

Memory Technologies and Data Recorder Design

Karl F Strauss, *Senior Member IEEE*

Jet Propulsion Laboratory

4800 Oak Grove Drive

Pasadena, California 91109 US

karl.f.strauss@jpl.nasa.gov

Abstract—Missions, both near Earth and deep space, are under consideration that will require data recorder capacities of such magnitude as to be unthinkable just a few years ago. Concepts requiring well over 16,000 GB of storage are being studied. To achieve this capacity via “normal means” was considered incredible as recently as 2004. This paper is presented in two parts. Part I describes the analysis of data recorder capacities for missions as far back as 35 years and provides a projection of data capacities required 20 years from now based upon missions either nearing launch, or in the planning stage.

The paper presents a similar projection of memory device capacities as baselined in the ITRS – the International Technology Roadmap for Semiconductors. Using known Total Ionizing Dose tolerance going back as far as a decade, a projection of total dose tolerance is made for two prime technologies out to the year 2028.

Based upon the two prime technologies, the design of a 130 Tb recorder is discussed in Part II. Further, it is noted that, for all the missions and technologies analyzed, the parameters of a recorder – mass, power, volume – remain constant despite ever-increasing capacity requirements.

Table of Contents

1. HISTORY OF DATA RECORDER CAPACITY	1
2. AN EXAMINATION OF SOLID STATE DEVICE HISTORY	3
3. SOLID-STATE TECHNOLOGY BASELINES – THE ITRS.	5
4. RADIATION EXPOSURE VERSUS CAPABILITY	7
5. A NEW MAXIM	9
6. HIGH CAPACITY SOLID STATE RECORDER.....	11
7. HIGH-CAPACITY SSR DESCRIPTION	12
8. CONCLUSION	15
DEFINITIONS.....	16
ADDITIONAL MATERIAL	16
REFERENCES.....	16
BIOGRAPHY	17

Part I

In this part of the paper, we examine the historical capacity of data recorders for deep space, Earth orbiting, and Lunar missions commencing from the year 1967. This information is augmented by the inclusion of missions still in the planning stage and the currently-allocated data recorder size.

1. HISTORY OF DATA RECORDER CAPACITY

Eighteen NASA, JPL, JAXA, and ESA deep space missions were analyzed as far back as 1973 – the series of Pioneer

¹978-1-4244-2622-5/09/\$25.00 ©2009 IEEE

¹ IEEEAC paper #0627, Version 1, Updated 22 Sept 2008

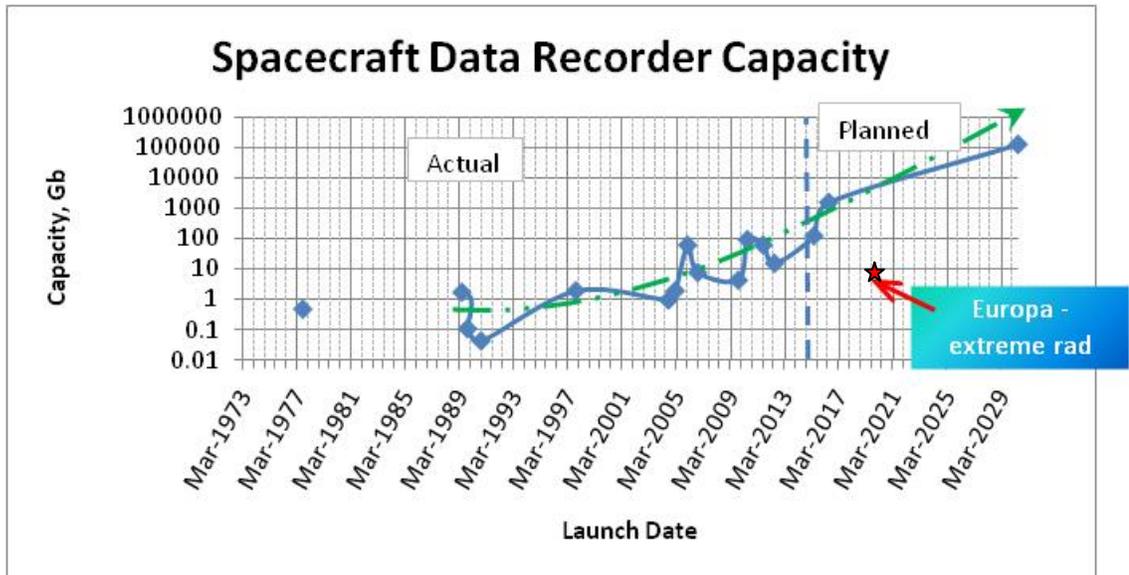


Figure 1 - Historical & Projected Spacecraft Data Recorder Capacity by Launch Year. Green line indicates Trend. (Extreme-radiation exposure mission to the Jovian moon Europa is an outlier and not included in the trend analysis)

missions – specifically Pioneer 11, launched April 1973, to study the survivability of the interplanetary asteroid belt and to assess the radiation belts surrounding Jupiter³. Presenting this information as a graph (Figure 1), it can be readily seen that the trend (shown in Green) of deep space missions is to have a doubling of on-board data capacity at a rate approximately every three years.

Tape Recorders – not without problems

Up through 1990, the spacecraft examined used multi-track tape recorders. The largest tape recorder (in terms of data capacity) launched was the 1.8Gb, 12-track behemoth used on the Magellan mission to Venus. The Magellan recorder featured over one-half mile of precision Mylar tape. These precision mechanisms were not without problems – culminating in the near loss of the Galileo mission in October, 1995, due to a position encoder error compounded with electrostatic build-up on the tape’s surface. From that point on, all missions (that is to say, all missions examined) have utilized solid state recorders (SSRs).

Solid State Recorders – A Watershed Moment

NASA’s first mission using a solid state data recorder was the Cassini mission launched in 1997. It features a 2 Gb SSR consisting of 640 specially screened DRAM devices and eight full custom ASICs. With the advent of the Cassini recorder, two trends emerged. First was the use of solid state devices for data storage and the second was the abandonment of the so-called “Mil-Grade” product by most

semiconductor manufacturers. The latter required the development of additional procedures to “up-screen⁴” commercial devices to what is termed *pseudo-mil*. Often this involved the development of a partnership with the device manufacturer to ensure that the devices chosen, screened, and accepted were manufactured at the same time, and often from the same wafer lot.

The use of commercial devices was both a bane and a blessing. The bane was two-fold. One, as memory devices were no longer made for specific space-“type” capability – such as temperature or radiation – a survey of many devices from many manufacturers was required. These surveys involved at least two types of expensive radiation testing – Total Ionizing Dose (TID) and Single Event Effects (SEE). Complicating the issue, more often than not, manufacturers would alter various manufacturing settings (such as ion implant depth) to achieve better commercial yields, and these changes, however minor, would completely invalidate previous Radiation Lot Acceptance Testing (RLAT). Therefore, this brought in the second complication: the partnering with manufacturers. This was rarely successful as semiconductor manufacturers are generally not set up to provide customers with lot-specific data. Therefore, so-called lifetime buys were instituted with random parts being examined internally for monitoring of certain on-die

³ Pioneer 11 ceased transmitting to Earth March 1997.

⁴ The term “up-screen” refers to a process wherein a set of commercial devices receives extra examination as far as quality and operational parameters are concerned often over a set of voltages and environments in excess of what the manufacturer states the device is designed for. The “up-screening” may include the use of X-Ray or Ultrasonic based inspections, additional radiation testing, and involves the destruction of a sub-set of the devices for physical examination at an Electron Microscopic level

manufacturing numbers and patterns. This added greatly to the cost of the recorder.

2. AN EXAMINATION OF SOLID STATE DEVICE HISTORY

The blessing is that as the semiconductor industry matured and developed new, aggressive products, the design of a recorder often went with it. Intel founder and chief scientist Gordon Moore – then working for Fairchild Semiconductor – predicted in 1965 that, essentially, the number of transistors per fixed die size would double roughly every two years to thirty months [1]. To date, the Industry hasn't proven him wrong [2].

Of the thirteen solid-state based data recorders analyzed, nine of them are based upon Dynamic Random Access Memory (DRAM) or Synchronous DRAM (SDRAM) technology and its variants. The remainder based upon non-volatile Flash. The predominance of DRAM can be attributed to the general ability of DRAM devices to have a higher total ionizing dose tolerance than Flash (discussed later in this paper).

MEMORY TECHNOLOGY

Of all the devices needed to operate a computer, especially a computer in a space environment, space-specific memory technologies have received the least amount of institutional development. While commercial technology is achieving density and power reductions unimaginable even a few years ago, space-specific memory is several generations behind. Commercial Flash technology embraces 4 and 8 Gbits per single die – the latest space-rated non-volatile technology is 4 megabits (Mb) – more than three orders of magnitude smaller.

Non-Volatile technologies

Non-Volatile Memory technology today consists of the following:

Rotating, Mechanical—Hard Disk Drive, Optical Platter (DVD, CD)

Hard Disk Drive— Current technology supports 500 GB per platter, Figure 2 [3]. Despite a successful research study undertaken for the Prometheus/JIMO program, due to the competitive nature of the industry, no supplier willing to support this effort has been located.



Figure 2 - Hard Disk Platter with Cantilever (courtesy Seagate)

Optical Platter—Current technology supports 50GB per platter [4]. Due to the competitive nature of the industry, no supplier willing to propose a rad-tolerant optical platter design has been located.

Solid State—Ferroelectric (FeRAM), Magnetic (MTJ MRAM), Ovonic (CRAM), Trapped Charge (ONO Flash, EEPROM), NROM (Nitride-oxide memory)

Ferroelectric (FeRAM)—In this technology, an Oxygen atom is positioned within a crystalline matrix usually consisting of Lead, Zirconium, and Tantalum. The position of the Oxygen atom is sensed as a positive or negative charge within the matrix, Figure 3.

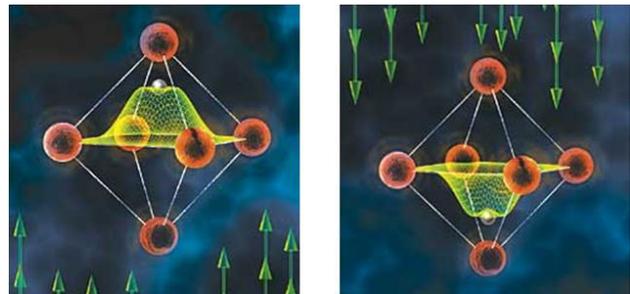


Figure 3 - Operation of FeRAM (courtesy Carnegie Mellon). Data is stored as a shift in the oxygen atom within the crystalline bond. As a result of the shift, the upper pole of the crystal is either positively or negatively charged.

A limiting factor in Ferroelectric is the destructive read out nature of the device. In this method, any knowledge stored at the cell must be disturbed during the read operation; it is then restored by register circuitry on the chip itself. The most common issues are data retention (relaxation of the position of the oxygen atom), and data imprint (a tendency of the covalent bonds to prefer that the oxygen atom be in one position of the crystal.)

The highest capacity devices are 4 Mb per die.

Magnetic MTJ (MRAM)—In this technology, a magnetic Spin is imposed into a Magnetic Layer by the summation of currents into the layer caused by Eddy currents flowing near its junction. (Figure 4)

Determination of the 1/0 state of a cell is achieved by the reluctance of the Magnetic Layer to lightly imposed current flows in the direction of – or in the direction opposite to – the stored spin.

A limiting factor in MTJ memory is the current required per bit to impose magnetic spin into the Free Layer. The most common issues are adjacent bit flip (field confinement) and electro-migration. As with the aforementioned FeRAM, these devices, too, are low density – also 4 Mb per die. Likewise, these were not considered for the HC-SSR.

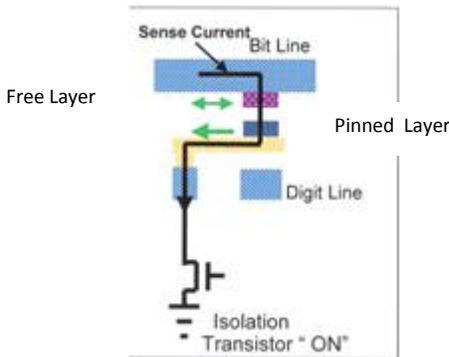


Figure 4 - Operational Principle of MTJ Device (courtesy EverSpin). A summation of magnetic lines of force between currents flowing in the Digit Line and the active Bit line imposes magnetic spin into the device’s ‘free layer.’ Resistance to current flow (read cycle) is determined by parallel or anti-parallel lines of force between the pinned layer and the free layer.

Phase Change—In this technology the physical state of a plastic-like material is changed from an amorphous to crystalline state and back by controlled, direct heating of the memory bit being addressed. The resistance of the material varies by many orders of magnitude (Figure 5). This large spread in resistance makes it easy to discern between a One and a Zero data state.

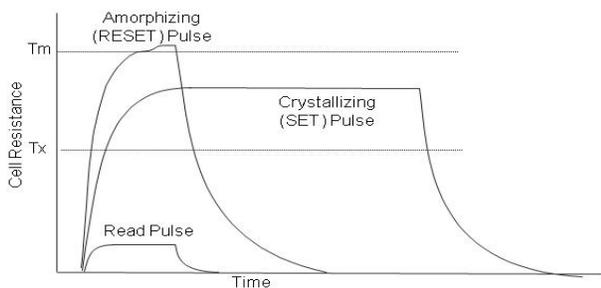


Figure 5 - Phase Change Resistance versus Temperature Profile (courtesy BAE)

Testing has shown that the data storage material is immune to strategic radiation, and therefore is a prime candidate for this application – as increases in radiation tolerance at the device level brings about a decrease in the amount of radiation shielding required.

A limiting factor in Phase change memory is the energy (heat) required in order to liquefy the memory cell. Typically this is one milliwatt per bit, with cell temperatures reaching 230 Celsius, Figure 6.

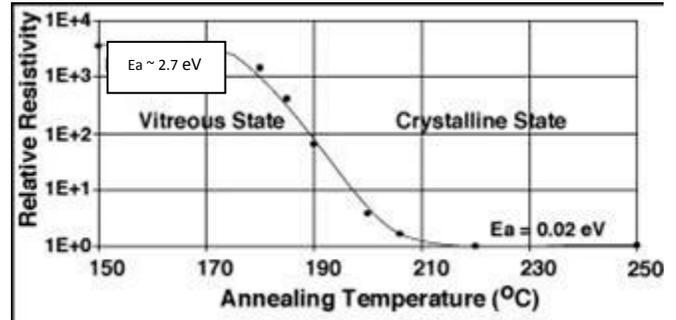


Figure 6 - Activation Energy versus State of an Ovonic Cell (derived from Ovonic)

The main issue with Phase Change memory is the large activation energy (EA) for a cell frozen in the amorphous (high resistance) state. EA is computed to be approximately 2.7 eV; an external application of heat may result in loss of data at the cell.

The densest available Phase Change device today is 4 Mb.

Trapped Charge—Several technologies are available under this category in which a charge (electron or proton) is trapped within the crystalline structure of an Oxide at an interface site. For this paper two technologies were investigated: Flash and NROM.

Floating Gate (Flash)—Non-volatile memory devices have traditionally relied on floating gate technology. A floating gate memory cell contains an electrically isolated gate, a floating gate – below the standard control gate and above the transistor channel. (Figure 7) The floating gate is composed of a conducting material, typically layer of polysilicon. The floating gate memory device stores information by holding electrical charge within the floating gate. Adding or removing charge from the floating gate changes the threshold voltage of the cell, thereby defining whether the memory cell is in a “programmed” or “erased” state [5].

A limiting feature in the Flash technology is the necessity to overcome large energy barriers to impose a charge or remove it (erase). In all devices on the market today, this is done by on-chip generation of a relatively large voltage

(typ. 25V) to bias the substrate and achieve Fowler-Nordheim tunneling.

This action stresses the oxides leading to an accumulation of defect sites that eventually renders the cell useless. Typical device lifetimes are reached after 100,000 erase cycles.

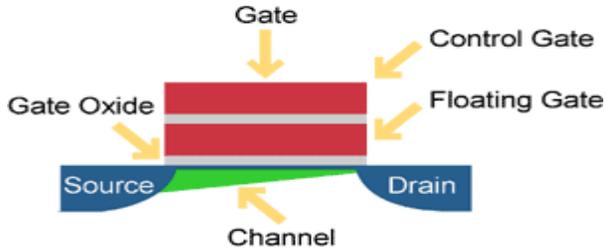


Figure 7 – Typical Floating Gate Memory Cell [courtesy Saifun Semiconductor]

The largest available Flash device today is 8 Gb per die.

Nitride (NROM)—NROM technology contains a nitride layer which traps or stores the charge, instead of a floating gate above the cell. The nitride layer is surrounded by two insulating silicon dioxide layers. (Figure 8) A charge may be accumulated and confined at each end of the nitride layer, effectively storing two separate and independent charges. Each charge can be maintained in one of two states, either “programmed” or “erased,” represented by the presence or absence of a pocket of trapped electrons.

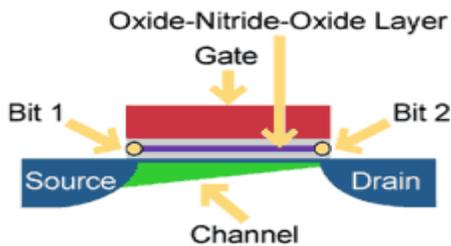


Figure 8 - Typical NROM Cell Construction [Saifun] Note that the device typically stores two-bits per site.

The limiting factors in NROM devices are, similar to Flash, large voltages (in this case 8V) necessary to achieve a write, and the very small number of electrons (sometimes as low as 6), used to signify data.

The largest NROM device available today is 1 Gbit per die; this is not a reflection on the capability of the technology, merely a reflection of the newness of it.

3. SOLID-STATE TECHNOLOGY BASELINES – THE ITRS

The International Technology Roadmap for Semiconductors (ITRS) is a compendium of exhaustive research into the technology trends for the semiconductor industry [6]. Updated yearly, the ITRS identifies emerging technologies as well as tracks technologies widely in use today. Technological advancements are predicted and identified on a year-by-year basis for nearly two decades hence.

A study of the ITRS for the memory technologies identified above reveals, as expected, a doubling of bits-per-die capacity every 30 or so months. The ITRS does not outright predict that a particular technology will double its capacity, rather, the prediction is borne out by an investigation of the feature size and cell-to-cell pitch. Further, some devices will be implemented with multiple bits per cell (example, Figure 8) – with four bits per cell being typical in 2009; this, too, being identified in the ITRS.

A very telling story of expected device capacity is shown in the Figure 9.

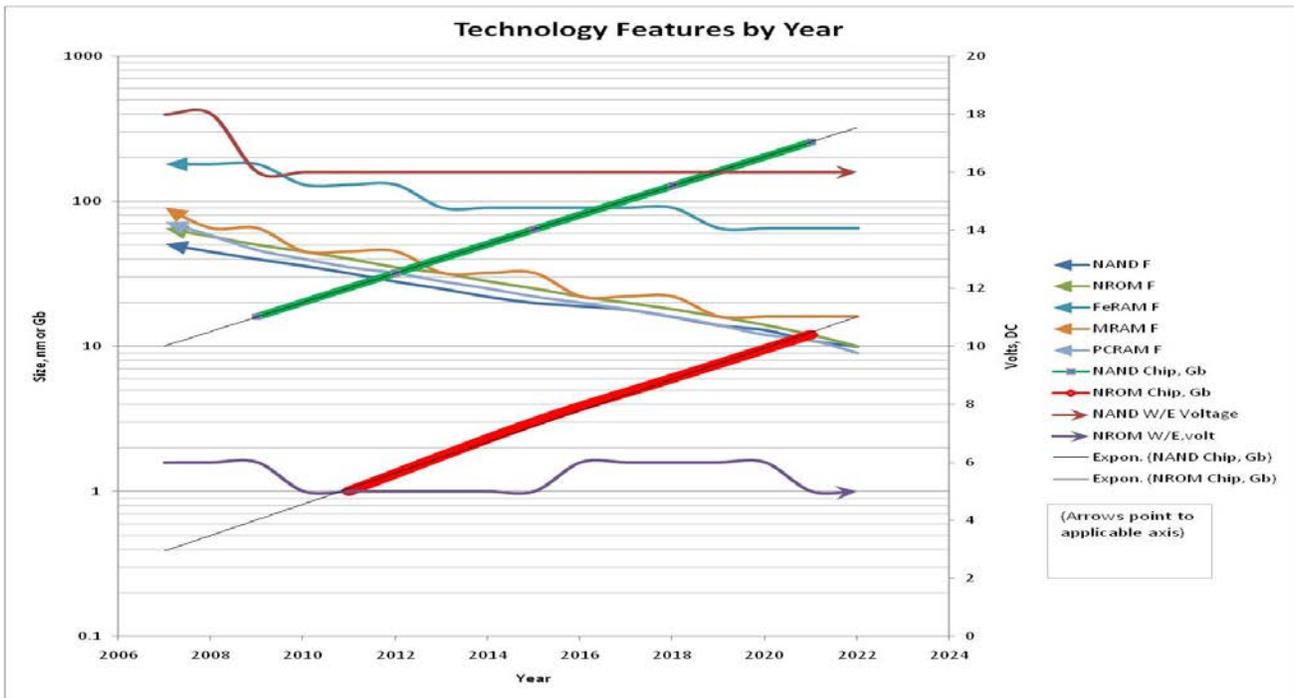


Figure 9 - Technology Features by Year [derived from ITRS] This graph details anticipated device feature size by year (left axis, arrows pointing to left), the Write/Erase Voltage per technology (right axis, arrows pointing to right) and the chip data capacity for NAND Flash and NROM technologies (green and red lines)

Figure 9 details the feature size, write/erase voltage, and chip data capacity for several non-volatile technologies as projected by the ITRS. The graph shows that the feature size for the listed technologies (NAND Flash, NROM, Ferroelectric, MTJ MRAM, Phase Change) reduces at an exponential rate, with a reduction by a factor of 2 approximately every 30 months. (Left axis, arrows pointing to Left). This follows Moore’s Law.

The graph also shows that the Write Erase Voltage (Right axis, arrows pointing to Right) reduces as well. This is an important factor as it indicates that the operating power for a system will be reduced from one generation to the next when considered on a per-bit basis.

The graph concludes with an estimation of per-die data capacity when considering feature size reduction and active die area. (heavy green and red lines, Left Axis).

Flash

Currently available single bit per cell Flash devices reach a capacity of 8 Gb per die. An increase to 16 Gb per die is expected in the 2011/2012 timeframe. Figure 10 plots Flash die capacity predicted by year for next two decades.

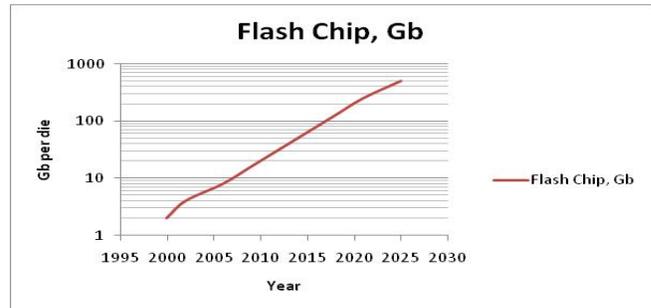


Figure 10 - Flash Capacity per die by Year [ITRS]

NROM

Current NROM technology is up to 1 Gbit per die. There are no radiation tolerant devices available today – other than those obtained by normal processing. Certain information contained herein is predicated upon successful development of devices according to a timeline obtained by private communiqué [7]. The advantage of NROM is that a single cell can be used to store two bits of data (Figure 8). The disadvantage is that a lesser number of electrons represent the state of any one bit; therefore, the assumed likelihood of data loss by heavy ion strike increases.

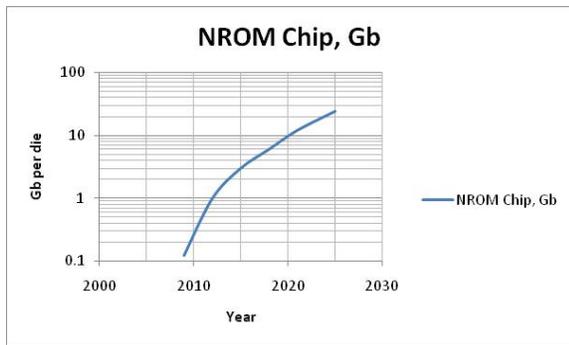


Figure 11 - NROM Capacity per die by Year [ITRS]

The important item to note here is that this is the only high-capacity technology in which a vendor has expressed an interest to develop Megarad level devices. It is conceivable to have a multi-gigabit device operable to at least 1000 krads within 48 months after receipt of funding [7].

Technology Comparison

A comparison of the two technologies (Flash and NROM) reveals this interesting phenomenon:

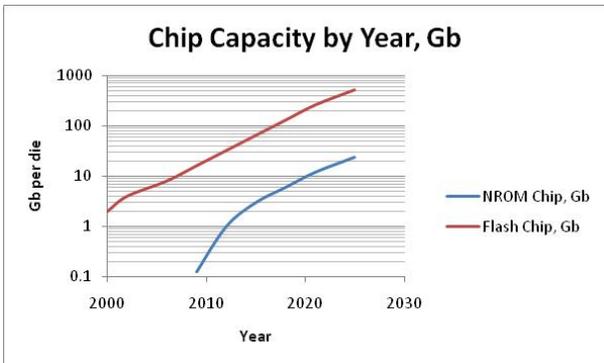


Figure 12 - A Comparison of Predicted bits-per-die by Year [derived from ITRS]

For all years applicable to this paper, the expected capacity of both Flash-based devices and NROM devices increases at the same rate, and follows Moore’s Law almost precisely.

Chalcogenide

For completeness, a discussion of Chalcogenide (CRAM®) technology is presented here. This technology is in its infancy, the largest available device today is 4 Mb. An analysis of the technology is presented:

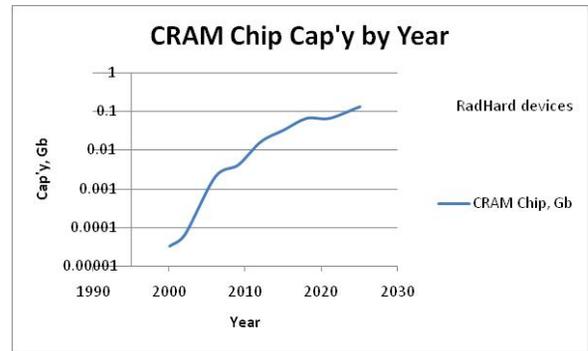


Figure 13 - CRAM® Device Capacity by Year [ITRS]

Note that even two decades from today, the CRAM technology – that is to say the Radiation Hardened CRAM technology – will achieve only 128 Mbits per die.

4. RADIATION EXPOSURE VERSUS CAPABILITY

A study was performed by the Jet Propulsion Laboratory Environmental Analysis Group on seven JPL-managed missions, already launched, soon to be launched, or planned. The total radiation exposure over the lifetime of the mission is shown in Table 1.

Table 1 - Radiation Exposure for Electronics by Mission

Launch Date (or Planned)	Total Dose Exposure (krad, Si)	Mission Objective
Nov 2003	1	Sun
Aug 2005	4	Mars
June 2006	5.5	Kuiper Belt
Sept 2009	1.5	Mars
Nov 2012	21	Venus
June 2016	17.5	Saturn
June 2020	2,670	Jupiter

This information is used in Figure 15, later in this paper.

Radiation Tolerance by Feature Size

Feature size is the semiconductor industry term for the general length of a source-drain channel in a single MOSFET transistor. In general, the total dose radiation tolerance of a semiconductor (or microchip) increases as feature size decreases [8]. This is mainly due the thinner oxides used to form the gate of a MOSFET transistor – the thinner gate results in dramatically fewer places for free electrons to become entrapped in crystalline defect sites. This, in turn, reduces the ability of the transistor to be controlled and, eventually, would lead to a circuit becoming inoperable as intended, usually with serious system consequences.

A previous development of JPL is the X2000 series of non-volatile memory cards. Each card contained twenty Samsung 128Mbit Flash memory devices. These devices were manufactured with 300nm design rules (feature size) and an oxide thickness of 90 Angstroms. Testing performed by JPL showed these devices maintained operation within specification to approximately 15 kilorads for the specific environment and operational cycles [9].

By comparison, testing a 4 Gbit Flash memory [10] manufactured at 90 nm feature size showed normal operation to 100 krad. From there , total dose tolerance has increased dramatically regularly achieving 200 krad [10] [11] for 65 nm.

A plot of Feature Size and Total Dose tolerance is included as Figure 14. One should note that the total dose tolerance of any device studied is dependent upon operational function prior, during, and after exposure. It is conceivable, and normal to have a single device, tested many times and many ways, to have varying total dose

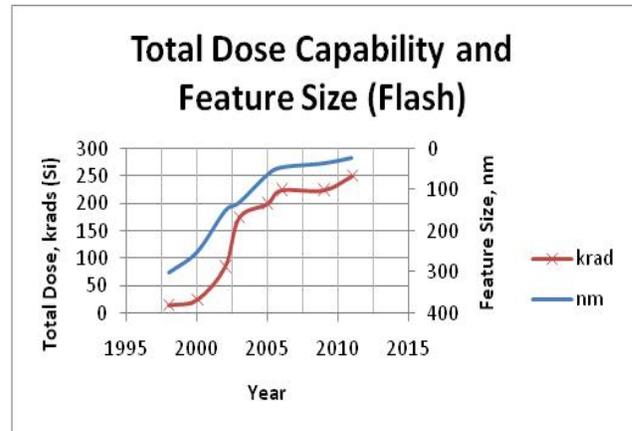


Figure 14 - Total Dose Capability and Feature Size (note CY2009 and beyond Speculative)
capability.

An excellent treatise on the control of radiation effects through design and processing can be found in [12].

A chart representing launches (and planned launches) versus measured and expected (for years 2009 and beyond) radiation tolerance for Flash and NROM technology devices is included as Figure 15. Note that NROM device technology flat-lines at 1 Mrad – this is based upon a business model by the supplier not to pursue capabilities in the multi-Megarad regime at this time.

This chart reveals that for all the missions studied --- save for one --, the use of storage devices manufactured without any special radiation processing is adequate. One must read this with a good amount of caution as the radiation capability of any technology beyond the year 2009 is highly speculative but is intended to indicate an increase in total dose tolerance as device features continue to shrink

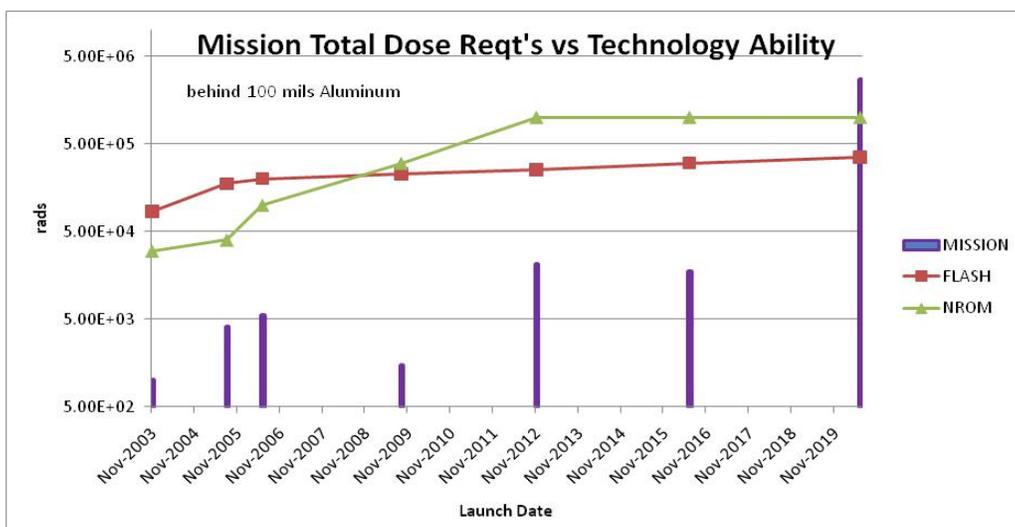


Figure 15 - Mission Radiation requirements and Technology Capability

according to Moore's Law.

5. A NEW MAXIM

Therefore, the selection of Flash and NROM devices will more than suit all mission needs. The sole outlier here is the Jupiter Europa Orbiter mission which requires devices capable of operating after megarads of exposure.

One will note that the evidence suggests a certain exponential growth exists for the required data volume of data recorders. Further, that by its very nature, the semiconductor memory industry is matching such needs with ever increasing density of its products. (Figure 12)

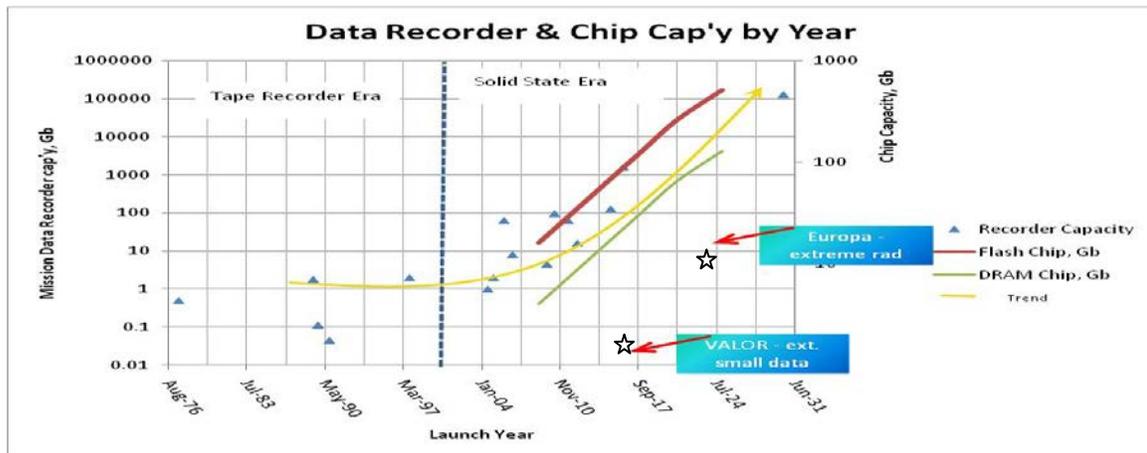


Figure 16 - Recorder and Device Capacity by Year. Note that the Trend for mission data recorder capacity (yellow) is increasing at the same rate as chip capacity for the two technologies analyzed on a year-by-year basis.

Flash

Due to decreased feature size, Flash devices of higher density will experience greater total ionizing dose tolerance and, conversely, experience a high upset rate of data or control information caused by the SEFI⁵ phenomena [13].

Current radiation tolerance is in the realm of 100 krad to 200 krad [14]. This is not a result of special design, but rather the normal increase of total dose capability due to processing, as discussed earlier in this paper.

NROM—NROM is currently the only technology supporting high capacity designs in which the vendor has agreed to develop devices that will perform to radiation immune levels. Testing on existing commercial samples reveals, much as Flash, tolerance to 200 – 300 krad; by design levels exceeding 1000 krad are possible [7].

CRAM—By design, this technology handily maintains normal operation in an environment exceeding 1000 krad [15]. For extreme environments, until other technologies are funded and move forward (e.g. NROM), there is no alternative.

⁵ Single Event Functional Interrupt – a loss or unexpected change of state of a device often by a charged particle upsetting a latch or register within the device's state machine.

Including the desired Recorder capacity, by Year, gives us the Figure 16. This series of curves is telling us that the required capacity of a recorder matches the data capacity afforded by any selected technology on a match-for-match basis year by year.

Impact of Technology on Power

Figure 17 identifies predicted Watts per Terabit for the memory technologies shown (Flash-red; DRAM-blue; NROM-yellow) on a generational basis. These lines exhibit the now familiar slope of Moore's Law. Therefore, this series of curves shows that for any recorder of a particular size and generation, it is easy to predict operational power based upon information available today.

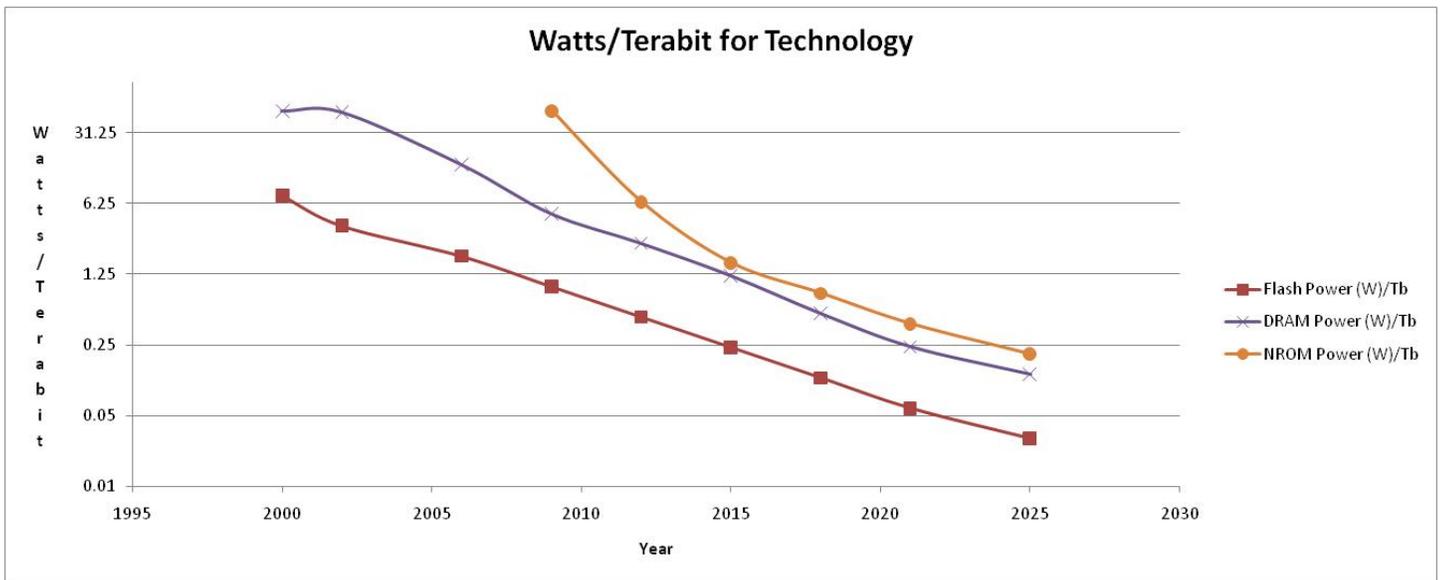


Figure 17 - Power Consumption (Watts per Terabit) for Flash, DRAM, and NROM technologies. (ITRS)

Table 2 - Memory Technology Comparison (Stoplight Chart)

	MEMORY TECHNOLOGY COMPARISON (valid Oct 2008)								FUTURE TECH	
	SRAM	DRAM	FLASH	Chalcogen ide	Magnetic (MRAM)	Ferro-electric (FRAM)	Hard Disk Drive	Tape Recorder^	Nanotube	Holographic
POWER	Red	Red	Green	Yellow	Red	Green	Red	Red	Green	Red
SPEED	Green	Green	Yellow	Red	Green	Green	Red	Red	Green	Green
DENSITY	Red	Green	Green	Red	Red	Red	Green	Green	Green	Green
RADIATION TOLERANCE	Green	Yellow	Yellow	Red	Green	Yellow	Red	Yellow	Green	Green
VOLTATILITY	Green	Green	Red	Red	Red	Red	Red	Red	Red	Red
ENDURANCE	Green	Green	Yellow	Yellow	Green	Green	Green	Yellow	Green	Green
AVAILABILITY										
-Commercial	Green	Green	Green	Red	Green	Green	Green	Red	Red	Red
-Industrial	Green	Green	Yellow	Red	Green	Green	Green	Red	Red	Red
-Mil	Green	Green	Red	Yellow	Green	Red	Green	Red	Red	Red
-Space-grade	Green	Red	Red	Yellow	Yellow	Red	Red	Red	Red	Red
							* redesign of custom ASIC required	^ No Longer mfg'd	Still in the Laboratory. 2011?	Still in the Laboratory. ???

The impact of such a simple statement is at the same time beautiful and difficult to realize. What this is telling the designer is that: “For any technology selected, mass and power of any data recorder design will remain relatively constant even as mission data demands increase from one generation to the next.”

What it Means

For all intents and purposes, the “fit” of the SSR can remain relatively constant throughout the years – that as the required capacity doubles, the media meets the challenge at the same pace. Therefore, assuming a recorder of Size *X* bits today and performing mass and volume predictions on that design, moving that recorder ahead several generations will result in the same mass and volume, even though the capacity of the data recorder increases exponentially.

Technology Comparison

Table 2 is a compendium of known, published technologies in use or in development today enabling a comparison of any one technology listed for specific attributes against another. In this stoplight chart, red indicates unfavorable attributes, green equals favorable, and yellow indicates cautionary acceptance. It can be easily seen that in terms of memory technology for space, the selection is limited – especially so when it comes to non-volatile radiation hardened devices.

Part II

6. HIGH CAPACITY SOLID STATE RECORDER

The impetus of this paper was the assignment to predict mass, power, volume, and cost for multi-terabit data recorders of the capacities shown in Table 3 with deliveries occurring in the year given [16].

Table 3 - Data Recorder Capacity Requirements

YEAR	DATA CAPACITY
2010	2.88 Tb
2020	30 Tb
2030	130 Tb

This part of the paper presents a design of a 130 Tb Non-Volatile data recorder which could be built for delivery in the year 2030, based on the information presented in Part I. In developing the SSR design, technology selection was limited to the two technologies examined in detail in Part I (Flash and NROM) and a technology freeze date of 36 months in advance of delivery was assumed.

SSR – Flash & NROM Based designs

Table 4 - ITRS predicted device capacity (Gb) by Year, for three technologies

Year	2015	2018	2021	2025
Technology				
CRAM	0.032	0.064	0.064	0.128
NROM	4	8	16	32
Flash	64	128	256	512

In this particular case, the year 2025 is selected as the freeze date and appropriate device capacity is used. The ITRS is predicting that by the year 2025, CRAM-based devices will have a capacity of 128 Mbit per die; NROM will have 32 Gb per die. Flash will be available at 512 Gb per die [6].

Table 5 highlights the number of die necessary of three discussed technologies to achieve a recorder of specified capacity in the timeframe given.

Table 5 - Minimum Device Count by Technology to achieve a Data Recorder of Specified Capacity in the Year shown. (Example, a 2.88 Tb recorder would require 720,000 CRAM devices, or 2,800 NROM devices, or 180 Flash devices)

Year	Recorder Capacity	Technology / Device Capacity /Chip Count		
		CRAM 4Mb	NROM 1Gb	Flash 16Gb
2010	2.88 Tb	720,000	2,880	180
2020	30 Tb	468,750	1,875	118
2030	130 Tb	1.02E+06	4,063	254
EDAC		Required	Not Required	Required

Therefore, a 130 Tb data recorder designed using devices predicted available in the year 2025, would require the following number of devices – not including any additional devices used for Error Detection and Correction (EDAC) or sparing and word control:

- CRAM: > 1 million devices
- NROM: approx 4060 devices
- Flash: approx 256 devices

By examination, it can be easily seen that developing a high capacity data recorder as proposed using CRAM technology is impractical.

Error Detection and Correction Codes

All technologies are subject to loss of information, be they the loss of a single bit or the loss of an entire chip, the inclusion of Error Detection and Correction is an appropriate action by any system designer.

Just by using EDAC, there is an associated bit-penalty ranging from a few percentile to over 25 per cent. The exact selection of most-appropriate Error Correction codes is highly dependent upon error phenomena not only for each technology [17], but also certain aspects of the exact part being considered as well as the environment to be encountered [18] [19]. For the recorder described below, a Reed-Solomon code was selected. It is important to note that this may not be the prime candidate, but was chosen for known characteristics.

Selection of a more appropriate EDAC code [19] would, most assuredly, bring about a reduction in recorder mass

and power.

One final, important, distinction between the two prime technologies is required: NROM devices have built-in an incredibly powerful error detection and correction machine, capable of detecting four bits in error in a single word, and correcting three bits per word [7]. With this method, the Bit Error Rate is calculated by the manufacturer to be less than $1.18E^{-19}$ errors/bit-day – equivalent to 3 unrecoverable bit errors every 550 years for the entire recorder; therefore, the package numbers presented here do not include extra devices required for EDAC other than that provided by the chip itself.

Table 6 - HC-SSR Basic Features

- Redundant SpaceWire interfaces
- Redundant Power Converters
- Redundant Controllers/Supervisors
- Multi-plane Non-volatile storage
- Error detection and correction
- Device isolation to the EDAC Word

7. HIGH-CAPACITY SSR DESCRIPTION

Block Diagram

A block diagram of a recorder meeting all requirements

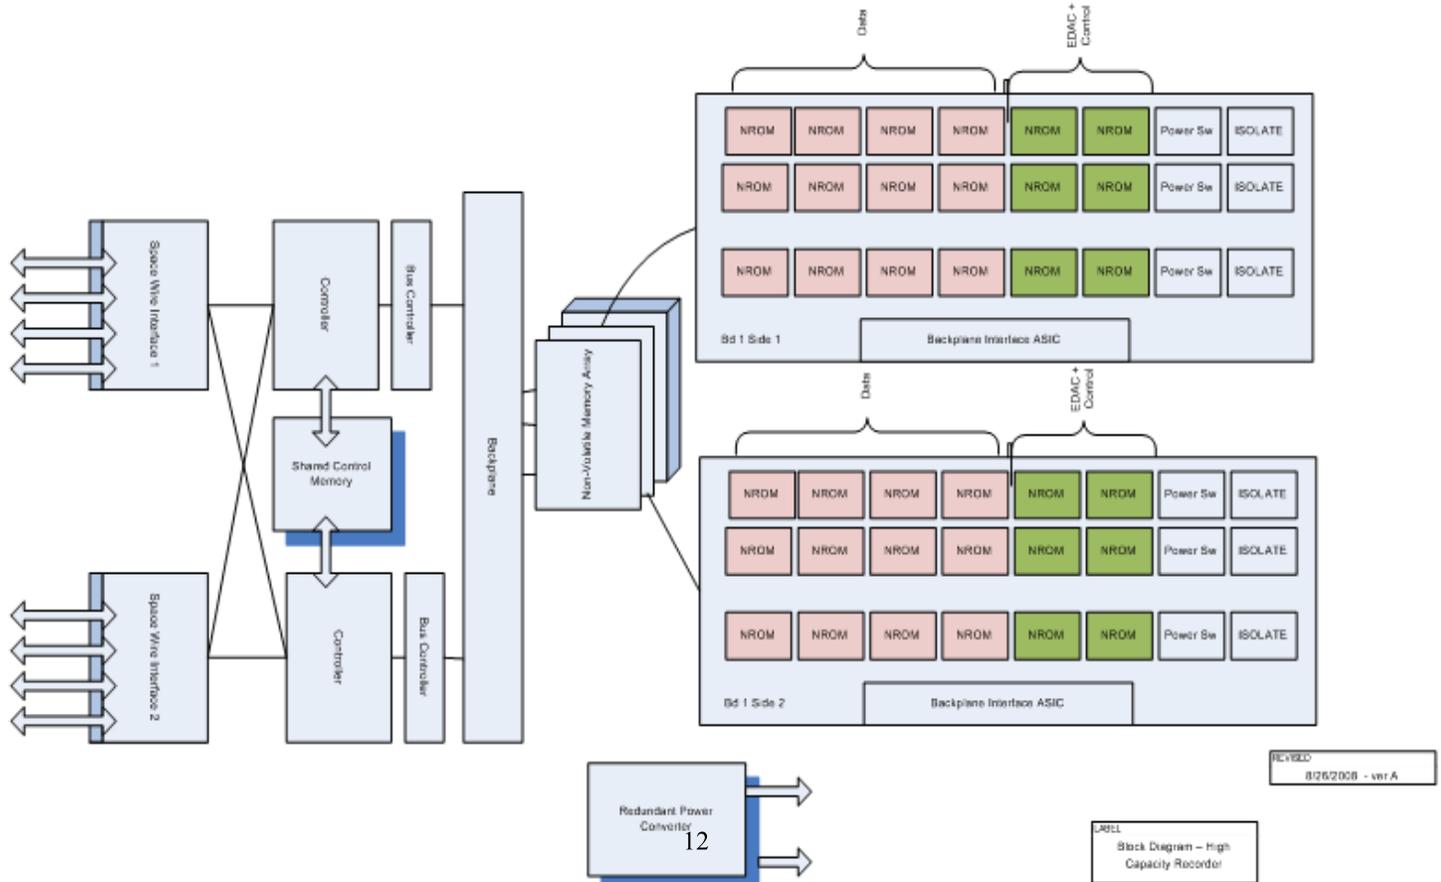


Figure 18 - HC-SSR Conceptual Block Diagram

listed in Table 6 is shown as Figure 18. The recorder features cross-coupled SpaceWire interfaces as well as redundancy down to the memory plane.

Cross-coupled controllers with shared memory ensure operation without the need to re-build memory allocation tables should a fault occur.

Recorder Summary

Using conservative packaging techniques, two recorder designs were developed: one using Flash memory and the other using NROM. A summary of their characteristics is provided in the tables below. A more detailed discussion of design and implementation follows.

Board Design

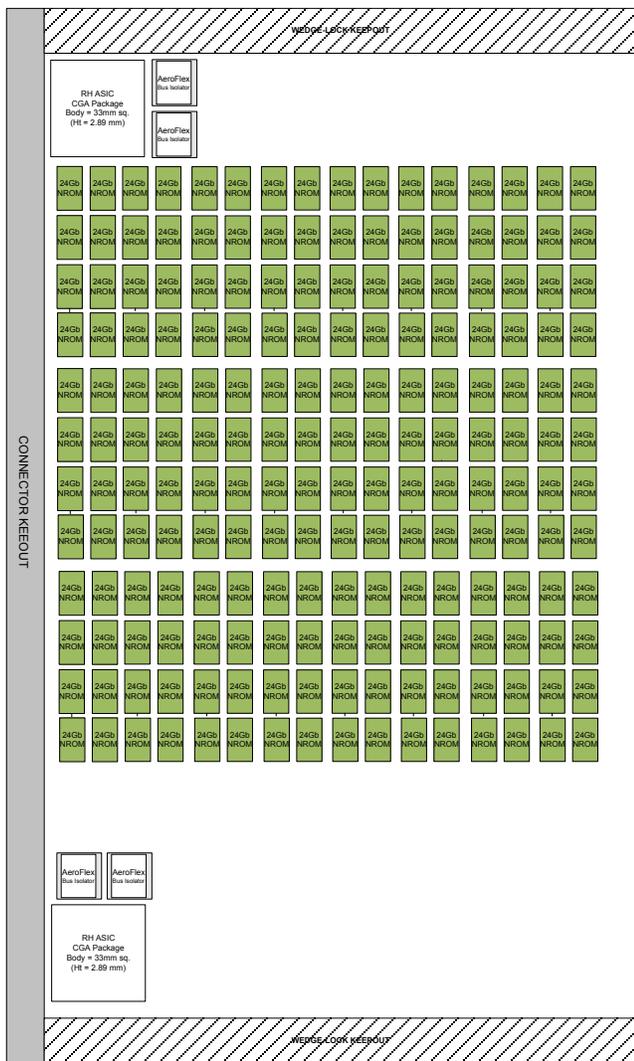


Figure 19 - One side of 8Tb Memory Slice using NROM Technology

The NROM-based solution is presented as Figure 19. In the figure, nearly two hundred 2-Gbit NROM devices are installed on each face of the 9U card. (The whitespace is preserved in this example for passive components and design margin.) At the time of this writing, the availability of chip-stacked NROM devices was uncertain and omitted from the trade space.

Design of a Flash-based memory card would be of similar ilk. Device density would permit the use of a 6U card. Compare the number of devices required for each design as given in Tables 7 and 8.

Table 7 - NROM-based HC-SSR Characteristics

ITEM	Quantity	Units
Size, 9U	Approx 40 x 33 x 30	cm (h, w, d)
Number of PWBs (slices)	14	ea.
Size, 9U ⁶	Approx 40 x 33 x 30	cm (h, w, d)
Mass	10	kg ⁷
Power	30	W, (not incl. uncertainty)
Interface	Dual Channel Space Wire	
Capacity	129.6	Tb (Sparing extra)
Devices	4060	32 Gb die

⁶ IEEE 1101.1-1998

⁷ Assumes worst case devices are operable to expected regime without shielding

Table 8 - Flash-based HC-SSR Characteristics

ITEM	Quantity	Units
Size, 6U	Approx 28 x12.7 x 25	cm (h, w, d)
Number of PWBs (slices)	4	ea.
Mass	3.5	kg
Power	6	W, (not incl. uncertainty)
Interface	Dual Channel SpaceWire	
Capacity	129.6	Tb (EDAC & Sparing extra)
Devices	254	512 Gb (not incl. EDAC)

Controller Board

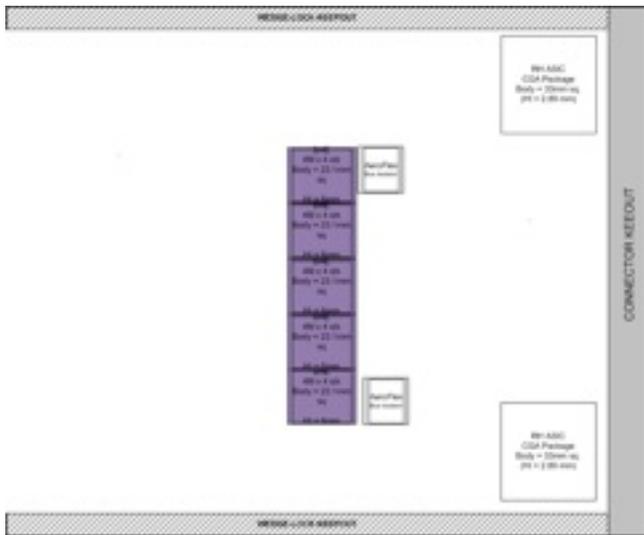


Figure 20 - Controller Card Layout (Violet = Shared Memory)

Figure 20 shows one of many possible methods of design for an IO controller slice. The IO controller slice provides redundant SpaceWire interfaces and redundant interfaces to the backplane. Although there are two controllers shown,

there is, in this concept, a single set of Shared Memory which is common between both controllers.

One of the main functions of the active controller is to build and maintain a directory of memory sector-to-physical address translation tables. Having the table stored in Shared Memory assures timely restart of the recorder function, should it become necessary to activate and swap to the alternate controller.

Interfaces

The Data Recorder interfaces with the Command Data Handling (CDH) subsystem via the standard SpaceWire interface⁸. Utilizing conservative design techniques, a sustained Write speed of at least 200 Mbps into memory is guaranteed. The recorder is capable of supporting simultaneous Writing and Reading of data subject to usual certain timing restrictions imposed by the memory devices themselves.

Standard 28 Volt bus power is applied to the SSR which features two power converter units for redundancy.

Redundancy & Fault Tolerance

The execution of good systems engineering practice is recommended and will, most likely result in a physical enlargement of the SSR in one dimension. This may consist of addition of spare memory to each memory slice, or the addition of one or more spare memory slices to the entire design.

With respect to Fault Tolerance, aside from the use of redundant interfaces and controllers, described above, the HC-SSR is designed such that individual blocks of words as determined by EDAC boundaries can be tagged as Do Not Use. An example of an EDAC word with embedded control bits is included as Figure 21.



Figure 21 - EDAC Word Bit-Level Assignments

File System

The HC-SSR contains built in File Management system supervised by Interface/Controller ASICs. File storage is compatible with the CCSDS File Delivery Protocol (CFDP), thereby reducing the need for processing and packetizing data externally.

⁸ ESA Standard ECSS-E-50-12A

The architecture permits random access to any EDAC word within the recorder. Simultaneous writing and reading is supported. Memory is divisible into blocks with an unlimited number of blocks forming a Sector. Each sector contains information that permits reconstruction of the File Structure independent of the CDH.

Sector Definition

As unique Data Units are identified by the CDH host, the SSR commences placement of information in the first available memory location. The information placed therein contains information identifying the Sector and relevant information with respect to data unit identification.

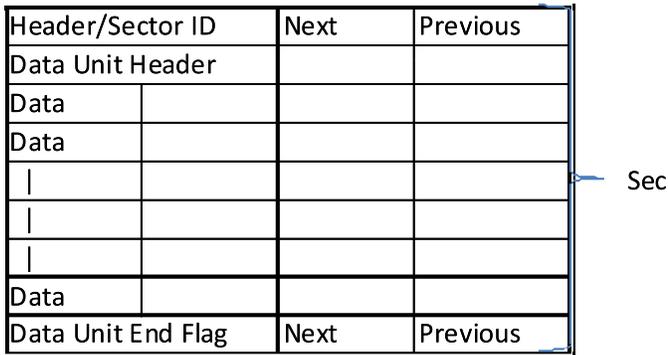


Figure 22 - Sector & Data Unit Structure

Each sector of information has contained within identification as to the next sector tag as well as the tag of the previous sector. In this way, should the file structure need to be rebuilt, a scan of memory *anywhere* will reveal not only the type and source of information encountered but also the logical placement of all relevant information before and after this data block.

Consumables

For Flash memory technology, the data is stored, and erased, by the imposition of high voltages onto the substrate and controlling gates. Therefore, this technology is stressful to the device and is rated to 100,000 write or erase cycles. NROM technology uses a different charge storage mechanism and is therefore not subject to write/erase degradation.

Stability

Testing has shown commercial NROM devices to retain data up to 300 krad, a simple refreshment following exposure permits a likewise gain in radiation tolerance. That is 300krads, refresh, 600 krad, etc [7].

Flash devices are prone to damaging charge-trapping and therefore do not benefit by the refreshment scheme outlined above. The users of Flash must take the surrounding radiation environment into consideration when declaring the number of write cycles a certain design can accommodate.

Chassis

Dimensions and mass for the HC-SSR assume the chassis is manufactured in 100 mils aluminum. One possible design is shown in Figure 23.

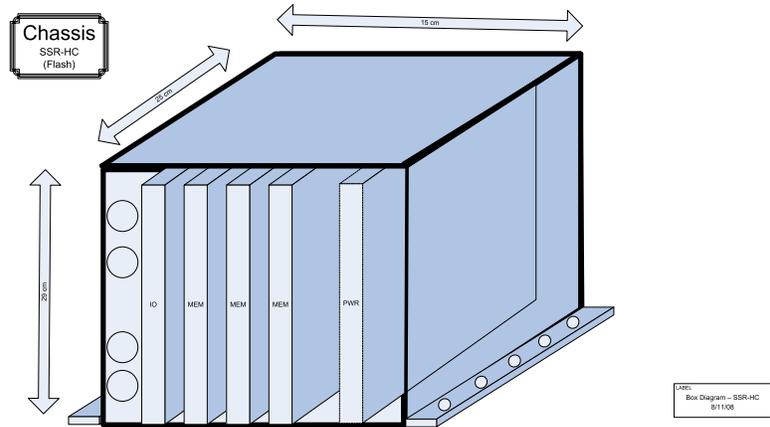


Figure 23 – Packaging Concept – HC-SSR

Power

Worst case power estimation for the Flash unit is 6W and 30W for an NROM-based design⁹. The power estimation includes both SpaceWire interfaces enabled and simultaneous read and write of information occurring. Unused non-volatile memory devices are unpowered and the non-addressed memory controller ASICs are in a Deep Sleep mode.

8. CONCLUSION

A review of data storage technologies was presented and evidence presented to affirm the assertion known as “Moore’s Law” that storage capacity of a single element will increase every (approx.) 30 months.

A prognostication of data recorder capacity requirements for future JPL mission and programs has been made. Based upon a past history spanning over thirty years, it is clearly evident that required capacity will progress at an exponential rate, effectively mimicking the rate provided in Moore’s Law. A new Axiom is coined: A reasonable accurate prediction of data recorder volume and power can

⁹ Power margin (uncertainty is not included in these figures)

be made when based upon today's technology and today's data storage requirements.

A High Capacity Solid State Recorder (HC-SSR) is suggested using devices predicated to be available for each decade. The design is robust and easily implemented using conservative design and manufacturing techniques.

DEFINITIONS

rad (radiation absorbed dose): the dose causing 0.01 joule of energy to be absorbed per kilogram of matter. As the absorption is greatly affected by the molecular structure of the material, citations should also indicate the material as a subscript to the term "rad", as in rad_(Si), indicating Silicon equivalency. For the purposes of this paper, radiation equivalency always assumes Silicon.

(For completeness, it should be noted that System International replaced the "rad" with the unit Gray (Gy), and having an equivalency of 100 rads = 1 Gy [21]. However, the use of rads, kilorads, megarads remains in the industry vernacular and is used in this document.)

Moore's Law Named after Fairchild Semiconductor technologist Gordon Moore, Moore's law was derived from empirical data which shows that the dimensions of basic memory cells will shrink by approximately 50% of the previous value every 30 to 36 months. It is Moore's Law, more or less, that forms the backbone of the ITRS examinations for memory devices.

ADDITIONAL MATERIAL

Standard Dose Rates for Various Orbits and Missions (per year)

Earth		
-	LEO	100 rad (protons)
-	MEO	100 krad (protons & electrons)
-	GEO	1 krad (electrons)
-	Transfer Orbit	10 krad (protons & electrons)
Mars		
-	Surface	2 krad (electrons)
-	Orbit	5 krad (protons)

-	Transit	5 krad (protons)
Jovian		
-	Transfer	100 Mrad (protons & electrons)

ACKNOWLEDGEMENT

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

- [1] G. Moore, "Cramming more Components onto Integrated Circuits," *Electronics*, vol. 38, no. 8, April 1965.
- [2] G. Moore, "No Exponential is Forever: But "Forever" Can Be Delayed," *Digest of Technical Papers, International Solid State Circuits Conference*, pp. 1.1-1 thru 1.1-19, 2003.
- [3] Seagate Technology Company. Seagate Technical Corporation. [Online]. http://www.seagate.com/docs/pdf/marrketing/Article_Perpendicular_Recording.pdf
- [4] Blu-Ray Disc Association. (2005, March) Blu-Ray Disc Technical Papers.
- [5] Saifun Semiconductor. (2008) NVM Technology Overview. [Online]. <http://www.saifun.com/content.asp?id=113>
- [6] International Technology Roadmap for Semiconductors. (2008, August) ITRS Home. [Online]. www.itrs.net
- [7] S. Tyson, "Private Communique," Albuquerque, NM, 2008.
- [8] N. Haddad and T. Scott, "Adapting Commercial Electronics to the Naturally Occurring Radiation Environment," in *IEEE Nuclear and Space Radiation Effects Conference Short Course*, Tucson, 1994, pp. iv-14.
- [9] D. R. Roth and et al, "SEU and TID Testing of the Samsung 128 Mbit and the Toshiba 256 Mbit flash memory," in *Radiation Effects Data Workshop*, Reno, 2000.

- [10] F. Irom and D. Nguyen, "Single Event Effect Characterization of High Density Commercial NAND and NOR Nonvolatile Flash Memories," , Honolulu, 2007.
- [11] C. Hafer, M. Lahey, and et_al, "Radiation Hardness Characterization of a 130nm Technology," in *Proceedings IEEE Nuclear and Space Radiation Effects Conference*, Honolulu, 2007.
- [12] T. R. Oldham, J. Friendlich, and et_al, "TID and SEE Response of an Advanced Samsung 4Gb NAND Flash Memory," , Honolulu, 2007.
- [13] R. C. Laco, "CMOS Scaling, Design Principles and Hardening-by-Design Methodologies," in *Nuclear and Space Radiation Effects Conference Short Course Notebook*, Monterey, 1993, pp. II-1 thru II-142.
- [14] J. Patterson and S. Guertin, "Emerging SEFI Modes and SEE Testing for Highly-Scaled NAND FLASH Devices," in *Proceedings 2005 Non-Volatile Memory Technology Symposium*, vol. CD-ROM, Dallas, TX, 2005, pp. G-3, Session G ; Paper 3.
- [15] D.. Nguyen and F. Irom, "Total Ionizing Dose (TID) Tests on Non-Volatile Memories: Flash and MRAM," in *2007 IEEE Radiation Effects Workshop*, vol. 0, Honolulu, 2007, pp. 194-198.
- [16] G. Noreen and et_al, "Low Cost Deep Space Hybrid Optical/RF Communications Architecture," , Big Sky, Montana, 2009, Pre-print.
- [17] BAE Systems North America, Microcircuit, CMOS, 3.3V, NVRAM 8406746, April 28, 2008, Rev A.
- [18] T. Sasada and S. Ichikawa, "Application of Solid State Recorders to Spacecraft," in *Proceedings, 54th International Astronautical Congress*, Bremen, 2003.
- [19] H. Kaneko, "Error Control Coding for Semiconductor Memory Systems in the Space Radiation Environment," in *Proceedings, 20th IEEE International Symposium in Defect and Fault Tolerance in VLSI Systems*, DFT2005, Monterey, 2005.
- [20] T. Sasada and H. Kaneko, "Development and Evaluation of Test Circuit for Spotty Byte Error Control Codes," in *Proceedings, 57th International Astronautical Congress*, Valencia, 2006.
- [21] Bureau International des Poids et Mesures. (2008, August) SI Base Units. [Online]. http://www.bipm.org/en/si/base_units/

BIOGRAPHY



Author, Karl Strauss, has been employed by the Jet Propulsion Laboratory for over 22 years. He has been in the Avionics Section from day One. He is considered JPL's memory technology expert with projects ranging from hand-woven core memory (for another employer) to high capacity solid state designs. He managed the development of NASA's first Solid State Recorder, a DRAM-based 2 Gb design currently in use by the Cassini mission to Saturn and the Chandra X-Ray observatory in Earth Orbit. Karl was the founder, and seven-time chair of the IEEE NonVolatile Memory Technology Symposium, NVMTS, deciding that the various symposia conducted until then were too focused on one technology. Karl is a Senior IEEE member and is active in the Nuclear and Plasma Science Society, the Electron Device Society and the Aerospace Electronic Systems Society.

Karl is also an active member of SAE.

Karl thanks his wonderful wife of 28 years, Janet, for raising a spectacular family: three sons, Justin, Jeremy, Jonathan. Karl's passion is trains and is developing a model railroad based upon a four-day rail journey across Australia's Northern Outback.