Next Generation of Embedded Computing Platforms

Written by: Hector Lin, Advantech Corporation, Industrial Automation Group

A Little Bit of History First

It all started with Moore's Law, of course. In the April 1965 issue of Electronics Magazine, Intel co-founder Gordon E. Moore described the doubling of electronic capabilities. "The complexity for minimum component costs has increased at a rate of roughly a factor of two per year ..." he wrote. "Certainly over the short term this rate can be expected to continue, if not to increase."

Even though there were and are pundits who believe that Moore's law will finally be exceeded, the cost and power of electronic products continues to follow his law. Costs drop by half, power increases by a factor of two, every two years. This has now been going on for almost 43 years. Every time it looks like there will be a slowdown, new processes are developed to continue to make more and more powerful electronics less expensively.

So, what has this meant? Manufacturing, indeed all of society, has been radically changed by the applications of Moore's Law. In 1965, manufacturing was done with paper routers and instructions. Machining was done by hand, to drawings. Drawings themselves were done with pencil or India ink using draftspersons that did nothing else all day. Engineers used slide rules.

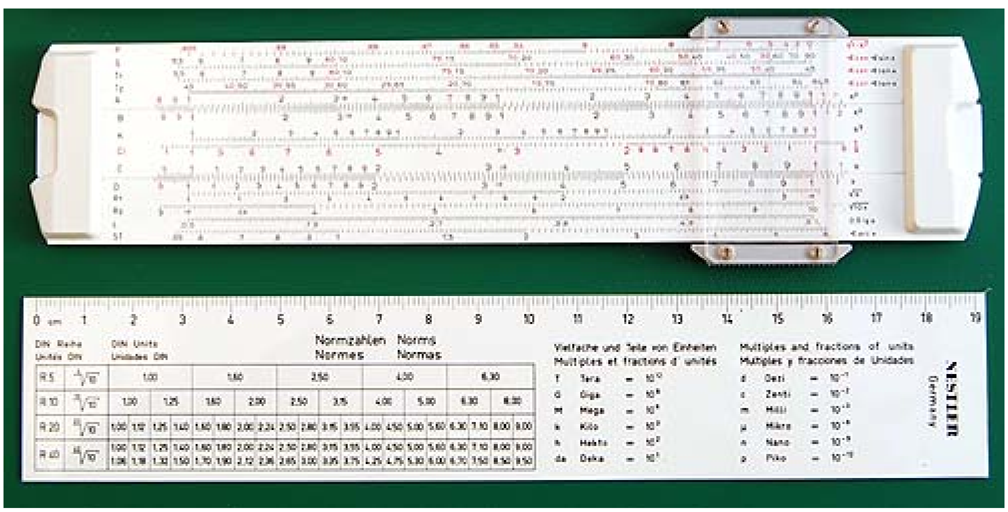

Machines were controlled by electromechanical relays and mechanical timers, and human operators. If a new product was required, the production lines needed to be shut down, re-designed, re-wired, and re-started, often at the cost of months of lost production.

The computers that put a man on the moon in 1969 had far less processing capability than the average inexpensive cell phone will in 2009.

In 1968, Dick Morley, of Bedford Associates, designed a computer-based replacement for hardwired relays that he called a Programmable Logic Controller, or PLC. The first generation PLCs had 4 kilobytes of memory, maximum. They revolutionized industrial production, both in the discrete and in the process fields.

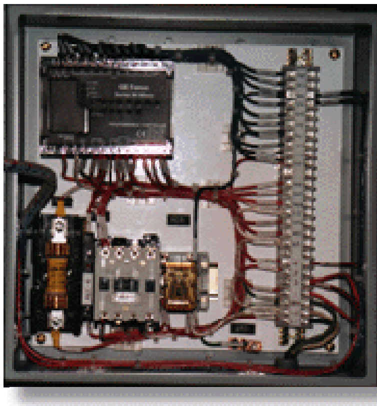

In 1976, Robert Metcalfe, of Xerox, and his assistant, David Boggs, published "Ethernet: Distributed Packet-Switching for Local Computer Networks." Most computers and PLCs today use some form of Ethernet to move data. The patents, interestingly, were not on the software but on the chips to produce hubs, routers, etc., which had become practical because of Moore's Law.

In 1981, Moore's Law permitted IBM to release their first Personal Computer, or PC. It ran with 16 kilobytes of RAM memory on an Intel 4.88 MHz 8088 chip. Each of the original PCs were more powerful than the triple modular redundant computers that still (in 2009) drive the Space Shuttle.

By 1983, Moore's Law had progressed to the point where a joint venture of Yamatake and Honeywell produced the first "smart transmitter." This was a field device: a pressure transmitter, which had an onboard microprocessor transmitter—a computer inside a field instrument, that could communicate digitally and be programmed like it was a computer. Other companies quickly followed suit.

In 1996, Fisher-Rosemount Inc., now Emerson Process Management, changed the definition of a Distributed Control System by combining a COTS (commercial-off-the-shelf) PC made by Dell with a proprietary field controller and a suite of integrated proprietary software, running over standard Ethernet networks and called it Delta V. This was only possible because Moore's Law had made the PC powerful enough to replace the “big iron” proprietary computers used in previous DCS designs, both from Fisher-Rosemount and other vendors.

And in 2002, Craig Resnick, an analyst with ARC Advisory Group, coined the name “Programmable Automation Controller” for an embedded PC running either a version of Windows or a proprietary RTOS (real time operating system).

From then until now, Moore’s law has continued to drive costs down and processing power and speed up with no sign of let up. This has led to the development of embedded computing platforms, such as the PAC.

Embedded computing platforms, of course, aren’t limited to PACs. They are used universally for many different tasks. They are used in Ethernet switches and routers. They are used in analytical instrumentation. They are used in machine control and vision control systems. They are used in washing machines, stoves, and other appliances in the consumer electronics industry, as well as in gaming machines, environmental controls, building automation, and sundry other applications. It is difficult, in fact, to think of an industry which does not use embedded computing platforms.

The Next Generation Embedded Computing Platform

What does the current iteration of Moore’s Law hold for embedded computing? Not only have speed and computational power gone up exponentially, and cost dropped the same way, other improvements have increased, too.

The original 8088-powered IBM PC of 1981 used over 150 watts to operate. The Pentium PCs of the early 2000s used over 300 watts to operate. They generated enough heat that many of those watts went to operate cooling fans. Because cooling fans pull dirt and debris into the enclosure, this made PCs a questionable tool on the plant floor. Making an embedded computer that can be operated without internal cooling apparatus is a critical advance.

The next generation of embedded computing platform has no moving parts, no fans, and has components that require low power. Low power computing also saves energy, and makes the embedded computing platform greener than computers that require fans for cooling.

Among those non-moving parts can be a “diskless” drive. Many of the remaining watts of an early computer went to run the hard disk drive. Recent advances in flash memory have made it possible to use flash memory as replacements for hard drives. Flash memory drives up to 16 Gigabyte are common, 32 Gigabyte are expensive but available, and according to Moore’s Law, by 2010, users should be seeing 64 Gigabyte flash drives in sizes and at prices that are economical enough to be used in embedded computing applications.

Where the original PC used a single core 8088 processor, the next generation of embedded computing platform uses high performance multi-core processors, with larger shared L2 caches, and delivers better power efficiency and performance than such systems did only a few short years ago. Moore’s Law strikes again. See for example the Core 2 Duo CPUs in Advantech’s UNO 2182 and 3182 embedded PCs, making them thousands of times more powerful than the original PC.

The fact that the next generation embedded computing platform has no moving parts, and is designed to be a low power, low heat dissipation system means that they can also be designed for a very wide operating temperature—wider, for example, than a rotating hard disk drive can tolerate.

Since they don’t need to allow space for cooling fans, or large heat sinks, or a hard disk drive, the footprint of the next generation embedded computing platform can be very small, and the entire system can be very light.

Fanless computing platforms can easily be used on more critical applications than traditional PCs with fans and rotating hard disks could not be used on in the past—power substations, trains, maritime applications, and hazardous areas in industrial applications.

In addition, like any PC, the next generation embedded computing platform is fully capable of running operating systems like Windows or Linux, or proprietary embedded OS, as the user chooses. The next generation of embedded computing platforms can be powerful enough to run full implementations of Windows XP, or Linux. They can also run embedded versions of Windows or Windows CE, and Embedded Linux.

Like any PC, the next generation of embedded computing platforms has many communication options that were not available to the original industrial PLCs. The next generation, thanks again to Moore’s Law, will even have Gigabit LAN capabilities. This means that they can be used for even the most data and bandwidth intense applications, like real-time vision systems, or motion control.

They will also have high speed graphics capabilities through their PCI Express (PCI-e) application interface. Introduced by Intel in 2004, PCI-e is intended to replace the original PCI local bus, the AGP graphics interface, and even the PCI-X high speed bus. This will make possible HD (high definition) video graphics on an embedded computer!

The very high speed capabilities of the PCI-e channels also give the embedded computing platform the ability to operate very high speed communications, such as 10 Gigabit Ethernet, and to perform in RAID (redundant array of independent disks) applications.

PCI-e isn't the only option for expansion in the next generation of embedded computing platforms, either. PC/104, PCI-104, PC/104+, PCI, and PCMCIA connections are also available. This provides an embedded system with openness to third party devices using a variety of standard bus interfaces.

In addition, some suppliers will integrate industrial I/O (including isolated digital I/O, RS232/422/485, and perhaps even analog I/O) directly on the board so the embedded computing platform can be used as a single board controller for industrial applications.

It is doubtful that the industrial controls user of 1965 would be able to imagine the changes that Moore's Law would bring. As Moore's Law continues to operate, its effects will continue to be felt on the next generation embedded computing platform, and the next, and the next.

###

