

MoSys Explains 1T-SRAM Technology

Unique Architecture Hides Refresh, Makes DRAM Work Like SRAM

by Peter N. Glaskowsky

For the first time, MoSys has offered a complete explanation of the operation of its unique 1T-SRAM technology (see MPR 8/3/98, p. 9), used in chips from MoSys or as a core that can be integrated into ASICs from various foundries. The 1T-SRAM can be made in pure logic processes for higher speed, or with an embedded DRAM process to achieve greater density.

Though MoSys (www.mosys.com) calls the design an SRAM, it is based on single-transistor DRAM cells. As with any other DRAM, the data in these cells must be periodically refreshed to prevent data loss. What makes the 1T-SRAM unique is that it offers a true SRAM-style interface that hides all refresh operations from the memory controller.

First, Start With a Fast DRAM

The key to the 1T-SRAM is its high speed. This speed comes from the use of many small banks, often just 32 Kbits each, to create 1T-SRAM arrays. The small banks have short word lines with little capacitance, eliminating a major source of delay in conventional DRAM banks.

Each independent macrocell on a 1T-SRAM device consists of a number of these banks, all connected by a 250-MHz 256-bit macrobus. Each chip has one or more macros. MoSys can customize the number of banks per macro, as well as the bank size, according to the intended application.

The address-decoding logic is also fast. MoSys says its 0.25-micron parts perform a complete read operation in every 250-MHz clock period, with one cycle of pipeline delay and no bus-turnaround delay between reads and writes. Flow-through devices have also been constructed. In a 0.18-micron process, 1T-SRAMs will operate at up to 400 MHz.

Design Hides DRAM-Specific Operations

The 1T-SRAM design works just like an SRAM, with no wait states or other overhead for precharging or refreshing its internal DRAM cells. MoSys hides these operations with a combination of raw speed and clever design. Precharging takes place during every access, overlapped with the end of the cycle and the decoding portion of the next cycle.

Hiding refresh is more difficult. Every row in every bank must be refreshed at least once during the refresh interval (typically 1–16 ms) to prevent data loss.

MoSys provides a separate refresh controller for every bank. In each clock cycle, all banks not involved in the current transaction can be refreshed invisibly. If accesses are guaranteed to be distributed among multiple banks, the plurality of refresh controllers solves the refresh problem.

When all accesses during a refresh interval are directed to one bank, as can happen in signal-processing applications, a more sophisticated approach is required to hide refresh. So-called pseudostatic SRAMs, sometimes used in these applications, require a memory controller that can hold off accesses when a refresh operation is needed. MoSys's unique qualitative advantage over these parts (in addition to quantitative improvements in clock speed and power consumption) is that the 1T-SRAM never needs to hold off accesses, and indeed it has no hold-off signal. The circuitry that gives MoSys this advantage is fairly simple but has not previously been disclosed.

Clever Caching Prevents Pauses

To handle applications with no predetermined access patterns, MoSys gives each 1T-SRAM macrocell an SRAM cache with the same capacity as one bank. As read accesses take place within the macro, the data is cached. Any read hit to the cache allows all the banks in the macro to be refreshed in parallel with the read response, thus solving the refresh problem.

The remaining obstacle to be overcome is an access pattern that directs all read accesses to a single bank during a refresh interval while leaving one or more of the rows in the bank untouched—and therefore unrefreshed.

This circumstance presents no problem for the 1T-SRAM, because the refresh interval divided by the number of rows (which MoSys calls the proper refresh period) is much longer than the number of cycles needed to access every row in a single bank. Even the worst-case access pattern will therefore result in at least one cache hit during the proper refresh period, allowing a row to be refreshed.

The cache, its tags, and the related circuitry impose a die-size penalty of less than 10%. This penalty is in addition to the extra area required by the multiple banks and macros of the 1T-SRAM design. Altogether, a 64-Mbit 1T-SRAM is about 10–15% larger than a commodity 64-Mbit SDRAM—but 70% smaller than an SRAM of the same capacity.

MoSys Builds Strategic Relationships

Perhaps the most valuable design win the 1T-SRAM will ever see was revealed recently when Nintendo announced it will use the MoSys design in the Flipper graphics chip for its forthcoming Dolphin video-game console (see MPR 5/31/99, p. 5), due out late next year. Nintendo expects to sell millions of the Dolphin consoles, making it a lucrative customer for MoSys.

Analog Devices, NEC, and TSMC have taken 1T-SRAM foundry licenses. MoSys says it has taped out 20 different 1T-SRAM designs at these and other foundries and expects additional announcements this year. These deals, and the Nintendo win, should give 1T-SRAM a bright future. □