

Refresh your memory

Designing drop in pseudostatic memories. By Jarrod Eliason.

Pseudostatic memories are designed to replace srams, even though their internal operation is not static. Two commercially available technologies are pseudostatic ram (psram) and ferroelectric ram (fram). Whilst psram targets slow sram applications and competes on a cost/bit basis, fram targets battery backed sram applications and competes on the basis of system cost and logistics. FRAM also targets non volatile data acquisition applications.

Address transition detection (ATD) – a particular feature of asynchronous sram – allows address pins to change constantly at any rate for any amount of time. The memory is then guaranteed to output the correct data within t_{AA} (address access time) of the address pins stabilising.

This flexibility allows designers to relax when it comes to controlling the relative timing of memory interface signals. A typical micro with a direct memory interface shows its chip select output and addresses as being driven from the same clock edge (see figure 1). Whilst the timing from CLK to /CS (t_{CS}) and CLK to A (t_{ADR}) are usually specified, the relative timing between /CS and A is not usually discussed, much less guaranteed.

For an sram, this doesn't matter; the only requirement is that t_{AA} is fast enough. The time allowed by the micro is $2T$ minus the address or chip select

propagation delay (the greater of t_{ADR} and t_{CS}) and the micro data setup time (t_{SU}). In most cases, t_{ADR} and t_{CS} are the same.

In an effort to control pin count, almost every port on modern micros serves more than one purpose. One side effect of function multiplexing is that each pin can have a different internal delay. Unless each delay is controlled specifically, it is likely that at least one address pin will be slightly slower than chip select.

External factors also complicate matters. A common system design technique involves sharing the address bus between memories and/or peripherals. However, chip selects are not usually shared between external devices and, as a result, even if the addresses are guaranteed to precede the chip select


signal, this relationship may be lost at the memory pins.

An asynchronous sram sees any address skew as a reduction in required access time, but this is not a threat to functionality. For pseudostatic memories, the critical parameter is the relative timing between the chip-select and the slowest address.

PSRAM operates internally as a dram; the memory must be refreshed constantly to preserve the memory state. Reading a psram cell degrades the signal in the cell and some time is required to sense the data state and restore the signal to its full level.

The act of reading a dram is considered destructive and, once an access to a dram cell has begun, it must be completed. Whilst dram cannot handle continually changing addresses in the same way as

Figure 1: Direct micro interface





sram, it can handle the finite timing uncertainty seen in most micro based systems. Whilst fram is inherently non volatile and does not require continuous refresh, it shares destructive reads in common with dram and common techniques may be applied to create usable pseudostatic fram.

For a pseudostatic memory to replace sram, it must handle:

- address skew from chip-select (negative address setup time)
- address skew from other addresses in ATD accesses.

The first scenario can be handled at the expense of an increase in access and cycle time. The solution for the second scenario comes as part of the pseudostatic ATD implementation. Figure 2 shows a simple asynchronous delay circuit that can be used internally to delay the falling edge of the chip-enable signal in order to allow the addresses to continue to change after the external chip-enable signal has fallen.

The magnitude of required negative setup time is rarely specified in datasheets and should be characterised for each system, but will be less than 5ns for most designs. The disadvantage is this delays the start of memory access and consequently increases access time.

The pseudostatic implementation of ATD is simpler and safer than the asynchronous ATD implemented by srams. Since any pseudostatic memory access must be completed and the negative address setup requirement has been addressed, ATD becomes a matter of comparing the address at the inputs when the internal read is complete to the address latched internally when the access began.

In an fram, the internal access is divided into read and write back phases for two reasons. First is sram's requirement for late-write compatibility – catering for the possibility that the /WE input falls after /CE and of data changing after /WE falls. Secondly, the internal data bus is usually wider than the external bus. If the internal bus is 64bit wide and the external bus width is 16bit, 48 additional bits must be read and

restored, even during write cycles.

To account for this, all fram accesses begin as reads and operate as reads until the data latched by the internal sense amps is connected to the external data path. For reads, sense amp data is driven onto the memory's data pins. For writes, data from the memory data pins is driven into the sense amp. The direction of data flow is the only difference between a read and a write. The restore or write-back operation is the same in either case, with data latched in the sense amp being driven back into the cell.

In older frams, the /CE pin alone controlled the transition between phases. The falling edge of /CE would begin the read and the rising edge of /CE would begin the restore. ATD requires that addresses can also initiate transitions between phases. Figure 3 shows how the internal version of the chip-enable signal (cebint) is modified to account for ATD.

The internal cebint signal is delayed slightly from the external /CE signal. As the access read phase begins, addresses


“FRAM ... targets non volatile data acquisition applications.” Jarrod Eliason, Ramtron

are latched internally. The latch control signal is labelled alatch and the latched addresses are labelled la[n:0]. Prior to read phase completion, any change in external address is ignored. Once the read phase is complete, as indicated by the data ready (datardy) signal, the ATD signal will go high if the external A[n:0] bus differs from the internal la[n:0] bus.

Once an address transition has been detected – as indicated by the ATD signal – cebint is forced high to begin the restore phase. Once restore is complete, control of cebint is passed back to the external /CE pin. If /CE is still low, a new access will begin at the address present when the restore phase is completed.

In this implementation, memory restore time allows for address skew. The fastest address will initiate ATD and the slowest address has the entire restore time

Figure 2: Falling edge delay



to arrive. Most pseudostatic memories can tolerate address to address skew times in excess of 10ns or even 20ns. ■

Author profile:

Jarrod Eliason is a senior design engineer with Ramtron.

Figure 3: Pseudostatic ATD

