Cahners
MICROPROCESSOR REPORT
www.MPRonline.com
THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

CHAMELEON CROSSES CPU, FPGA

Reconfigurable Processor Targets Communications

By Peter N. Glaskowsky {6/12/00-01}

Three simple steps are essential to any successful product—identify a problem, solve the problem, deliver the solution. Too often, media-processor companies have created solutions without first defining the problem they're trying to solve. Others have failed to deliver a complete solution—or sometimes anything at all.

Chameleon Systems seems to have avoided these common pitfalls in the design of its new CS2112 reconfigurable communications processor (RCP), the first chip in a family of RCPs Chameleon expects to introduce over the next year. As its name suggests, the RCP was designed for one specific class of applications: communications processing. Eventually, Chameleon's new architecture may be used in other ways, but the company has wisely decided to focus on one niche market before expanding into other areas.

The CS2112 includes a conventional 32-bit ARC RISC microprocessor core, a 64-bit memory controller, and a 32-bit PCI interface, but these are merely supporting actors. The real star of the show is a reconfigurable processing fabric (RPF) that consists of a two-dimensional array of 32-bit integer datapath units (DPUs) with associated memory. Figure 1 shows a top-level block diagram of the CS2112, which includes four "slices," each with three "tiles."

Figure 2 shows one of the 12 tiles in the RPF. Each tile contains seven DPUs, two integer multipliers, four 32-bit x 128-entry four-port static RAMs, and a control/logic unit (CLU) that manages data movement and processing. Each SRAM block can perform one read and one write per cycle, for an aggregate on-chip bandwidth of 48GB/s at 125MHz.

Using local routes, each DPU can receive operands from any DPU within eight locations above or seven locations below in each slice. In Figures 1 and 2, DPU 2 in tile B in slice 0 can read from the output register of DPU 2 in tile A, DPU 1 in tile C, or any DPU in between. DPUs can also read and write to the LSMs within a similar distance; all these routes wrap around between the top and bottom of each slice.

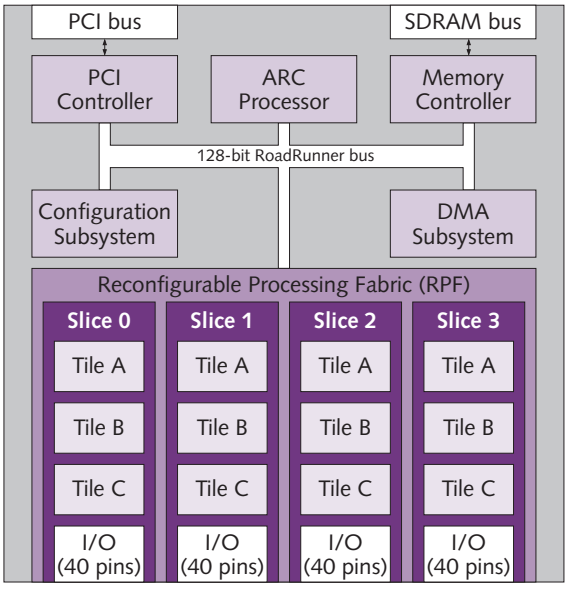

Figure 1. The CS2112 includes an ARC CPU core, PCI and memory interfaces, and a reconfigurable processing fabric (RPF) with a total of 84 32-bit ALUs, 24 integer multipliers, and 24K of distributed SRAM. Future family members will support more or fewer slices.

Additional intraslice routing resources allow DPUs to communicate directly with all other DPUs and LSMs within the same slice with no time penalty. Interslice routing resources connect each DPU with all the other DPUs and LSMs on the chip using one added clock cycle of latency. A limited number of these intra- and interslice routes are available, but Chameleon does not expect this limit to pose a problem for developers, since most connections use the dedicated local routes.

DPUs Perform Multiple Operations Per Cycle

Figure 3 shows a single DPU. A DPU is not a full processor but is more like a reconfigurable ALU. A DPU contains an instruction buffer that defines up to eight different 56-bit microcode instructions. On each clock cycle, the CLU in the parent tile sends three bits to each DPU to select one of the eight instructions. The instructions specify the data sources (selected by input-routing multiplexers in each DPU) and the operations to be performed.

Each DPU can perform multiple operations per cycle. A barrel shifter performs the expected shift operations and also handles word and byte swapping plus word duplication (but not byte duplication). The barrel shifter can also be configured to generate 5-bit constants.

The register and mask units resynchronize data before entering it into the operator unit and can also perform 32-bit AND/OR mask operations. The operator unit itself performs all classic ALU operations, signed and unsigned shift and mask functions, minimum and maximum selection operations, priority encoding, and saturating adds. The operator unit can process data as 8-, 16-, or 32-bit words; it also supports a SIMD mode, where 32-bit inputs are operated on as two independent 16-bit values. Output data and condition flags are made available on the DPU output registers.

Originally, Chameleon planned to have eight DPUs per tile, but discussions with communications-equipment makers led Chameleon to swap one DPU for two multiplier units.

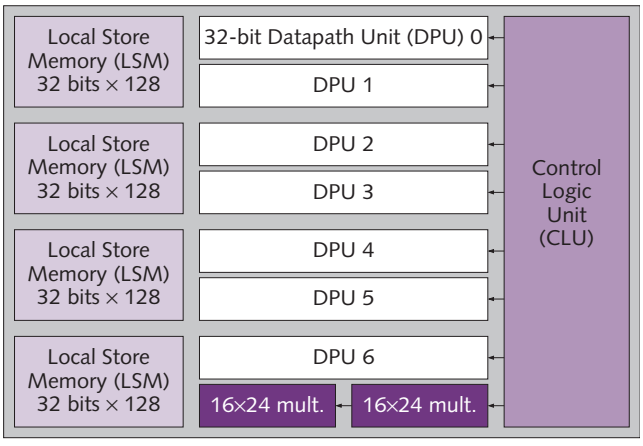

Figure 2. Each RCP tile contains memory, processing resources, and control logic.

These units can perform 16-bit x 16-bit = 32-bit signed multiplies and can also support 16-bit x 24-bit multiplies, where the eight LSBs are truncated to create a 32-bit signed result. With a total of 24 such multipliers, a 125MHz CS2112 can perform three billion multiply-accumulate operations per second—five times faster than TI's 300MHz TMS320C62x.

With all DPUs and multipliers busy, the theoretical performance of the CS2112 is 24 billion 16-bit integer operations per second, or 45GOPS if shift operations are also considered—a number that compares well with the peak-throughput claims of other application-specific processor vendors. These claims are of little practical value to those vendors' customers, but they can be effective marketing tools nonetheless.

With no floating-point hardware or hardwired function units for video-codec functions such as motion compensation or discrete-cosine transforms, the CS2112 is not worth considering as a general-purpose media processor. The DPU bit-masking functions and high degree of parallelism could make the CS2112 a good network processor, however, though Chameleon is not currently pursuing that application.

Programmability Is Limited But Adequate

Though just eight of 36 possible instructions can be selected on each cycle, Chameleon believes this approach provides more than ample flexibility for its target application. Chameleon says its customers report that no more than four different instructions are needed for any application.

Chameleon developed this unique architecture to match the normal design processes used to create signal-processing software. In this type of software, one data set is processed to create another data set by a sequence of operations that is both sequential (has little conditional branching) and readily parallelized. These programs may be represented by flowcharts in which operations are represented by a bubble, usually with two inputs and one output. Data flows down through the chart and eventually produces results at the bottom.

This type of flowchart may be converted by a Chameleon-savvy programmer into a CS2112 slice configuration, where data flows down through the slice until it

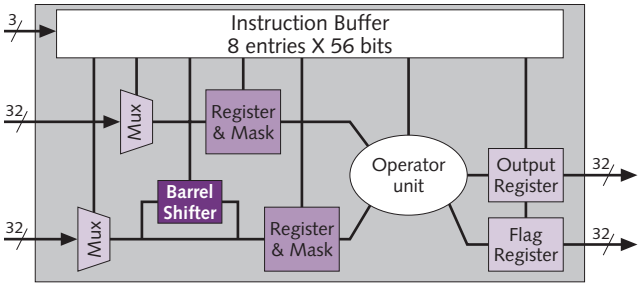

Figure 3. Each DPU contains routing multiplexers that select data sources, a barrel-shift unit, registers with AND/OR masking, an operator unit, and output registers for data and flags.

produces a final result that is saved to local or off-chip memory. By changing instructions from one clock cycle to the next, up to eight different operations can be performed by a single DPU over successive clock periods before the data proceeds down the stack. The data flow can be modified in flight according to the flags generated by each DPU operator unit.

Other programming techniques can be used, of course, and the precise configuration of processing, memory, and routing resources in the CS2112 may occasionally force the programmer to work around some limitation. For the most part, however, Chameleon believes its architecture is a natural fit for the way signal-processing software is developed and says its customers agree.

It's possible that Chameleon's architecture will be easier to manage than arrays of independent processors such as Cradle's Universal Microsystem (see *MPR 10/6/99-05*, "Cradle Chip Does Anything") or Silicon Magic's DVine (see *MPR 3/27/00-02*, "Silicon Magic: DVine-ly Inspired?"). The Chameleon approach involves just one program at a time managing dozens of parallel execution units, whereas these other chips host many independent programs. It will certainly be much more difficult to develop a complete configuration for a Chameleon chip than for any single element of a conventional multiprocessor chip; however, reducing, or in some cases eliminating, interprocessor synchronization overhead may more than make up for this difficulty.

Improving the CS2112's flexibility is its ability to hold a second complete set of configuration bits and switch between the active and the alternate configuration in just one cycle. All configuration data is stored in SRAM cells, so it can be reprogrammed at any time. It takes just three microseconds to load a new slice configuration into the backup set while the chip continues to run on the primary set. The whole chip can be reconfigured in 12 microseconds.

This capability allows the chip to be devoted entirely to one task at a time and reconfigured quickly when the chip is ready to perform the next required task. For example, the company describes an application for the CS2112 as a chip-rate processor in a third-generation (3G) cdma2000 cell-phone base station. (A CDMA "chip" is the code used to identify each channel sharing a single frequency—hence the name code-division multiple access.) In this application, the CS2112 must perform a sequence of four basic tasks—pseudorandom noise generation, demodulation, and two different search functions—every 1,250 microseconds. Four different configurations, one for each task, can be prepared ahead of time and the chip reconfigured between tasks with no lost cycles. In this application, Chameleon says a single CS2112 can process 50 channels of data, compared with just one or two channels on the conventional DSPs used today.

The company has also provided two other benchmark numbers. Chameleon says the CS2112 can perform a 1,024-point complex FFT in 10 microseconds or run a 48-tap symmetric finite impulse-response filter at 125 million samples per second.

The final element of each slice is a 40-pin configurable I/O subsystem. Eight of the pins are reserved for control, while 32 pins handle data transfers at up to 125MHz, giving the CS2112 a total of 2GB/s of peak I/O bandwidth. These subsystems are not as powerful or configurable as that found in Cradle's UMS, however, leaving more work for the processing fabric to handle.

CPU, Peripherals Round Out Chip Design

A 32-bit ARC core running at 125MHz is included to handle higher-level control tasks such as processing-fabric configuration. The core includes 4K of instruction cache and 4K of local data memory. All the configuration bits, local-memory arrays, and registers in the part can be read and written by the ARC core to simplify setup and debugging operations.

Additional I/O is available through hardwired PCI and memory controllers. The 32-bit, 33MHz PCI controller handles master and slave cycles and is compliant with version 2.2 of the PCI spec. The CS2112, however, is not able to act as a PCI system controller; it is meant to be used as a peripheral in a larger PCI-based system.

The memory controller, with its 64-bit, 125MHz memory bus, supports synchronous SRAM, SDRAM, and

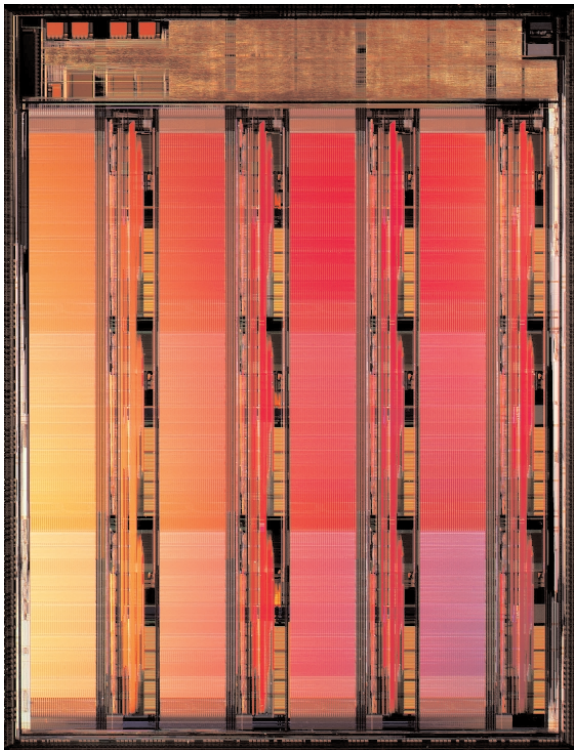

Figure 4. The CS2112 die is 300mm^2 in a 0.25-micron, four-layer-metal process. The die is dominated by the four slices of the reconfigurable processor array, each of which contains three tiles. The area at the top of the die contains the ARC CPU core, the PCI and SDRAM interfaces, and the global bus.

Price & Availability

Chameleon's CS2112, a 12-tile implementation of the CS2000 architecture, will be available in sample quantities in 3Q00 for $295 each in 100-unit quantities. The price of the chip is expected to decline to less than $100 by 2H01. The company also plans to introduce CS2103 and CS2106 derivatives at lower prices. For more information, visit Chameleon's Web site at *www.chameleonsystems.com*.

flash memory. Peak bandwidth is 1GB/s for SRAM and SDRAM; the flash mode supports narrower bus widths of 8 and 16 bits that reduce bandwidth but allow simpler, less-expensive implementations.

The ARC core, PCI and memory controllers, configuration manager, and RPF are connected by a 128-bit, 125MHz pathway Chameleon calls the RoadRunner bus. A DMA controller manages 16 simultaneous channels of DMA activity among the various modules on this bus.

Silicon Is on the Way

Chameleon has seen first silicon of the CS2112. Figure 4 is a photo of the CS2112 die, which is 300mm^2 in a 0.25-micron, four-layer-metal process. Though this chip is priced at $295 in sample quantities, Chameleon expects to sell it for less than $100 in volume by the second half of 2001. The next two members of this family, the CS2103 and CS2106, will have one and two slices (three and six tiles), respectively.

Along with silicon, Chameleon will deliver a software-development environment with a C compiler and a Verilog synthesizer, a chip-level simulator, a debugger, and a development board. The C language is used for the ARC core, while Verilog must be used to accomplish RPF configuration design. Chameleon provides a software layer to simplify the ARC core's management of the RPF.

A bit- and cycle-accurate simulator, while slow, permits full visibility into the internal operation of the chip. The development board allows native execution of code at higher speeds, but with reduced visibility. The full package is priced at $25,000 per seat.

As with other such products, Chameleon's software-development tools will make or break its hardware architecture. It is too easy for these products to run afoul of the 80/20 rule, where the last 20% of the problem requires 80% of the total work. Customers are all too often seduced by the ease with which the first 80% of the problem can be solved and don't always recognize that the true challenge lies in what remains.

With an architecture that makes it easier to break problems into independent tasks, each of which can be solved by a different chip configuration, Chameleon may help its customers avoid the 80/20 pitfall. The company's focus on just one application, and its early work in adapting the CS2112 architecture to this application and developing pieces of the needed software, make it more likely that Chameleon can turn its new architecture into a profitable business. Other applications and other chip designs may follow, but Chameleon's success depends on its ability to deliver complete solutions for each application before moving on to the next.