

Bigger Jobs in the Same Space

The increase in CPU power, memory and interconnect bandwidth in step with Moore's Law and implemented as PCIe/104, continues to multiply the computing power within the PC/104 form factor. And this one still has room to grow.

by Martin Mayer, Advanced Digital Logic

The PC/104 envelope continues to benefit from Moore's Law, as CPU performance has reached a new plateau of 20 GIPS starting with Intel's T7400 CPU and continued with the SP9300. Feeding such a powerhouse requires fast memory, and with DDR3, over 5.7 Gbyte/s of memory bandwidth is available to support compute-intensive floating-point and integer tasks while simultaneously pumping 180 Mbyte/s out the video port.

While such an arrangement may be ideal for parallel computation on large arrays of numbers, ray tracing, animation, gaming and fractal generation, none of these tasks are core to embedded applications, which thrive on data collection and processing. Networks of sensors or direct image processing of digital camera feeds are but two bandwidth-intensive applications that are served by the introduction of PCI Express into the PC/104 embedded space. These limited physical-volume systems may now address the challenges of sipping gently from a fire hose.

PCIe provides the local data highway necessary to connect high-bandwidth peripherals to multicore CPU systems. PCIe



FIGURE 1

These cubes offer a visual impression of the data in Table 1 as a function of CPU performance, bus bandwidth and memory.

in the PC/104 embedded space increases I/O bandwidth and throughput, expanding the realm of applications that may be addressed by modular PC/104-based systems. Modularity permits the PC/104 embedded space to morph and tailor the peripheral mix to suit the application, while providing the infrastructure necessary to sustain multiple product offerings and expanded performance levels in the future. As more processing power is brought to

bear on the flood of I/O data delivered via PCI Express, new embedded applications in high-definition video, signal processing, high-volume upper tier data communications, management and encryption are beginning to emerge.

The serialization of the PCI bus into the Express variant changes the fundamental topology from a multipoint simplex bus to a dedicated star, which permits simultaneous full-duplex communication to each endpoint. Capable of driving up to six endpoint peripherals, PCIe/104 is composed of four x1 lanes at 250 Mbyte/s, and a multi-lane portion that typically drives a single x16 endpoint at 4 Gbyte/s, two x8 endpoints at 2 Gbyte/s or two x4 endpoints at 1 Gbyte/s.

Using dedicated transmit and receive signal paths, PCI Express leaps forward as a communications bus. Bandwidth capabilities suggest that a single PCIe/104 CPU could handle the task of inserting selective security into a full speed OC-192 data link. Similarly, a suitably fast CPU and memory could multiplex and demultiplex eight Gigabit Ethernet links into an OC-192 data link, again, adding value within the stream.

The x16 Express interface offered by PCIe/104 was originally devised as a high-bandwidth interconnect for an external



Get Connected

with companies mentioned in this article.

www.rtc magazine.com/getconnected

graphics processor. Emergence of the Compute Unified Device Architecture (CUDA) has tasked graphics processors with alternate roles as vector processors and physics engines. Market requests indicate some desire to exploit this functionality, which waits only for the development of silicon capable of meeting the rigorous environmental characteristics of the PC/104 embedded space.

Adaptation of existing x4, x8 and x16 PCI Express designs are currently being driven by new market demand. Ultra definition and high-speed image capture are two applications that can leverage 4 Gbyte/s x16 interface, as inputs to the system rather than outputs. The increased signal processing power of the SSE4 instruction set makes compression and signal processing tasks much easier to handle “in-core” as opposed to designing specialized hardware specifically for these tasks.

When coupled with sufficient processor and memory resources, sometimes referred to as the processing footprint, the I/O bus can be graphed in a third dimension to express the computational volume of a system. With memory represented in the diagonal Z-dimension, Figure 1 compares current bus champions, for ISA, PCI and PCIe.

I/O bus bandwidth is vertical and processor speed is horizontal. The smallest shape represents a 133 MHz Elan SC520, 128 Mbyte SDRAM, PC/104 CPU. The medium is a 1.8 GHz Pentium-M, 1 Gbyte DDR, PCI/104 CPU. The large cube represents a 2.26 GHz Intel SP9300 Core 2 Duo, 4 Gbyte DDR-3, PCIe/104 CPU. All of the CPUs chosen have the ability to execute one or more instructions per clock cycle, making their integer performance meet or exceed the CPU clock frequency.

Note that the shape of the PCI Express computational volume is more cubic than flat, wide and deep like the other two volumes. This is the first indication that the combination of bus, CPU and memory has not reached the bus maximum.

With PCI Express, total memory bandwidth must exceed four times the total bus bandwidth before full-duplex store-and-forward operation can maximize bus utilization. Increasing the memory bandwidth may require a wider and faster channel, and this is one of the

CPU	MHz	BUS	MB/sec	MB RAM	MMv	$\sqrt[3]{MMv}$	$\log_2(MMv)$
SC520	133	ISA	12.72	128	216460.16	60.04	17.72
PM745	1800	PCI	127.16	1024	234374997.66	616.55	27.8
SP9300	2260	PCIe	4768.37	4096	88281250000	4452.69	36.36

TABLE 1

Numerical representation of Figure 1. Note that the cube root column for the PCIe system shows there is still room to grow to a potential of around 6200 before the PCIe interconnect fabric is saturated.

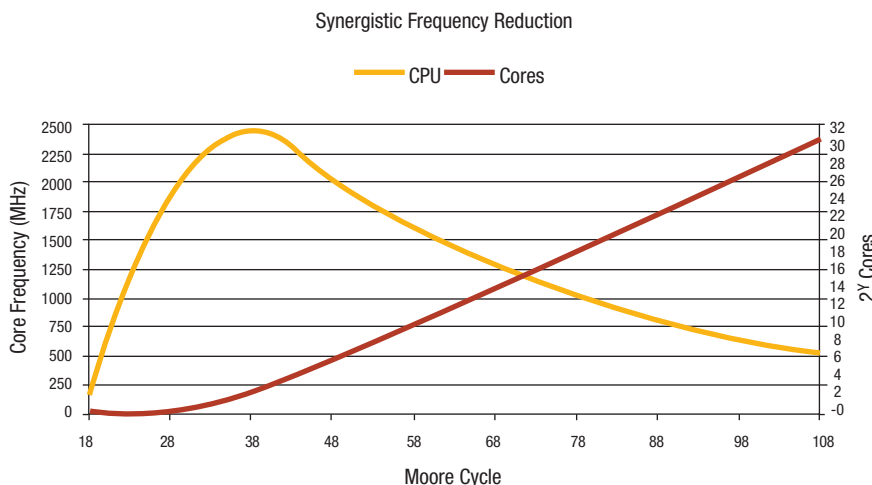


FIGURE 2

The synergistic effect on clock frequency caused by adding CPU cores while continuing to advance through future Moore cycles.

fundamental architectural challenges that must be addressed before PCI Express, in even its basic form, is challenged to the fullest extent. The PCI Express bus is likely to remain useful for several Moore Cycles past saturation.

Technical limitations of the PCIe/104 implementation will hold the current per-lane data rate at its current level, as the increased bandwidth of second-generation PCI signaling will cross the boundaries of reliability on some systems. Knowing when the next I/O interconnect solution must be deployed becomes critical to the continued longevity of the physical PC/104 envelope.

The cubes in Figure 1 represent Moore Volumes, which are the product of CPU Frequency (MHz), memory size (Mbyte) and bus bandwidth (Mbyte/s-1). One MMv has the implicit units of MHzB²s⁻¹. It should be noted that a system of 1 MHz by 1 Mbyte/s by 1 Mbyte, such as the IBM PC/XT would have a rating of 1 MMv. It is worth noting that the base-10 bus speeds were converted to MiB, so that a correct Moore Volume is computed.

This preserves the scale implied by the fundamental HzB²s⁻¹ unit.

Table 1 reveals the dimensions of the three cubes. The first two rows are for the saturated volumes, and we see that there is a rough factor of 10 in the cube-root of the Moore Volumes, and an increase of 10 in the base-2 logarithm of the Moore Volume. For the cube root column, we would expect the PCIe generation to have a number on the order of 6200, revealing that this number can grow. Likewise, the log₂ column seems to suggest that 1.5 more cycles remain before saturation.

While the length of a Moore Cycle is no longer fixed at 18 months, the log₂ column still shows that ten doublings of Moore Volume occur between new bus saturations in PC/104. Considering that there were many Moore cycles that occurred before the PC/XT, it is reasonable to class these shifts in I/O paradigm as occurring every Moore Decade. There is considerable overlap from one bus to the next, as two stackable buses at most are currently supported on a PC/104 CPU. This says nothing of other ubiqui-

tous I/O offerings that bristle on various connectors.

With Mega-Moore generations 18 and 28 completed, we are currently racing toward the completion of generation 38, which is witnessing the vast deployment of multicore CPUs. Many studies have been performed on twin-core units, showing that many applications can be unrolled and paralleled to nearly double the performance. Of course, we know that

tem with >12 km radius and a 3D display running at 1600x1200 resolution. Once memory bandwidth exceeds 16 Gbytes⁻¹, the market will begin clamoring for a faster embedded bus and more cores. What will it take to satisfy such a voracious appetite?

Imagining the 48th Moore Cycle and beyond, there is limited usefulness in expecting processes to be continuously splittable. However, this is not out of the

sure on core clock rates, the number of processor cores per CPU must increase 16-fold every 10 Moore Cycles in order to saturate the next-generation interconnect fabric. Given this explosive rate, one might be inclined to look at each process as a vector that is at a right-angle to another vector, with the result being a mutual vector with a magnitude of $\sqrt{2}$, representing an average synergistic effect of each new processor core added to the CPU. Figure 2 expresses this possible future.

Note that the secondary Y-axis is logarithmic, with the maximum value representing 2³⁰ cores in a single processing unit. Memory and bus bandwidth must grow to feed a billion core unit. Table 2 lists upper bound values for each decade in Figure 2.

While only time will answer the precise I/O bandwidth beyond the 38th Moore Cycle, the expectation that bus speed continues to advance 10x for each new bus may prove different than what is predicted. Many other I/O solutions may be invented between now and then, yet the factor expressed in the table represents the minimum payload that must be delivered before a new I/O interconnect is considered for incorporation into the PC/104 embedded space.

Equally tenuous is the rate of memory expansion. While it is certainly reasonable to expect that memory density will continue to increase, the question remains as to when, if ever, it may achieve a density that permits it to serve in the confines of the PC/104 embedded space.

As we focus on the tasks necessary to finalize the transition into the age of the Giga-Moore, and wonder what innovations will bring the 48th Moore Cycle to fruition, there is certainly no schedule implied in this application of Moore's Law and the concept of the computational volume. Rather, milestones are indicated, and when, not if, they are reached, we will know that it is time to explore all that technology has to offer, selecting the best of breed to keep the PC/104 embedded space a growing and thriving ecosystem. ■

**Advanced Digital Logic
San Diego, CA.
(858) 490-0597.
[www.adl-usa.com].**

Historical Comparisons for Perspective

1) The Cray-3, which was implemented to ¼ of its designed capacity before being abandoned, was rumored to have run as a dual-core vector processor for some time after being decommissioned. The theoretical throughput of a 16-core Cray-3 was 15.17 GFLOPS. The Intel SP9300 CPU has been clocked in excess of 16 GFLOPS using SSE3 instructions and SiSandra 2007 and Windows XP-SP3. The Cray-3 consumed 88 kW in 1995 and was bath cooled in Freon. The 2010 PCI/104-Express SP9300 needs less than 30W and can be cooled radiatively, with no moving parts, at an ambient temperature of +85°C.

2) In a tight loop, the 6502 CPU could execute 250,000 instructions per second, with a more typical performance of 125,000 IPS. With an installed memory of 64 kB and a 1.023 MHz 2-phase clock, the I/O expansion bus would move 0.975 Mbytes⁻¹. The computed Moore Volume for an Apple][with these specifications is 0.00762 *MMv* typical and 0.0152 *MMv* peak. Compared to the IBM PC/XT's rating of 1*MMv*, the Apple][appears at -6.0356 to -7.0356 Moore Cycles.

No. Cores	Moore cycle	No. Cores (Log ₂)	f(Core) (MHz)	RAM Size	I/O Bandwidth	Moore Unit
1	18	0	133	128 MB	12.7 MBs ⁻¹	<i>MMv</i>
1	28	0	1800	1 GB	127 MBs ⁻¹	
4	38	2	2488	8 GB	4.66 GBs ⁻¹	<i>GMv</i>
64	48	6	1990	64 GB	46.6 GBs ⁻¹	
1024	58	10	1592	512 GB	466 GBs ⁻¹	
16384	68	14	1274	4 TB	4.55 TBs ⁻¹	<i>TMv</i>
262144	78	18	1019	32 TB	45.5 TBs ⁻¹	
4194304	88	22	815	256 TB	455 TBs ⁻¹	
67108864	98	26	652	2 PB	4.44 PBs ⁻¹	
1073741824	108	30	522	16 PB	44.4 PBs ⁻¹	<i>PMv</i>

TABLE 2

Moore Volume Parameters: M, G, T and P correspond to 2²⁰, 2³⁰, 2⁴⁰ and 2⁵⁰ respectively.

there are limits to synergistic processing, rendering the simple product of cores and clock frequency less than useful for more than about 4 cores.

It can be predicted that a quad core system operating at a nominal 2.5 GHz with 8 Gbytes of system memory will be able to saturate the existing PCIe/104 deployment, perhaps in the form of a real-time mobile Doppler weather-radar sys-

question as the complexity of the tasks to be accomplished will continue to increase in response to available embedded system processing power. At the 48th Moore Cycle we may see full advance weather prediction, with HD-3D display, running an advanced WRF 24 km weather model, fed by the previous generation's radar unit hanging off of a USB-3 port.

To apply continued downward pres-