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

StarCore SC140 Gets Baby Brother

Half the Core Size, Half the Power, Less Computational Ability

By Steve Leibson {9/18/00-02}

StarCore, the Motorola-Lucent DSP-core alliance, has produced its second core design. The new SC110 is the baby brother of the previously announced SC140 (see *MPR 5/10/99-03* "StarCore Reveals Its First DSP") and bears a strong family resemblance. In fact, as Figure 1 shows, the major difference between the SC110 and the SC140 is that the SC110 has in its data ALU only one-fourth the number of computational units in the SC140.

StarCore created the SC110 variant because the SC140 is unnecessarily large and power hungry for many DSP applications that do not need the SC140's processing power, such as handheld 3G voice-only terminals. Eliminating three-fourths of the execution units present in the SC140's data ALU reduces the SC110's peak MAC rate by 75%. Overall computational power is reduced to roughly 50% of the SC140's, because the SC110 retains all the SC140's registers and non-data ALU faculties (two address-arithmetic units and a bit-manipulation unit). The SC110's data buses and dispatch-issue width are also half the size of the SC140's. Because of the narrower data buses, the SC140's quad 16-bit move instructions, dual 32-bit move instructions, and quad 16-bit Viterbi move

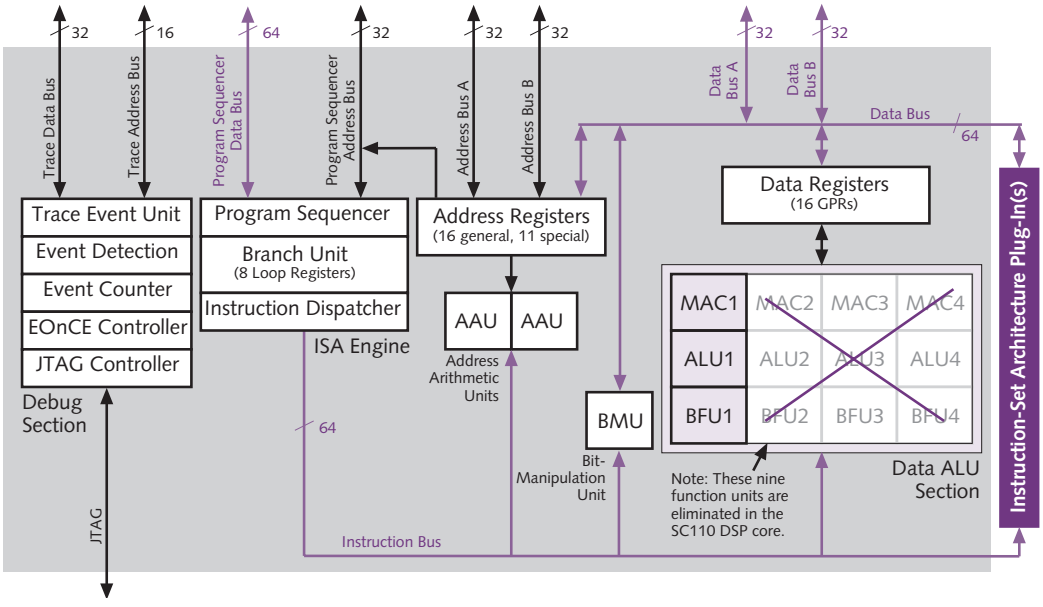

Figure 1. The major hardware difference between StarCore's SC140 DSP core and the new SC110 core is the elimination of three-fourths of the execution units in the data ALU. In addition, several buses are now 50% smaller than they are in the SC140, and the instruction-issue width is also cut by half. (Differences between the SC110 and the SC140 are shown in purple.)

Price & Availability

The SC110 is a DSP core that will be used by Lucent and Motorola to create chips. No such chips based on the SC110 have yet been announced. Information about the SC100 DSP- core family is available at *www.motorola.com/SPS/DSP/documentation/MSC8100.html* and *www.lucent.com/micro/starcore/doc.html*.

instructions are not present in the SC110 instruction set; these instructions require the wider 64-bit bus width. What remains after all the chopping and hacking is a DSP core that's upwardly code compatible with the SC140, executes 1 MMAC/MHz, runs at half the operating power of the SC140 core, and requires half the silicon.

That statement is not meant to imply that the SC140 is a power hog. Peak power for the SC140 is 239mW running at a core voltage of 1.5V and a core operating frequency of 300MHz. It's a mere 34mW at a core voltage of 0.9V and a core operating frequency of 120MHz. At the same operating frequencies and core voltages, the SC110 power dissipation numbers are 103mW and 15mW, respectively. These are power-dissipation numbers for the cores fabricated in a 0.18-micron process, not for chips based on these cores.

The SC100 family is designed to be programmed largely through a high-level language. The core design has a large, C-friendly register bank that's identical in the SC110 and the SC140. A new compiler switch allows software developers to create code for the SC110 and to run that code on existing chips based on the SC140 for prototype purposes until SC110-based silicon becomes available. Using the SC140 as a prototyping platform for the SC110 produces the same bit-exact results and execution times (as long as program and data address ranges are separate and there is no memory contention). The SC110 retains the SC140's VLES (variable-length execution set) architecture, which combines CISC-like, prefixed, variable-length instructions with a VLIW hardware architecture. The compiler seeks parallelism in the source code and groups instructions into bundles to exploit that parallelism.

The fewer computational resources and smaller buses in the SC110 do reduce the core's computational ability. In addition, the SC110's instruction bundles are limited to a maximum of four instruction words each (versus eight for the SC140), and the SC110 can issue only three instructions per clock (versus six for the SC140). Because the SC100 family's instructions range from 16 to 48 bits, the SC110's compiler will clearly need to use somewhat different packing algorithms for the SC110's 64-bit bundle size than those used for the SC140's much more generous 128 bits. For example, two 48-bit VLES instructions will fit in an SC140 bundle but not in an SC110 bundle. Again, this limitation means that code compiled for the SC110 will run on the SC140, but not necessarily vice versa.

In one way, the SC110's reduction in execution resources actually makes life easier for programmers. The SC140's quadruple set of data-ALU execution units is unquestionably harder to program in assembly language, as is any highly parallel DSP architecture. Manually scheduling the SC140's many resources will challenge even the most skilled programmer. However, the SC110 has far fewer computing resources, making the smaller core easier than the SC140 to program in assembly language.

StarCore plans to hand off the SC110 design to Lucent and Motorola in November. According to StarCore, the SC110 is intended for chip designs with manufacturing costs as low as $5–$10 in 0.18-micron process technology. Once the client companies receive the SC110 core, they will begin to develop system chips. Only after these chips have been designed and fabricated will the SC110 be available in silicon.