



Tech Info Library

PowerPC Business Issues: Competitive Analysis (3/95)

Article Created: 27 January 1994

Article Reviewed/Updated: 17 March 1995

TOPIC -----

This article takes a look at a number of practical business reasons why PowerPC will be the better choice of a microprocessor than Pentium. The business issues boil down to the fact that as an older architecture, Pentium is more expensive to build, will continue to be more difficult to improve, and suffers from the same shortcomings as any sole-source part. PowerPC is less expensive to produce, easier to advance, and has the advantage of being designed and manufactured by multiple sources.

DISCUSSION -----

To best understand where the two microprocessor families stand in today's market it is necessary to take a look at the primary design criteria of both. PowerPC is a new architecture which was designed specifically to meet the computing needs of this decade and beyond. Those needs are very high performance at low cost. Intel acknowledges that Pentium was designed with backward compatibility with the installed base of DOS applications as its number one design criteria.

PowerPC's roots lie in the powerful IBM POWER (Performance Optimization With Enhanced RISC) architecture introduced in 1990. The POWER architecture is a very high performance RISC architecture designed initially for use in high end workstations and servers. PowerPC is a cost reduced adaptation of this architecture by the joint design teams of Apple, IBM and Motorola. PowerPC starting with its first its generation has always been intended for use in high performance computers, be they desktop, server, or portable.

While PowerPC's roots are in high performance computing, the x86 architecture's roots lie in the simple 4004 calculator chip which Intel built in the early 1970s. The 4004 led to the 8008 and then 8080 the immediate predecessor to the 8086. The x86 architecture was essentially defined by its namesake 8086, the grandfather of all succeeding x86 generations including Pentium. Intel had not intended to make the 8086 the basis for a long succession of chips so it probably didn't put as much effort into designing these early-generation parts as it would have had it intended the x86 to be the foundation for many generations to come. In fact, according to John Wharton, author of many articles about Intel's x86 processors and a former Intel employee, the first generation of the 8086 was designed by three engineers in a hotel room over the course of a couple of weeks. Intel was putting its serious microprocessor

design efforts into its ill fated APX432 processor.

Because the IBM Power architecture was designed with very high clock speeds and modern implementation techniques like superscalar execution and deep pipelining in mind, it was optimized for the ability to run at highs speeds and to take advantage of these implementation techniques through the use of such RISC architectural features as fixed size instructions, large register sets, and minimal number of memory addressing modes. Because the x86 was defined long before today's advanced implementation techniques were available to microprocessors (and at a time when 5 MHz was considered fast) it was designed as a complex CISC architecture which must be squeezed into these modern implementations.

Intel's effort to bring an older architecture closer to modern performance standards has resulted in a processor with a large and expensive die size, thermal output levels which has brought a new meaning to the term "a hot PC," and ultimately delivered considerably less than industry-leading performance.

Lets look at each of these factors from a business perspective:

Die size

CISC architecture microprocessors, like Pentium, typically have a higher transistor count than comparable RISC processors. And larger number of transistors will generally translate to a larger die size and higher thermal output. There has already been a considerable amount written in the press about the larger die size of Pentium relative to PowerPC but most computer buyers probably remain unaware of why the relative die size or transistor count is important to them. These characteristics directly affect the cost to produce a microprocessor. And if a manufacturer must pay more to produce a part it is likely that ultimately they will have to charge more for the part. Additionally, over the long term it will be increasingly difficult for the more expensive processor architecture to maintain its place in the market. (A more complete comparison of the transistor count and die size of Pentium and PowerPC are covered later in this paper.)

The higher thermal output of Pentium compared to PowerPC will lead to higher system costs because of the need for complex heat sinks or larger and bigger cooling fans. But the additional cost of these fans and heat sinks is not very great. The big financial impact will be on those individuals and organizations that had intended to upgrade their current 486 based systems to a Pentium based system but are unable to do so because of the significantly increased cooling requirements of Pentium processors were not considered by the manufacturers of their supposedly "Pentium upgradeable" 486 based system. (See the Appendix "Pentium Upgrade Hoax" for more details).

Development issues

Developing succeeding generations of an older architecture is likely to cost considerably more and take far longer than developing new generations of more modern microprocessors. The reasons for this again lie in the fact that much effort (both time and money) must be devoted to sections of the processor that exist primarily to maintain complete compatibility with previous generations of

the processor, generations for example that were designed when the fastest anticipated clock speed was 12 MHz rather than the 120 or 200 MHz expected by the mid 1990s.

In the past, when Intel didn't face as much competition in the microprocessor business, long expensive design cycles for successive generations of processors wasn't a significant issue. But now as the microprocessor market advances rapidly, long and expensive design cycles could put the Pentium and its successors at a significant competitive disadvantage since it should take Intel longer and cost it considerably more money to design future generations of its x86 processors than it will for the PowerPC alliance members to develop succeeding processor generations. In fact, while the next generations of the PowerPC processor will ship in 1994, experts are not expecting to see the next generation of processor from Intel until 1995. Motorola and IBM have already shown samples of the next versions of PowerPC, the 603.

Problems due to sole sourcing

When a key component, like Pentium, is available from only one supplier the supplier may make decisions not in the best interest of its customers. Traditionally, electronics manufactures have been unable to sell their parts if they were the only source for those parts. The reason for this is that having a single company as the only supplier of such a key component as a microprocessor leads to a number of conditions which are less than favorable for that part's customers and the end users who are the consumers of the products.

One of the key business reasons for the ultimate success of PowerPC will be the fact that it is being manufactured and sold by multiple suppliers. In fact one of the most important reasons for the formation of the PowerPC alliance was to ensure that PowerPC would be available from more than one supplier, and that the suppliers would have to compete with each other on pricing and other single supplier issues. This section looks at some of the numerous reasons why a sole source for a microprocessor, like Pentium, is a less than ideal situation for computer manufacturers and more importantly the customers of the computer manufacturers, the party that must ultimately pay for a single sourced part.

Pricing

It doesn't take a professor of economics to figure out that if only one company is supplying a critical component that company will charge an exceptionally high price for that part. Intel has made no secret of the fact that it has fought to maintain its sole source position in x86 parts because as a sole source supplier it can charge much more for its chips than it could if there were alternate suppliers. So long as a microprocessor is available from only a single supplier (especially a supplier that spends many millions of dollars in court fighting to remain a sole source for the part) that part will be priced much higher than one that has to be priced according to market forces.

For example, fear of cutting into its very profitable 486 business is cited as one of the reasons for the high price of Pentium. Although Pentium is very expensive to manufacture, at a selling price of over \$800 it should still be profitable for Intel. But, apparently because Pentium yields are low, producing more Pentium parts would require Intel to give up too much of its very

profitable 486 capacity so it is slow to bring Pentium to market in the numbers its customers require.

Those who are depending on Intel for the latest generation of microprocessor lose out. They must either pay the very high chip price, or make do with older generation parts. If an alternate Pentium supplier existed Intel could not keep Pentium prices high in order to protect its 486 profits. Both experience and common sense dictate that when multiple suppliers offer similar parts, as is the case with PowerPC, lower prices are inevitable.

Availability

Shortages of the latest generation parts can easily occur in the case of a sole supplier since any glitch in that single manufacturer's process could cause a halt in production or lead to production of parts with an unacceptable bug in the part. All production of the part would have to be slowed or stopped till the sole manufacturer solved all of the problems. Intel's problems supplying bug free versions of both the 386 and 486 in the volumes and time period originally announced have been well documented.

Although it easy to see why having a single supplier for the latest generation of a processor could lead to a short supply, a sole source situation often leads to availability problems with older and less expensive parts as well. Even parts that have already gone through the learning curve problems of new chips can be subject to these shortages When a chip maker can charge much more for its latest chip than it can for older chips but it has limited production capacity for all chips, it will devote its production to building the most profitable chips. And because semiconductor manufacturing capacity is so costly and takes so long to build, a sole source manufacturer often becomes capacity limited causing it to put its customers on allocation.

Many observers believed that the acute shortage of 386 parts in 1989 and the shortage of 486SX parts in 1992 and early 1993 was due to Intel's devoting most of its fabs to building the more profitable 486DX parts. Customers that had become dependent on Intel for their processors were forced to pay more for the available higher priced parts or simply wait until the less expensive parts became available.

Multiple sources for the same chip doesn't absolutely guarantee better availability of all generations of parts regardless of their price range, but industry experience has generally shown that it does.

Design advances

A sole source supplier whose customers have become dependent on it for a key component has less incentive to rapidly advance its products or create different versions of the products. Designing advances to a processor or creating new versions is an expensive process. It requires a large investment in time and money, and without competitive market forces a company has little incentive to make such investments until absolutely necessary.

Intel provided an excellent example of this behavior with its 486 microprocessor. Prior to the entry of AMD and Cyrix into the x86 market with

386 compatible parts (parts which Intel continues to contest in court and parts which may ultimately be determined to be illegal) Intel felt little pressure to offer different varieties of its processors. Once it felt competition from these companies though it started to design different versions of the 486. Some of these versions provided significant benefits to computer buying customers, yet it seems likely that many of these versions would never have been produced by Intel had its 486 monopoly not been breached by AMD and Cyrix.

Not only does direct competition between manufacturers spur advances but the combined resources of multiple companies makes advances possible that could be too costly or too risky for a single company to undertake. IBM, Apple and Motorola, have a combined R&D budget of over \$8 billion while Intel's total R&D spending is about one tenth as much at about \$800 million. The combined financial, engineering, design and manufacturing resources of the PowerPC alliance companies far exceeds Intel's resources.

Process limitations

A single manufacture is limited to using only those production processes for which it has experience or is able to obtain licenses. If a particular processor architecture can only be produced by a certain process or if a certain process is necessary for the production of certain speeds of a processor a sole source producer (and its dependent customers) would be at a significant disadvantage. The first generation of PowerPC, the PowerPC 601, is available from both Motorola and IBM even though it is being produced using IBM's advanced .6 micron technology.

Because Intel does not yet have sufficient .6 micron production capacity and it has not licensed any second sources, the Pentium processor is currently being produced with Intel's less-advanced .8 micron technology. But Intel, and its customers, would be better off if Pentium was currently being produced with the more advanced .6 micron technology. By the time Intel is able to produce significant volumes of Pentium with .6 micron technology IBM and Motorola should already be producing the next generation of PowerPC in .5 micron technology.

The combined technology and production resources of two of the best semiconductor manufacturers in the world should insure that PowerPC customers will always have the best process technology available.

For further competitive analysis on PowerPC versus Pentium, please refer to the following Tech Info Library articles:

- 1) PowerPC Technical Issues: Competitive Analysis
- 2) RISC and CISC, Why the Difference: Competitive Analysis

Our thanks to Stephen Dougherty of Apple Competitive Analysis for his permission to add this article to the Tech Info Library.

Article Change History:

- 17 Mar 1995 - Corrected PowerPC wording.
- 11 Mar 1994 - Reclassified for public access.

Support Information Services

Copyright 1994-95, Apple Computer, Inc.

Keywords: kppc

=====

This information is from the Apple Technical Information Library.

19960215 11:05:19.00

Tech Info Library Article Number: 14605