



Chip designers move back from the bleeding edge

By Ron Wilson
August 18, 2003

SAN MATEO, Calif. — The Hot Chips conference, which opens at Stanford University Monday (Aug. 18), has long been the showcase for processors that hit blistering speeds by way of advanced architecture, bleeding-edge processes and full-custom design. But this year two papers will offer a different perspective, touting unconventional semi-custom design methodologies that wring the same kind of speed from familiar architectures and ordinary processes.

The devices in question are an SPI-4 switch chip from Fulcrum Microsystems (Calabasas Hills, Calif.) and a vector signal processor from Telairity Semiconductor (Sunnyvale, Calif.). The Fulcrum chip achieves 1-GHz performance in a standard Taiwan Semiconductor Manufacturing Co. 130-nanometer process through a pragmatic blend of hard intellectual property (IP) and asynchronous design. The Telairity core is aiming for 600-MHz peak performance across a variety of 130-nm processes through use of a novel hierarchical library and methodology that implement circuits based on successive construction rather than synthesis.


Fulcrum began the journey as a spin-off from an asynchronous-design project at CalTech. But after several years of learning the bitter taste of evangelism, the company came to some conclusions. First, it was better to mix asynchronous design into existing flows than to attempt to convert the synchronous; and second, it was easier to sell chips than concepts.

Fulcrum's design flow today shows little of the zeal that characterized early advocates of asynchronous techniques, said vice president of product development Uri Cummings. The asynchronous tools—which Fulcrum is still discussing with Cadence Design Systems—can emit Verilog, and can be used in a variety of ways. Least doctrinaire of these is to simply use the Fulcrum tools to create multistage functional blocks that are internally asynchronous but are bounded by latches. To everyone except the block designer, the block appears conventionally synchronous. But internally, it runs very fast, dissipates little power and shows almost complete resistance to process, voltage and temperature variations.

Another, more aggressive blend of synchronous and asynchronous techniques produced the PivotPoint switch chip the company will describe this week. This six-port SPI-4 switch is built around the company's Nexus 16-port crossbar switch, which is fully asynchronous, non-blocking and exhibits a transit latency around 3 ns.

Fulcrum surrounded the Nexus switch with six full-duplex SPI-4 interface blocks. Fulcrum, with limited design resources, was aware that a fully asynchronous SPI-4 design would have significant benefits. But it would also require resources and interface expertise that weren't available, and it would entail a long, arduous cycle of compliance and interoperability testing. So the company opted to license the cores from Silicon Logic Engineering (Eau Claire, Wis.).

The result was an excellent set of interfaces and a chip that was working quickly, along with some understandable second thoughts. "The interface blocks behave beautifully," Cummings said. "But you look at



them and realize that a lot of design work went into making a synchronous design cover a huge number of corner cases that just never come up in real life. The synchronous blocks end up consuming most of the power in a 32 million-transistor chip. You just feel that extending the asynchronous design a little further into the interfaces could have saved a lot."

But Fulcrum had to ask not just if an asynchronous design would be better, but if doing the interfaces itself would be a strategic use of its engineers. "We look at that a lot," marketing vice president Mike Zeile acknowledged. "I think that in the future an efficient serializer-deserializer design implemented in a widely used interface like PCI Express might be considered strategic. Otherwise the economics probably won't justify the effort."

Telairity's paper, meanwhile, will take an entirely different route to a similar place. The startup was founded on a principle just as contrarian as Fulcrum's: that synthesis was not the best approach to digital logic design. Telairity's methodology rests on the successive construction of larger and larger functional blocks from a handcrafted cell library. The company assembles the basic cells into functional blocks, typically about 1,000 gates in size, that are hand-placed and routed in dedicated metal-1 through metal-3 routing. There are presently over 200 such blocks in Telairity's library, from data path and control elements to memory sections.

Just the right size

Telairity's key concept is simple: The hard blocks are just large enough that assembling them gives nearly all the optimization that would be possible from a fully custom design. But they are just small enough to not severely limit the flexibility available to the designer. Using something with a striking resemblance to schematic capture to put the blocks together, the digital designer can get much better than synthesis results, the story goes.

Telairity's initial aim was to launch a fabless ASIC company based on the block library. This proved less than practical in an environment of plummeting ASIC starts, shrinking design teams and epidemic conservatism. One of the company's demo projects was an ambitious vector signal processor. That design will now emerge at Hot Chips as a licensable hard macro portable to any reasonable 130-nm process.

Architecturally the design is aggressive, but not out of the mainstream. A 32-bit scalar RISC processor is paired on a high-speed bus with a bank of four 16-bit vector processors that support both conventional and chained vector operations from 32-element vector registers.

With all that horsepower, the potential throughput of the TVP400 is enormous. Consequently, so is its demand on memory bandwidth. Founder Howard Sachs estimates the peak memory bandwidth demand of the processing units at 38.4 Gbytes/second. The only way to support this appetite is to embed large memory structures next to the processing units.

That is where the Telairity methodology really comes into its own, Sachs said. A conventional ASIC methodology would couple synthesized logic with compiled memories to produce the processor core. But Telairity feeds its CPU with 600-MHz caches, and its vector units with a combination of a 12-port SRAM and a



crossbar switch, also operating at 600 MHz. Those memory structures are beyond the range of memory compiler techniques.

By constructing the TVP400 hard macro from its retargetable block library, Telairity can offer a piece of approximately 600-MHz IP that can be fabricated in a range of 130-nm logic processes. The company is estimating that the performance will be 600 MHz at 85 degrees C, dropping to 440 MHz at the slow/slow process corner. In the presence of large process variations, a conservative design could drop back to 400 MHz, Sachs suggested, leaving enough headroom to all but obviate signal integrity issues.

At this speed, and with the level of concurrency the core provides, the device should be able to handle a number of applications—such as HDTV set-top boxes—that are beyond the range of any signal-processing IP available today, Sachs maintained. It is an offering made possible by a novel design methodology enabling very fast multiport memories and fast data paths.