

26.7 A 4.8GHz Fully Pipelined Embedded SRAM in the Streaming Processor of a CELL Processor

Sang H. Dhong¹, Osamu Takahashi¹, Michael White¹, Toru Asano², Takaaki Nakazato², Joel Silberman³, Atsushi Kawasumi⁴, Hiroshi Yoshihara⁵

¹IBM, Austin, TX

²IBM, Yasu, Japan

³IBM, Yorktown Heights, NY

⁴Toshiba, Austin, TX

⁵Sony, Austin, TX

The Synergistic Processing Element (SPE) is a processor designed to accelerate media and streaming workloads [1]. The Local Store (LS) unit in the SPE is a local memory comprised of several macros performing: (1) load/stores, (2) transactions for DMA, and (3) instruction fetches into an instruction line buffer (ILB). Because the LS occupies one-third of the SPE floorplan, area, power and yield are as important as performance.

The LS consists of a sum addressed decoder (memdec), four 64kB memory arrays (mem64k), write accumulation buffers (wacc and wtB), and read accumulation buffers (rdb1, rdb2, and rdb41), distributed throughout the SPE (Fig. 26.7.1). It takes four cycles to complete a write and six cycles to perform a read. The 6-cycle read path is shown in Fig. 26.7.1. The numbered latch points in the pipeline diagram also appear in the physical image to indicate the latch locations. A memory access starts by setting address operand registers in the memdec macro. The memdec adds the operands while pre-decoding bit groups of the sum with the results latched at the macro output at the end of the first cycle. In the 2nd cycle, the pre-decoded indices are distributed to the four copies of the mem64k memory block. The decoding of the address is completed within mem64k in the 3rd cycle which ends with the word line selection latched at the WL driver. The sub-array access cycle, the 4th, starts with WL activation. A write operation is completed during this cycle. For a read, the sense amplifier (SA) senses the BL differential signal and holds the value until it is captured by the read latch (RL) at the beginning of the 5th cycle. The remainder of the 5th cycle transfers the data read from the arrays to either the read buffers (rdb1 and rdb2 in Fig. 26.7.1) or to a 4-ported latch (rdb41). The fetched data is forwarded to the execution units in the 6th cycle. Macro placement, number of inversions, pin location, balanced wire length, bus interleaving, hostile/quiet neighbor wire assignment, choice of metal layer, metal width and spacing are engineered to achieve the 6-cycle 11-FO4 path.

A new sum addressed decoder (Fig. 26.7.2) is based on "add by shift/rotate". Input registers A and B in memdec receive 18b. The A and B inputs in the bit groups shown in Fig. 26.7.2 are decoded and the decoded value of A for each group is rotated by the B amount producing a decoded representation of the sum of A and B within the group of bits. This add-decode ignores the possibility of a carry into the group from the sum of the lower order bits of the address. Carries into each group are calculated separately; a final conditional rotation of the pre-decoded signals by one position occurs if a carry input is generated. Absorption of generated carries in the pre-decoder eliminates the need for a carry-based late select.

The final decode is done in mem64k based on the index signals. The sum addressed pre-decoder allows 1/8 partial activation by highest-order 3b, which is shown as the hatched area in the physical image (Fig. 26.7.2). By f<0:7> indices, one-to-four mem64ks are activated providing a quad word or a line access. When executing a small graphics kernel, 80% of memory access is accounted for by quad words. A line access activates 16 of 128 sub-arrays.

Quad word memory accesses activate only four out of 128 sub-arrays, resulting in substantial array power saving.

The mem64k block (Fig. 26.7.3), occupies three cycles of the access path. It is comprised of 32 sub-arrays, with 128 word lines plus four redundant word lines per sub-array. Additionally each mem64k includes two redundant BLs. Each BL spans 64 or 68 0.99 μ m² six-transistor SRAM cells. One SA is shared by four BL pairs. All sub-array input signals are latched with the exception of write data. Each sub-array receives the global clock and generates timing locally for a 120 μ m \times 260 μ m area. Timing skew between near-end and far-end is a maximum of 3ps. The sub-array timing generator is activated only when addressed. Inputs to the sub-array are captured by transparent latches (L1) and launched by clocked NANDs. Critical timing signals, SA set signal and wordline reset signals are generated by a timing generator with external margin control. The L1 of the wordline latches drives two NANDs for odd and even WLs, saving area and power.

The SA to RL path, an 11FO4 path, is one of the most challenging. The SA (Fig. 26.7.4) is based upon a conventional gate-terminal input sense and latch circuit [2]. Cross-coupled PMOSTs (P2, P3) are added to reduce the sensitivity to Vt mistracking between both N0-N1 and N2-N3. P2 and P3, directly inject BL/BL_B signals to the SA internal nodes, net0 and net1, during pre-charge before SA enable (SAE) is asserted. Compared to a conventional design, feed-forward PMOSTs allow 54% larger N0-N1 Vt mistracking. Simulations show the largest improvement at 3 σ mismatched SA and 4 σ weak cells and are supported by hardware data. The timing diagram (Fig. 26.7.5) shows key signals in the SA and RL. In the "writeA" period in the diagram, write data is written into a cell. In "readB", a read access happens in some sub-array and the data shows up at qwdout in the next cycle.

Due to BL and SA output isolation, BL pre-charge starts once SA is set. By keeping SAE high, the SA output is extended to the next cycle. The SAE, in turn, is gated by address information so that only the SAs for the required bits of data are enabled. This selection, coupled with the pulsed behavior of the SA output, allows an inverter and dynamic NOR to collect data from eight sub-arrays in RL without a separate multiplexer.

The measured SPE AC power is 3W, excluding clock distribution and leakage when executing a small graphics kernel at 1.0V/3.2GHz. Both LS and SPE power simulations showed LS power is 8.6% of SPE power. Therefore the interpolated LS unit and mem64k power are 258mW and 24mW, respectively.

A 256kB Local Store designed utilizing a 0.99 μ m² cell with SA achieves the 11-FO4 cycle time. A 6-cycle read access path is obtained by carefully balancing the micro architecture, circuit, and physical design issues. A sum addressed pre-decoder is important to reduce address generation latency while providing the means to selectively enable only the required portions of the array. An SA with reduced Vt mismatch sensitivity and gate-isolated inputs facilitates cycling the array at 11-FO4. A 256kB all-good LS unit in 90nm SOI technology has achieved 5.4GHz at 1.3V/52°C (Fig. 26.7.6). A small graphics kernel runs on SPE using a 16kB subset of LS with ECC correction at 5.6GHz/1.4V/56°C. The interpolated power of the LS unit based on hardware measurement is 258mW.

Acknowledgement:

P. Hofstee, R. Cook, S. Yong, S. Cottier, H. Kim, T. Wagner, L. Maurice, M. Maurice, D. Heidel, S. Wu, D. Hathaway, L. Smith, B. Hoang

Reference:

- [1] B. Flachs et al., "A Streaming Processor Unit for a CELL Processor," *ISSCC Dig. Tech. Papers*, Paper 7.4, pp. 134-135, Feb., 2005.
- [2] T. Kobayashi et al., "A Current-mode Latch Sense Amplifier and a Static Power Saving Input Buffer for Low-power Architecture," *Symp. VLSI Circuits*, pp. 28-29, 1992.

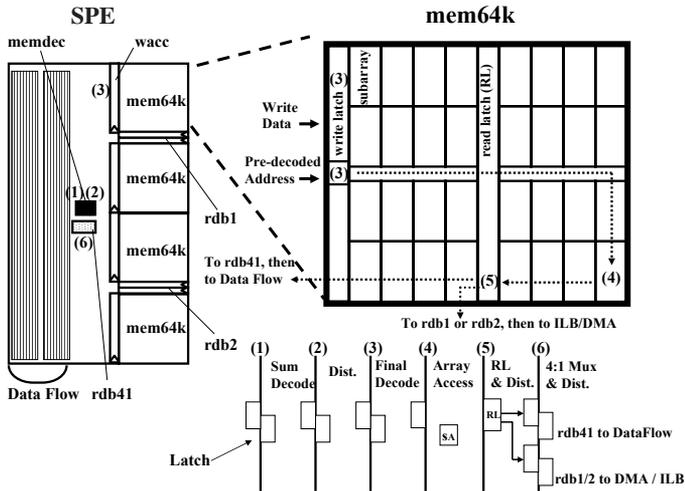


Figure 26.7.1: Local store macros in SPE.

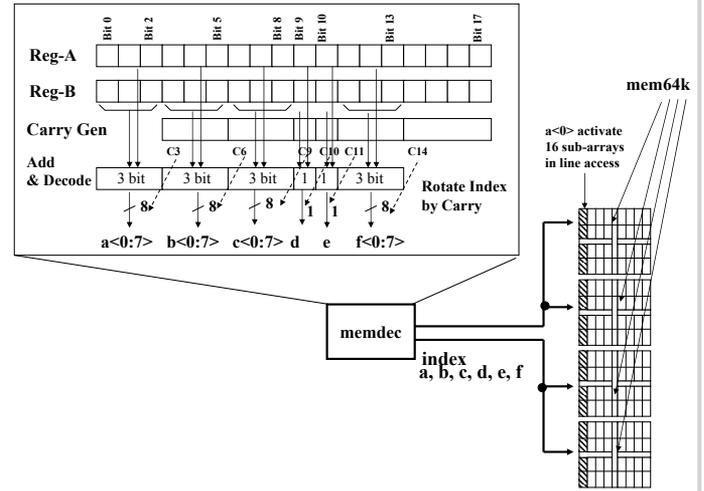


Figure 26.7.2: Sum address decode.

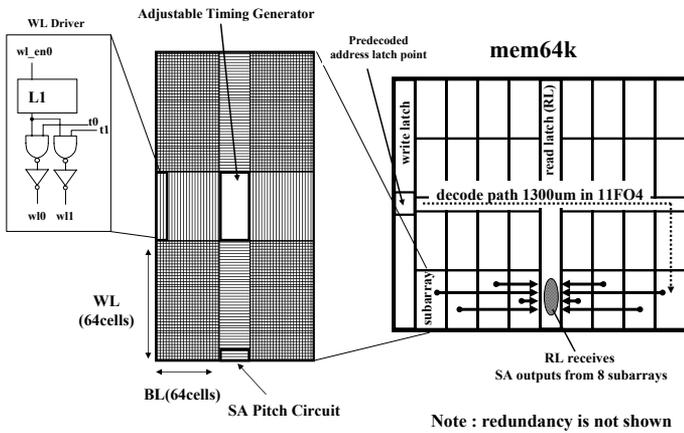


Figure 26.7.3: Sum addressed decoder (mem64k).

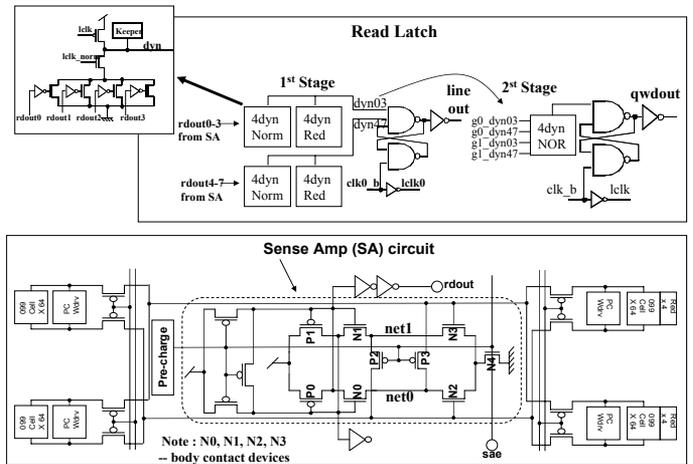


Figure 26.7.4: Read out scheme.

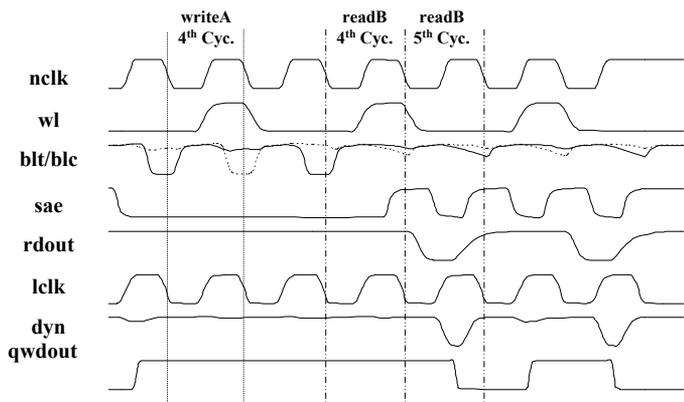


Figure 26.7.5: Simulation waveform.

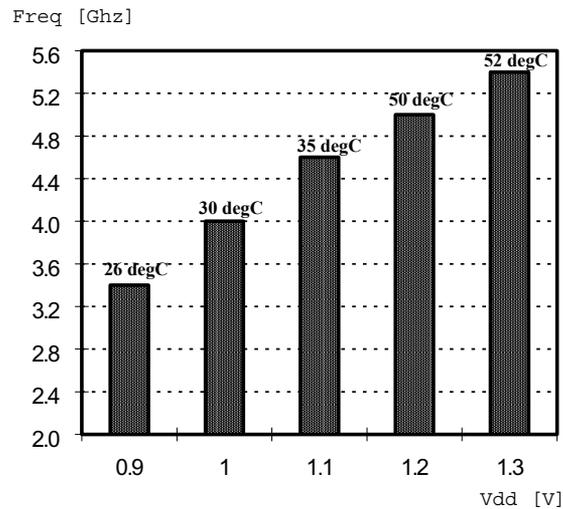
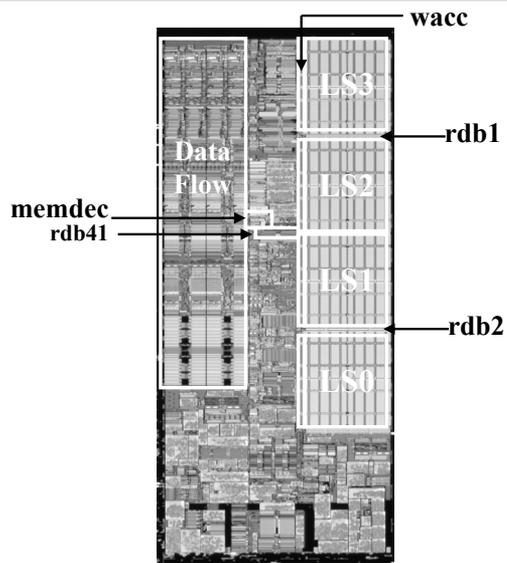


Figure 26.7.6: Frequency - voltage SHMOO.

Continued on Page 612



26.7: Die Micrograph.