Moore's Law and Energy and Operations Savings in the Evolution of Optical Transport Platforms

Sunan Han, Fujitsu Network Communications Inc.

ABSTRACT

As a telecommunications equipment vendor, we have witnessed a tremendous increase in the capacity of transport equipment over the past 16 years. Platforms have evolved from the first-generation add/drop multiplexer to the next generation multiservice provisioning platform and today's packet optical transport platform. This article discusses our findings in the application of Moore's Law to these optical transport platforms in capacity, power, space, and capital cost. It also discusses our findings in using Moore's Law as a presumption to explore the operations cost savings in power consumption, office floor space, and maintenance. Payback periods for the capital cost of new technology from energy and operations savings are investigated. The conclusions justify the replacement of the network technology for every other generation.

INTRODUCTION

Moore's Law [1] states that while technology advances, the number of transistors that can fit on an integrated circuit (IC) of the same size doubles every 18 months. Moore's Law has since been extended to many other related areas, such as computing speed, media capacity, IC chip cost, and even power consumption [2].

In a broader sense there are three basic elements in Moore's Law: a change that happens at a pace that always doubles, a space where the change takes place, and a time period or cycle over which the change repeats. In the original Moore's Law the change is the number of transistors, the space is the IC chip and the cycle is every 18 months. When these elements are applied to the computing speed of a processor, the change is the frequency of the processor clock, which dominates the computation of each micro step during the execution of a machine instruction, the space is the processor chip, and the cycle is the number of months. This is an example of extending the change element in Moore's Law from the number of transistors to the processing speed. The two other elements — space and time period — can also vary. The space can be a system chassis hosting a number of ICs, and the number of months can be 24 or more, for example.

According to Intel [3], the speed of personal computers (PCs) has increased 417 times, from 6 MHz to 2.5 GHz, in 18 years. If we assume this increase has followed the broader Moore's Law for each of the individual cycles, the cycle can easily be calculated and is equivalent to 24.8 months, which implies the performance of a PC, as a system box, doubles approximately every two years on average over the mentioned 18-year period.

While there are other observations or predictions for bandwidth growth in telecommunications that are based on Moore's Law [4, 5], in the following section we focus on formulating Moore's Law in a broader sense for optical transport platforms based on capacity, power consumption, space, and cost changes in the three generations of our products. These are the add/drop multiplexer (ADM), multiservice provisioning platform (MSPP), and packet optical transport platform (POTP). We use Moore's Law as a presumption to analyze the network operations cost savings from replacing ADM with POTP to determine how fast the investment is paid back by savings in operations. We contend then that if Moore's Law holds for future generations of optical transport platforms, updating network technologies as frequently as at least every other generation makes perfect sense because the capital cost is paid back from energy and operations savings in a surprisingly short period of time.

MOORE'S LAW FOR OPTICAL TRANSPORT PLATFORMS

The second element in the broader Moore's Law — space — is a constant, such as a fixed IC size measured by square inches or square centimeters. It can always be scaled up or down as long as it remains constant during the change. Our space in the analysis of an optical transport platform is a rack that holds three shelves of functional equipment that work together to provide a certain amount of bandwidth in gigabits per second and transport a certain quantity of communication signals in optical fiber over a distance.

The first generation of optical transport platforms, originally manufactured and deployed in the early 1990s, is an ADM of one OC-48 throughput with a maximum capacity of 10 Gb/s when all ports are protected. It utilizes one high-speed shelf for OC-48 1+1 or ring network transport, and two tributary shelves for add/drop or termination traffic for DS-3, OC-3, and/or OC-12. It fits into a standard rack in a service provider's central office.

Twelve years later, service providers started deploying a second-generation optical transport platform — MSPP — in their metro networks. A single-shelf version possesses 100 Gb/s capacity. It has the flexibility to configure any ports to be network or add/drop. It operates at a maximum capacity of 300 Gb/s when its configuration extends to all three shelves in a rack.

In early 2008 a third-generation optical transport platform — the POTP — became generally available. It offers 480 Gb/s capacity in a single shelf that utilizes 1/3 of a rack. This makes the rack capacity 1440 Gb/s.

These three generations of optical transport platform represent a 144-fold capacity increase over the past 16 years, from ADM to MSPP to POTP, in the same three-shelf rack space (Fig. 1).

This rate of capacity increase in a fixed space generates an average cycle of 26.8 months for the capacity to double each time in the past 16 years. To restate this in the form of Moore's Law, the optical transport capacity doubles in a three-shelf system rack space about every two years. It is also interesting to observe the power consumption in the same system space. Figure 2 depicts the increase in power consumption from the three generations of transport equipment housed in a single rack.

We see a much flatter trend line in Fig. 2 for power consumed by the ADM, MSPP, and POTP. Using Moore's formula to describe the power increase, the power consumption for telecommunication transport equipment doubles in a rack every 42.5 months. This flatter trend reflects the fact that the power consumption for each transported gigabit per second has decreased over the past 16 years, from 84 W/Gb/s for an ADM to 6.7 W/Gb/s for a POTP, a 12-fold improvement. Similarly, the cost increase for a service provider to own a rack of transport equipment follows a trend of doubling the cost every 42.7 months. In other words, the per-Gb/s transport cost is now about 1/12 what it was 16 years ago. See Fig. 3 for an illustration with normalized cost.

In the present mode of operation, the high-speed side of an ADM takes in one OC-48, and its lower-speed side terminates primarily DS-3s. The great majority of the DS-3s between ADMs or OC-48s are crossconnected through patch panels to achieve aggregation and grooming for the best utilization of network bandwidth. In addition to significant increase in capacity, one important benefit of POTP is that when multiple OC-48s are consolidated into one such platform, the DS-3 crossconnect among OC-48s is accomplished within the switch fabric of a POTP, saving the terminating ports originally necessary on ADM. We then assume in the following analysis that an OC-48 port on a POTP replaces an entire ADM if all ADMs in a CO can be replaced by one POTP.

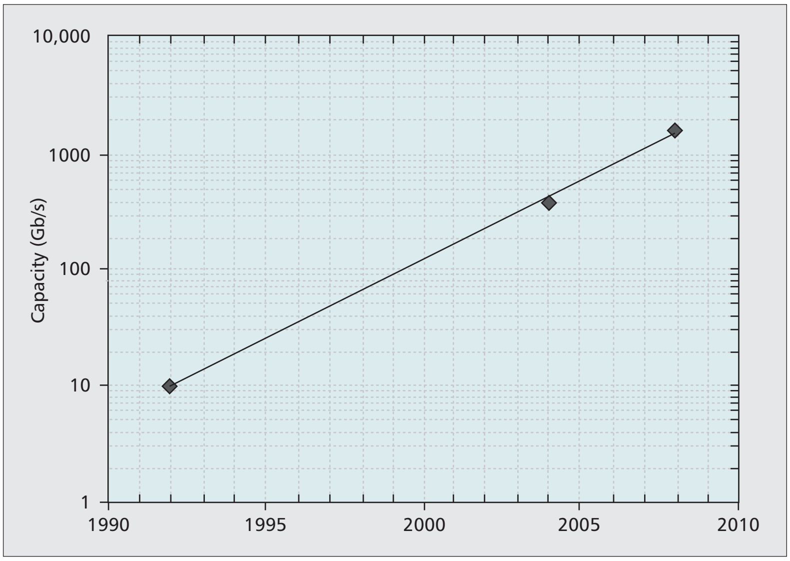

Figure 1. *Optical transport platform capacity increase in a system rack.*

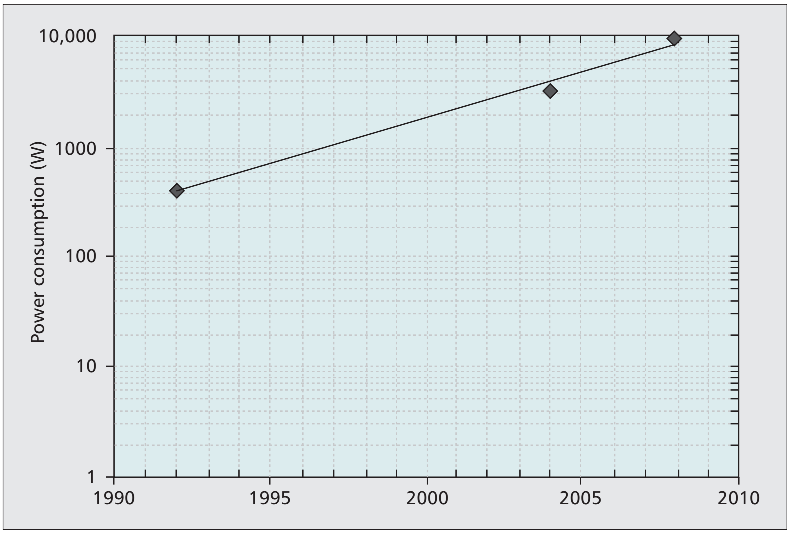

Figure 2. *Optical transport platform power consumption increase in a system rack.*

How Power Savings from New Technology Can Pay Off Capital Cost

In evolving their transport networks, telecommunication service providers have been updating the transport platforms with new technologies for three major reasons: capacity demand increase, service additions, and operations savings. Each time an update occurs, new capital is invested. In this section we try to justify the new capital from power savings implied by Moore's Law for transport equipment power consumption.

With rising energy costs, power consumption is becoming one of the major expenses in operating today's transport networks. Furthermore, the heat dissipation that results from every watt of electricity consumed by central office (CO) equipment requires 2 W of electricity for an air conditioning system to cool it down.

As indicated in the previous section, power consumption and cost increases in Moore's Law occur more slowly than capacity increases. This discrepancy brings about power savings based on the same bandwidth requirement in a CO between the new generation of transport platforms and an old one. Our first finding is that when replacing the first-generation transport platform with the third-generation one, OC-48 for OC-48, in a CO of moderate size (16 to 20 OC-48s on 16 to 20 ADMs), the power savings will offset the POTP capital cost in less than three and a half years. For a CO of 40 or more ADMs of OC-48 capacity, the payback period is less than three years. In Fig. 4 we assume a payback period window of five years (Y-axis). The electricity cost is at 11.5¢/kWh) in the first year with an annual increase rate of 5 percent from year two through year five.

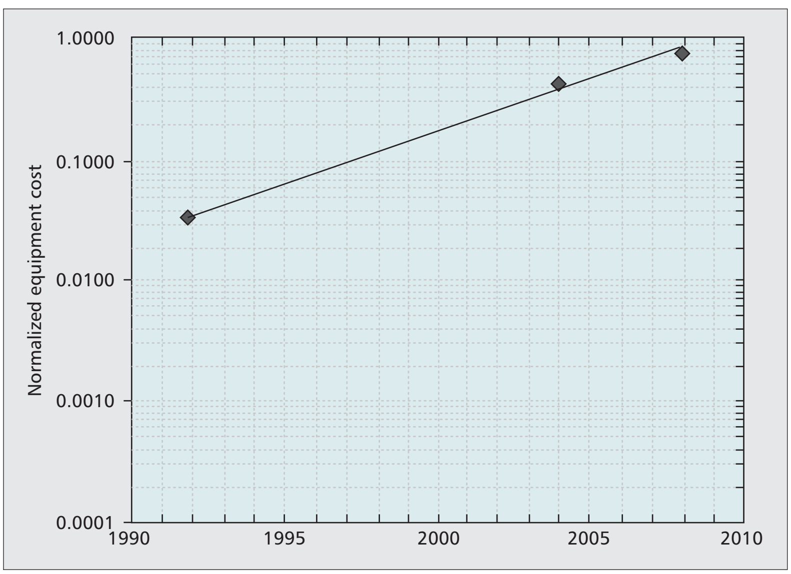

Figure 3. *Optical transport platform cost increase in a system rack.*

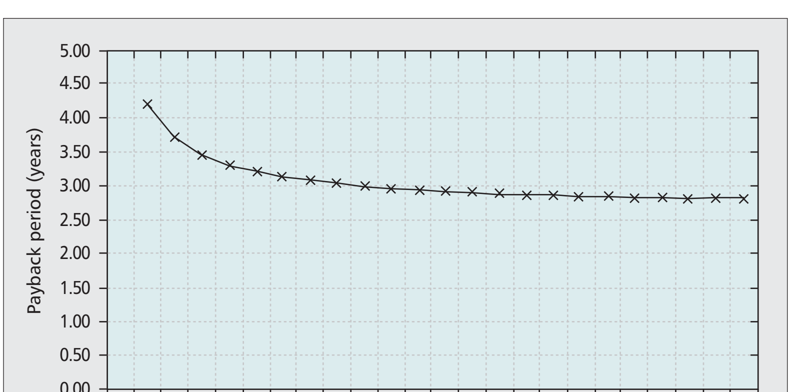

Figure 4. *Payback period analysis — power savings vs. capital cost.*

To apply the analytical result in Fig. 4 to each CO in a carrier's network, locate the number of deployed first-generation OC-48 ADMs using the X-axis.

Central Office Floor Space Savings

One of the operational costs of running a metro network is that of leasing floor space to host the optical transport equipment. In Moore's Law we assumed a constant space — a system rack — that holds three shelves of optical transport equipment. In terms of capacity, one rack of POTP equipment replaces 288 OC-48 ADM racks. This saves a tremendous amount of office space. In real-world deployments, for heat dissipation considerations, we place one POTP shelf in a single rack and leave five times the rack footprint space to create air space around the rack. For OC-48 ADMs a maximum of two such ADMs (six shelves) can be placed in a rack with three times the footprint space around the rack. Both platforms use the same standard size racks.

We assume a \$10/ft^2/mo lease cost for the office space. We then add the savings from floor space and those from power together. We then obtain a new payback-period curve in Fig. 5 based on cumulative savings from both power and floor space.

The reduction of payback time is significant. If replacing a CO of 16 to 20 OC-48 ADMs with a single POTP, the capital cost will be paid back in about two and half years.

Network Maintenance Savings

We surprisingly found that the numbers of active circuit packs in the three-shelf configurations for the ADM, MSPP, and POTP platforms are quite close, in the range of 76 to 85. While providing the same amount of bandwidth, the reliability of one POTP shelf vs. 288 ADM shelves (96 ADMs) of transport equipment components will be quite different.

The cost of maintaining the network mostly comprises replacement packs and the labor required to physically replace the failed packs in a shelf, which may be located anywhere in the network covering a given geographical area. Assuming a mean time to repair (MTTR) of 2.5 h and a per hour labor cost of \$50, we use calculated or predicted failures in time (FIT) for each of the circuit packs. FIT is the number of failures in 10^9 h, information each equipment vendor is required to provide to customers. The calculated FIT is usually higher than the real FIT observed in the field. We divide the FIT by 10 for each ADM component and divide the FIT by 5 for each POTP component to compensate for the fact that an ADM is a well established product, while a POTP is a new product with higher circuit density.

Based on these assumptions, the number of failures occurring annually on a three-shelf ADM is 27 percent, as many as those on a one-shelf POTP. However, in the case of a single POTP replacing multiple ADMs, the number of failures expected on a POTP is, for example, only 4 percent as many as for 96 ADMs. Thus, there is significant cost savings associated with maintaining a single-shelf POTP vs. maintaining multiple ADM shelves that provide the same amount of bandwidth. Figure 6 shows the payback periods resulted from a combination of power, floor space and maintenance savings.

Amazingly, with all the cost saved from power, floor space, and maintenance, the capital cost of deploying a third-generation optical transport platform in place of first-generation equipment is, in most cases, just as much as one and a half year's operations cost spent on the outdated first-generation ADMs.

Other Considerations

Of course, to update an existing network with new equipment is a very complex process. Among other things, such as service continuity, there is cost associated specifically with this process. Meanwhile, the equipment being decommissioned from the network is still of value in the so-called secondary market. We found that, even with a price of 5–10 percent of the original purchase price, the total value of ADMs being replaced may cover most of the reengineering and reinstallation costs for the new equipment.

The study so far has not been comprehensive, and there are other cost and operations savings, such as the removal of broadband crossconnects and DS3 patch panels, as well as benefits such as tax reductions for energy savings. This article has confined itself to tangible operational benefits from the newest generation of optical transport platform, the POTP. There are other potential benefits or savings from the new technology that are hard to quantify, such as efficiency, simplicity, and convenience, and new revenue opportunities with multiservice delivery and new offerings such as Ethernet aggregation.

We would also expect that recent endeavors in developing new and advanced technologies for energy efficiency will soon be kicked into Moore's Law to even flatten the curve against capacity increase. The consequence is a quicker payback period for the investment in the upgrade of network technologies.

Conclusion

This article has given a broader definition to Moore's Law and extended it to telecommunication optical transport platforms based on the evolution of our products. The article has also analyzed the network operations savings from energy, floor space, and maintenance, and thereby proved that the operations savings can pay back the new capital investment in a fairly short period of time when the network technology update happens every other generation. It is worth mentioning that all the network operations activity has significant impact, directly or indirectly, on energy utilization.

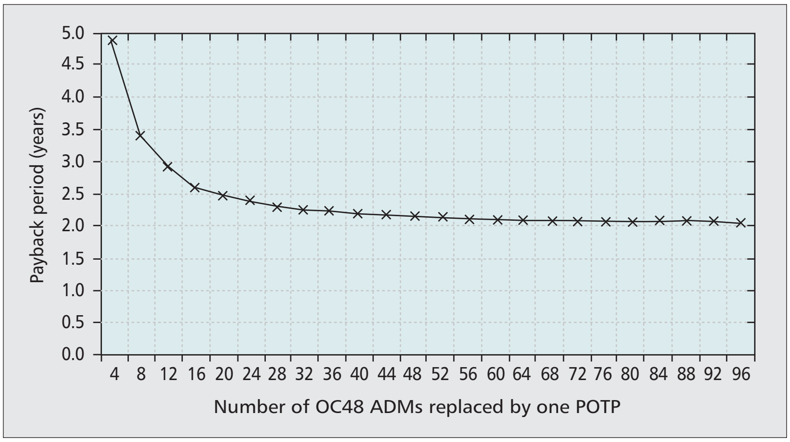

Figure 5. *Payback period analysis — power and floor space savings vs. capital cost.*

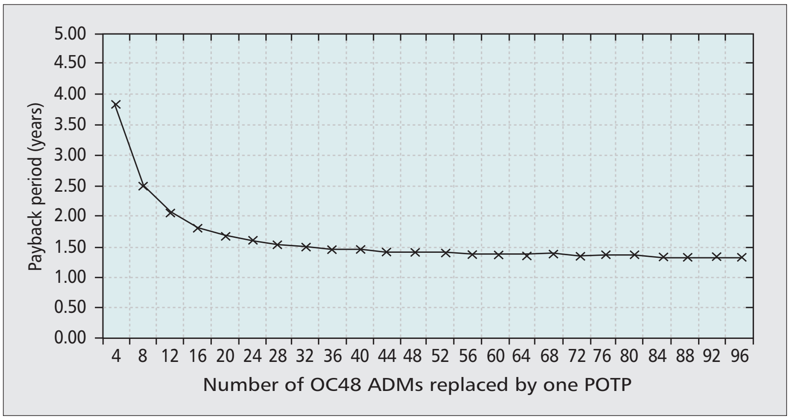

Figure 6. *Payback period analysis — power, floor space, and maintenance savings vs. capital cost.*

References

[1] G. E. Moore, "Cramming More Components onto Integrated Circuits," *Electronics Mag.*, vol. 38, no. 8, Apr. 19, 1965.

[2] W.-C. Feng, "Making a Case for Efficient Supercomputing," *ACM Queue*, vol. 1, no. 7, Oct. 2003.

[3] Intel, "PC Energy Efficiency Trends and Technologies," 2002; http://cache-www.intel.com/cd/00/00/10/27/102727_ar024103.pdf

[4] C. Eldering *et al.*, "Is There a Moore's Law for Bandwidth," *IEEE Commun. Mag.*, Oct. 1999, pp. 117–21.

[5] S. Cherry, "Edholm's Law of Bandwidth," *IEEE Spectrum*, vol. 41, no. 7, July 2004, pp. 58–60.

Biography

Sunan Han (sunanhan@gmail.com) graduated from Beijing University with a B.S. and Beijing University of Aeronautics and Astronautics with an M.S., both in computer engineering. He received his Ph.D. degree in computer science at the University of Texas at Dallas in 1992. From 1992 to 1996 he was an assistant professor of computer science at Adams State College. He joined a systems engineering group at Nortel in 1996, first as a member of scientific staff and then as manager of a voice and data network design team. He was a principal planner and architect with Fujitsu Network Communications from 2001 to 2009. His research interests include SONET, DWDM and packet network design, modeling and analysis, voice, video, data and transport network architectures, network optimization, economic and competitive analysis, and algorithm design and analysis.