

Excerpts from

Inside the StarCore SC110



A Technical Evaluation
by the staff of
Berkeley Design Technology, Inc.

The following are excerpts and abridged text from BDTI's report, *Inside the StarCore SC110*.

Contents include:

- Introduction
- Scope
- The SC110 Core
- Benchmark Performance:
 - Sample Execution Time Results
 - Sample Memory Usage Results
- Summary

The complete report may be ordered from BDTI. Details are on page 4.

Introduction

In 1998, Agere (then Lucent) and Motorola announced the formation of a joint development center dubbed "StarCore." StarCore's primary goal is to develop next-generation DSP processor cores for use in Agere and Motorola chip-level products.

In 1999, StarCore announced its first architecture: the SC100. StarCore planned to introduce a family of upward compatible cores based on the SC100 architecture. Each family member would offer different levels of parallelism, performance, and cost.

StarCore's first implementation of the SC100 architecture was the SC140, a VLIW processor notable for its extremely high level of parallelism (even in comparison to other VLIW-based processors) combined with very low power consumption. The SC140 core targets high-performance applications, such as cellular base stations and gateways, and portable applications such as cellular terminal devices. At 300 MHz, the SC140

is one of the fastest general-purpose DSPs to be demonstrated in silicon.

Following the introduction of SC140-based devices such as the Motorola MSC8101, StarCore has introduced the SC110, a scaled-down version of the SC140 core. The SC110 is binary upward-compatible with the SC140, and offers lower cost and power consumption by including fewer execution units and less on-chip memory bandwidth. First announced in September 2000, StarCore plans to target the SC110 at applications requiring moderate DSP performance, low cost, and very high energy efficiency.

Inside the StarCore SC110 evaluates the DSP performance of the SC110 core. In this report, the technical staff of BDTI explore the key differences between the SC110 and the SC140. This report evaluates how the SC110 architecture addresses the needs of DSP applications; it includes both a detailed qualitative analysis of the SC110's architecture and

About BDTI

Berkeley Design Technology, Inc. (BDTI) was founded in 1991 to assist companies in creating, selecting, and using DSP technology. The technical staff of BDTI has extensive experience in the development of DSP-intensive software and hardware for commercial applications. BDTI offers a variety of technical products and services, including:

- *Published reports on DSP processors and technology*
- *DSP software development services*
- *Technical advisory services*
- *Training*

a quantitative evaluation of the SC110's performance on a series of DSP benchmarks developed by BDTI.

As of October 2001, neither Agere nor Motorola have announced chips based on the SC110 core. However, the initial products based on this core are expected to operate at 300 MHz using a 1.5-volt power supply.

The SC140 is covered in BDTI's companion report, *Inside the StarCore SC140*.

Scope

Inside the StarCore SC110 is intended for anyone interested in understanding the DSP performance and capabilities of the SC110 core or SC110-based products. It presumes a basic knowledge of DSP processor concepts and terms, which are covered in BDTI's *DSP Processor Fundamentals*. *Inside the StarCore SC110* is especially useful for electronic system designers, hardware and software engineers, processor designers, engineering managers, and product marketing managers. This report will aid in the assessment of the SC110's suitability for a given application, and it will allow system designers to make informed decisions when considering the SC110 for their latest designs.

For the purpose of comparison, this report includes brief analyses of several other processors: the Motorola DSP5685x, the Texas Instruments TMS320C55xx, and the StarCore SC140. These processors have been included to give the reader insight into

how the SC110 compares to other well-known DSP architectures.

The SC110 Core

The SC110 is a VLIW architecture that can execute up to three instructions at a time. Starcore refers to instructions grouped for parallel execution as an “execution set.” Instructions are scheduled for execution at compile time by either code-generation tools or the assembly language programmer.

The SC110 contains a single 16-bit data path, the “Data ALU” (DALU), which consists of a register set and a combined ALU/MAC/BFU. The BFU (bit-field unit) contains a 40-bit barrel shifter. The SC110 also includes an address generation unit that contains two address arithmetic units and one bit mask unit.

Because the hardware elements inside the SC110’s DALU are not independent, the SC110 can only execute a single DALU instruction per cycle. It is not possible, for example, to execute a MAC and a BFU instruction simultaneously. For this reason, in each group of three instructions executed in parallel, only one can use the DALU. In contrast, the SC140’s data path consists of four combined ALU/MAC/BFUs and can execute four DALU instructions simultaneously. This is the primary difference between the SC110 and the SC140. The remaining two instructions in an SC110 execution set can use the address generation unit to perform data moves, pointer arithmetic, or bit mask operations; or they can specify program flow-control instructions.

The SC110 DALU performs single-cycle 16×16 multiplications with 32-bit results, with or without accumulation. The multiplier supports all combinations of signed and unsigned inputs and also supports fractional and integer data (however, both source operands must be integer or fractional).

Like the SC140, the SC110’s DALU supports SIMD dual addition and subtraction by treating values in registers as packed pairs of 16-bit data operands. For example, using SIMD operations, the SC110 can perform two 16-bit additions

per instruction cycle (compared to eight on the SC140).

Memory System

The SC110 has three 32-bit address buses supporting a 64-bit program memory data bus and two 32-bit data memory data buses. Program and data memory are unified; any address can contain either instructions or data. The SC110 can perform two data reads, two writes, or one read and one write per instruction cycle.

On each of the SC110’s data memory buses, reads or writes can access contiguous groups of data up to 32 bits wide. With both data memory buses active, the SC110 can access up to eight 8-bit, four 16-bit, or two 32-bit words per cycle. At 300 MHz, the maximum on-core data memory bandwidth on the SC110 is 1.2 billion 16-bit words/sec.

In contrast, the SC140 contains a 128-bit program data bus and two 64-bit data memory buses. Therefore the SC110 can achieve only half the on-chip memory bandwidth of the SC140 at the same clock speed. However, because the SC110 contains a quarter of the SC140’s 16-bit execution units, the ratio of memory bandwidth to execution units is much more generous on the SC110 than the SC140. With few exceptions, the on-chip memory bandwidth of the SC110 is more than sufficient to keep its execution units fed, as illustrated by the SC110’s BDTI Benchmark™ results.

Instruction Set

The SC110 fetches four 16-bit instruction words at a time and can execute up to three instructions in parallel. (In contrast to the SC110, the SC140 fetches eight instruction words and can execute up to six instructions in parallel.)

The SC110 uses two different methods for specifying which instructions will be included in an execution set: serial grouping and prefix grouping

Serial grouping uses the two most significant bits in an instruction word to determine the end of an execution set. If these two bits are not zeros, the instruction is considered to be the last instruction in the execution set.

Prefix grouping uses a one- or two-word prefix for an execution set. A one-word prefix defines how many instructions are included in the execution set and also contains information used for conditional execution and hardware looping. Prefix grouping must be used if instructions are to be executed conditionally. Two prefix words are added to an execution set if it contains instructions that use registers in the higher-numbered halves of the register banks (D8-D15, R8-R15).

The choice of prefix or serial grouping is not specified by the assembly programmer; the assembler determines the grouping method. Serial grouping is used whenever possible to minimize instruction memory usage.

SC110 binary executable programs can be executed by the SC140 without modification, providing an upward migration path to a more powerful processor. However, while it is possible to modify SC140 code to run on the SC110, due to its smaller complement of execution units and further limitations in instruction grouping, the SC110 cannot directly execute binaries generated for the SC140.

The high level of orthogonality and relatively simple pipeline make the SC110 simpler to program in assembly language than many DSP processors, and the processor lends itself well to compiler code generation.

Pipeline

The SC110 uses a partially interlocked five-stage pipeline consisting of pre-fetch, fetch, dispatch, address generation, and execute stages. The first three stages are implemented in the program sequencer unit; the last two stages are implemented in the data path.

Despite the SC110’s five-stage pipeline, which is relatively short in comparison to DSPs such as the DSP5685x and TMS320C55xx, the SC110 spends many cycles on branches and change-of-flow instructions. However, the SC110’s ability to perform conditional instruction execution often allows it to avoid branches and jumps, and the SC110’s support for delayed instruction execution also reduces wasted cycles.

The SC110's five-stage pipeline is benign; pipeline hazards and branch cycle penalties can generally be avoided with very little programming effort.

Parallel Move Support

The SC110 supports operand-unrelated parallel moves by allowing up to two data moves to occur in parallel with other instructions. These data moves can be either one load and one store, two loads, or two stores. The access width of each parallel move can be a byte, a word (16 bits), or a long word (32 bits).

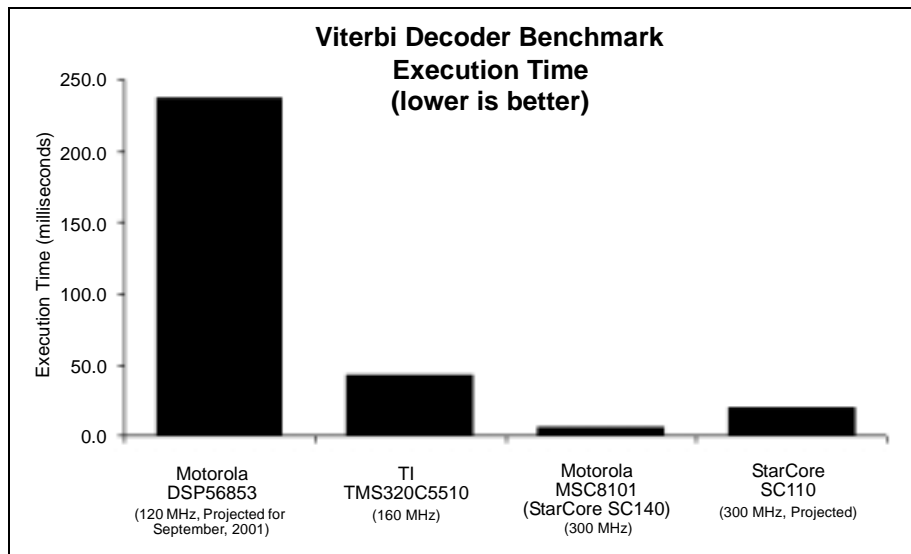
Benchmark Performance

Inside the StarCore SC110 includes extensive benchmark results, used to quantitatively evaluate processor DSP performance. For each benchmark, BDTI reports cycle counts, execution times, energy consumption, and memory usage. BDTI also provides extensive analysis of why the processors perform as they do. In this section, we present sample execution time and memory usage results taken from the complete set of results in the report.

Execution Time

To determine the execution time of a particular benchmark on a given processor, the number of instruction cycles the processor requires to execute the benchmark is multiplied by the processor's instruction cycle time. *Inside the StarCore SC110* includes tables and charts illustrating the number of cycles required by each processor to execute each benchmark and uses these results to generate corresponding tables and charts for execution times. The execution time results for the SC110 were obtained using a projected core clock speed of 300 MHz, which StarCore anticipates initial SC110-based chips to run at. Sample Benchmark Results

The execution time results for BDTI's Viterbi Decoder benchmark are shown in the figure above. As illustrated in this figure, the StarCore SC110 has a faster result on this benchmark than the Motorola DSP56853 and the TI TMS320C5510. The SC110 core has instructions dedicated to Viterbi decod-



ing (described in more detail in the full report); the SC110 is able to process 8 trellis butterflies in the Viterbi add-compare-select stage in 30 cycles by pipelining these special Viterbi instructions.

The TMS320C5510 also has special functions for Viterbi decoding that simplify the trace-back computations of the Viterbi algorithm. However, the architectural efficiency of the SC110 allows it to achieve a lower cycle count than the TMS320C5510 on this benchmark. The SC110's efficiency, combined with its 300 MHz clock rate, yield an execution-time result on BDTI's Viterbi Decoder benchmark that is over two times faster than that of the TMS320C5510.

About the BDTI Benchmarks™

The BDTI Benchmarks are a set of DSP software functions that BDTI has independently designed to provide an objective basis for comparing processor performance characteristics such as speed and memory use for DSP applications. The BDTI Benchmark functions are implemented in assembly language to allow a realistic assessment of processor DSP performance. The resulting software is then verified for functional correctness, optimality, and adherence to the BDTI Benchmark specifications. Benchmark performance results are obtained either through manual analysis and careful, detailed simulation, or by measurement on sample devices.

The highly parallel data path of the MSC8101 allows this processor to achieve a significantly faster result than that of the SC110 on the Viterbi Decoder benchmark.

Memory Use

Execution performance is often the primary metric used to compare processors. However, a processor's memory usage is also important. For example, the memory requirements of an application can have a significant impact on overall system cost. In addition, processors may experience significant performance degradation when application code and data do not fit in on-chip memory. Because of these and other factors, memory efficiency is an important metric in processor selection. For each of the BDTI Benchmarks, BDTI reports each processor's program, constant data, non-constant data, and total memory use.

Control Benchmark

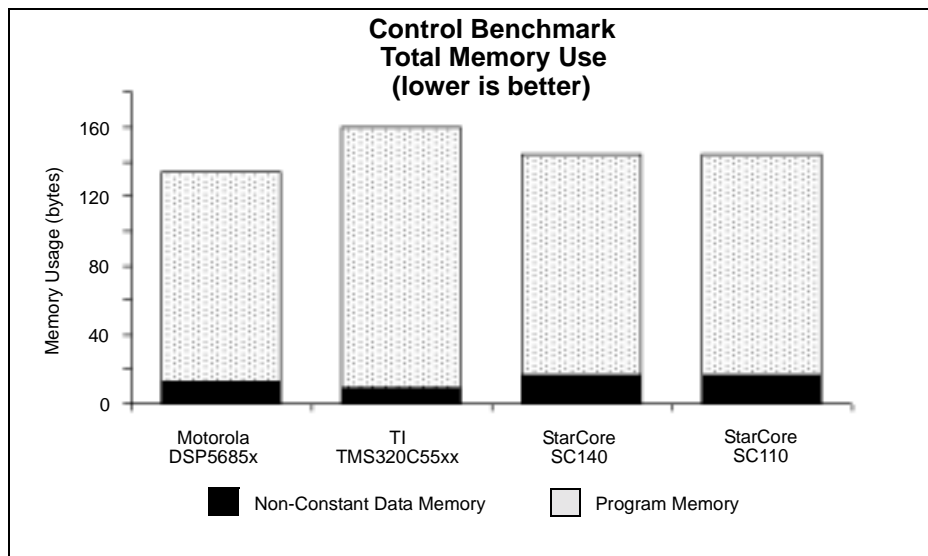
The BDTI Benchmarks™ include one benchmark function specifically designed to evaluate memory use for control-oriented software. Control-oriented tasks usually constitute the bulk of an application's program memory requirements, but only a fraction of the application processing time. Thus, in control-oriented tasks, memory use is usually a more serious concern than execution speed.

BDTI's Control benchmark is designed to represent control-oriented code. While most of the BDTI Bench-

marks™ are optimized primarily for maximum speed, BDTI's Control benchmark is optimized for minimum memory usage. This optimization hierarchy mirrors the approach generally followed by control-code programmers. Note that memory usage results on the Control benchmark are not necessarily indicative of processor memory use in signal-processing-intensive code.

Sample Benchmark Results

The memory usage results for BDTI's Control benchmark are shown in the figure at right. The SC110 and SC140 have identical memory usage on this benchmark. This is not surprising—the SC110 is binary upward compatible with the SC140 and thus should achieve the same memory efficiency on control-oriented software. Although the Control benchmark memory usage of the SC110 and SC140 is roughly average among the processors included in this report, their memory usage is notable considering that VLIW processors often have significantly higher memory usage than conventional DSP processors. The SC110 and SC140 achieve good memory efficiency through the use of short 16-bit instruction words, flexible conditional instruction execution, and their ability to use prefix words when needed.



Summary

Based on the SC110's results on the BDTI Benchmark™ suite, the SC110 is expected to be an unusually strong competitor in the world of low-cost and low-power DSPs. Power consumption estimates for the SC110 suggest that this processor will achieve excellent energy efficiency on DSP applications, even better than most conventional DSPs. Despite having only a single MAC unit and half of the SC140's on-chip memory bandwidth, the SC110's execution-time results on the BDTI benchmarks show that the SC140 is only twice as fast as the SC110, and the SC110 is more than two

times faster than TI's dual-MAC TMS320C5510.

The SC110 seems poised to suit applications where the SC140's very high performance (and presumably higher price) may be overkill. However, one question yet to be answered is where the SC110 fits into the product plans of Agere and Motorola. Surprisingly, nearly a year after the SC110's introduction, neither company has announced a product based on this core.

Order Form

Inside the StarCore SC110: A BDTI Technical Evaluation

Mail this form along with a check or fax with purchase order to:

Berkeley Design Technology, Inc. Tel: +1 (510) 665-1600
 2107 Dwight Way, Second Floor Fax: +1 (510) 665-1680
 Berkeley, CA 94704 USA Email: info@BDTI.com

Name
 Title, Division
 Company
 Address
 City, State, Zip, Country
 Tel: Fax:
 Email:

<u>Description</u>	<u>Qty</u>	<u>Price</u>	
First copy	1	x \$1500	=
Additional	x \$650	=
Tax (for CA orders)			=
International orders add \$75 for shipping & handling			=
TOTAL			=

Payment

- International orders must be prepaid in US dollars.
 Contact BDTI at info@bdti.com for volume discounts
- Check enclosed, payable to Berkeley Design Technology, Inc.
 - Purchase order, copy attached
 - Credit Card (contact BDTI for instructions)