

M I C R O P R O C E S S O R

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

MASSIVELY PARALLEL DIGITAL VIDEO

Fabless-Semi Startup Connex Reveals New Processor Architecture

By Tom R. Halfhill {1/9/06-01}

Three things in life seem certain: death, taxes, and new microprocessor architectures. Unlike the first two things, new architectures aren't necessarily bad, but they are becoming even more expensive. These days, processor architects need good reasons for instigating the

years of engineering and \$50 million or so investment required to bring a new chip to market.

Almost all new architectures introduced in recent years are unconventional and designed for relatively narrow classes of applications. (One reason is that almost all the conventional, general-purpose architectures have already been invented.) New architectures also tend to emphasize data parallelism, often on a massive scale, because architects have nearly exhausted the ideas for optimizing instruction-level parallelism in everyday software.

So perhaps it's no surprise that the latest new microprocessor architecture to emerge is unconventional, massively parallel, and optimized for the narrow domain of high-definition (HD) digital video. Although Connex Technology's architecture is applicable to other purposes—such as pattern-matching filters in security processing—digital video is the largest potential market offering an opportunity for a profitable return on investment.

One twist is that the chief scientist behind the new architecture is Dr. George Stefan, a professor of electrical engineering at the Polytechnica University of Bucharest in Romania. Stefan worked on his idea for 20 years before finding backers to help launch Connex in 2002. (The company's original name was Gemicer, an amalgam of "Gemini" and "Cancer"—the astrological signs of one cofounder's children.) Among Stefan's backers was John H. Sununu, former governor of New Hampshire and White House chief of staff under President George H.W. Bush.

Stefan has invented a concept he calls the integral parallel architecture. It orchestrates a massive array of simple processor cores to execute data-intensive algorithms. (The first Connex test chip, produced in December 2003, had more than 4,000 processor cores.) These cores have simplified instruction sets and can work together in different software-controlled configurations, adapting to the requirements of different algorithms and workloads. Programmers write the software using a modified version of C that adds a new vector datatype for parallel operations on video datastreams. Another essential element is an on-chip memory controller that vectorizes the serial datastream for the parallel-processor array.

Connex, a fabless semiconductor company, is now headquartered in Silicon Valley and emerging from stealth mode. Connex has disclosed its new architecture to *Microprocessor Report* but won't officially announce a chip until the first working silicon returns from the fab—perhaps as soon as this quarter (1Q06).

Enough Horsepower for Dual-Channel HD

Connex has defined three goals for video processors based on its new architecture. First, the processors must be competitive with the performance and cost of ASICs. (Specifically, the processors must be able to encode, decode, and transcode multiple channels of HDTV and require little or no additional silicon than a custom chip would.) Second, the processors must be fully programmable, so they can support different HDTV standards around the world and adapt to

evolving standards. Third, the processors must exploit the data parallelism inherent in digital video while preserving an easy-to-use sequential programming model in high-level C.

Achieving all three goals without making significant compromises is the challenge. Although nobody can change the laws of semiconductor physics, Connex can slightly bend the laws by narrowing the application domain. A processor that needs to do nothing but execute a few algorithms critical to digital video can discard some excess baggage found in general-purpose processors and DSPs. Software-development tools for such a processor can be highly optimized, because they needn't target a broad range of applications. (See the sidebar, "The Key to Massive Parallelism: Think Small.")

The first processor from Connex, known for now as the Connex Machine, will target HDTV and standard-definition (SD) video. Future Connex processors may target low-power portable products, but the initial chip is strictly a high-performance device for AC-powered systems. It will be able to encode, decode, or transcode HD-quality video using the MPEG-2, H.264/VC-1, and Windows Media 9 (WM9) codecs while performing various scaling and filtering operations. It can simultaneously decode two channels of HD video, which allows picture-in-picture viewing or concurrent dual-channel viewing and recording. The Connex Machine can also transcode an MPEG-2 bitstream into an H.264 stream, allowing a system to redistribute HDTV programming from room to room over a home network. Of course, a processor with that much horsepower can easily handle SD video, too. It can decode four SD video streams simultaneously or decode two SD streams while encoding or transcoding two additional SD streams.

All together, those capabilities will rank the Connex Machine among the elite of video processors. Yet it should be

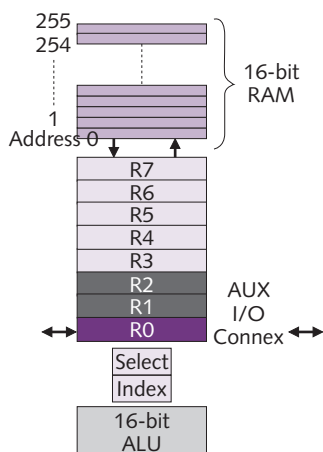


Figure 1. The Connex integral parallel architecture is based on a massively parallel array of processor cores known as processor elements (PE). In the first commercial implementation, each PE is a 16-bit processor core with an ALU, register file, and local memory. One unusual feature is a Boolean unit and select-index register that allow programs to specify which PEs work together during a vector operation.

an economical chip, because the company's foundry partner (undisclosed for now) will manufacture the device in a common 0.13-micron CMOS process. To deliver such high performance while retaining high-level programmability, the Connex architecture focuses exclusively on data parallelism, using a massive array of processor cores. The first commercial chip will have 1,024 cores.

As Figure 1 shows, each 16-bit processor core—or processor element (PE), in Connex terminology—has an ALU, eight 16-bit-wide registers, and 256 words of local memory (SRAM). Each PE also has a special select-index register for a Boolean unit, which allows programs to control which PEs will join together to perform a particular vector operation on a chunk of data. (More on this later.) All PEs in the array are identical and have a stripped-down instruction set of about 70 instructions, mostly 16-bit integer and Boolean operations.

Missing from the PEs are any features deemed unnecessary for the parallel operations of video decoding. Most notably, they have 16-bit datapaths and registers, not the 32-bit structures found in most other processors adapted for digital video. Unlike video decoders based on general-purpose processors, the Connex Machine has no legacy software to support. The PEs also lack floating-point hardware, as well as dedicated function units or hardware acceleration for multiply-accumulate (MAC) instructions.

Omitting the MAC hardware is particularly daring. Connex acknowledges that about 15% of the operations in video decoding and about 10% of the operations in video encoding require MACs. However, the company says its PEs can duplicate the functions of a MAC unit and achieve high performance by using Booth's algorithm, which requires only a few shifts and adds to multiply signed integers.

Processor Array Weaves a Simple Fabric

Another simplification of the Connex Machine is the on-chip network that permits all 1,024 PEs to communicate with each other and with the integrated memory controller. Most other massively parallel processors have mesh-like interconnect fabrics that allow each processor core to communicate with all neighboring cores in various ways. Examples include the massively parallel processors recently introduced by ClearSpeed (see *MPR 1/12/04-02*, "ClearSpeed Hits Design Targets"), Elixent (see *MPR 6/27/05-02*, "Elixent Improves D-Fabrix"), IBM Microelectronics (see *MPR 10/11/04-01*, "IBM Makes Designer Genes"), and PicoChip (see *MPR 10/14/03-03*, "PicoChip Makes a Big MAC").

Mesh fabrics are flexible but require complex webs of wiring. The wiring can quickly get out of hand, especially with fully synthesized designs relying heavily on electronic design automation (EDA) tools. Indeed, the wiring may occupy more die area than the processor cores do. Signal-propagation delays often restrict the size of such fabrics, limiting the amount of processing power the chip can bring to bear on a particular task. After taking a hard look at the complexity of a mesh fabric with 1,024 PEs, and after evaluating

the algorithms targeted by a digital-video processor, Connex steered a different course.

The Connex Machine’s interconnect fabric allows the PEs to communicate with each other along an “east-west” axis only—in other words, in a serial chain. Each PE can interact directly with only two other PEs. As Figure 2 shows, the first implementation of the Connex Machine arranges the PEs in a two-dimensional array that is 16 PEs wide by 64 PEs deep. Arrays of many other sizes and configurations are possible with this architecture, depending on the design target; the first test chip had a 64 × 64 array of 4,096 PEs. Connex settled on a 16 × 64 array for the Connex Machine because it’s sufficient for the most demanding task of simultaneously decoding two channels of H.264 video at HD resolution.

Alongside the array is the on-chip I/O controller, which provides the interface to external memory. In the first implementation of the Connex Machine, the controller has a 64-bit DDR-DRAM interface providing 3.2GB/s of peak bandwidth. The interface between the controller and the array is buffered to prevent the PEs from stalling on I/O dependencies. During every clock cycle, the I/O controller can read or write to local memory in 64 PEs—one column of PEs in each of the 64 rows of the array. Because the 16 PEs in each row are chained together serially, data must traverse the row to move among the PEs. A signal takes 16 cycles to propagate completely through a row of the array.

Connex has carefully matched the PE array’s size and the I/O controller’s bandwidth to the target algorithms. As Figure 3 shows, the Connex array can store 256 vectors, each with 1,024 values of 16 bits, in the local memories of the PEs. A simple stack-machine controller called the sequencer (seen in Figure 2) directs the PEs to execute algorithms on this data. In parallel with these operations, the I/O controller transfers data to and from local memories by assembling 1,024-component vectors fetched from external memory.

The I/O controller assembles the vectors by storing each 16-bit value in register R7 of each PE. When the PEs finish operating on the data, the I/O controller transfers the results back to external memory. Because Connex has matched the amount of external memory bandwidth to the size of the array, the chip isn’t I/O bound, except when simultaneously decoding two channels of H.264 video in HD. Usually, the I/O controller handles memory transfers in the background, without stalling the PEs.

A Version of C You May Never See

To supervise the vast array of PEs, the Connex Machine uses the sequencer seen at the top of Figure 2. The sequencer is a simple stack-machine processor running under software control, and it activates the PEs that will participate in a particular vector operation. Some operations don’t require the services of every PE. A program can activate only the PEs needed during each clock cycle, idling the other PEs to save power.

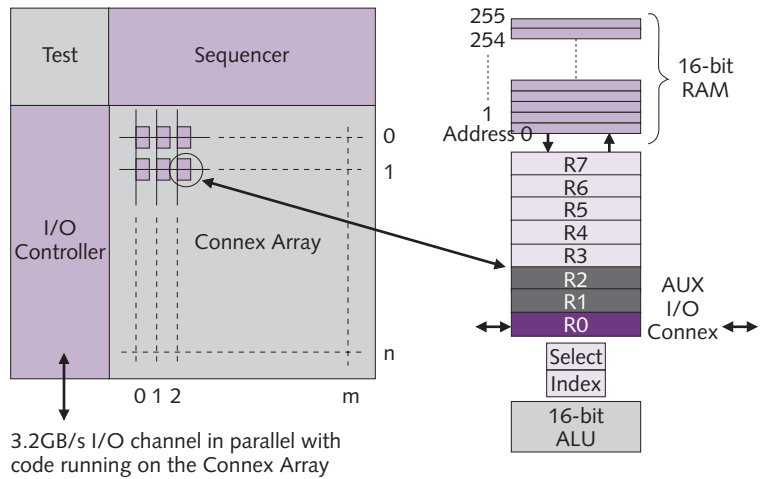


Figure 2. The Connex architecture arranges numerous PEs in a two-dimensional array, much like other massively parallel architectures, but it simplifies the on-chip interconnect fabric by severely limiting the connections among the PEs. Each PE can communicate directly with only two other PEs along the horizontal axis of the array. An integrated memory I/O controller keeps the PEs fed with data. The sequencer at top is a simple stack-machine controller for supervising the array. The array’s size and the width of the memory interface can vary from one chip design to another, according to the performance required for a particular application.

This feature is very flexible, but it requires using a special version of C—albeit one that few customers will ever see.

Connex Programming Language (CPL) is a proprietary extension of ANSI C. Among other things, it defines a new datatype (vector) and a new command (sel) for operating on vector data. The sel command tells the sequencer which PEs to activate for a particular vector operation. For example, the

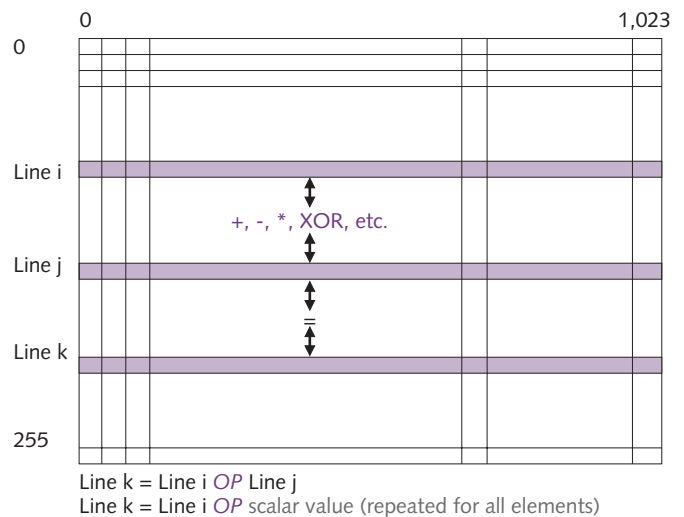


Figure 3. The Connex Machine can operate on 1,024 words of data simultaneously. These 16-bit words are arranged in a single-dimensional array or vector. The Connex Machine can perform most basic integer and Boolean operations on these vectors in a single clock cycle. For instance, it can add all the elements in two vectors and store the results in a third vector, or it can apply a scalar mask for a Boolean operation across an entire vector.

statement `sel = r3 > 100` selects all PEs having a value greater than 100 in register 3. The statement `sel = !sel` reverses the previous selection—in this case, it would select all PEs having a value of 100 or less in register 3 (equivalent to `sel = r3 <= 100`). The `sel` command is flexible enough to select any individual PE or almost any combination of PEs in the array.

Figure 4 shows a snippet of code written in CPL. The first section defines four vectors (`x`, `y`, `min`, and `max`) of 16-bit integers. Essentially, these are one-dimensional arrays, except the programmer needn't specify the dimensions, as is necessary when creating conventional data arrays. The CPL compiler adjusts the dimensions of the vectors to fit the processor's array of PEs. The next section of code selects all the PEs (`sel = all`) for the vector operations that follow. Finally, the vector operations find the minimum and maximum elements in the `min` and `max` vectors. These vector operations execute in only four clock cycles, because the processor evaluates an entire vector at a time.

CPL is an interesting variation of C, but few customers need worry about it. Connex says it will provide industry-standard audio and video software codecs for the Connex Machine, so customers will have little or no reason to interact directly with the processor. Initially, Connex doesn't plan to ship the CPL compiler or any other software-development tools with the processor, unless a customer needs to port a proprietary codec that Connex doesn't offer. For those cases, Connex offers a software-development kit. Most developers will continue writing their application software in conventional C, then link their code to the precompiled binaries of the standard codecs written in CPL. This model makes the Connex Machine transparent to programmers and eliminates a major barrier to adopting an unfamiliar, massively parallel microprocessor architecture.

Specialization Boosts and Limits Performance

The Connex integral parallel architecture, at least as implemented in the first Connex Machine, is highly optimized for

```

{
  ...
  const short OFFSET = 15;
  ...
  short vector x, y;
  short vector min, max;
  ...
  sel = all;
  x += OFFSET;
  ...
  min = x;
  max = x;
  min = (min > y)? y; /* min = min(x, y) */
  max = (max < y)? y; /* max = max(x, y) */
  ...
}

```

Figure 4. The Connex Machine can use its array of PEs to find the minimum and maximum values in two vectors of 16-bit integers by examining entire vectors at a time. Connex created a proprietary version of ANSI C, known as Connex Programming Language (CPL), that adds new vector datatypes and commands.

some algorithms but inefficient for others. This is a common characteristic of extreme architectures. Obviously, a massively parallel processor isn't the best choice for running the branch-ridden sequential code of business productivity software. However, Connex is carrying specialization to an even higher level with a massively parallel architecture that severely restricts the way the PEs communicate through their interconnect fabric. Connex readily acknowledges that some tasks, even within the domain of digital video, are unsuitable for the Connex Machine.

One such task is Context-based Adaptive Binary Arithmetic Coding (CABAC), a compute-intensive compression algorithm integral to the H.264 codec. CABAC bitstreams are sequential streams of data offering few or no opportunities for parallelism. That's why some video processors offload CABAC decoding to a separate stream processor. The Connex Machine does this, too, by integrating a stream processor (a group of eight PEs) with a licensed 32-bit RISC core on chip. In fact, the Connex Machine will integrate multiple RISC processors for operations that don't map well onto the PE array.

Beyond digital video, the Connex Machine has only a few applications. One possibility is running pattern-matching algorithms on high-bandwidth datastreams for network security or electronic surveillance. Security processors often have special hardware for scanning data packets, which may carry telltale signs of viruses and other malware. In surveillance applications, a processor uses pattern-matching algorithms to watch for certain character strings or other clues.

Connex has sold a few evaluation boards to In-Q-Tel, a government-sponsored venture capital fund that promotes research and development for the Central Intelligence Agency. Each board has four Connex test chips—the early devices with larger arrays of 4,096 PEs. Presumably, the CIA is interested in these powerful processors for electronic surveillance, code breaking, or some equally mysterious mission. However, Connex doesn't view the deal as a major business opportunity. The spooks pay well, but their purchase volumes are too low for Connex to earn a significant return on investment.

Digital video offers a better opportunity. Falling prices of flat-panel displays, coupled with government mandates to phase out analog TV broadcasting, will drive a massive turnover of TVs over the next several years. (Congress keeps vacillating, but the deadline for the transition to digital TV will probably come in 2009.) When analog broadcasts fall silent, even the holdouts who keep their old TVs will have to buy digital converter boxes—possibly subsidized by the government. Of course, Connex isn't the only company to recognize this huge opportunity. Competitors are flocking to the digital-video market from all points on the compass.

Everyone Wants Their HDTV

The competition breaks down roughly into three classes of devices: programmable video processors like the Connex Machine, with varying degrees of integration; fixed-function ASICs, or ASICs with limited programmability; and

The Key to Massive Parallelism: Think Small

If “Connex Machine” sounds vaguely familiar, some readers may be remembering an earlier attempt at massively parallel processing called the Connection Machine. The Connection Machine was a massively parallel supercomputer from a Massachusetts-based company called Thinking Machines, founded in 1982.

The largest Connection Machine had more than 64,000 simple processors. A later model with 128,000 processors was designed but never finished. The grand ambition for these supercomputers was to mimic the thought processes of the human brain by applying massively parallel processing to a wide range of applications. At first, Thinking Machines believed its technology could be applied to aeronautics, finance, genetics, medicine, physics, data mining, digital imaging, and more. Unfortunately, programming the Connection Machine was more difficult than anticipated, and the company was forced to narrow its vision. Thinking Machines declared bankruptcy in 1994, and some of its technology ended up at Sun Microsystems.

There’s no direct connection between the Connection Machine and the Connex Machine, except that Connex has evidently learned a few things from the woes of its predecessors. One lesson—recognized by other companies as well—is that massively parallel machines are more practical when they focus on smaller, well-defined problems. By concentrating on data-parallel algorithms in digital-video codecs, Connex can succeed where Thinking Machines failed.

Among other companies reaching similar conclusions are IBM Microelectronics, Elixent, and PicoChip. IBM’s

BlueGene/L supercomputer, with 65,536 dual-core PowerPC processors, matches the scale of the biggest Connection Machine ever conceived. (See *MPR 10/11/04-01*, “IBM Makes Designer Genes.”) Lawrence Livermore National Laboratory is using BlueGene/L mainly for simulating tests of nuclear weapons. Other BlueGene supercomputers will be used for very specific aspects of medical research.

Elixent’s massively parallel D-Fabrix processor cores have been adopted by Matsushita Electric and Toshiba, mainly for accelerating data-intensive multimedia and communications algorithms. (See *MPR 6/27/05-02*, “Elixent Improves D-Fabrix.”) Extracting data parallelism from those applications is much easier than finding instruction-level parallelism in other types of software.

PicoChip designed its massively parallel PC101 and PC102 processors to replace multiple chips in next-generation cellular and wireless networks by executing data-intensive communications algorithms. The PC102 integrates 344 processor cores, including 260 multiply-accumulate (MAC) processors. (See *MPR 10/14/03-03*, “PicoChip Makes a Big MAC.”)

None of these massively parallel machines is as easy to program as a uniprocessor system. However, they are capable of delivering very high performance within their relatively narrow application domains, and they are fully programmable, unlike fixed-function ASICs. The lesson seems to be that massively parallel processing is more successful when it thinks small, not big.

programmable processor cores sold as licensable intellectual property (IP) for integration in custom SoCs.

MPR has recently covered some of the licensable IP in detailed articles, which include surveys of the competitive landscape. (See *MPR 11/28/05-01*, “Tensilica Previews Video Engine”; *MPR 11/21/05-01*, “ARC Shows SIMD Extensions”; and *MPR 11/7/05-01*, “Videantis Chases Digital Video.”)

The configurable processors from ARC International and Tensilica are general-purpose 32-bit RISC cores enhanced with custom extensions for digital-video decoding. Although they deliver high performance for general-purpose architectures, these particular configurations can’t match the Connex Machine’s multichannel decoding and encoding at HD resolution. The ARC and Tensilica cores are better suited for low-power portable systems. The same is true for the new TriMedia TM3270 video-processor core recently announced by Philips Semiconductors. (See *MPR 12/5/05-01*, “Philips TriMedia Goes Mobile.”)

Videantis, a German startup, is more specialized. Its programmable v-MP2 core is based on a clean-slate architecture

designed solely for video processing. Videantis licenses two video coprocessor modules using the v-MP2 core: the v-MP2000M, a single-core design for mobile systems, and the v-MP2000HD, a triple-core design for HD systems. The v-MP2000HD can decode H.264, MPEG-2, and WM9 video at HD resolutions, but it’s not powerful enough to decode two HD channels simultaneously or to perform HD encoding and transcoding.

Customers preferring to buy their chips off the shelf turn most often to Broadcom, which has a large catalog of video processors, audio processors, and support chips. Two examples of dual-channel HD video decoders from Broadcom are the BCM7038, based on a 64-bit MIPS core, and the BCM7411, which decodes H.264 as well as MPEG-2. Direct comparisons with the Connex Machine must wait until later, when Connex publicly discloses all the features of its chip—particularly the integrated peripherals and I/O interfaces. However, Connex has privately disclosed some information to *MPR*, indicating that the Connex Machine will have a fairly high degree of integration. In 2007, Connex

For More Information

Connex Technology will officially announce the first commercial chip based on its new architecture after receiving working silicon from its foundry partner, which could be as early as this quarter (1Q06). Details about the chip's features and volume pricing will be announced at that time. For more information about Connex, visit www.connextechnology.com.

plans to introduce a second-generation video processor with considerably more integration. Of course, Broadcom won't be standing still, either.

ASICs are another obstacle for Connex. Volumes in the consumer market are certainly large enough to justify developing ASICs. Even with today's soaring costs of mask

sets and nonrecurring engineering (NRE), some leading consumer-electronics companies prefer to spin their own silicon. However, fixed-function ASICs or those with limited programmability are less versatile than fully programmable chips like the Connex Machine, especially if the product developers hope to accommodate the smorgasbord of digital-video standards throughout the world. If the Connex Machine can match the performance and economy of large-volume ASICs—a challenging goal—programmability will tilt the balance in its favor.

In three years, Connex has intelligently molded a university professor's pet project to fit an application that leverages the architecture's capacity for massive data parallelism. And Connex has targeted an application that has lucrative business potential. If Connex can take the next step—deliver a commercially viable implementation that fulfills the promise of the architecture—the market will have a formidable new competitor. ♦

To subscribe to Microprocessor Report, phone 480.483.4441 or visit www.MPRonline.com