W(h)ither Intel

Lack of Demand for Higher Performance to Pummel Core Business

by Nick Tredennick

Frank Gill, Les Vadasz, Arthur Rock, Max Palevsky, Gerhard Parker. What do these people have in common? They are Intel executives or early investors, and they have all recently registered to sell large amounts of Intel stock. (To us working saps, a large amount of stock might be 100 shares. Here, the smallest amount we are talking about is 50,000 shares.) Is this because Intel executives must sell stock to keep food on the table, or because they know something we don't know? Is this unusual in Intel's history? I don't know, but it made me think about where the PC market might be headed and why.

Intel Money Machine Driven by Moore's Law

Intel's business model is based on high-volume production coupled with clockwork improvements in its semiconductor process technology. Intel reduces prices quarterly. Each quarter, a new high-end, high-margin microprocessor enters at the top price bracket, and an old part drops out the bottom. In general, Intel stays about two speed grades ahead of its competitors with its enormous investments in semiconductor process improvements. By leading the x86 pack on this front and in production capacity, Intel commands high margins.

Intel's business model rests on the microprocessor's historical doubling in performance every 18 months. Part of this doubling is from semiconductor process improvements. The black curve in Figure 1 shows that, at a fixed cost, performance doubles every two years. Alternatively, the cost to implement a fixed amount of performance halves every two years. We're all familiar with this Moore's Law performance curve, which is the supply curve for the microprocessor industry.

The electronics business has been in such a continuous boom since the invention of the IC that we forget there's a complementary demand curve. Intel's business model assumes persistent demand for increased performance.

The performance-demand curve should look similar to the performance-supply curve, but it started at a higher performance level and could have a different compound growth rate. We can overlay these two curves and speculate about what will happen. There are two possible scenarios.

• Demand increases faster than supply (Figure 1). This is the world as Intel would have it. In this scenario, I never get as much performance as I need. If I buy a new system every few years, I get a big jump in performance with each new system. But over that interval, my computing requirements have grown even faster, so I never get the performance I need.

• Supply increases faster than demand (Figure 2). Supply may have started well below demand, but some day the curves cross. The demand curve is a hard thing to know. It tells us how much people will buy. What a mess! To take the present example, how much PC performance does a buyer demand? The answer depends on what software is popular

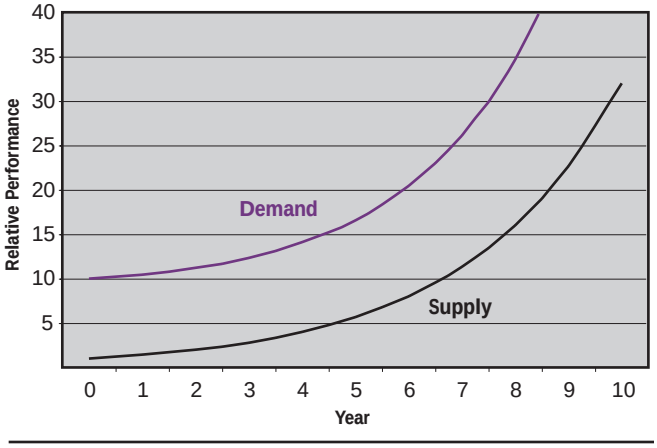

Figure 1. The world as Intel would have it. The demand for performance (purple) starts above the supply curve (black) and has a higher compound growth rate. The supply curve, driven by Moore's Law, is shown doubling every two years.

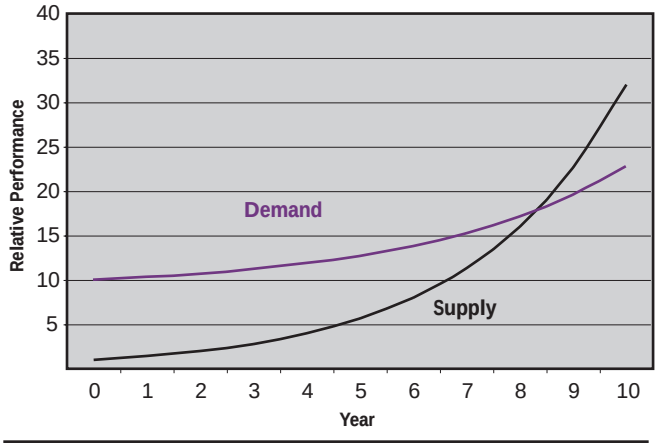

Figure 2. The more probable scenario. The performance supply curve (black) starts below demand but, driven by Moore's Law, overtakes the demand curve (purple) and puts Intel processors under severe price pressure.

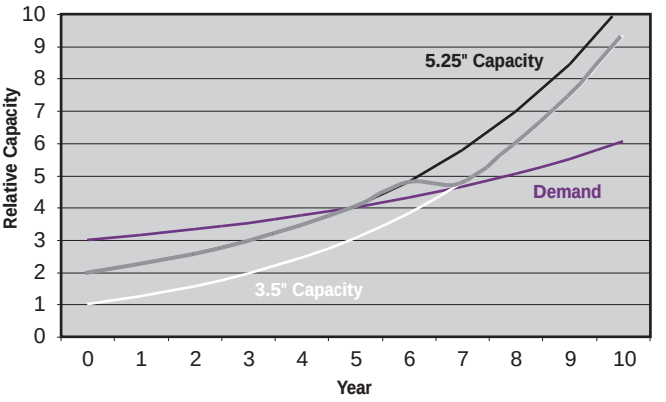

Figure 3. In an analogous historical case, the supply of 5.25" hard disk capacity (black curve) grew faster than demand (purple curve). When the curves crossed, the market shifted (bold gray curve) to 3.5" drives (white curve). (Source: *Harvard Business Review*)

and on how its complexity changes with time. To make matters worse, the composition of the consumer base varies with time. At first, PC owners were hobbyists, then businesses, then households, and now game players. As PCs become cheaper and more capable, the consumer base changes. Demand may be an unknowable function of many variables, but we know from experience that it has powerful effects.

Before the supply and demand performance curves cross, average selling prices (ASPs) are determined by what buyers can afford. This is because buyers can't get as much performance as they need. After the curves cross, some buyers get the performance they need. As the number of satisfied buyers increases, ASPs decline. Some analysts believe the PC business will be driven by nonconverging supply and demand performance curves for a long time.

We've seen supply and demand curves cross before. In disk drives, for example, it helped precipitate the change from 5.25" to 3.5" drives. Figure 3 shows how the market belonged to 5.25" drives until the storage-capability curve for 3.5" drives crossed the curve for storage demand. The market shifted to 3.5" drives, opting for lower cost over higher capacity (the "performance" metric in diskland).

Since there is no PC-equivalent of a 3.5" disk, when the PC-performance supply curve crosses the performance demand curve, we will see ASPs decline. Figure 4 shows that it's already happening.

As long as the demand for performance is greater than the supply, people will pay a premium for it. Figure 5 shows the range of Intel's products from its low-end Celeron processors through its high-end Pentium II processors. In 3Q98, Celeron-266 was priced at $86, and the high-end Pentium II-400 was priced at $589. Looking at performance, the NT SYSmark32 for the Celeron-266 is 185, while the corresponding SYSmark32 for the Pentium II-400 is 368. Across these microprocessors there is a two-to-one ratio in performance, but almost a seven-to-one ratio in price!

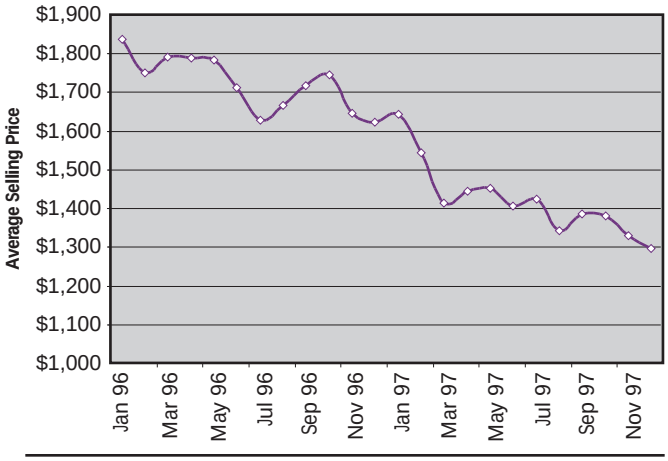

Figure 4. PC system ASPs have been declining steadily over the past two years. (Source: Hal Hardenbergh)

Historically, the price ratio has been even higher, more like eight to one, and this higher ratio is likely to return when the 500-MHz Katmai enters the product line early next year. Buyers are indeed willing to pay a premium for performance—as long as demand exceeds supply.

Intel's competitors offer nearly equivalent performance at one or two speed grades below Intel's fastest parts, and they offer them at about 40% lower prices. Instead of buying the fastest Intel CPU, a buyer might choose to step back one or two speed grades, giving up a small difference in SYSmark rating. With the savings, this buyer can purchase more memory and a better display card to recover the performance. As the demand for the highest performance processors weakens, Intel's sales are pushed down to where its competitors are better positioned.

Lower PC ASPs Threaten Intel's Margins

Lower PC ASPs and decreased demand for leading-edge microprocessors both threaten Intel's margins. There are four things Intel is doing to head off this threat:

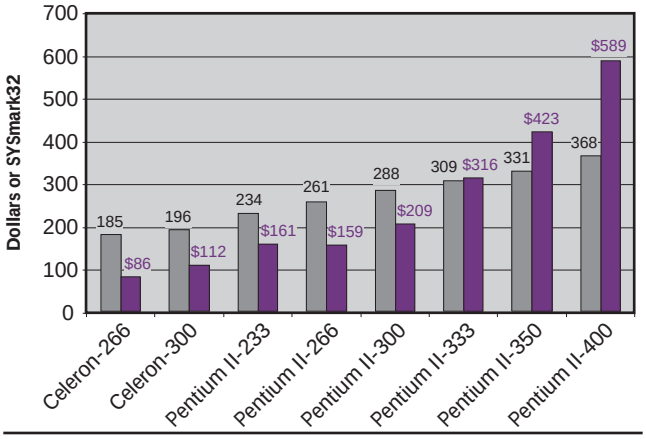

Figure 5. In 3Q98, the price of Intel microprocessors (purple) has a range of nearly six to one from the high end to the low end of the product line, while the performance range (gray) is only two to one over the same processors. (Source: Intel)

- It is diversifying from microprocessors to other system elements—such as network interface cards, motherboards, and graphics chips—to collect margins no matter where they sprout in the computer system.
- It is doing what it can to increase the demand for higher performance.
- It is attempting to change the PC world from Socket 7 to its proprietary Slot 1. For desktop systems, the Pentium II is available only in a Slot 1 module. The Pentium and all of Intel's competitors' microprocessors are available only in packages for Socket 7.
- It will try to move the market from the x86 architecture to the new IA-64 architecture.

Intel's attempt to move the PC market from Socket 7 packages to Slot 1 modules reminds me of IBM's attempt to move the PC market from the ISA bus to the Micro Channel. If IBM had succeeded, it would have regained control of the PC market. If Intel can move the market to its proprietary Slot 1, it will control margins. I don't expect buyer resistance to Slot 1 to be significant, because consumers make decisions independently of thoughts about their aggregate effects.

The major advantage of Slot 1 over Socket 7 is the dedicated backside L2 cache bus. Today's Socket 7-based processors must access the L2 cache using the system memory bus. This increases cache latency and limits cache bandwidth to the system bus speed (66–100 MHz). In contrast, today's L2 cache in a Slot 1 module operates at half the CPU speed (117–200 MHz). Intel is able to do this because it can precisely control the electrical parameters inside the module.

Intel's leadership in chip manufacturing and its dominance of the chip-set and motherboard markets means Intel is well positioned to make the Slot 1 conversion happen. The critical time lies with the next generation of x86 microprocessors from Intel's competitors.

The next-generation part from AMD, the K6-3, will have on-chip L2 cache that could turn Intel's Slot 1 advantage into a Slot 1 millstone. An on-chip L2 should work at CPU speed, giving Socket 7 a performance advantage over Slot 1 until Intel also integrates the L2 cache. But when Intel integrates the L2 in Mendocino (see MPR 8/24/98, p. 1), there will be no reason for Slot 1, as Intel has recently acknowledged (see MPR 7/13/98, p. 4).

Today, Intel enjoys high margins and the lion's share of the market. If ASPs continue to decline, or if Slot 1 is slow to displace Socket 7, Intel can lower margins and maintain market share, or it can maintain margins and lose market share. Intel has built an enormous production infrastructure. It will be motivated to keep its manufacturing plants operating at capacity, which means Intel will give up margins to maintain market share.

Intel Chooses to Switch Rather Than Fight

For the future, I believe Intel's thinking is that if it cannot maintain high margins on its IA-32 (x86) microprocessors, perhaps it can move the market to its new 64-bit IA-64 architecture (see MPR 6/22/98, p. 1). With Merced, the first implementation of IA-64, Intel has an opportunity to break the hold of its x86 legacy and to benefit from recent advances in computer architecture. But the primary benefit to Intel of switching to Merced is that Intel would have exclusive control of the market.

It is almost a foregone conclusion that Merced will dominate the workstation market. Sun remains the lone holdout. But when Merced is introduced, Sun's offerings will probably lag in performance. This is a familiar, and seemingly comfortable, competitive position for Sun. Sun's contrary efforts notwithstanding, the workstation market will consolidate around Merced and Windows NT. Any workstation applications that are not already available on x86 will move onto x86 concurrently with their move to Merced/Windows NT.

Intel's introduction of Merced at the high end of the desktop market is a monumental change in strategy. In its battle for the desktop against RISC-based workstations, Intel used high-volume production as the way to get increased performance. Intel's RISC competitors built for performance, on the assumption that leading-edge performance would grab market share and lead to volume production. Intel's volume-based strategy thoroughly defeated the performance-based strategy. As if it wishes to prove it can win using either strategy, Intel is introducing Merced with a performance-based strategy.

It's not going to work. Merced will dominate the workstation market, but workstation unit volumes are only 1% of desktop-system volumes. Meanwhile, Intel's x86 competitors will treat Merced as a free "concept car" for new architectural ideas. Some of the ideas will fail, and some will be clear winners. AMD, Cyrix, IDT, and possibly even Intel will adopt ideas from Merced in future x86 microprocessors.

Time for PC Geeks to Move On

Intel's business model is built on microprocessor speed and on the assumption of infinite demand for more performance. But processor improvements don't translate directly to improvements in system speed, because the bottleneck isn't in the processor, and the (performance) supply and demand curves have crossed. If Intel converts the market to Slot 1, Intel's margins will decrease, since the curves have crossed. If Intel doesn't convert the market to Slot 1, Intel's market share will decrease, since Intel has a high-cost solution and the curves have crossed.

That's a losing strategy.

As the performance needs of more buyers are satisfied, the PC is becoming a commodity. It's time for electronic-gadget enthusiasts to move on. The short-wave radio, the citizens-band radio, and the handheld calculator preceded the PC as enthusiasts' gadgets. The PC has had a good run, but now it's a consumer-electronics item, so enthusiasts will find a new gadget. What will it be? I don't know—perhaps it's time for PDAs to have their chance. M