	2	3	4	5
Board Density: (Gb/s)/ in^2	1	2	3	4	5
Energy Density: (Gb/s)/ Watt	1	2	3	4	5
Cost Effectiveness: (Gb/s) per dollar	1	2	3	4	5
Other:	1	2	3	4	5
Other:	1	2	3	4	5

Characterizing the Dominant Design

10. For your largest volume optical interconnects, what do you see as the dominant design regarding the selected light source?

a. Will the light source within interconnects be **VCSEL** or **Edge emitting**?

All Interconnects VCSEL	Most Interconnects VCSEL	Equally Edge and VCSEL	Most Interconnects Edge emitting	All Interconnects Edge emitting
1	2	3	4	5

b. Will the light source be **Multimode** or **Single mode?**

Entirely Multimode	Mostly Multimode	Equally Multimode or Single mode	Mostly Single mode	Entirely Single mode
1	2	3	4	5

c. Will the light be **Externally modulated** or **Directly modulated?**

Entirely Externally modulated	Mostly Externally modulated	Equally Externally and Directly	Mostly Directly modulated	Entirely Directly modulated
1	2	3	4	5

11. Why is the chosen light source preferred in the design? And what are the shortcomings of using that light source?

Integrating Optics in the Printed Wire Board

12. Do you see any value in creating a fully integrated solution (integrating optical waveguides in the PWB)? If so, where do you see value in terminating a fully integrated solution and why?

    i. Board edge? ii. On the Board? iii. In the Board?

    a. When you envision optical waveguides in an integrated board solution, do you see a change in the dominant design of the components?

    b. When would you imagine needing this solution? When do you imagine incorporating this technology?

13. What problems could optical waveguides integrated with PWB solve?

    a. Density? b. Thermal Management? c. Cost? d. Other?

Defining a Standard Interface

14. Should the industry adopt a standard interface? If yes, where in the optical link should a standard interface exist?

    a. Board WG at board edge? b. Chip/Module to board edge? c. Chip-to-Chip?

    d. Module-to-module? e. Other?

15. What interface might solve the bulk of needs?

    a. "One size fits all"? (80% of performance @ 20% cost)?

    b. "Substantially Differentiated" (2-5 sizes fits all)?

    c. "Heavily Differentiated" (Customized to get 100% performance @ cost premium)?

    d. Other?

B.2 HPC Questionnaire Version 2

Optical Components – Compute Market Segment Questionnaires

The Microphotonics Consortium at M.I.T. is conducting research on the adoption of optoelectronic components in high performance computing market segments. As an element of this research, this interview attempts to ascertain the current and future requirements for interconnects used in each of the four market segments (HPC Box-to-Box, Backplane-to-Board, Board-to-Chip, Chip-to-Chip). Please provide responses for **TWO** of the four market segments on this page. Also, please provide responses to the **TWO** corresponding questionnaires.

Market / Technology Attributes	HPC Box-to-Box	Backplane-to-Board	Board-to-Chip	Chip-to-Chip
Bit Error Rate Response Type: Rate Range: 10 E-3 to 10 E-20	10E-12	10E-12	10E-12	10E-15
Wavelength Response Type: nanometers Range: 400 nm to 1550 nm	850 nm	850 nm	850 nm	850 nm
Link Types Response Types: Freespace, Waveguide, POF to MM GOF to SM GOF, OTHER	MM GOF	MM GOF	MM GOF	MM GOF
Link Lengths Response Type: mm to Km	1m – 100 m	<1m to 2m	<1m	<1m
Temperature Response Type: Degrees Celsius Range: -55 C to 150 C	Controlled environment ~ RT	Controlled environment ~ RT	Controlled environment ~ RT	Controlled environment ~ RT
Standards Response Type: ITU, IEEE, OIF, Others	IEEE HSSG, IEC	IEEE HSSG, IEC	IEEE HSSG, IEC	IEEE HSSG, IEC
Reliability Response Type: Failures in Time	10	10	3-5	.3

HPC Box-to-Box Scenario

Define **link** (We assume that a link is a dedicated transmit or receive signal on a single fiber).

Data Transceiver Properties	2007	2010	2013	2016
Data Rate Per Link Response Type: Gbps Response Range: 5-100 Gbps				
Energy Density Per Link Response Type: Watts per Gbps Response Range: 5 mW – 50 mW **Supply Power Per Link** Response Type: Watts Per Link Response Range: 0W – 5W				
Optical Connector Type Response Type: Direct to LC to SNAP 12				
Electrical Connector Type Response Type: Differential LVDS, SFP, XAUI, XFI, SFP+				
What cost would need to be achieved by the time of adoption? Response Type: $ per Gbps Response Range: $0.50 - $5.00				
Link Properties	**2007**	**2010**	**2013**	**2016**
Number of Links (link = 1 fiber) Response Type: Per blade or drawer Response Range: 100-10000				
Channels per Link (wavelengths/fiber) Response Type: Ribbon fibers or Wavelengths				
Physical and Economic Properties	**2007**	**2010**	**2013**	**2016**
# fiber/sq. cm Response Type: # fibers Response Range: 1 - 96				
Transceiver on PWB or Active cable assembly? Response Type: PWB or ACA				
Cost Target for an entire Link (end to end) Response Type: $ per individual link Response Range: 50 - 300				
# links/Transceiver Response Type: # Response Range: 1- 96				
Estimated Annual Unit Volumes of Transceivers Response Type: In terns of Transceiver Volumes Response Range: 0 – N				

Board-to-Backplane Scenario

Define **link** (We assume that a link is a dedicated transmit or receive signal).

Data Transceiver Properties	2007	2010	2013	2016
Data Rate Per Link Response Type: Gbps Response Range: 10 - 240 Gbps				
Energy Density Per Link Response Type: Watts per Gbps Response Range: 2.5 mW – 50 mW **Supply Power Per Link** Response Type: Watts Per Link Response Range: 0W – 5W				
Optical Connector Type Response Type: Direct to LC to SNAP 12				
Electrical Connector Type Response Type: Differential LVDS, SFP, XAUI, XFI, SFP+				
What cost would need to be achieved by the time of adoption? Response Type: $ per Gbps Response Range: $0.25 - $5.00				
Link Properties				
Number of Links Response Type: Per Board Response Range: 1 – 10 K				
Channels per Link Response Type: Ribbon fibers or Wavelengths				
Physical and Economic Properties				
# fiber/sq. cm Response Type: # fibers Response Range: 1 - 96				
Transceiver on PWB or Active cable assembly? Response Type: PWB or ACA				
Cost Target for an entire Link (end to end) Response Type: $ per individual link Response Range: 50 - 300				
# links/Transceiver Response Type: # Response Range: 1- 96				
Estimated Annual Unit Volumes of Transceivers Response Type: In terns of Transceiver Volumes Response Range: 0 – N				

Chip-to-Chip Scenario

Define **link** (We assume that a link is a dedicated transmit or receive signal).

Data Transceiver Properties	**2007**	**2010**	**2013**	**2016**
Data Rate Per Link Response Type: Gbps Response Range: 10 - 240 Gbps				
Energy Density Per Link Response Type: Watts per Gbps Response Range: 2.5 mW – 50 mW **Supply Power Per Link** Response Type: Watts Per Link Response Range: 0W – 5W				
Optical Connector Type Response Type: Direct to LC to SNAP 12				
Electrical Connector Type Response Type: Differential LVDS, SFP, XAUI, XFI, SFP+				
What cost would need to be achieved by the time of adoption? Response Type: $ per Gbps Response Range: $0.01 - $1.00				
Link Properties	**2007**	**2010**	**2013**	**2016**
Number of Links Response Type: Per Chip-Chip Response Range: 0 - 100				
Channels per Link Response Type: Ribbon fibers or Wavelengths				
Physical and Economic Properties	**2007**	**2010**	**2013**	**2016**
Faceplate Area (per port) Response Type: Per Port Response Range: 1 - 4 sq. cm				
Board Area Response Type: footprint of transceiver and connector Response Range: 12 - 24 sq. cm.				
Cost Target for an entire Link (end to end) Response Type: $ per individual link Response Range: 0 - 10				
Annual Unit Volumes (High) Response Type: In terms of Transceiver Volumes Response Range: 0 – 1000K				
Annual Unit Volumes (Low) Response Type: In terns of Transceiver Volumes Response Range: 0 – 20K				

Appendix C – Model Code

HP DL 360 G5 Server CPU Utilization – Power Consumption Data

Server CPU Utilization	Power Consumed
0 %	180 W
0.099 %	190 W
0.199 %	200 W
0.299 %	210 W
0.399 %	221 W
0.499 %	234 W
0.595 %	247 W
0.697 %	258 W
0.803 %	270 W
0.907 %	281 W
0.994 %	288 W

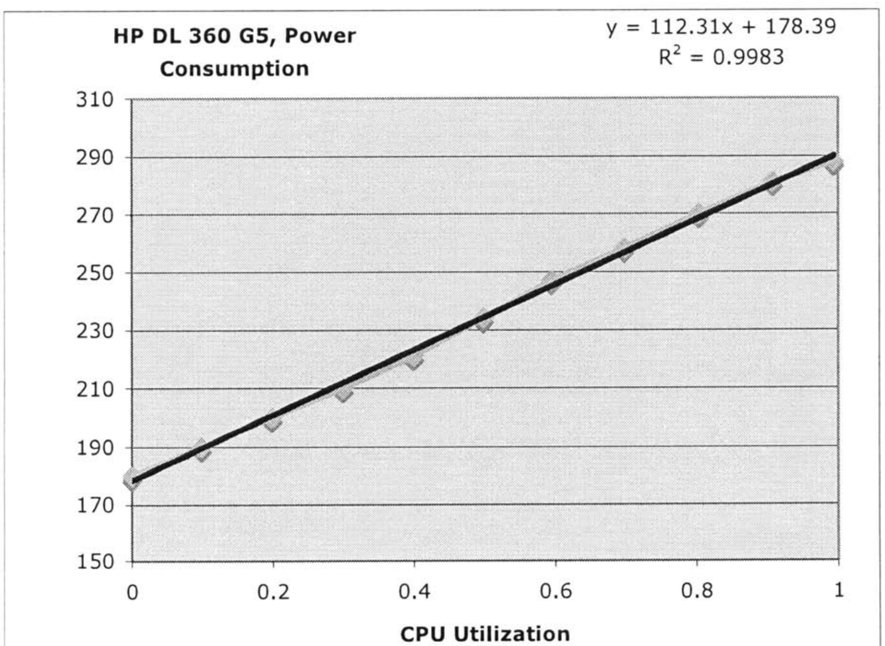

CRAC Sensible Cooling Data [17, 18]

Exhaust Temperature	Sensible Cooling
95 F	99000 BTU
90 F	102000 BTU
85 F	104990 BTU
80 F	102000 BTU
75 F	99000 BTU
72 F	97000 BTU

Sensible Cooling Per CRAC Unit (>= 85F)	
Constant	12743
Exponent	0.4747
Sensible Cooling Per CRAC Unit (< 85 F)	
Constant	1000000
Exponent	-0.5279

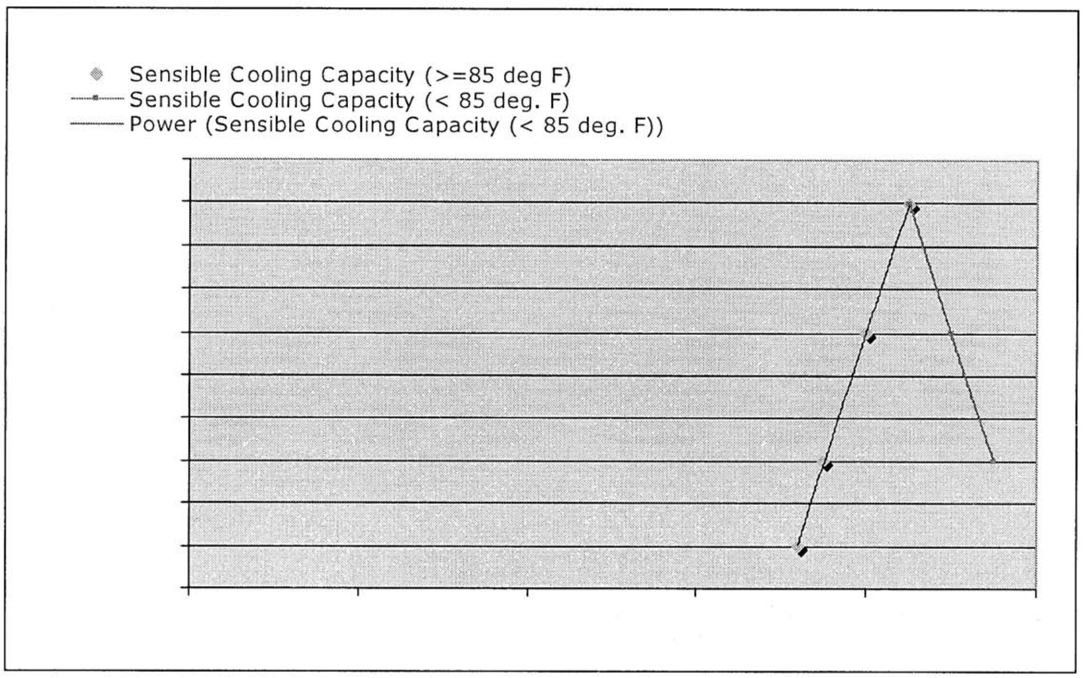

```
Public Function TreeNetworkSwitches(ClusterSize As Integer, PortsPerSwitch As
Integer, NetworkSwitchUtil As Double) As Double
  'Function to calculate the number of switches required
  'for a cluster of servers in a tree configuration.
  'Author: Johnathan Lindsey
  'Date: 1/25/09
  'Unpublished Work © 2009, Johnathan Lindsey

  'Declare variables used within the function
   Dim temp1 As Double
   Dim temp2 As Double
   Dim PortsUsed As Integer

  'Initialize temp1, temp2 to 0
  'temp1 = counter for servers per level
  'temp2 = counter for servers in the entire cluster
   temp1 = 0
   temp2 = 0

  'Initialize PortUsed as the number of ports utilized in a switch.
  'Note that one port is automatically reserved to connect to a storage device
   PortsUsed = PortsPerSwitch * NetworkSwitchUtil - 1

   If (ClusterSize <= PortsUsed) Then
  'If the number of ports available is greater than the number of servers in a cluster
  'then only one switch is needed, the function terminates and return 1 to spreadsheet
      TreeNetworkSwitches = 1
   Else
  'If more than one switch needed, then multiple levels of switches are needed.

   'First, temp1 = the number of switches needed on the base level
      temp1 = ClusterSize / (PortsUsed - 1)

   'If-Else-End If statement essentially roundup the number of switches
      If (temp1 > Round(temp1, 0)) Then
      'Round down then add 1
         temp1 = Round(temp1, 0) + 1
      Else
      'Round up
         temp1 = Round(temp1, 0)
      End If

     'Add the number of switches on the base level to the total number of switches
      temp2 = temp1

      If (temp2 <= (PortsUsed - 1)) Then
```

```
    'If the number of ports available per switch is enough to accommodate the number
    'of switches on the base level, then add a head node switch and return the total
    'number of switches to the spreadsheet.
      TreeNetworkSwitches = temp2 + 1
    Else
    'If not, enter a do-while loop that counts the number of switches needed
    'per level as long as the number of switches per level is greater than
    'the number of ports that a SINGLE switch can accommodate.

      Do While (temp1 > (PortsUsed - 1))
        temp1 = temp1 / (PortsUsed - 1)
        If (temp1 > Round(temp1, 0)) Then
          temp1 = Round(temp1, 0) + 1
        Else
          temp1 = Round(temp1, 0)
        End If

    'At the end of each iteration, the number of switches per level is added to the
    'total number of switches per cluster.
        temp2 = temp2 + temp1
      Loop 'End Loop
    'If the function does not terminate earlier, the function (once do-while loop
    'terminates) returns the total number of switches + 1 head node switch.
      TreeNetworkSwitches = temp2 + 1
    End If
  End If
End Function
```

Public Function ROI(Investment As Double, AnnualReturn As Double, DiscountRate As Double, EquipLife As Integer) As Double

```
  'Function to calculate the ROI of an investment
  'Author: Johnathan Lindsey
  'Date: 1/30/09
  'Unpublished Work © 2009, Johnathan Lindsey

  'Declare variables used within the function
  Dim Gain As Double
  Dim temp As Double

  'Initialize variables to zero, Counter to the equipment lifetime
  Gain = 0
  temp = 0
```

```
    'Calculates the present value of gains over the life of IT equipment
    'Using Excel's PV function, annual gains must be entered as a negative number
    Gain = PV((DiscountRate / temp), (EquipLife * temp), ((-1) * AnnualReturn), ,0))

    'Calculates the present value of annual investments over the life of IT equipment
    'Using Excel's PV function, annual gains must be entered as a negative number
    Investment = PV((DiscountRate / temp), (EquipLife * temp), ((-1) *
    AnnualInvestment), ,0))

    'Return ROI value to main excel workbook.
    ROI = (Gain - Investment) / Investment

End Function
```

Appendix D – Default Model Inputs

Global Inputs	**Units**	**Totals**
E.A. Space Utilization		**2**
Floor Space Utilization	%	**50%**
Total Rack Space	U	**45**
Rack Area	sf	7.46
Rack Depth	ft	**3.27**
Rack Width	ft	**2.28**
CRAC Area	sf / unit	**22.8**
Utilized Rack Space	%	**75%**
Auxilary Load Factor	%	**35%**
Maximum Conditioned Load Req.	Kw	**100**
IT Load Utilization		**75%**
Cluster Size	# of Server Nodes	**256**
Network Switch Size	Ports / Switch	**48**
# of Ports Per Row in a Switch	Ports / Row /Switch	**24**
Network Switch Utilization	% of Ports/Switch	**100%**
Network Switches Per Cluster	# of Switches	6
Uplink Ports Per Storage Device	# of Ports	**2**
Storage Devices Per Clusters	# of Storage	3
Air Temperature (from CRAC)	deg F	**60**
CRAC Power Requirement	W	**39000**
Industrial Power Costs	$ / Kwh	**$0.08**
Rack Costs	$ / Rack	**$3,000**
External Hardwired Connections	$/rack	**$5,000**
Rack Management HW Cost	$/rack	**$3,000**
Software Licensing Costs	$ M	**0**
Raised Floor Cost	$/ft^2	**$220**
CRAC Cost	$/Unit	**$25,000**
Cost of Land	$ / sf	**$45.20**
Equipment Lifetime	Yr	**3**
Facility Lifetime	Yr	**15**
Weighted Average Cost of Capital	%	**0.18**
Revenue Per Server Employed	$ Per Server	**$8,045**

Copper / Optical Transceiver Characteristics

Energy Penalty		
Fiber	Add'l Watts	-0.44
Copper	Add'l Watts	0
Cable Radius		
Fiber Cable Radius	ft	**0.006**
Copper Cable Radius	ft	**0.015**
Cable Bend Radius		
Fiber Bend Radius	ft	**0.230**
Copper Bend Radius	ft	**0.417**
Transceiver Height		
1GB Optical Interconnect Height	ft	**0.028**
1GB Copper Interconnect Height	ft	**0.046**
Cable Management Utilization	%	**70%**
Transceiver Pwr Consumption		
1000Base-T Copper SFP Transceiver	W	**1.1**
1000Base-SX Optical SFP Transceiver	W	**0.66**
Transceiver Thermal Output		
1000Base-T Copper SFP Transceiver	BTU	**3.74**
1000Base-SX Optical SFP Transceiver	BTU	**2.25**
Transceiver Cost Premium		
1000Base-T Copper SFP Transceiver	$	**0**
1000Base-SX Optical SFP Transceiver	$	**100**

Appendix E – Gains and Investment Results From Sensitivity Analysis

A. Transceiver Height

Transceiver Height	Mesh Gain	Mesh Investment	Tree Gain	Tree Investment
0.005	862.902	230.8957913	830.6102	252.9778
0.01	862.902	230.8957913	830.6102	252.9778
0.015	862.902	230.8957913	830.6102	252.9778
0.02	862.902	230.8957913	830.6102	252.9778
0.025	862.902	230.8957913	830.6102	252.9778
0.03	862.902	230.8957913	830.6102	252.9778
0.035	197.6079	246.6454757	260.4208	271.7279
0.04	197.6079	246.6454757	260.4208	271.7279
0.045	197.6079	246.6454757	260.4208	271.7279
0.05	197.6079	246.6454757	260.4208	271.7279

Transceiver Height	Mesh-CHG in Total Costs / Yr	Tree-CHG in Total Costs / Yr
0.005	$152.64	$133.70
0.01	$152.64	$133.70
0.015	$152.64	$133.70
0.02	$152.64	$133.70
0.025	$152.64	$133.70
0.03	$152.64	$133.70
0.035	-$32.83	-$25.07
0.04	-$32.83	-$25.07
0.045	-$32.83	-$25.07
0.05	-$32.83	-$25.07

B. Transceiver Power

Transceiver Power	Mesh Gain	Mesh Investment	Tree Gain	Tree Investment
0.1	971.4941	228.33099	939.5588	249.642
0.2	918.7513	229.6112884	938.5879	249.642
0.3	917.9071	229.6112884	885.8834	251.3071
0.4	917.0546	229.6112884	884.914	251.3071
0.5	916.1963	229.6112884	883.9379	251.3071
0.6	915.3684	229.6112884	882.9976	251.3071
0.7	862.5564	230.8957913	830.2171	252.9778

0.8	196.4197	246.6454757	206.5708	273.4671
0.9	195.5883	246.6454757	205.6258	273.4671
1.0	194.7317	246.6454757	204.6508	273.4671
1.1	141.1179	247.9865772	151.1415	275.2122
1.2	140.2571	247.9865772	150.1615	275.2122
1.3	139.4245	247.9865772	149.215	275.2122
1.4	138.57	247.9865772	148.2425	275.2122
1.5	84.91241	249.3321866	94.69872	276.9633
1.6	84.07843	249.3321866	93.7504	276.9633
1.7	83.22435	249.3321866	92.77835	276.9633
1.8	82.34823	249.3321866	91.78027	276.9633
1.9	81.52751	249.3321866	38.17189	278.7203
2.0	67.72471	250.6823266	37.21977	278.7203

Transceiver Power	Mesh-CHG in Total Costs / Yr	Tree-CHG in Total Costs / Yr
0.1	$182.25	$164.67
0.2	$168.15	$164.24
0.3	$167.77	$149.55
0.4	$167.39	$149.12
0.5	$167.01	$148.68
0.6	$166.63	$148.26
0.7	$152.49	$133.52
0.8	-$33.37	-$40.54
0.9	-$33.75	-$40.98
1.0	-$34.14	-$41.42
1.1	-$48.90	-$56.77
1.2	-$49.29	-$57.22
1.3	-$49.67	-$57.65
1.4	-$50.06	-$58.10
1.5	-$64.86	-$73.49
1.6	-$65.25	-$73.92
1.7	-$65.64	-$74.37
1.8	-$66.04	-$74.82
1.9	-$66.41	-$90.27
2.0	-$62.94	-$90.71

C. Cable Radius

Cable Radius	Mesh Gain	Mesh Investment	Tree Gain	Tree Investment
0.005	872.4348	230.8957913	841.3704	252.9778
0.006	862.902	230.8957913	830.6102	252.9778
0.007	853.3692	230.8957913	819.85	252.9778
0.008	843.8364	230.8957913	809.0898	252.9778
0.009	834.3036	230.8957913	798.3296	252.9778
0.010	824.7708	230.8957913	787.5694	252.9778
0.011	815.2381	230.8957913	776.8092	252.9778
0.012	805.7053	230.8957913	766.049	252.9778
0.013	796.1725	230.8957913	755.2887	252.9778
0.014	786.6397	230.8957913	744.5285	252.9778
0.015	777.1069	230.8957913	733.7683	252.9778

D. Cable Bend Radius

Cable Bend Radius	Mesh Gain	Mesh Investment	Tree Gain	Tree Investment
0.04	912.9431	230.8957913	883.8277	252.9778
0.08	902.4081	230.8957913	872.624	252.9778
0.12	891.8731	230.8957913	861.4203	252.9778
0.16	881.3382	230.8957913	850.2167	252.9778
0.20	870.8032	230.8957913	839.013	252.9778
0.24	860.2683	230.8957913	827.8093	252.9778
0.28	184.4332	246.6454757	247.0019	271.7279
0.32	172.7826	246.6454757	182.2698	273.4671
0.36	162.2429	246.6454757	171.5347	273.4671
0.40	151.7032	246.6454757	160.7995	273.4671
0.44	141.1634	246.6454757	150.0643	273.4671

E. Network Switch Utilization

Network Switch Utilization	Mesh Gain	Mesh Investment	Tree Gain	Tree Investment
0.4	961.8742	416.5981017	926.3803	491.6451
0.5	936.1577	344.2389647	1014.982	389.1594
0.6	936.5517	321.1302682	936.1577	344.239
0.7	881.7296	274.8362601	896.2635	297.2255
0.8	830.6102	252.9777651	896.2635	297.2255
0.9	830.6102	252.9777651	881.7296	274.8363
1.0	862.902	230.8957913	830.6102	252.9778

F. CPU Utilization

CPU Utilization	Mesh Gain	Mesh Investment	Tree Gain	Tree Investment
0.4	188.4648	239.6788297	203.1476	265.1403
0.5	862.902	230.8957913	830.6102	252.9778
0.6	899.601	236.6923919	865.4918	259.7162
0.7	202.3925	258.7890447	275.4813	285.9567
0.72	941.6431	243.2424581	954.485	267.0577
0.74	943.1307	243.4713726	955.9904	267.3225
0.76	207.8339	263.1957735	224.7283	292.922
0.78	208.2059	263.4982674	174.2689	293.2833
0.80	975.9661	245.3225465	976.1159	271.3951
0.82	207.3191	267.5352439	284.0463	296.2051
0.84	207.6998	267.8540529	225.0532	298.4872
0.86	208.082	268.1740774	225.4715	298.8697
0.88	518.1539	512.9375903	345.4133	623.0282
0.9	131.7496	214.84375	212.1098	233.8695
1.0	199.223	217.2579023	212.5063	238.1706

References

[1] Speerschneider, Michael J. "Technology and Policy Drivers for Standardization: Consequences for the Optical Components Industry." (Sc. M Thesis, Massachusetts Institute of Technology, 2004). 9-10.

[2] Kawai, Shigeru. Handbook of Optical Interconnects. Boca Raton: Taylor and Francis Group, 2005. Pgs. 3-4.

[3] Jerry Bautista, et al. "Silicon Microphotonics." Microphotonics: Hardware for the Information Age. 2005. Pgs. 7-10.

[4] Christensen, Clayton. "Value Networks and the Impetus to Innovate." The Innovator's Dilemma. Cambridge: Harvard Business School, 1997. Pg. XV.

[5] Christensen, Clayton. "Value Networks and the Impetus to Innovate." The Innovator's Dilemma. Cambridge: Harvard Business School, 1997. Pgs. 44-47.

[6] Utterback, James and Happy Acee. "Disruptive Technologies: An Expanded View." International Journal of Management. March 2005. Pgs. 1-17.

[7] Pew Research Center. "Americans and Their Cars: Is the Romance on the Skids." August 2006. <http://pewresearch.org/pubs/316/americans-and-their-cars-is-the-romance-on-the-skids>. June 2008.

[8] Oliver Wyman, Inc. "Electronics Are Driving the Development of the Automobile Industry." Dec. 2006. <http://www.oliverwyman.com/ow/pdf_files/1_en_PR_Automotive_Elektronics.pdf>. July 2008.

[9] Frost and Sullivan. Stretegic Analysis of North American In-Vehicle Network Technology and Protocols. Dec. 2005. <http://www.frost.com.libproxy.mit.edu>. June 2008.

[10] Parnell, Karen. "Put the Right Bus in Your Car." Xcell Journal. Winter 2004. Pg.5 <www.xilinx.com/publications/xcellonline/xcell_48/xc_pdf/xc_autobus48.pdf.> May 2008.

[11] Muyshondt, Henry. MOST – Entertainment and Information Systems Infrastructure in Automobiles. Dec. 2002. <http://www.cptwg.org/Assets/Presentations/MOST%20CPTWG%20Pres%202002-12.pdf>. April 2008.

[12] National Research Council. Getting Up To Speed: The Future of Supercomputing. Washington, DC: National Academies Press, 2005. Pgs. 20, 44-53.

[13] Top 500. Top 500 Processor Performance Development. June 1993 – Nov. 2007. <http://www.top500.org/stats/list/30/processor>.

[14] Alan Benner, et al. "Exploitation of Optical Interconnects in Future Server Architectures." IBM Journal of Research and Development. July 2005. <http://www.research.ibm.com/journal/rd/494/benneaut.html>. 25 March 2008. Pgs. 758-759.

[15] Bergh, Arpad. "Today's High Speed Interconnects." Future Opportunities in Optoelectronics. 7 July 2006. <www.cns.cornell.edu/documents/Arpad BerghPresidentEMOIDA.pdf>. April 2008.

[16] Hopkins, Ken and Richard Pitwon. "Pluggable Optical Backplane Technology." Xyratex White Papers. October 2005. <http://www.xyratex.com/pdfs/ whitepapersXyratex_white_paper_Optical_Backplane.pdf>. June 2008.

[17] Hopkins, Ken et al. "An Optical Backplane Connection System With Pluggable Active Board Interfaces." Xyratex White Papers. July 2006. <Xyratex_white_paper_Optical_Backplane_Connection_System_2-0.pdf>. June 2008.

[18] Best, Jo. "Analysis: What is a Smart Phone?" Mobile and Wireless. Feb. 2006. <http://networks.silicon.com/mobile/0,39024665,39156391,00.htm>. June 2008.

[19] Yu, Julia. "Optical Interconnects for Mobile Applications." [Presentation to the Microphotonics Consortium Center Spring 2007 Meeting]. April 2007. <http://mph-roadmap.mit.edu/project/documents/motorola.pdf>. June 2008.

[20] Frost and Sullivan. "Advances in Embedded Memories – Technical Insights." March 2007. <http://www.frost.com.libproxy.mit.edu/.> June 2008.

[21] Gartner DataQuest. Mobile Phones Quarterly Statistics Worldwide. March 2007. <http://www.gartner.com/press_releases. 2000-2007>.

[22] Karkkainen, Leo and Ari Tervonen. "Opportunities in Optics for Mobile Devices." [Presentation to the Microphotonics Consortium Center Spring 2007 Meeting]. April 2007. <http://mph-roadmap.mit.edu/project/documents/nokia.pdf>. June 2008.

[23] Miao, Tremont. "Re: Phone Teardowns Reports." Email to CTR Fellow. 12 Dec. 2006.

[24] Schumacher, Robert W. et al. "Multimedia Entertainment: Vehicle Technology and Service Business Trends." Proceedings of the 2002 Society of Automotive Engineers International Congress on Transportation Electronics. Warrendale: SAE, October 2002.

[25] D'Avello, Bob and Jim Van Bosch. "Portable and Embedded Wireless Devices as Conduit for Telematics Applications." Proceedings of the 2002 Society of Automotive Engineers International Congress on Transportation Electronics. Warrendale: SAE, October 2002.

[26] McGuckin, Paul. "2007 Data Center Conference Poll Results for Power and Cooling Issues." Gartner. March 2008.

[27] Snevely, Rob. Enterprise Data Center Design and Methoology. Pg. 188. Palo Alto: Sun Microsystems Inc., 2001.

[28] United States Environmental Protection Agency. Report to Congress on Server and Data Center Energy Efficiency. Aug. 2007. <http://www.energystar.gov/index.cfm?c=prod_development.server_efficiency>. Aug. 2007.

[29] Barroso, Luiz A., Xiaobo Fan and Wolf-Dietrich Weber. "Power Provisioning for a Warehouse-sized Computer." Proceedings of the ACM International Symposium on Computer Architecture. San Diego: ACM, 2007.

[30] Patel, Chandrakant D. Cost Model for Planning, Development and Operation of a Data Center. 9 June 2005.<http://www.hpl.hp.com/techreports/2005/HPL-2005-107R1.html>. 1 Nov. 2007.

[31] Snevely, Rob. Enterprise Data Center Design and Methodology. Pgs. 38-40. Palo Alto: Sun Microsystems Inc., 2001.

[32] Koomey, Jonathan. Estimating Total Power Consumption by Servers in the US and the World. 15 Feb. 2007. 1 July 2007. <http://enterprise.amd.com/Downloads/svrpwrusecompletefinal.pdf>.

[33] Energy Information Agency. Worldwide Electricity Consumption, 1980-2005. 10 Oct. 2007. 15 March 2008. <http://www.eia.doe.gov/>.

[34] Gartner DataQuest. Servers Quarterly Statistics Worldwide. 1 June 2007. <http://www.gartner.com/press_releases. 2000-2007>.

[35] AFCOM Data Center Institute. Five Bold Predictions for the Data Center Industry [Keynote Presentation]. March 2007. <http://www.afcom.com/files/pdf/dci_keynote_final.pdf>. April 2008.

[36] Computer Industry Almanac. Worldwide Internet Usage – Press Releases, 1995-2006. Feb. 2007. 25 March 2008. < http://www.c-i-a.com/pr_info.htm>.

[37] Wilson, Dominic and Roopa Purushothaman. "Dreaming with BRICs: The Path to 2050." Goldman Sachs Global Economics Papers. 1 Oct. 2003. <http://www2.goldmansachs.com/ideas/brics/brics-dream.html>. 15 March 2008.

[38] Crosman, Penny. "The High-Speed Arms Race on Wall Street Is Leading Firms to Tap High-Performance Computing." Wall Street and Technology 19 May 2007. 29 June 2007. <http://www.wallstreetandtech.com/showArticle.jhtml?articleID=198001925>.

[39] Hewitt, Jeffrey, et. al. Server Virtualization Produces a Shift in Server Shipments. Gartner Group Research. April 6, 2007.

[40] Strohmaier, Erich. 20 Years Supercomputer Market Analysis. May 2005. 1 April 2008. < http://tclark.ittc.ku.edu/eecs739fall2007/papers/strohmaierSC2005.pdf>.

[41] Top 500. Top 500 Architecture Share. June 1993 – Nov. 2007. <http://www.top500.org/stats/list/30/archtype>.

[42] Coleman, Doug. "Optical Connectivity a Good Choice in the Data Center." Cable Installation and Maintenance. Sept. 2007. <http://cim.pennnet.com/display_article/305141/27/ARTCL/none/none/1/Optical-connectivity-a-good-choice-in-the-data-center/.> July 2008.

[43] "Blade Server." <http://www.searchdatacenter.com>. July 2008.

[44] Patel, Chandrakant D. Cost Model for Planning, Development and Operation of a Data Center. 9 June 2005.<http://www.hpl.hp.com/techreports/2005/HPL-2005-107R1.html>. 1 Nov. 2007.

[45] HP DL 360 G5 Server Technical Specifications. <http://h18004.www1.hp.com/products/quickspecs/12476_na/12476_na.html>

[46] Nortlel 5510 Technical Specifications <http://products.nortel.com/go/product_assoc.jsp?segId=0&parId=0&catId=-9245&rend_id=19784&contOid=100175856&prod_id=42140&locale=en-US>

[47] Sun Storage J4500 Array Technical Specifications. <http://www.sun.com/storagetek/disk_systems/expansion/4500/specs.xml>

[48] Emcore. Emcore Connects Cables: Product Brief. June 2007. <http://www.emcore.com/assets/fiber/DDR_Product_Brief_06092008.pdf> April 2008.

[49] Turner, Pitt W. "Industry Standard Tier Classifications Define Site Infrastructure Performance." <http://www.datacenterdynamics.com/Media/DocumentLibrary/Tier_Classification.pdf>

[50] Stallings, William. Data and Computer Communications. River: Prentice Hall, 2004.

[51] Snevly, Rob. Enterprise Data Center Design and Methodology. Palo Alto: Sun Microsystems Inc., 2001.

[52] HP DL 360 G5 Server Technical Specifications. <http://h18004.www1.hp.com/products/quickspecs/12476_na/12476_na.html>

[53] Nortel 5510 Technical Specifications <http://products.nortel.com/go/product_assoc.jsp?segId=0&parId=0&catId=-9245&rend_id=19784&contOid=100175856&prod_id=42140&locale=en-US>

[54] Sun Storage J4500 Array Technical Specifications. <http://www.sun.com/storagetek/disk_systems/expansion/4500/specs.xml>

[55] Hydraulic Diameter of Rectangular Tubes. <http://www.engineeringtoolbox.com/hydraulic-equivalent-diameter-d_458.html>

[56] Heat Transfer. <http://www.drphysics.com/convection/convection.html>

[57] Nusselt Number for Rectangular Tube. <http://www.uofaweb.ualberta.ca/mece//pdfs/ecE370_7_External_Forced_ Convection.pdf>

[58] Properties of Air. <http://www.engineeringtoolbox.com/air-properties-d_156.html>

[59] Cisco 1000Base-T SFP. <http://www.cisco.com/en/US/prod/collateral/ modules/ps5455/product_data_sheet0900aecd8040bd01.pdf>

[60] Cisco 1000Base-SX SFP. <http://www.cisco.com/en/US/prod/collateral/ modules/ps5455/ps6577/product_data_sheet0900aecd8033f885.pdf>

[61] Daniels, Kristin. Customer Service at Cisco Systems - Phone Interview. 9 February 2009.

[62] Turner, Pitt W. "Industry Standard Tier Classifications Define Site Infrastructure Performance." <http://www.datacenterdynamics.com/Media/ DocumentLibrary/Tier_Classification.pdf>

[63] Patel, Chandrakant D. Cost Model for Planning, Development and Operation of a Data Center. 9 June 2005.<http://www.hpl.hp.com/techreports/2005/HPL-2005-107R1.html>. 1 Nov. 2007.

[64] Prof. Pindyck's Lecture Notes on "Production, Cost, and Demand" Sections 4.1, 4.2 on Open Course Ware – Fall 2004 15.010

[65] Pratt, Shannon. Cost of Capital: Estimation and Application. 2002.

[66] Liebert. Industrial Cooling Series: Technical Data Manual. Columbus: Emerson, 1994.

[67] Sanford, James. Phone Interview. 2 February 2009.

[68] United States Environmental Protection Agency. Report to Congress on Server and Data Center Energy Efficiency. Pgs. 90-95. Aug. 2007. <http://www.energystar.gov/index.cfm?c=prod_development.server_efficiency>. Aug. 2007.

[69] ASHRAE 90.1-2001 Power Densities Using Space-by-Space Method. <http://www.lightingtaxdeduction.org/standard.html>. May 2009.

[70] Energy Policy Act of 2005. Public Law 109-58. 8 Aug. 2005. <http://www.epa.gov/oust/fedlaws/publ_109-058.pdf>. May 2009.

[71] Counting Kilowatts. <http://www.savetheearthnc.org/PDF-ARCHIVE/Counting_Kilowatts-3.pdf>. April 2009.

[72] Pacific Gas and Electric Demand Side Programs. <http://www.pge.com/mybusiness/energysavingsrebates/demandresponse/>. May 2009.

[73] Commercial Rebates. <http://www.cpsenergy.com /Commercial/Commercial_Rebates_2009/>. May 2009.

[74] Energy Efficiency Programs for Texas. <http://www.tnpeefficiency.com/>. May 2009.

[75] Austin Energy Commercial/Business Rebates and Incentives. <http://www.austinenergy.com/Energy%20Efficiency/Programs/Rebates/Commercial/Commercial%20Energy/index.htm>. May 2009.