

Start-up Unveils Flexible Supercomputing Approach



A computing pioneer has developed an innovative approach that would make supercomputers easily programmable and useful for multiple tasks, instead of just one as has traditionally been the case with high-performance machines.

Convey Computer cofounder and chief scientist Steven J. Wallach has developed the HC-1 supercomputer, which uses microprocessors, field-programmable gate arrays (FPGAs), Linux, and the open source Open64 compiler. Its price range starts at about \$32,000.

Convey built its computer for use in a broad range of scientific and engineering applications, particularly those that are easily vectorizable and parallelizable, according to Wallach.

Typically, he explained, supercomputer software is designed specifically to solve a single class of complex problems very quickly and thus is generally inefficient at other types of tasks.

Some designers have tried to build machines with enough performance to solve multiple types of problems by simply adding more microprocessors.

Some supercomputers also use FPGAs, graphics processing units, and other types of chips for performance acceleration. However, because these resources aren't tightly integrated in typical supercomputers, the machines require extra compilers and additional management resources to handle tasks such as cache coherency. In

addition, they require complex programming techniques.

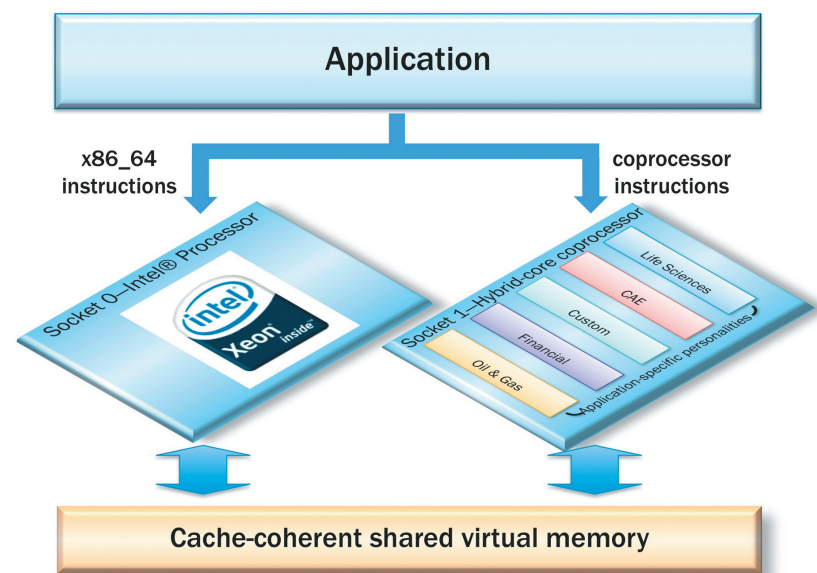
According to Wallach, his goal was to design a flexible supercomputer that would offer a standard programming environment that appears to developers as a single x86 instruction set, which is not the case with typical supercomputers. Convey accomplished this by tightly integrating the microprocessor with the FPGAs on each node.

The single environment makes the HC-1 relatively easy and inexpensive to program. The easiest-to-program systems always win in the marketplace, Wallach said.

The Convey computer works with multiple nodes. In a single node, the machine uses an x86 microprocessor and, for flexibility, 18 Xilinx FPGAs, which are programmable logic devices composed of an array of configurable logic cells. The programmability lets the system adapt easily to new applications.

A bus links the cache memories of the microprocessors and the FPGAs, which provides the system with fast access to cached data.

According to Wallach, a single Convey processor node is 10 to 100 times faster than nodes used in typical supercomputers.



Convey Computer has developed an innovative supercomputer that would be easily programmable and useful for multiple tasks, instead of just one as has traditionally been the case with high-performance machines. The HC-1 computer uses an architecture that includes a standard x86 chip and multiple field-programmable gate arrays designed for specific functions on each tightly integrated node. The FPGAs add flexibility and easy programmability to the HC-1.

Larry Getman, Xilinx's senior director of corporate development, said his company has backed Convey because it wants to invest in companies with innovative applications for the firm's FPGAs. Intel has also invested in Convey.

So far, supercomputers have made limited use of FPGAs' programmatic

flexibility, according to Rajesh Gupta, chair of the University of California, San Diego's Department of Computer Science and Engineering.

Convey has taken an innovative approach to utilizing this flexibility to improve supercomputer performance, particularly in allowing developers to make maximum use of FPGAs' pro-

grammability, Gupta said.

The system could also help increase FPGA use in mainstream computing, he added.

Wallach has been designing high-performance computers since the early 1970s and has been active in several US government studies on supercomputing. **■**

New API Helps with Multicore-Chip Communications

An industry association has developed a standardized API designed to help developers of software that will run on multicore chips efficiently build applications that can run on multiple vendors' multicore architectures.

The Multicore Association—which creates industry-standard approaches for multicore-processor-related technology and certifies products as conforming to standards—has completed and released its Multicore Communications Application Programming Interface.

Proponents say MCAPI is important because there is concern over the ability of application vendors to write software to run over the increasingly popular multicore processors.

MCAPI would make it easier for vendors to write applications for multicore systems, said Eric Heikkila, industry analyst with VDC Research Group. The availability of more applications would encourage increased adoption of the systems, he explained.

The new standard provides a common way to implement the communications necessary to run an

application on a multicore processor, regardless of the chip manufacturer's own intercore API.

The technology would scale easily and address communications issues such as sending and receiving data, as well as synchronizing and managing processing elements, threads, and end points.

Today's multicore-chip vendors offer proprietary communications APIs for their processors. This forces application developers to spend time and money rewriting their software for each multicore processor on which they want their programs to run.

With MCAPI, application developers only have to write to the common API, regardless of the underlying processor, said Multicore Association president Markus Levy.

And by letting developers write code only once, the standardized communications API also makes it easier for them to try different mappings of an application across multiple cores for optimal parallelization and performance, explained Sven Brehmer, president of PolyCore Software and chair of the Multicore Association's MCAPI Working Group.

According to Brehmer, MCAPI probably will be implemented primarily in multicore processors. However, it could also handle communications for multiple processors on a board.

The technology would run efficiently and could be used in many computing scenarios, including embedded applications and streaming approaches such as those used in signal, network-packet, and multimedia processing.

Efficiency would be critical for applications running on processors that have tight execution and memory constraints and that require reliable on-chip interconnects and high system throughput.

The Multicore Association designed MCAPI with help from various member companies, including Freescale Semiconductor, Intel, and Wind River.

Now, chip makers and software developers will have to adopt the standard for it to be successful, said John Carbone, vice president of marketing for Express Logic, a vendor of real-time operating systems and development tools for embedded systems.

So far, PolyCore and Express Logic have supported it. **■**