CompactPCI ™ and Digital's Alpha

A balanced architecture? Close to perfection…

by Glen Lowry

The benefits of the 64-bit CompactPCI architecture together with Digital's newest generation of Alpha microprocessors makes possible a "balanced" open architecture which benefits many real-world, performance-driven applications. Alta selected the Alpha microprocessor for its CompactPCI Single Board Computer (SBC) products based upon the Alpha's clear lead as the most powerful off-the-shelf microprocessor, providing the widest memory bandwidth in the business, and providing easy access to the 64-bit peripheral I/O of the CompactPCI standard. This combination would make it a clear winner for embedded and near real-time applications, but matching application expectations with performance claims isn't as easy as it may have once been.

Measuring one particular feature of a system and staking "worlds best" claims for that feature is not all that is required when selecting a system architecture. In addition to the software and development tools issues, today's high-performance requirements are more complex than those of last year. Here, several "old standby's" for performance measurement are reconsidered for the industry's fastest new entries into embedded and real-time systems. Comparisons with the Pentium are meant to provide a reference point which readers may find useful. Ultimately, the price/performance of the total solution, including software and support costs must be combined with the total performance picture.

The actual performance is "application specific". Specific benchmarks for graphics engines, database servers and SPEC performance all try to evaluate some performance characteristic of a particular architecture for a particular application (usually by manufacturers of products like graphics engines, database engines, and microprocessors.) Benchmarks with little relation to the problem are often the only available tools for performance evaluation. Real time and embedded users are left to apply these benchmarks to their *specific application* (great weasel words for salesmen). The best architecture for a particular application may not be known until well after the project nears completion – and the bangs and bucks of performance are moaned about or glorified.

In years past, the solution was to "get the biggest and fastest" as measured by *"application specific"* benchmarks often at a higher price than current budgets allow. But for the typical real-time or embedded systems designer, the choices have been limited to a few selections - once again with *"application specific"* results. Today's wider range of processors, busses, peripherals, and software makes the decision even more complex. With the newest products to be introduced, it still may be useful to go back to some very basic measurements and to determine their relationship to the real world.

For the measurements *du jour*, let's select processor performance, memory bandwidth, and high bandwidth access to I/O peripherals. The relationship between each of these measurements is – you guessed it: *"application specific"*. Higher processor performance generally requires higher memory bandwidth to keep busy; the I/O required for real-world problems may impact the memory bandwidth. The I/O to computation delay may be too slow. The design objective for a particular problem may begin to sound Dilbert-like: "Eliminate the bottlenecks before you design the system (right!)". Or more scientific: anticipate the bottleneck of your system and attack it with the most powerful feature of a vendor architecture. Back to the biggest and fastest rule.

Why did Alta select the Alpha for CompactPCI?

The easy answer is the Alpha's blazing fast processor speed. But processor speed is more easily described with facts than forecasted results for a real-world solution. Using the processor speed (500 MHz vs. 200 MHz, for example) doesn't always prove the only reliable criteria for decision making. Factors such as pipelining, on-chip cache utilization, scalar/floating point mix, and the proficiency of the compiler writer and programmer all make major differences to useful processor performance. Customer experience with all of these factors have shown the Alpha's exceptional performance in areas of high computation.

Table 1 - SPECmark Comparisons

PROCESSOR	SPECint95	SPECfp95
500 MHz Alpha 21164	15.40	21.10
233 MHz Pentium II	7.12	5.21

As shown in Table 1, SPECmarks, like other standard benchmarks, provide a helpful indicator of processing power for a very specific combination of integer or floating point functions. The 500 MHz Alpha SPECmark performance shows a doubling of the 233 MHz Pentium for SPECint95. But this still provides only a partial indication as to the Alpha's benefits. Real-world applications with floating point components have the potential for even greater improvements – more than four times the floating point performance. In general, for applications which have high computation requirements, the Alpha produces a more practical, high end solution than any other off-the-shelf processor technology. Future versions of the Alpha promise even higher performance.[1 2]

But processor speed is only part of the equation...

In addition to processor speed, memory bandwidth must be sufficient to feed the processor and the I/O devices. For applications with high memory bandwidth requirements, the Alpha exceeds any other architecture which we considered. The CompactPCI 21164 Alpha implementation uses a 256-bit wide bus and up to 1024 MBytes of inexpensive DRAM DIMM modules. Compared with synchronous memory for a 64-bit bus, the potential difference in memory performance is very evident: 60% higher bandwidth. Other variables of the memory performance equation deliver equally impressive results – with three levels of memory cache and a low memory latency. The higher memory bandwidth is important not only to the processor but is a requirement for the higher I/O bandwidth requirements of a 64-bit PCI bus.

The tough one: I/O bandwidth

The third bottleneck is often the most difficult to solve: high speed access to I/O and peripheral devices. In an open architecture, the I/O bus (such as VME or PCI) generally becomes the limit on bandwidth into memory. The current 33 MHz bus speed of the PCI bus limits practical transfer rates to under 100 Mbytes/sec. for a 32-bit peripheral and under 200 Mbytes/sec for 64-bit transfers. The 66 MHz standard will roughly double those numbers to 200 and 400 Mbytes/sec. Compared with maximum 64-bit VME transfers of only 50 Mbytes/sec, the PCI bus provides a growth path for embedded or near real-time applications.

The Alpha's native 64-bit PCI bus support is a natural for 64-bit PCI. In a CompactPCI backplane, all 7 slots (or more with bridge chips) can be 64-bit peripherals. But 64-bit PCI makes promises that peripheral chip manufacturers have been slow to keep. There is a great amount of inertia for 32-bit PCI peripherals and interface chips but relatively few chips with the 64-bit PCI interface – with more coming in the next year. For those applications which need the higher bandwidth, Alta and others will create custom 64-bit CompactPCI interfaces.

The real problem with a PCI bottleneck may be conflicting requirements of I/O and processor for memory bandwidth. Here again, the wide path and low latency of the Alpha's support chips provide more than adequate bandwidth for both I/O data and memory intensive applications. When selecting the Alpha for CompactPCI applications, Alta's customers wanted 64-bit I/O to flow uninterrupted together with memory-hogging applications on the processor.

What about software?

Off-the shelf Alpha products run on the CompactPCI platform, providing developers with Windows NT, Digital UNIX, VxWorks, OpenNT and Caldera (Linux) tools. This software compatibility also allows software development to be performed on less expensive Alpha workstations and moved to the CompactPCI system.

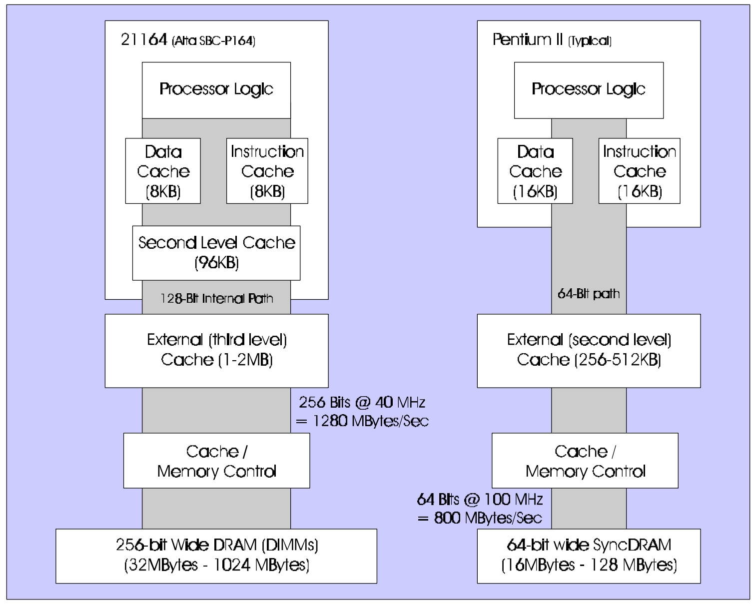

When choosing the Alpha, many of Alta's customers required the ability to use clustered or distributed processors in an off-the-shelf C, C++ or FORTRAN environment. The exceptional processor speed of the Alpha family (soon to be over 600 MHz) suggested to our engineers that many computationally intensive tasks that were being performed by DSPs could easily be performed by Alphas without a significant learning curve for the customer. Using high-speed fibre channel links, the Alpha is easily adapted to the clustered or distributed environments required in many DSP environments.

Nice package...

The 6U CompactPCI chassis is a natural fit to leverage all of the industry's VME knowledge. After addressing the obvious problems for an Alpha implementation (such as oversized heat sinks for the 500 MHz processor) the CompactPCI specification overcomes the power and bandwidth deficiencies of VME architectures. Multiple CompactPCI backplanes in a 19" rack provides users with multiprocessing capability – and split backplanes can be used to combine legacy VME I/O with CompactPCI backplanes and peripherals. Bridge chips and complex backplanes provide access to up to 28 (32-bit) 3U PCI devices. Then, add PMC modules for an increasing number of functions from IEEE-1394 (FireWire™) to old standbys like RS422, and the CompactPCI becomes a tremendous growth path for many existing applications.

There will always be bottlenecks – and benchmarks to measure them. Today's solutions will be replaced by even newer technology. But for today, the CompactPCI Alpha SBC hits most of the high-end problems with a powerful solution – the industry has a sure bet on it!

Glen Lowry is the president and co-founder of www.altatech.com

[1] See *Digital 21264 Sets New Standard* by Linley Gwennap, © MicroDesign Resources, October 28, 1996 found on http://www.digital.com/semiconductor/microrep/digital2.htm