# A Comparison of Space-Grade Memory Technologies

Real-time as well as Store & Forward on-board processing applications are increasingly requiring larger amounts of fast on-board storage, and the choice of memory technology has a major impact on capacity, physical size, power-consumption, speed, reliability and mission lifetime.

Conventional SRAM stores each bit within a latch, typically realised using four or six transistors, and space-grade devices offer capacities up to 32 Mb and access times of 12 ns.

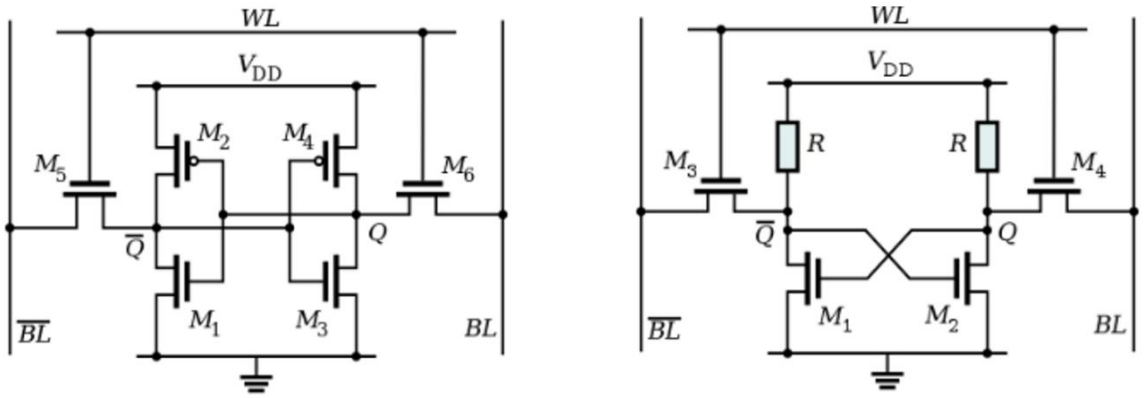


**Figure 1 : Six and four-transistor SRAM cells: WL = word line, BL = bit line.**

The fastest and largest semiconductor memory is SDRAM which is organised as a logical array of cells, with each comprising a capacitor and a FET as the control gate. Every cell stores one bit and a simple 4-bit memory is illustrated below. The transistor opens or closes based on the voltage on the row, charging or draining respectively the capacitor attached to it. After the desired 'word-line' is charged, the column selector is switched to access the required capacitor for the up-coming read/write operation. Due to its natural discharge, the cells have to refreshed periodically, which involves reading and then re-writing the data back.

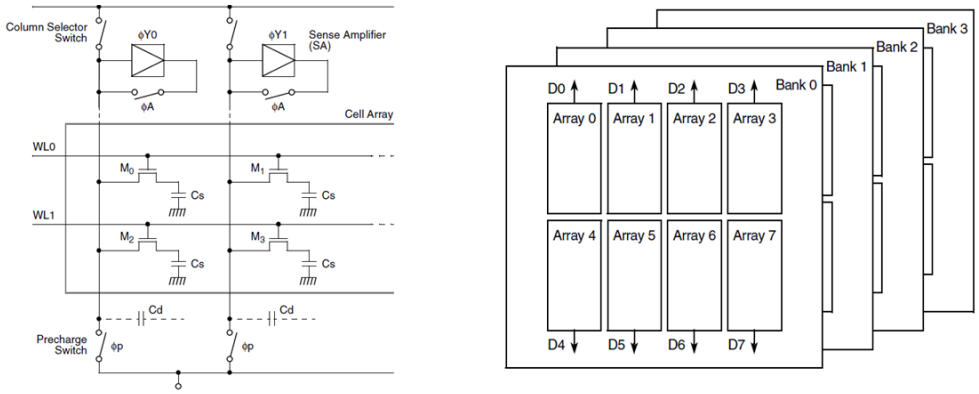


**Figure 2 : SDRAM bit cells and the organisation of an SDRAM chip.**

SDRAM architecture comprises memory cells organised into a two-dimensional array of rows and columns. To select a particular bit, it is first necessary to address the required row and then the specific column. Once the desired row is open, it is possible to access multiple columns, hence improve speed and reduce latency through successive read/write bursts.

To increase the word size, the memory has multiple arrays which means when a read/write access is requested, the memory only requires one address to access 1 bit from each array.

To increase overall memory capacity, banks are added to the internal structure of SDRAM as illustrated above. Bank interleaving further increases performance and each can be addressed individually.

To perform a read or write, the ACTIVE command first has to be issued by the memory controller to activate the required row and bank. Once the desired operation has completed, the PRECHARGE command closes a particular row in one or more banks. A new row cannot be opened until the previous one has been closed.

SDRAM operation is achieved using its control signals, Chip Select (CS), Data Mask, (DQM), Write Enable (WE), Row Address Strobe (RAS) and Column Address Strobe (CAS), with the last three determining which command is issued as listed below:

| Command | $\overline{CS}$ | $\overline{RAS}$ | $\overline{CAS}$ | $\overline{WE}$ | ADDR |
|---|---|---|---|---|---|
| Command inhibit (NOP) | H | X | X | X | X |
| No operation (NOP) | L | H | H | H | X |
| Active (select bank and activate row) | L | L | H | H | Bank/row |
| READ (select bank and column, and start READ burst) | L | H | L | H | Bank/column |
| WRITE (select bank and column, and start WRITE burst) | L | H | L | L | Bank/column |
| PRECHARGE (deactivate row in bank or banks) | L | L | H | L | X |
| AUTO REFRESH or SELF REFRESH (enter self refresh mode) | L | L | L | H | X |

**Table 1 : The SDRAM command truth table.**

SDRAM has evolved significantly since its release in 1992: the initial version was Single Data Rate (SDR) SDRAM was has the same internal clock frequency and I/O rate. SDR SDRAM can only read or write once in a clock cycle and has to wait for the completion of the current operation before starting the next.

Double Data Rate (DDR) SDRAM achieves greater bandwidth without increasing the clock frequency by transferring data on both clock edges, doubling the I/O transfer speed without increasing the clock frequency. This is accomplished by using a 2*n*-prefetch architecture where the internal data-path is twice the width of the external bus, allowing the internal frequency to be half the external transfer speed. For each single read access, two external words are fetched and for a write operation, two external data words are combined internally and written in one cycle. DDR1 is a true source-synchronous design, where the data is captured twice per clock period using a bi-directional data strobe.

DDR2 SDRAM operates the external bus twice as fast as DDR1 doubling the I/O transfer rate. This is achieved by using a 4*n*-prefetch buffer where the internal data-path is four times the width of the external data bus. DDR2 can operate at half the clock frequency of DDR1 and achieve the same transfer speed, or at the same rate with double the information bandwidth.

DDR3 SDRAM operates the external bus twice as fast as DDR2 doubling the I/O transfer rate by using an 8*n*-prefetch architecture. The width of its internal data-path is eight bits compared to DDR2's four. DDR3 can operate at half the clock frequency of DDR2 and achieve the same transfer speed, or at the same rate with double the information bandwidth.

Table 2 summarises the space-grade SDRAM options which are currently available to satellite and spacecraft manufacturers.

| SDRAM | Capacity (Gb) | Data Bus (bits) | Clock Frequency (MHz) | Transfer Speed (MT/s) | Bandwidth (Gbps) |
|---|---|---|---|---|---|
| | | | | | |
| SDR | 0.5 to 4 | up to 72 | 133 | 133 | up to 9.6 |
| DDR | 8 | up to 72 | 100 to 200 | 200 to 400 | up to 28.8 |
| DDR2 | 8 | up to 72 | 200 to 400 | 400 to 800 | up to 64 |
| DDR3 | 24 | up to 72 | 300 to 666 | 600 to 1,332 | up to 95.9 |
| DDR4 | 48 | up to 72 | 933 to 1200 | 1,866 to 2,400 | up to 153.6 |

**Table 2 : Current, space-grade SDRAM capability.**

NAND and NOR flash use floating-gate transistors to produce non-volatile memory: NAND has physically smaller cells resulting in high densities with the latest devices offering Tbs of storage. Due to yield-related issues, NAND flash is fabricated with a known number of erroneous blocks, typically 2% of the overall capacity at beginning-of-life. Furthermore, this technology degrades with time due to repeated use of the memory cells and there is a specified reliability limit regarding the total number of bytes that can be written (TBW) over the lifetime of devices. NAND flash contains intrinsic error-correction spare bits, is intended for linear, contiguous storage and can be considered a solid-state hard drive. An ASIC/FPGA controller is required to manage its interface, the location of bad blocks and additional deterioration over the duration of a mission.

NOR flash has physically larger cells resulting in less density with the latest devices offering a maximum capacity of 2 Gb. Compared to NAND, NOR flash provides address lines to map the entire range, allowing random access and faster initial and sequential read times, making it better for code execution or FPGA configuration. NOR flash does not suffer from bad blocks or lifetime endurance issues, but has slower write and erase speeds due to its architecture and how the floating-gate transistors are connected.

NAND flash is typically used for data storage and is organised and functions differently to conventional semiconductor static and dynamic memory. NAND flash memory is organised into pages, blocks, planes and logical units (LUNs) as illustrated below. A target is the amount of memory accessed by a chip enable signal and can contain one or more die, which is the minimum that can independently execute commands and report status. According to the Open Nand Flash Interface (ONFI) specification, a die is referred to as a LUN.

For the example shown below: each page contains 8640 bytes, comprising 8192 bytes of data and 448 for error correction and wear-levelling, *i.e.* ensuring blocks are exercised uniformly.

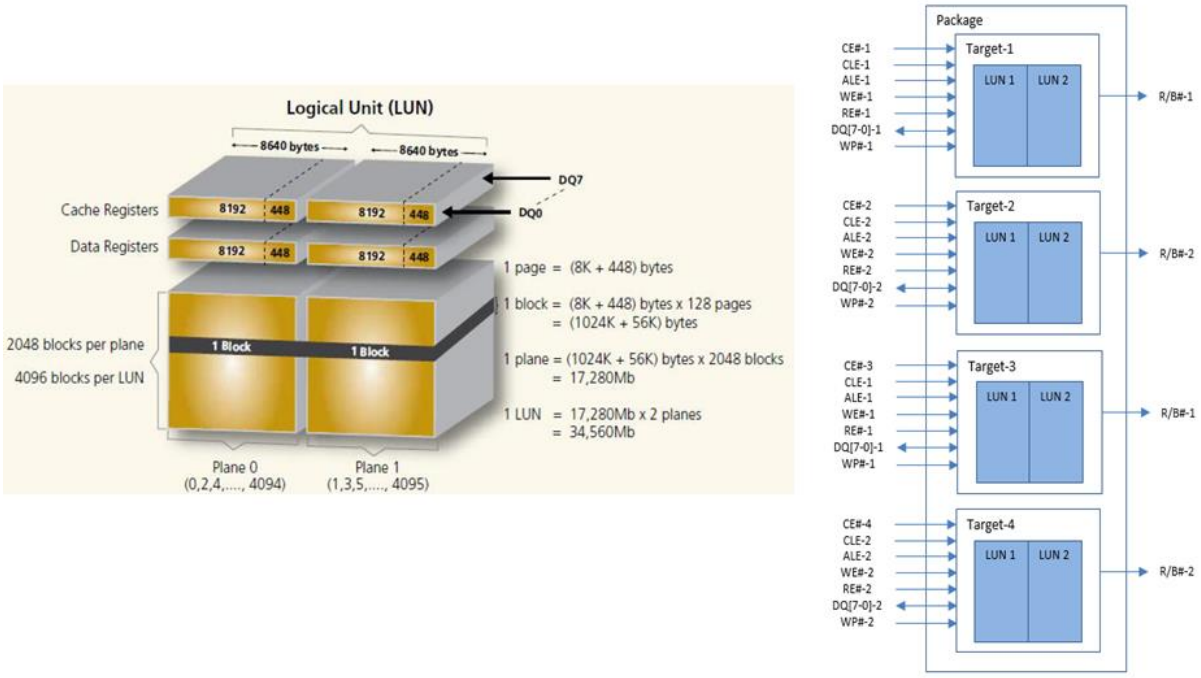


**Figure 2 : Organisation of NAND flash memory.**

To reduce pin count, data, commands and addresses are multiplexed onto the same pins and are received by I/O control circuits. Commands are latched into a register and transferred to logic for generating internal signals to manage device operations. The addresses are latched and sent to a row/column decoder to select the desired location within the NAND flash memory array.

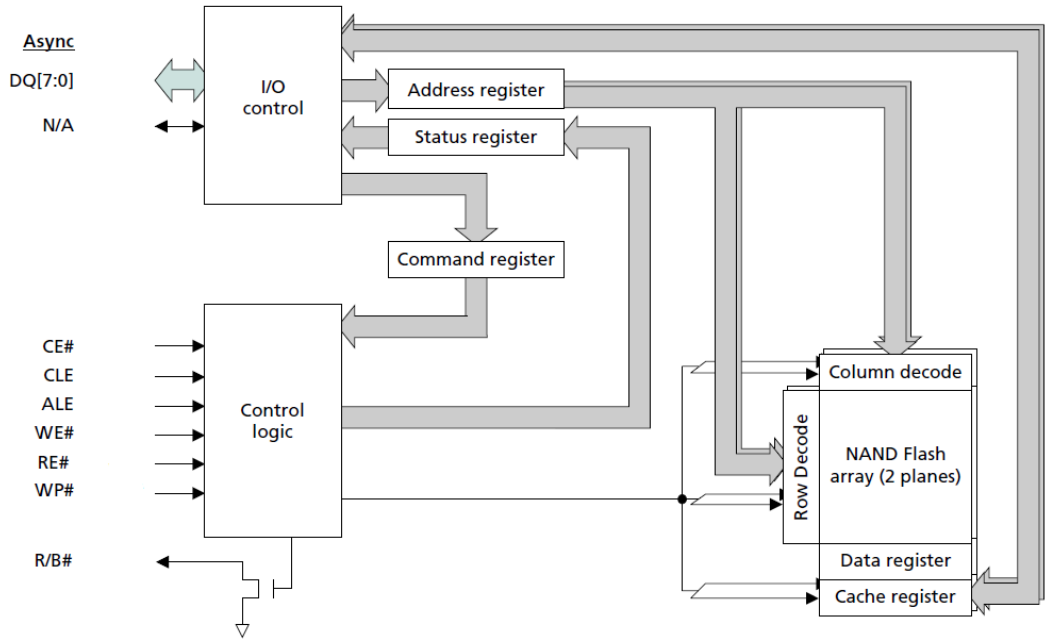


**Figure 3 : The LUN functional block diagram.**

The NAND Flash memory array is programmed and read using page-based operations and is erased using block-based instructions. In normal page mode, the data and cache registers act as one. During cache operations, the data and cache registers operate independently to increase data throughput, *i.e.* you can access data from the cache register while array data is being transferred to the data register.

Previously Spacechips was asked to develop a 1 Tb mass-memory unit which had certain speed requirements: we considered SDRAM and NAND flash and the differences in overall physical size, cost and power consumption are staggering:

| | DDR3 16Gb x 16 | NAND 256Gb x 16 |
|---|---|---|
| **Number of Devices** | 64 | 4 |
| **Minimum Real Estate ($mm^3$)** | 152,320 | 16,796 |
| **Power Dissipation (static/dynamic)** | 17W / 17W | 13.2mW / 5.2W |
| **Approximate Cost ($)** | 650k | 30k |
| **Storage Rate, 16-bit bus (Mbytes/s)** | 2,667 | 100 |
| **Total Byte Writes (TBW)** | Unlimited | 7530 TB |

**Table 3 : Realising 1 Tb of on-board storage.**

NAND flash has endurance limits on the maximum number of write operations that can be performed over its lifetime before impacting device reliability. This requires a detailed understanding of your mission's *on-board* storage requirements and operational duty cycle before selecting appropriate parts, *e.g.* how many writes per orbit, per day, per year and over the duration of the complete mission.

If a NAND flash part has a specified lifetime endurance of 40,000 cycles, this denotes the maximum number of times blocks can be written or erased. If the mission duration is three years, then each block can be written to or erased 13,333 times annually, 36 times per day or once per hour. The choice of manufacturer will depend on the amount of data to be stored per orbit and how this reconciles with device capacity and/or memory bandwidth.

Many of our Earth-Observation clients have very specific on-board storage requirements, *e.g.* one customer needs to write 100 Gb every orbit while another needs to store 1.08 Tb during eight-minute captures. A 256 Gb part with a specified lifetime endurance of 40,000 cycles results in a TBW of 1,280. If a mission needs to write 2.16 TB per day or 788.4 TB annually, then this device will limit the mission duration to 1.6 years. If a three-year LEO lifetime is required, then the options are to use a more endurant part, reduce the amount of data to be written or use multiple memory chips controlling how often data is stored to each.

Compared to commercial-grade devices, space-qualified NAND flash is slow, operating in asynchronous mode with a maximum I/O speed of 50 MHz, with the **RE#/WE#** control input transitioning read and write transfers respectively. The maximum time to write one page (8640 bytes) of data to the cache is 173 µs, and 560 µs to transfer this from the cache to the memory array. The subsequent time to write one byte is 64.8 ns and this must occur within the clock period of 1/50 MHz = 20 ns.

To increase sustained array (page program) bandwidth, within a device, multi-plane mode allows both planes to be written simultaneously halving the above time to 32.4 ns. Each plane contains independent cache and data registers and as shown in Figure 2, there are two die per chip enable allowing multi-plane, multi-LUN mode permitting multiple die to be accessed at the same time, *e.g.* by using two die, the write (page program) time of 64.8 ns can be reduced to 64.8 / 4 = 17.4 ns. By simultaneously interleaving data to four planes on two die, you can fit the time required to write one byte within the 20 ns clock period so the array page programming time is now faster than the I/O time. The data will not be stored contiguously, but spread over two die, and a memory controller within a FPGA/ASIC must manage this process.

To improve yield and data retention, some manufacturers of NAND flash exploit SONOS (silicon-oxide-nitride-oxide-silicon) technology which uses an insulating layer to trap or store charge. SONOS FETs have higher endurance than conventional, floating-gate transistors and have been successfully integrated into existing CMOS process flows.

All of the above memory technologies use electron charge to store data, whereas, MRAM uses magnetic elements to store information, providing inherent radiation immunity, offering the speed of SRAM, a density approaching that of DRAM, the non-volatility of flash memory, unlimited read/write endurance and high-power efficiency at all times. MRAM uses magnetic states and polarisation of a ferromagnetic material for writing data and magneto-resistance for reading back bits.

MRAM uses magnetic states and polarisation of a ferromagnetic material for writing data and magneto-resistance for reading back. Each storage element consists of a single transistor, magnetic tunnel junction (MTJ) memory cell as shown below. The MTJ comprises a fixed magnetic layer, a dielectric tunnel barrier and a free magnetic layer which can have its direction changed using a magnetic field or by applying a polarised current.

The insulating barrier is as thin as a few atomic layers and when an external bias is applied to the MTJ, electrons tunnel through the normally insulating material changing its resistance. A logical zero is stored when both layers have the same direction, *i.e.* the MTJ has low resistance, and a logic one when both layers have different directions, *i.e.* a high resistance.

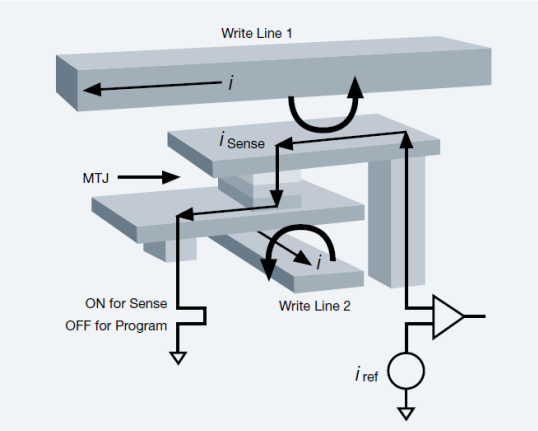


**Figure 4 : Schematic of a single transistor, MTJ cell [Freescale].**

Data is stored as a magnetic state rather than a charge and sensed by measuring its resistance without disturbing its polarization. This offers a number of major benefits for space applications:

1. The magnetic state does not leak away with time so your information remains when power is turned off.

2. Switching between the two states does not involve the actual movement of electrons or atoms, thus no wear-out mechanism exists.

3. MRAM is immune to bit flips (SEUs) due to radiation effects, latch-up resistant up to 85.4 MeV/cm$^2$/mg and a measured total-dose tolerance up to 1 Mrad(Si). Additional RHBD

practices are required for support circuitry at a device level which differ by vendor, impacting actual radiation hardness.

Space-grade MRAM devices are now available with parallel and serial interfaces and my previous post provided examples of rad-tolerant parts.

While space-grade NAND flash offers GBs and Tbs of storage capacity, its speed, reliability, data retention and the need to manage erroneous bits limits its use to missions with specific lifetimes and operational duty cycles. MRAM and NOR flash provide Gbs of fast, non-volatile storage without endurance limits.

The choice of memory technology has a major impact on overall storage capacity, physical size, power-consumption, speed, reliability and mission lifetime. How have you managed your on-board storage needs? Please leave your comments below and the best answer will win a Courses for Rocket Scientists World Tour tee-shirt. Congratulations to Anna from Sweden, the first to answer the riddle from my previous post.

*Dr. Rajan Bedi is the CEO and founder of Spacechips, which designs and builds a range of advanced, VHF to K-band, re-configurable, multi-band, ultra high-throughput on-board processors, transponders, as well as Edge-based OBCs and memory-mass units for telecommunication, Earth-Observation, navigation, internet, 5G and M2M/IoT satellites. The company also offers Space-Electronics Design-Consultancy, Avionics Testing, Technical-Marketing, Business-Intelligence and Training Services. (www.spacechips.co.uk). Rajan can also be contacted on Twitter to discuss your space-electronics' needs: https://twitter.com/DrRajanBedi*

*Spacechips' Design-Consultancy Services develop bespoke satellite and spacecraft sub-systems, as well as advising customers how to use, select and procure the right components, how to design, test, assemble and manufacture space electronics.*