Tech Info Library

RISC and CISC, Why the Difference: Competitive Analysis (3/94)

Revised: 3/11/94
Security: Everyone

RISC and CISC, Why the Difference: Competitive Analysis (3/94)

==========

Article Created: 28 January 1994
Article Reviewed/Updated: 11 March 1994

TOPIC

This article describes some of the differences between RISC and CISC computer architectures and the implications for the PowerPC and Pentium processors.

DISCUSSION

RISC computer architectures take advantage of features such as simple memory addressing and fixed length instructions to achieve higher performance levels than CISC processors which use variable length instructions and multiple memory addressing modes. To understand why the designers of RISC processors decided to use less complex (reduced) instructions than CISC processor designers it's necessary to understand something about the systems for which these processors were designed. The following is but one example of the type of design trade-off decisions made in developing a processor architecture.

CISC computer architectures were originally developed in the 1960s and 1970s, an era when a computer's available memory was both very limited and quite expensive, and many of the processor design decisions made at the time were based on minimizing the program memory requirements. One way to reduce program memory requirements was to simplify software as much as possible by building more complexity into the processor itself. Although increasing the complexity of the processor had a negative impact on processor performance, it was a reasonable trade off at the time. After all, it didn't matter how fast your processor could run if you didn't have enough memory left to load your data and run your programs. But as the price and availability of memory improved there was a need to reexamine the performance versus complexity trade offs. An architecture that made sense in the 1970s when a computer might have only 16 or 32 kiloytes of memory might not be an optimum architecture in the 1990s when 4 or 8 megabytes of memory would be standard.

The goal of a RISC architecture is to allow instructions to be executed as fast as possible. One of the ways of accomplishing this was to simplify the number and type of instructions. The smaller and simpler instructions of a RISC processor could run faster than the longer and more complex instructions of a CISC processor. Of course, this would require moving some complexity into the computer's software programs which, in turn, would require more system memory and better compiler technology. By the mid-1980s with compiler technology quite advanced and memory more readily available, although still expensive, RISC processors made sense for high performance workstations where the cost of memory was not an important issue. And now in the 1990s with very advanced compilers widely available and 8 megabytes or more of memory becoming common, high performance of RISC computing makes sense even for personal computers.

The market conditions that made CISC technology the right choice for personal computers in the 1970s, limited system memory and modest performance requirements, have been replaced by new conditions in the 1990s, inexpensive memory and heavy performance demands. And these requirements of the 1990s are the very conditions which RISC processors were designed for.

Floating point

Intel's engineers and marketing people seem to be at odds over the importance of fast floating point performance. Intel engineers devoted a considerable amount of resources to improving the poor floating point performance of the earlier x86 processors. This effort is reflected in the relatively large amount of space devoted to floating point execution on the Pentium processor. Yet the baggage of the x86 architecture was so great that the Intel engineers were only partially successful in improving floating point performance. Pentium is faster than the 486 in floating point but significantly slower than PowerPC and most current generation processors. As a result Intel marketing people now go out of their way to de-emphasize the importance of floating point.

The truth is floating point performance is not critical for some of today's typical desktop computing tasks such as word processing or database programs, so Pentium's lack of good floating point performance is not too critical for these applications (or at least for the versions of these applications that have been written thus far.) But for the large number of professionals that depend on their computers for financial calculations, high-end graphics or digital imaging, the lack of sufficient floating point performance could make them less competitive. And in the future, fast floating point performance will enable new and far better versions of many existing applications as well as enable new classes of applications that require fast floating point. To the extent that users will want to run these newer applications fast floating point will become essential.

Applications are getting cheaper

The DOS/Windows installed base has far more invested in its data than it

does in its applications. Because the price of applications continues to drop, and because data is the easiest thing to move between platforms, customers will have less and less reason to stay with the declining architecture.

Number of fans

It used to be that the latest x86 based PC would brag about its high clock speed or MIPS rating. Now many of the current press releases and ads for upcoming Pentium systems talk about the number of fans a system includes. It is quite likely that in the near future PC trade magazines will start counting the number of fans a system includes as part of the system's overall rating. The need to have multiple cooling fans is not a particularly desirable trait for a microprocessor. Fans, which are mechanical devices, have a much higher failure rate than any of the electronic components used in computers. So a computer that depends on fans for cooling the CPU will have a reliability rating (MTBF) similar to the relatively low MTBF of its fans. Multiple cooling fans also generate considerable noise. While this noise may not be an issue for file servers in closets it can be quite irritating in a desktop computer.

Intel's reaction

Intel's reaction to the introduction of PowerPC has been unusually vigorous. This is no doubt due to the fact that possibly for the first time it faces a significant threat to its high-end microprocessor business. Intel apparently feels so threatened by PowerPC that it has started running ads in Macintosh oriented publications extolling the virtues of Pentium and the x86 architecture. But MacUser in commenting on the veracity of Intel claims noted that a close look "revealed more obfuscation than clarification".

For further competitive analysis on PowerPC versus Pentium, please refer to the following Tech Info Library articles:
1) PowerPC Business Issues: Competitive Analysis
1) PowerPC Technical Issues: 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:
11 March 1994 - reclassified for public access.

Support Information Services

Tech Info Library Article Number:14607