

TECHNICAL ARTICLE

“Evolutionizing” the DRAM Business: EDO and Burst EDO DRAMs

This paper was originally presented at the 1995 Silicon Valley Personal Computer Design Conference (SVPC), San Jose, California, USA.

Written by Bob Fusco
Product Marketing Engineer
Micron Technology, Inc.
8000 S. Federal Way
Boise, ID 83707-0006 (PCTP-103)

ABSTRACT

There has been an explosion in the variety of DRAMs being offered in the market in recent years. Extended Data Out DRAM (EDO DRAM), Synchronous DRAM (SDRAM), Rambus DRAM (RDRAM), and Enhanced DRAM (EDRAM) are among the variety of proposals intended to speed up the basic, vanilla DRAM. Each may have its technical merits, but which will make good on claims to dominate the highest volume application, PC main memory?

INTRODUCTION

In 1995, a sleeper has made significant inroads into the PC main memory segment. This is the EDO DRAM, occasionally referred to as hyper page mode. Through most of 1994, EDO DRAM was practically ignored by many DRAM manufacturers, anxious to promote more exotic offerings. How is it that this relatively modest performance boost is finding success while higher performance DRAMs are turning into marketing disappointments?

EDO DRAM

The most appealing feature of the EDO DRAM is its similarity to the commodity Fast Page Mode (FPM) DRAM. Though DRAM marketers report that cheap synchronous DRAMs are just around the corner, their factories keep churning out asynchronous DRAMs. PC OEMs and chipset designers have discovered that EDO DRAM offers just the right blend of cost vs. performance, with the emphasis where it must be, on cost.

EDO DRAM merely separates the two functions of the $\overline{\text{CAS}}$ pin. In FPM DRAM, the $\overline{\text{CAS}}$ HIGH-to-LOW transition latches the column address and the LOW-to-HIGH transition turns off the output buffer.¹

With EDO, the LOW-to-HIGH transition of $\overline{\text{CAS}}$ no longer turns off the output buffer. This change provides an ex-

tended data valid time, hence its name. EDO allows the microprocessor to sample that output data even while the address is set up in advance for the next cycle. Additionally, it makes the system relatively insensitive to the rising edge of $\overline{\text{CAS}}$. After satisfying a minimum $\overline{\text{CAS}}$ LOW pulse width, the $\overline{\text{CAS}}$ pin may be taken HIGH at whatever system clock edge is convenient while still being assured that the output data remains valid.

FPM transition times require precise control of an irregular $\overline{\text{CAS}}$ waveform to achieve their rated cycle time. Further, t_{PC} is merely the sum of $\overline{\text{CAS}}$ LOW and $\overline{\text{CAS}}$ HIGH (precharge), which isn't a useful metric of a FPM DRAM's bandwidth. Due to EDO DRAMs' flexible CAS duty cycle, t_{PC} is an accurate measure of its usable bandwidth.

WHAT CHANGES ARE REQUIRED TO USE EDO DRAMs?

A contributing factor to the emerging success of EDO DRAMs is that little or no change is necessary to the system or the DRAM itself. Perhaps a more appropriate question is “What changes are not required to use EDO DRAMs?” The new $\overline{\text{CAS}}$ logic on the EDO DRAM is so trivial that it can be incorporated without die size penalty. Because the pinout and package are the same as a FPM DRAM, it is most often implemented as a bond option on the same die. EDO DRAMs also can be used on the existing standard 72-pin SIMM modules.

The importance of a new DRAM being offered as an option of the current high volume standard product cannot be overstated. Such an approach, and only such an approach, ensures that a new DRAM alternative quickly moves down the cost learning curve in tandem with the high volume DRAM (FPM) it is intended to replace. New process shrinks will be implemented quickly. There won't be temporary gaps in pricing while the lower volume die (such as SDRAM) “catches-up” on processing and die size. Additionally, the assembly tooling already exists for SOJ and TSOP packages which you're using in your FPM application.

IS EDO BACKWARD COMPATIBLE?

This question is often raised, and in most cases the answer is “Yes”. While $\overline{\text{CAS}}$ high does not tristate the

1. The $\overline{\text{OE}}$ pin provides another method of disabling the output on both FPM and EDO DRAMs.

output buffer, a high level on the $\overline{\text{RAS}}$ input does turn off the output buffer. Users ask, “Are there situations where $\overline{\text{RAS}}$ is LOW, but $\overline{\text{CAS}}$ is HIGH resulting in “old data” still being driven on the bus?”

Such a condition could occur between read cycles if the system is using two way interleave with a shared $\overline{\text{CAS}}$ input. If the system is relying on $\overline{\text{CAS}}$ and $\overline{\text{CAS}}$ to “ping-pong” between two banks, use of EDO DRAMs would result in bus contention. Fortunately, few systems actually perform this sort of interleave because of the standard SIMM architecture. Even FPM prohibits keeping two $\overline{\text{RAS}}$ “banks” active at the same time. This is because the $\overline{\text{CAS}}$ pins are used to select 8-bit (or 9-bit) banks, and hence are shared between banks. An easy way to check to see if a $\overline{\text{CAS}}$ interleave is used is to identify the number of $\overline{\text{CAS}}$ pins

on the controller. If there are four $\overline{\text{CAS}}$ drivers on a 32-bit chipset or eight $\overline{\text{CAS}}$ drivers on a 64-bit chipset, EDO DRAMs are likely to be a drop-in replacement.

If the EDO DRAMs do cause a system error, it can be detected immediately upon booting the system, much the same as the system would detect if 80ns FPM DRAMs were installed when the machine required a faster speed grade. In the event the controller isn’t sufficiently intelligent to detect the memory error and reconfigure itself to operate without interleave, at least the user would be immediately aware that the system’s failure to boot was due to the newly installed memory.

But what about the situation where multiple bytes of data are to be read, and the controller holds $\overline{\text{RAS}}$ LOW, keeping the page open? If the controller were to follow this with read

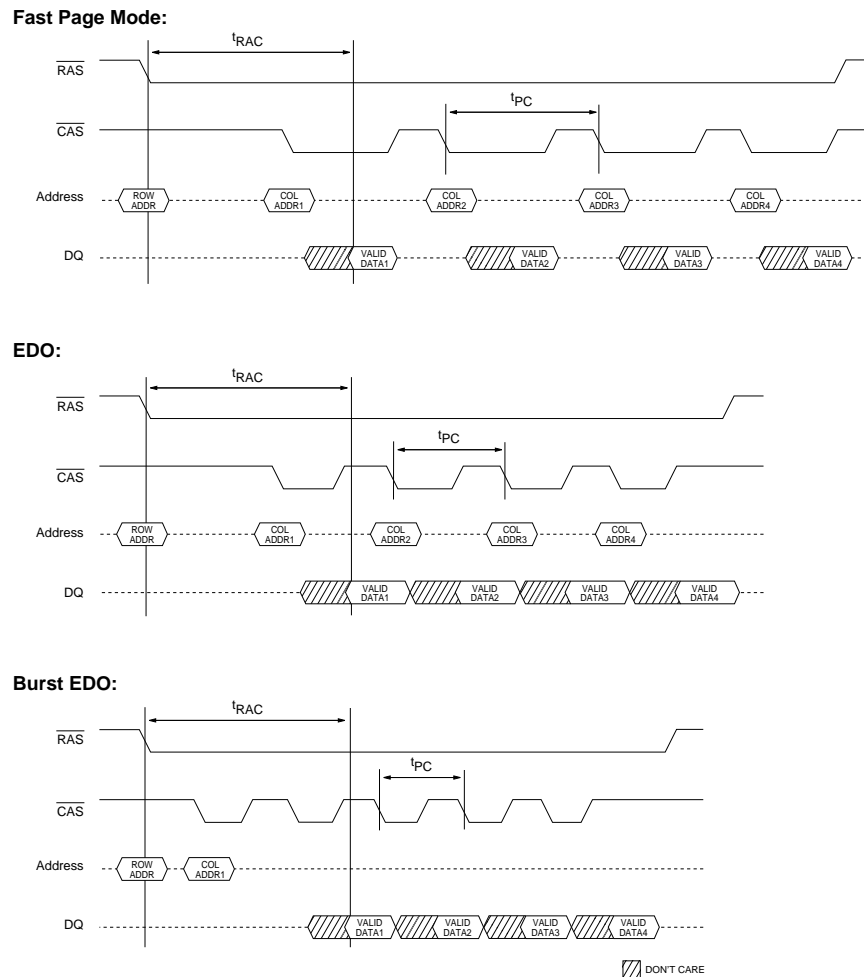


Figure 1

This figure illustrates the faster page cycle times (t_{PC}) possible with EDO and Burst EDO. Note the “dead” time between valid data is reduced with each enhancement to the operating mode.

data, from the same page or any other page, the old data will disappear and allow the “new data” to be driven on the bus. But what if a write operation to the same page were to follow? Wouldn't the “new” write data be in contention with the old data? The answer is “No”, because when \overline{WE} goes low, as it must for any \overline{WE} operation, the old data is tristated clearing the data bus and allowing for the new data to be written.

WHAT IS BURST EDO?

Burst EDO (BEDO) takes the good idea used in EDO DRAM (leaving data valid even after \overline{CAS} goes HIGH) and makes it better. Because most computer applications access DRAMs in four cycle bursts to fill cache memory, once the first address is known, the following three addresses can easily and quickly be provided by the DRAM itself.

Generating the subsequent addresses for BEDO DRAMs requires only the addition of a 2-bit counter. A few simple logic gates are added to provide for the two different burst orders used in common microprocessors. At the system level, this alleviates the tight timing requirements of address transitions every cycle on the heavily and variably loaded address bus.

The other performance enhancing feature of BEDO is the inclusion of a pipeline stage. This pipeline allows the page access cycle to be divided into two components. The first component accesses the data from the array to the output stage, while the second component drives the data bus to the appropriate logic level. Because the data is already available in the output stage when \overline{CAS} falls, a faster access time from \overline{CAS} (t_{CAC}) is achieved.

Figure 1 compares the control signals required for a four cycle burst for FPM, EDO, and BEDO DRAMs. The same RAS access time (t_{RAC}) is shown for all three options. However the burst is completed faster with BEDO.

Table 1 compares the maximum operating frequency ($1 / t_{PC}$) of each. Actual performance advantages of EDO and BEDO are even greater than shown, because the FPM cycle times cannot be achieved without an asymmetric duty cycle on \overline{CAS} .

Table 1

Minor changes in operating mode have boosted performance of asynchronous DRAMs by up to 100%.

DRAM SPEED (t_{RAC})	DRAM TYPE		
	FAST PAGE MODE	EDO	BURST EDO
70ns	25 MHz	33 MHz	50 MHz
60ns	28 MHz	40 MHz	60 MHz
50ns	33 MHz	50 MHz	66 MHz

Different microprocessors expect data in different burst sequences. The two alternative schemes are interleaved and linear bursts. Since BEDO DRAM was designed to work on standard 72-pin SIMM modules that can be installed and interchanged between systems with a variety of microprocessors, the DRAM burst order must be programmable. Programming is accomplished in the same way as SDRAM, with a WCBR cycle. This is a \overline{CAS} -before- \overline{RAS} refresh with \overline{WE} held low.

Because the BEDO DRAM will read its address inputs during a WCBR cycle, the controller provides the appropriate address bits (20H for linear, 21H for interleave) to configure the DRAM for the proper burst sequence. Though this sequence can be changed with another WCBR cycle at any time, most system designs will select the burst sequence only on initial power up. In instances where it is desired to switch from interleave to linear burst, such as when a DMA access isn't aligned on an even boundary, it is faster to start the interleaved burst on the odd address, terminate it, and restart on the next even address. This typically adds only two cycles, which is negligible in a long DMA access. Reprogramming the burst type takes 6 clocks or more, and would have to be performed twice, including once to return to interleaved burst mode.

BEDO resulted from an effort to reduce the cost and complexity of Synchronous DRAM. BEDO achieves a system's design objectives: reasonable bandwidth at low cost; while eliminating the least useful features of the SDRAM. Although SDRAM started out simple, a “design by committee” approach burdened it with excessive options. Its attempt to satisfy the full spectrum of applications from, graphics memory and workstations, to low cost PC applications result in added cost which will limit its market acceptance (see Table 2).

TRANSITION TO 3.3V OPERATION

BEDO DRAMs preserves many of the best ideas of SDRAM, such as programmable burst sequences and a pipelined output stage. However it avoids the changes in pinout and increased pin count for SDRAM which dictate new packaging in a costlier TSOP. Another good idea taken from SDRAM was its definition as a 3.3V-only part, bypassing the remainder of the 5V generation. 5-volt tolerant I/Os are provided on BEDO DRAMs, allowing for use on 3.3V or 5V busses. 3.3V BEDO DRAMs can be used on 5V SIMMs with the addition of a small, inexpensive three-terminal regulator. As users become more comfortable with 3.3V availability, the module can be powered directly with 3.3V, eliminating the external regulator. Figure 2 illustrates how the migration path for PC main memory can be achieved with 3.3V while meeting the wider and faster data bus requirements.

Table 2

This table compares some of the functional differences between BEDO DRAM and SDRAM.

	SDRAM	BEDO DRAM
# of Banks	1 or 2	1
Read Latency	1, 2, or 3	2
Burst Length	2, 4, 8, 512	4
Burst Sequence	Linear, Interleave	Linear, Interleave
Programmable by:	WCBR	WCBR
Burst Advance	CLK	CAS
RAS Control	Pulsed	Level*
Byte Control	New DQM	CAS*
x32 Module	None	72 pin*
x64 Module	168 or 200 pin	168 pin
Supply Voltage (Vcc)	3.3 V	3.3 V (5 V tolerant)
Relative Die Size	1.03 ~ 1.05	1.0*
Defined by:	Committee	PC architects, chipset designers, and Micron

* = Same as Fast Page and EDO

Part proliferation always leads to higher production costs, so cost is reduced by streamlining this new operating mode to only the new voltage range. As BEDO DRAM ramps, the market won't be needlessly divided into two smaller 3.3V and 5V camps. Additionally, a 3.3V only strategy ensures the end user will have a forward migration path to 64 Meg DRAMs and third generation (i.e., low cost) 16 Meg DRAMs using the 0.35 micron (64 Meg) process.

The new 8-byte (x64) DIMM modules are offered with a mechanical key to identify 3.3V operation. The 64-bit data path allows use of a single module when upgrading memory in fifth generation CPUs (see Figure 2). Changing to a new module also eliminates the restrictive burden of backward compatibility. Switching to BEDO DRAM, 8-byte DIMMs, and 3.3V operation all at once is a course many system suppliers are pursuing.

SUMMARY: EVOLUTION vs. REVOLUTION

While the efforts of many DRAM suppliers are aimed at “revolutionizing” the DRAM business, its history clearly indicates that an evolutionary approach is more likely to be successful. Despite clear technical advantages, even the simplest changes have often failed to overcome the inertia of the DRAM market. Backward compatibility is critical, and new controllers will allow for a mix of FPM, EDO, and BEDO modules in different sockets within the same system.

EDO is enjoying success not because of technical elegance or dazzling performance. It is successful because it is easy and cost effective to adopt. BEDO also is headed for widespread adoption, because it too uses an evolutionary approach.




Figure 2

BEDO DRAMs provide a migration path to 3.3V process technology, higher bus speeds and wider data paths of fifth generation microprocessors.