

Semiconductor Technology from A to Z

Digital technology

The DRAM Cell	3
<i>A Capacitor Instead of Six Transistors</i>	3
<i>Access Transistor, Word Line, and Bit Line</i>	3
<i>Leakage and Charge Decay</i>	4
<i>The Refresh Cycle</i>	4
<i>The Sense Amplifier</i>	5
<i>DRAM versus SRAM</i>	5

The DRAM Cell

A Capacitor Instead of Six Transistors

Where the SRAM cell holds its bit as a stable switching state of two cross-coupled inverters, the DRAM cell (Dynamic Random Access Memory) takes a fundamentally different approach: it stores information as electric charge on a single capacitor. A single NMOS transistor is enough to selectively connect this capacitor to the bit line or isolate it from it – hence the name 1T1C cell (one transistor, one capacitor).


This radical reduction to two components instead of six makes DRAM significantly denser and cheaper per bit than SRAM – main memory is therefore built almost exclusively from DRAM. The price for this: the stored charge is not stable. It slowly leaks away through parasitic paths, which is why every cell must be refreshed periodically to avoid losing its content – hence “dynamic”.

Access Transistor, Word Line, and Bit Line

As with the SRAM cell, the word line (WL) and bit line (BL) handle addressing – but with a single access transistor instead of two. Its gate connects to the word line, and its source-drain path connects the storage capacitor to the bit line. When WL is low, the transistor is fully off, and the capacitor is electrically isolated – aside from leakage, the stored charge has nowhere to go. Only when WL is activated does the access transistor open and connect the capacitor to the bit line.

To write, the bit line is driven low-impedance to the desired level from outside while WL is active – the capacitor charges or discharges accordingly through the open access transistor. Once WL falls low again, the capacitor is isolated once more and holds its charge – for now.

The DRAM Cell (1T1C)



Leakage and Charge Decay

The storage capacitor of a DRAM cell is tiny – typically just a few femtofarads. Even the smallest leakage currents through the off-state access transistor, through the capacitor dielectric, or through the pn junction to the substrate are enough to noticeably drain the stored charge within a few milliseconds. Without countermeasures, the stored value would simply vanish – a state no circuit, however clever, can tolerate.

This charge problem is exactly what fundamentally distinguishes DRAM from SRAM: while the cross-coupled inverters of an SRAM cell actively defend their state against disturbances, a DRAM cell is purely passive – it cannot maintain its own content.

The Refresh Cycle

To compensate for charge loss, the memory controller reads each cell at regular intervals and immediately writes the detected value back – a process known as refresh. Typical DRAM devices must refresh every cell roughly every 64 milliseconds, with some modern, denser cells needing it even more often. Since a single read access already activates an entire row of the array at once, refresh can be performed efficiently row by row rather than cell by cell.

This refresh costs time and energy: while a row is being refreshed, it is unavailable for regular read or write access. In very large memory devices with billions of cells, refresh overhead accounts for a noticeable share of total

bandwidth and power consumption.

The Sense Amplifier

During a read, the tiny storage capacitor briefly discharges onto the bit line, which was previously precharged to a mid-level voltage – the resulting voltage change is minimal, often just a few tens of millivolts. A sense amplifier at the end of each bit line detects this tiny signal and amplifies it into a clean digital level.

This amplification serves two purposes at once: it not only delivers the read value to the outside world, but also drives the original capacitor, via the activated word line, back to its full level. Every read access is therefore automatically also a refresh of the currently active row – unlike SRAM, where reading barely affects the stored state, reading a DRAM cell is destructive and makes this “rewrite” mandatory.

DRAM versus SRAM

The fundamentally different storage principles lead to very different characteristics, each suited to different applications:

	SRAM	DRAM
Transistors per bit	6	1 (+ 1 capacitor)
Cell size	large	very small
Refresh required	No	Yes (~every 64 ms)
Access speed	very fast	slower
Cost per bit	high	low
Typical use	cache, registers	main memory (RAM)

In practice, the two complement each other: fast but expensive SRAM cells in the cache levels close to the processor core, and dense, inexpensive DRAM cells for high-volume main memory.