TMT ANALYTICS

4DS Memory Ltd.

Next-gen Cloud and DC storage technology

4DS Memory Ltd (4DS) is a semiconductor development company with facilities in Silicon Valley, USA. It is developing an emerging memory technology for a new memory class, which has some major advantages over traditional memories, such as NAND Flash memory, used to store data in mobile phones, laptops and data centers. Due to progress in semiconductor manufacturing, allowing for ever-smaller chips, NAND Flash is expected to run into physical limits within the next seven years. The increasing amount of video content is the secular driver of the insatiable demand for cloud and data center storage capacity.

Addressing fast growing cloud and data center storage markets

4DS is developing an Interface Switching Resistive Random Access Memory (ReRAM) technology, which has distinct advantages over other emerging storage technologies, including scalability of cells, low power consumption and high endurance levels. The company will need to prove technical and commercial viability of the technology at high resolutions (small circuitry linewidths) with the primary goal to scale below 3D NAND Flash to attract licensees. Given the specific characteristics of Interface Switching ReRAM, the cloud and data center storage markets, that are growing at double digits, are key targets.

3rd party validation by data storage leader Western Digital

For the past two years, 4DS has had a strategic partnership with HGST, a subsidiary of NASDAQ-listed Western Digital (WDC, EV ~US$20 BN). The joint development agreement (JDA) was renewed for another 12 months in July 2016. As part of this agreement, HGST received an option to a non-exclusive license of the 4DS technology. Western Digital has publicly stated that 3D ReRAM is their technology of choice to address the Storage Class Memory market of the future.

End game for 4DS may be take-out by strategic player: BUY

Our initial, near term, price target of A$ 0.05 per share is based on the valuation of other ASX-listed memory companies as well as the value of the HGST JDA and invested capital to date. However, we believe there is substantial upside, to A$ 0.12 per share, if and when 4DS can scale its technology to resolutions of 45 nanometers (nm) and below. We believe this will draw substantial strategic and competitive interest in the company, including that of its current JDA partner HGST. Based on its development track record to date, we expect 4DS will be successful in doing so in the near term, thereby unlocking the share price upside. Hence, we start coverage of 4DS with a BUY rating and a PT of A$ 0.05.

4DS Memory Limited	
Number of shares (m)	659.2
Number of shares FD (m)	833.0
Market capitalisation (A$ m)	15.8
Free Float (%)	83%
12 month high/low A$	0,046/0,018
Average daily volume (k)	1,014

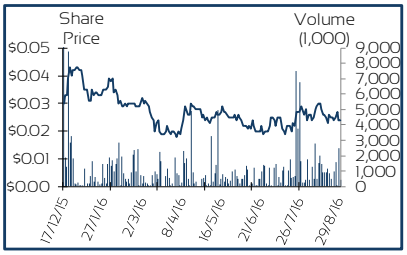

Readers should be aware that TMT Analytics has been engaged by the company covered in this report for ongoing research coverage. Please refer to the final page of this report for the General Advice Warning, disclaimer and full disclosures.

Company report

4DS.ASX

Semiconductors & Semiconductor Equipment

Australia

Risk: High

4DS Memory Limited (4DS.ASX) is a semiconductor development company aiming to provide an enterprise grade storage memory for cloud and data center storage markets. The company is developing a proprietary Interface Switching ReRAM technology leveraging expertise from a strategic partnership with a leading data storage player. 4DS plans to demonstrate commercial viability in the 2017/2018 timeframe.

BUY

Current price: A$ 0.024

Price target: A$ 0.05

1 September 2016

Share prices are closing prices of 31 August 2016

Analyst: Marc Kennis

marc.kennis@tmt-analytics.com.au

+61 (0)4 3483 8134

Summary and Conclusion

Interface Switching ReRAM has key advantages over other emerging technologies

The next-generation computer memory cells 4DS is developing are based on changes to the resistive properties of the used materials in order to store information. These cells use less energy than other emerging technologies. Additionally, because of the lower operating currents, the memory cells degrade substantially slower than these other technologies, resulting in very high endurance levels. Furthermore, due to the specific method in which Interface Switching ReRAM cells are switched, 4DS expects it will be able to scale the technology down to narrower linewidths compared to what competing technologies are expected to achieve over time. This is important because computer chips are becoming ever smaller, facilitated by ongoing advancements in semiconductor manufacturing technologies, and memory cells also need to keep offering more storage capacity at lower costs. So being able to scale down the cells is extremely important for a new memory technology, not just for the first commercially available iteration of the technology, but for many years to come. This ability to substantially shrink the memory cells in the next 10 to 20 years is not obvious for some of the other emerging alternatives to today's NAND Flash memory because of the high electrical currents required.

Addressing a data storage market that is growing at double digits

Because of these specific characteristics, Interface Switching ReRAM technology (also referred to as Non-Filamentary ReRAM) is very well suited to high capacity data storage applications (Storage Class Memory, SCM), such as cloud and data center storage, that require low energy consumption and high endurance rates. The cloud and data center storage market has been growing very fast for a number of years on the back of strong growth in cloud computing and online video content, such as Netflix, YouTube and video in Social Media. Global growth in data center capacity is projected to be in excess of 10% CAGR in the foreseeable future.

4DS will need to demonstrate GB silicon storage viability in the next few quarters

In order to attract licensees and/or potential strategic acquirers, 4DS will need to demonstrate so-called GB silicon storage viability, which will require the company to demonstrate cell scalability below 50nm. Additionally, the company will need to characterize cycle endurance, read/write speeds and data retention across a statistically significant number of memory cells. We expect 4DS to be able to take this technical hurdle in the near term, assisted by its valuable development partner HGST, which is owned by global storage giant Western Digital (WD) that publicly endorsed ReRAM as its Next-Gen SCM of choice.

Strategic end game for 4DS may be an acquisition by industry player

Once GB silicon storage viability is demonstrated, we expect the company will be in a position to start discussions with potential licensees, such as Samsung, Micron, SK Hynix etc, who would need several more years to commercialize the technology, i.e. market entry around 2019-2020. HGST has an option for a non-exclusive license with 4DS. However, based on the acquisitive nature of semiconductor manufacturers when it comes to unique Intellectual Property (IP), we believe an outright acquisition of 4DS may be more likely than a scenario of licensing the IP to multiple licensees.

We expect 4DS to rerate to valuations of ASX-listed semiconductor companies

ASX-listed computer memory companies are valued substantially higher than 4DS, even though 4DS seems to be well ahead technologically. Other non-memory ASX-listed semiconductor companies are valued even higher, based on executed license and technology evaluation agreements. Hence, we anticipate a substantial rerating of 4DS' shares if and when the company succeeds in demonstrating GB silicon storage viability, which would be the prelude to licensing deals or an outright take-over. We start our coverage of 4DS with a BUY rating and an initial price target of A$ 0.05.

4DS working on Resistive Random-Access Memory (ReRAM)

As the name partially suggest, ReRAM is based on the resistive characteristics of certain materials in order to store data. Specifically, ReRAM uses an electric current to change the resistance level of the materials inside the memory cell, with a low resistance level associated with a value of 1 and a high resistance level associated with the value of 0.

In a ReRAM cell a dielectric, or non-conductive, material is sandwiched between two metals. When a positive or negative voltage is applied to one of the electrodes, the electrical resistance level of the sandwiched material in the middle changes to high or low, depending on the polarity of the voltage (+ or -).

ReRAM has several key advantages over NAND Flash

Because ReRAM doesn't use an electrical charge to store information, as NAND Flash does, but instead uses a low voltage to change a cell's resistance level, ReRAM is much more energy efficient, which is important in many applications, such as storage in data centers and in mobile devices. Also, the switching speed of ReRAM can be 100 nanoseconds or faster, depending on the voltage used, which is substantially faster than NAND Flash, which typically switches at speeds of more than 1 microsecond (1 microsecond = 1,000 nanoseconds).

Furthermore, ReRAM has a far superior endurance with the number of P/E cycles potentially orders of magnitude higher compared to today's MLC NAND Flash. Additionally, because the manufacturing process for ReRAM cells is simpler and involves fewer process steps, ReRAM should be cheaper per unit of storage once it's in volume production.

> *Most importantly though, ReRAM cells can be scaled to much smaller geometries than NAND Flash, which allows for far superior densities and substantially more runway when it comes to taking advantage of future technological developments, especially in lithography.*

Different approaches to switching a ReRAM cell: filament and interface switching

There are different ways to alter resistance levels in ReRAM cells. One way is to create a conductive filament, or conductive channel, from one electrode to the other. As illustrated in Figure 1, when a voltage is applied to the top, copper (Cu) electrode, positively charged copper ions travel through the non-conductive dielectric material towards the bottom electrode made of platinum. The copper ions start to form atoms at the bottom electrode and eventually form a conductive filament connecting the top and bottom electrode. In this state the cell exhibits a low electric resistance level representing a value of 1. Applying a reversed voltage, the filament is destroyed, changing the material's resistance back to high, representing a value of 0.

FIGURE 1: THE FORMING OF FILAMENT IN RERAM

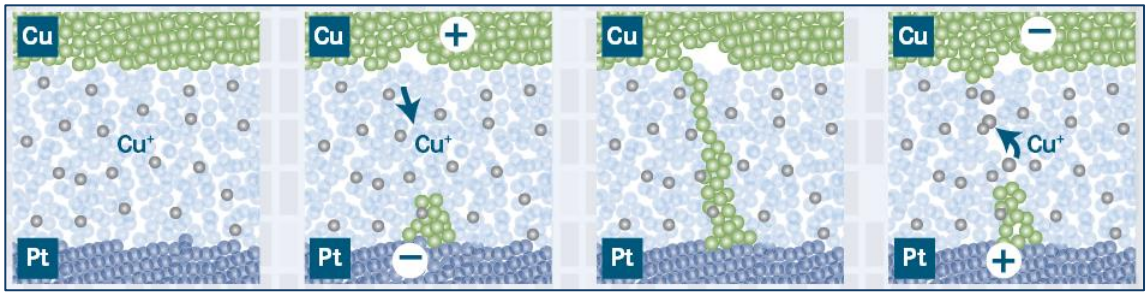

Source: Forschungszentrum Jülich, TMT Analytics

Filament-based ReRAM cells have been developed by companies such as Panasonic (in collaboration with Belgian research institute IMEC), Crossbar and Adesto Technologies (Conductive Bridge RAM, CBRAM) and are commercially available in low volumes for small capacity storage applications in high end applications, such as certain medical devices, wearables and IoT applications. ASX-listed Weebit Nano (WBT) is developing filamentary ReRAM technology using Silicon Oxide (SiOx) as the switching material, i.e. the material in which the filament is formed.

4DS is developing an interface switching ReRAM cell

Another way to manipulate the resistance level of the switching material in a ReRAM cell is based on interface switching. Rather than using materials that stimulate the forming of a filament between the two electrodes when an electrical current is applied, in interface switching memory the resistivity of the entire dielectric layer in the middle is manipulated.

Work that MOHJO!

The heart of 4DS' technology consists of a Metal Oxide Hetero Junction Operation (MOHJO). In semiconductor manufacturing, a hetero junction simply refers to the interface between any two solid-state materials. As illustrated in Figure 2, a material called Praseodymium Calcium Manganese Oxide (PCMO) is deposited onto the bottom electrode made of an inert metal. PCMO, which is a so-called Perovskite material, has certain properties making it highly suitable for this specific type of application. An ion-conducting insulator and a reactive metal, that can partially oxidize, are subsequently deposited onto the PCMO layer, followed by the top electrode, which is also an inert metal.

When a positive voltage is applied to the top electrode, oxygen ions move from the top of the PCMO layer, across the ion conducting insulator to the reactive metal. This substantially increases the resistance level of the PCMO layer to which the value of 0 is attributed. Alternatively, when a negative voltage is applied to the top electrode, oxygen ions move in the opposite direction to the top of the PCMO layer, which becomes conductive again. A value of 1 is attributed to this state of the PCMO layer.

In this memory cell architecture the PCMO layer is effectively the cell switch, meaning that if the correct electric voltage is applied, it will switch from a low resistive state (LRS) to a high resistive state (HRS) or vice versa. By attributing a value of 1 to a LRS and 0 to a HRS, for example, the device can be used to store binary information, i.e. zero or one, similar to other forms of RAM.

FIGURE 2: DEVICE STRUCTURE OF INTERFACE SWITCHING RERAM

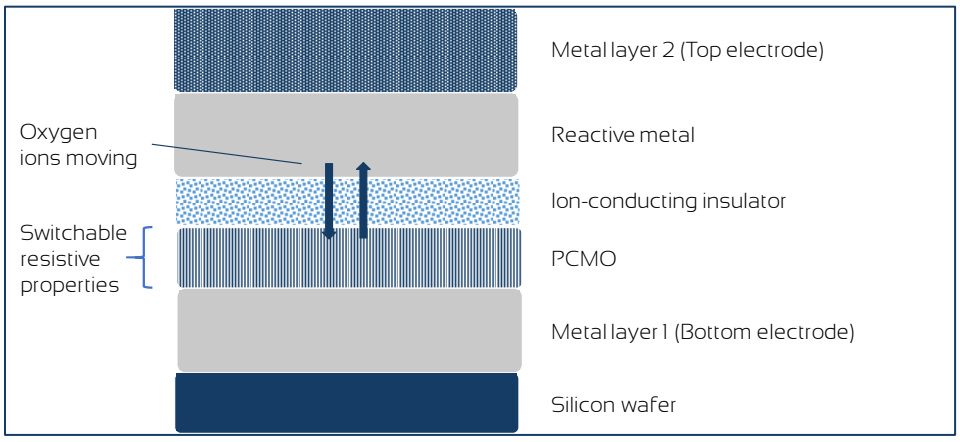

Source: US Patent # 8,378,345 B2, 4DS, TMT Analytics

Filamentary versus Interface Switching solutions

One of the key challenges for the filamentary approach to ReRAM is the scaling problem. Memory cells will continue to become ever smaller on the back of technological advancements in semiconductor manufacturing over time. However, the electrical currents required to read and write filamentary ReRAM cells, i.e. to change the resistivity level, will stay roughly the same because switching current is not related to the size of the cell or to the materials used.

However, there is a maximum current that can be applied to any given semiconductor circuit. This so-called current density limit (maximum electrical current per square nanometer) essentially dictates how narrow the lines in a chip can become. If the same amount of electrical current needs to travel through a thinner wire, this higher current density will likely lead to metal migration in the wire and eventually in wire breakdown.

More importantly, every memory cell needs an access device to select the proper cell to write to or read from and the access device needs to be a certain size to carry the filament switching current. The actual size that determines the memory density is the largest of (1) the memory cell size, (2) the minimum wire width to carry the current, and (3) the minimum transistor size needed to control the current.

Roadmap to commercialization

Scalability to higher resolutions is key

The name of the game in semiconductor manufacturing is scalability, i.e. being able to reduce feature sizes (improve the resolution) while improving functionality with each new version of a computer chip and maintaining the performance characteristics achieved in earlier iterations.

Technical validation process similar for all storage technologies

In order to validate a particular storage technology, all memory developers follow a very specific development process. During this multi-year process, memory cells are scaled down towards a certain target resolution and then characterized for cycling endurance, access speed and data retention. The cell architecture and manufacturing process is then further optimized to achieve the targeted goals for these metrics.

Memory device performance is measured on the following key metrics:

Cycling endurance: The number of times a memory cell can be switched from a LRS (a value of 1) to a HRS (value of 0) and back is called cycling endurance.

4DS' Class 1 performance shares, expiring on 31 December 2018, are based on the company achieving a minimum number of switching cycles (from 0 to 1 and back).

Data retention: The amount of time a memory cell can stay in a LRS or HRS, and thus represent a value of 1 or 0, is an important metric for memory manufacturers. Data retention will typically improve if higher energy levels are used to switch the cell. However, this may result in the cell degrading faster. So a balance between the two needs to be found, again depending on the specific end application.

Access speed: The speed at which data can be written and read is called access speed. This needs to be sufficiently fast to make the technology suitable for various applications.

Step-by-step optimization process

After memory developers optimize each metric at a given resolution, they then need to revalidate all previous metrics at that resolution, as illustrated in Figure 3. For instance, when a change is made to optimize access speed at 50nm, the developers have to access whether cell behavior, endurance and data retention are the same or better than before the change to optimize the access speed was made. Only when all these metrics are validated at a particular resolution, does the process move on the next scaling step (smaller linewidths, higher resolution).

FIGURE 3: SILICON STORAGE TECHNOLOGY VALIDATION CYCLE

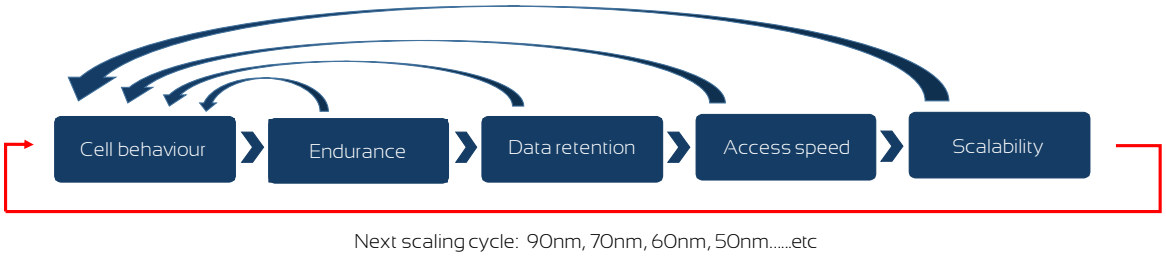

Source: Company, TMT Analytics

Statistical yield analysis to demonstrate overall performance consistency

When working to further shrink the memory cell size, it is important to demonstrate that the metrics above are stable or better compared to the previous iteration of the cell. It is equally important to demonstrate this for a statistically significant number of memory cells, across different wafers and in different batches (groups of wafers being processed simultaneously). Statistical yield analysis of the collected data is then used to derive memory chip performance levels and functional specifications that can be guaranteed to licensees of the technology at that resolution.

4DS has progressed down to 50nm resolutions, moving towards 45nm and below

4DS has been working on its Interface Switching ReRAM technology for the last ten years and has been granted a range of US patents around ReRAM cell structures and the associated manufacturing processes. While the company's initial experimental memory devices measured 1,000nm in resolution (1 micron), ten years of R&D have resulted in memory cells featuring a resolution of 50nm today. When 4DS commenced the JDA with HGST in July 2014, the company had 800nm cells and made significant progress by demonstrating scalability to 50nm within 19 months.

FIGURE 4: EVOLUTION OF 4DS' MEMORY CELL SIZES

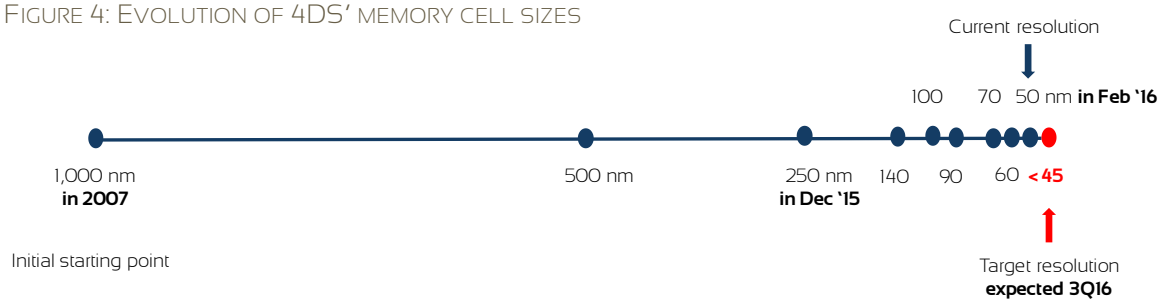

Source: Company, TMT Analytics

Next step: Scale down to 45nm and lower in the near term

For 4DS, the technology validation process described above means further improving the resolution of the memory circuitry down to 45nm and potentially beyond to allow for higher densities when the memory chip will go into volume production in due course. This is done by adjusting cell architecture, the material thickness, tuning the fabrication process and refining the read/write controls. At the same time, device performance needs to be maintained, which will be key in commercializing the technology.

> *While the sub-45nm resolution will likely need to be further improved upon before 4DS' memory cell will move into commercial production, we believe a resolution between 40nm and 45nm should be sufficient to demonstrate the technical and commercial viability of the technology for Storage Class Memory applications to potential licensees.*

Focus has shifted to improving performance and durability

While 4DS' development efforts, in collaboration with HGST, were initially focussed on manufacturing consistency at higher resolutions, this focus shifted to characterizing cycling endurance and access speed. Two years ago 4DS partnered with HGST, a subsidiary of Western Digital, to expedite this development process.

Joint Development Agreement with HGST has been instrumental

HGST is a provider of hard disk drives (HDD), solid-state drives (SSD), and enterprise optimization software. The company is a subsidiary of Western Digital (WD), one of the world's key players in data storage devices. WD achieved revenues of US$14.6BN in FY15. This was down 4% year-on-year due to an 8% decline in shipments of HDD's. Revenues from SSD's, on the other hand, grew by nearly 60%, hence the company's interest to invest in future SSD storage technologies.

Access to HGST's expertise has limited 4DS' cash burn

We believe 4DS' Joint Development Agreement (JDA) with HGST, which was recently renewed and entered its third year, is focussed on cell scaling (i.e. improving the resolution/shrinking the memory cells) and characterization of cycling endurance and data retention. In JDA's such as these it is customary for both parties to share resources and technical expertise while each party carries their own costs and each party retains ownership of its background intellectual property (IP). 4DS wholly owns all of the background IP. In the case of HGST, we believe 4DS is taking full advantage of HGST's expertise, which has enabled 4DS to operate at a cash burn rate of around A$600k per quarter, which is very low for a development company in the semiconductor space at these high resolutions.

Given 4DS' limited financial resources, we believe the JDA with HGST in the past two years has been instrumental in driving the technology to where it is today.

Under the terms of the JDA, HGST has the option to purchase a non-exclusive license to use 4DS's technology for up to 20 years. Additionally, HGST has a right of notification of any third party acquisition or financing proposal.

> *Assuming that a third-party proposal to acquire 4DS would only materialize once the company's Interface Switching ReRAM technology has been proven to be scalable and commercially viable, i.e. with HGST's assistance under the JDA, we expect HGST will want to pre-empt any such third party approach.*

Furthermore, during the recent Flash Memory Summit held in Santa Clara, California, Western Digital publicly endorsed 3D ReRAM as their technology of choice for Next-Gen Storage Class Memory. While WD didn't specify whether they meant Filamentary or Interface Switching ReRAM, HGST's involvement with 4DS makes a strategic tie-up with 4DS for Interface Switching ReRAM very logical, in our view.

Beyond 2016: Commercial deal or strategic action most likely

From the current state of technological development, we essentially see three scenarios for 4DS going forward; firstly, a scenario in which the company succeeds in finalizing the minimum required development of its technology to the point where prospective licensees step up to the plate. In this scenario 4DS will need to demonstrate scalability of the Interface Switching ReRAM memory cell down to sub-45nm resolutions with consistent cell behavior using a scalable manufacturing process.

Still several years of development required by licensees

Any licensee will need to further develop the technology in the next several years to the point where high density, Interface Switching ReRAM memory chips can be manufactured in existing fabs, which we would expect around 2019-2020. This development process will require tens of millions of dollars, possibly up to US$100M, in our view, which is why 4DS will likely not embark on this journey, at least not by itself.

Any license agreement that is non-exclusive would see 4DS receiving multiple up-front, one-off license payments, which could amount to several millions of dollars each. Additionally, 4DS would receive royalties per chip sold once the technology goes into commercial production. Memory chip royalties typically amount to a single digit percentage of the sales price of the chip.

An exclusive license to use 4DS' technology would likely require the licensee to pay a substantially larger up-front license fee, potentially several tens of millions of dollars, in our view, in addition to royalties once the chip goes into commercial production.

Potential strategic action ahead of any license deal

The second scenario would be an acquisition by an established manufacturer in the data storage space. In this scenario as well, we believe 4DS would first need to complete the minimum required development of its technology in order to prove commercial viability of Interface Switching ReRAM. Any acquirer will then need to further develop the technology to commercial insertion into the market, similar to the first scenario.

Given 4DS' development agreement with HGST and the acquisitive nature of its parent company, we believe Western Digital would be a very likely acquirer in this scenario. Additionally, companies like SK Hynix, Micron and Samsung might have a keen interest in securing Interface Switching Memory technology to gradually take over from 3D NAND Flash in a few years' time.

The third and final scenario would see 4DS' technology not being rolled out commercially in due time, either due to insurmountable and/or overly expensive technological issues with Interface Switching ReRAM, or the emergence/dominance of another non-volatile memory technology, e.g. filamentary ReRAM, MRAM etc. While the likelihood of this scenario seems relatively small at this stage given the rapid technological progress being made together with HGST, it can never be excluded in the dynamic semiconductor industry. As history of this

sector has shown, many promising technologies have never made it to the finish line for a variety of reasons.

> *However, if 4DS is able to execute on its technology roadmap, we believe a commercial deal and strategic action are the most likely scenarios.*

Commercial/strategic outcome likely to be very different from R&D endpoint

The above implies that the endpoints for commercial roll out of 4DS' technology into the market and return on investment for shareholders are two very different things. If technological progress follows the development roadmap, it will still take a number of years for Interface Switching ReRAM to hit the market commercially. However, an IP license agreement or an all-out acquisition of 4DS are likely to happen well before then, if the company hits its development milestones.

We see four main risks for 4DS

Technical hurdles proving too difficult or too expensive to overcome

First and foremost, there is a risk that 4DS will not be able to develop its Interface Switching ReRAM technology to the point where it becomes a technically and commercially viable alternative to existing and other emerging storage technologies. For instance, technological hurdles may prove to be too high, resulting in limited scalability, lower than expected cell endurance and slower than planned read and write speeds. Alternatively, the company may be able to overcome such hurdles through more expensive manufacturing processes, rendering the end product too expensive to effectively compete in the storage market.

Extension of current technology prolonging lifespan of NAND Flash

The introduction of 3D NAND, such as Samsung's VNAND and 3D XPoint by Intel and Micron, has prolonged the lifespan of NAND Flash by a number of years. In addition, MLC and TLC have increased memory cell capacity by enabling storage of more bits per cell. It is not beyond the realm of possibilities that the industry will develop additional techniques that will further prolong the economic life of NAND Flash, pushing out the need for alternative technologies by a few more years.

Other emerging technologies

In addition to Interface Switching ReRAM, several alternative emerging technologies are being developed, including Filamentary ReRAM, MRAM, PCM, STT-RAM, FeRAM, and CeRAM etc, some of which should be considered viable technologies. We would expect at least one or two of these technologies to be developed to the point where they may be commercially introduced to the market. And while the sweet spot of those technologies may be different from 4DS' SCM sweet spot, they would still pose a potential threat to Interface Switching ReRAM technology.

Termination of Joint Development Agreement with HGST

Currently in its third year, the JDA with HGST has proven to be extremely valuable for 4DS in driving technological progress. Without this JDA, we believe 4DS would not be where it is today, i.e. having demonstrable Interface Switching ReRAM cells at 50nm resolution. Should the JDA with HGST be terminated for whatever reason, we expect the speed of development

to slow down significantly, potentially resulting in Interface Switching ReRAM not coming to market at all. However, given the unique IP and advanced development progress of the technology, there is a possibility that 4DS could find another partner.

Traditional semiconductor IP licensing model may not apply

Once 4DS' Interface Switching ReRAM technology is fully developed, the company will not be manufacturing memory chips itself, but could aim to license the technology to existing manufacturers of memory chips that have their own multi-billion dollar fabs, such as Micron and Samsung. When it comes to licensing semiconductor IP, typical revenue models include a one-off license fee for use of the technology and recurring royalties based on the volumes of chips sold by the licensee, i.e. a percentage of the chips' sales prices. Additionally, 4DS may be able to charge non-recurring engineering fees, related to integrating its technology into customers' existing chipsets.

However, we believe it may not come to this for 4DS' technology.

Semiconductor industry has a history of IP acquisition rather than licensing

There are several factors that lead us to believe that 4DS' technology will be acquired rather than licensed. Firstly, the number of potential Interface Switching ReRAM licensees/customers is very limited, i.e. between five and eight, including the companies in Figure 14 and companies such as Apple that consume large amounts of non-volatile memory.

All of these companies have development capabilities of their own and most are working on next-gen technologies, targeted for succession of NAND Flash. We believe it is highly unlikely that if one of these prospects were to license 4DS' technology, they would allow competing memory manufacturers to license the technology as well, leaving open the possibility that 4DS' Interface Switching ReRAM technology would be competing directly with them.

Additionally, because 4DS' technology will most likely be tailored to the high capacity storage market, i.e. for data centers and cloud storage, we don't anticipate licensing agreements for low volume applications and embedded memory by smaller manufacturers.

> *In other words, we expect any license agreement 4DS may sign to be an exclusive one with a large player, which would be a de facto acquisition of the company.*

4DS' valuation is a three-stage rocket

Assuming successful completion of product development into a viable memory module, commercial roll out of storage products based on 4DS' Interface Switching ReRAM will still be at least three years out. If we were to model 4DS' revenues from 2019 onwards, this would lead to a highly uncertain financial model, i.e. having to make assumptions around how many companies would license the IP, at which royalty percentages, the commercial selling price per GB and which volumes licensees could sell beyond 2020 etc. More importantly, we believe a scenario with multiple licensees is not the most obvious one.

For these reasons we have looked at industry valuations under a scenario which we believe is more realistic, i.e. an acquisition of the company if and when it succeeds in developing Interface Switching ReRAM technology to the point where it is considered viable for GB silicon storage.

Western Digital on an aggressive expansion strategy

sTEC Inc. acquired by HGST in 2013

sTEC provides enterprise-grade SSD's, flash cards and flash modules for use in high-performance enterprise applications and high-density DRAM modules for networking, communications and industrial applications, amongst other products. The company was acquired by 4DS' joint development partner HGST in 2013 for US$ 321M. Corrected for US$ 116M net cash, sTEC's Enterprise Value amounted to US$ 205M. Trailing twelve-month revenues just prior to acquisition amounted to US$ 122M, implying an Enterprise Value/Sales multiple of 1.7x.

Virident also acquired by HGST

Also in 2013, Western Digital's HGST subsidiary acquired Virident Inc., a provider of PCIe-based enterprise flash storage, controllers and SSD solutions for database, Web 2.0 and virtual desktop infrastructure environments. Clearly Virident is highly geared towards cloud storage. WD paid US$ 685M in cash for Virident (EV of US$ 645M) on reported revenues of US$ 20.7M, implying a EV/Sales multiple of 30.7x.

Seagate acquired LSI's Flash business from Avago

Seagate acquired LSI Logic's Accelerated Solutions Division (ASD) and Flash Components Division (FCD) from Avago for US$ 450M in 2014. Avago earlier acquired the whole of LSI Logic. Seagate was particularly interested in LSI's enterprise PCIe Flash and SSD controller capabilities. The ASD and FCD units combined had revenues of approximately US$ 150M, implying a P/Sales multiple of 3x.

Other acquisitions in the SSD space include the US$ 57.5M acquisition of Shannon Systems by Silicon Motion, the US$ 645M acquisition of storage software player Dot Hill by Seagate in 2015 and the acquisition of Proximal Data (caching software) by Samsung in 2014, which focused on software-based efficiency gains in SSD storage.

In our view, WD is currently the most aggressive player when it comes to expansion into the SSD market through acquisitions. Players such as Samsung and Micron, already established in the SSD market, have taken the autonomous growth route, i.e. through proprietary R&D. All of the acquisitions mentioned above involved established companies addressing the existing NAND Flash market, underscoring the appetite of the industry to expand into the SSD segment. Acquisition prices for the smaller SSD companies roughly ranged from US$ 300M (A$ 400M) to US$ 650M (A$ 870M).

All of these companies generated revenues when they were acquired, making a comparison with 4DS based on revenues and earnings impossible. However, 4DS is addressing a next generation storage technology, rather than the existing NAND Flash market, which we believe puts a high strategic value on the company for potential acquirers.

HDD giant Western Digital has become SSD challenger after SanDisk acquisition

Through the acquisition of SanDisk, which closed in May 2016, Western Digital (WD) has set its sights on the SSD market. The company is currently the world's largest player in HDD with a 44% market share. However, driven by the decline in HDD's in favor of SSD's, the company has strategically moved into the SSD market, in which it had a small presence through HGST.

The fact that WD was willing to pay US$16BN to acquire SanDisk illustrates WD's commitment to the SSD storage market, in our view, which bodes well for 4DS' joint development of Interface Switching ReRAM with HGST, HGST having been a fully owned subsidiary of WD since 2012.

While all of the market participants in Figure 14 could be potential licensees of 4DS' technology, we believe WD, through the HGST connection, is the most likely candidate to be the first licensee of the technology or, alternatively, is best positioned to acquire 4DS.

ASX-listed memory companies imply ~100% upside to 4DS' share price

Despite 4DS' rapid technological development in recent years and its highly valuable JDA with HGST, the company is valued well below ASX-listed sector companies Weebit Nano and Strategic Elements (Figure 5). Given the aggressive roadmap of 4DS' JDA partner and its parent company WD, we believe 4DS is actually in a better position than these companies to monetize its technology in the medium term, e.g. through an exclusive license or a full-blown acquisition of the company. This assumes 4DS' development roadmap in the next twelve months will be met.

FIGURE 5: PRE-REVENUE ASX-LISTED SEMICONDUCTOR COMPANY VALUATIONS

Company	Code	Semiconductor sub sector	Share price	Fully diluted market cap (A$ M)
Weebit Nano	WBT	Filamentary ReRAM	0.039	49.1
Strategic Elements	SOR	Printable memory ink	0.125	30.6
XPED	XPE	IoT communications protocol IP	0.035	78.4
BluGlass	BLG	Semiconductor equipment	0.315	114.6
Brainchip	BRN	Artificial Neural Networks	0.12	102.0
4DS Limited	**4DS**	**Interfacial ReRAM**	**0.024**	**20.0**

Source: FactSet, TMT Analytics

First stage of the valuation rocket is A$ 0.05 per share

On the basis of valuations of other ASX-listed memory companies alone, we believe 4DS should be valued roughly between A$ 31M and A$ 49M, or between A$ 0.037 and A$ 0.059 per share (fully diluted).

Another way to look at this is to take into account total development costs and the intangible benefit of working together with HGST to develop Interface Switching ReRAM. So far 4DS has spent US$ 11M (approximately A$ 15M) to develop its technology, while we estimate HGST has spent at least that amount during the past two years of joint development with 4DS. Adding the current, third, year of the JDA to that puts the cash valuation in the same valuation range.

However, this does not take into account the high probability that 4DS and HGST will continue their JDA in the foreseeable future or the intangible value of the expedited development track 4DS can now follow.

The above valuation range implies roughly a doubling of the current share price if we take the mid-point of that range, which translates into A$ 0.05 per share (fully diluted).

Rerating in second stage based on demonstration of sub-45nm resolution

As some of the other non-memory ASX-listed companies in Figure 5 have illustrated, valuations can expand quite quickly if companies are able to move closer to commercialization, e.g. through securing license or technology evaluation agreements. In the case of 4DS, we would expect such an event to be preceded by announcements regarding progress in scaling below 50nm.

> *In our view, such an announcement in itself should result in a substantial rerating of the company as we consider this the last remaining, significant, hurdle in the initial validation of the technology for GB storage solutions.*

As illustrated by the valuations of other pre-revenue semiconductor companies listed on the ASX, we believe 4DS' valuation has very substantial upside if and when that hurdle is taken. In other words, the single most important near term share price catalyst for 4DS, in our view, is confirmation of scalability to 45nm and below.

In this second stage of the company's valuation, we would expect 4DS' valuation to expand towards levels that are realized by other pre-revenue semiconductor companies, i.e. around A$ 100M and higher. This would imply a value of at least A$ 0.12 per share.

Third stage of the valuation rocket takes long-term strategic value into account

Even though more development work will need to be done post GB storage validation, we believe 4DS taking the sub-45nm scaling hurdle will put the company on the M&A short list of established storage players. If and when 4DS engages in meaningful strategic discussions with potential acquirers, we would expect the company to be able to incorporate the long-term strategic value of its Next-Generation SSD storage technology into the eventual acquisition price.

While it would be premature to attach a price tag to 4DS in such a scenario, given the many uncertainties surrounding it and the speculative nature of pre-revenue technology companies, recent acquisition prices have shown what the consolidating semiconductor industry is willing to pay for high valued assets and technologies.

Initiating coverage with a BUY rating and A$ 0.05 price target

Near term price target of A$ 0.05 based on memory companies' valuation

While our initial, near term, price target of A$ 0.05 per share is based on the valuation of other ASX-listed memory companies, we believe there should be substantial upside towards A$ 0.12 per share if and when 4DS can take the sub-45nm scaling hurdle, i.e. when the second stage of the valuation rocket kicks in.

Based on its technological progress to date and our expectation that 4DS will be able to scale cell resolutions to 45nm and potentially lower before year-end 2016, we start coverage of 4DS with a BUY rating.

Storage technology: A walk down Memory Lane

Computer memory chips come in many forms and are used in many different devices, including PC's, mobile phones, tablets, cars, data centers, kitchen appliances, satellites in space, consumer electronics in general, robots and a whole range of other devices. Basically, any electronic device that is required to store information for a certain period of time will require some sort of memory device. In many cases, electronic devices require more than one type of memory.

Volatile versus non-volatile memory

A desktop computer, for instance, has two main types of memory. Firstly, working memory or system memory that is used for temporary storage of information, such as a Word document that the user is working on, the Word application itself when it is opened by the user and information about the processes being run on the PC when its operating. When the user saves the Word document to the computer's permanent memory (see below) and closes MS Word, all this temporary data will be lost. Also, when the PC loses power for some reason, all temporary information will be lost as well. This type of memory is referred to as volatile memory. It requires a continuous electric power to retain the stored information.

The second type of memory stores information, such as the user's Word document, more permanently. This type of memory will retain stored information, even when the power is turned off, unlike temporary memory. This permanent type of memory, or storage, is called non-volatile memory as it doesn't require a steady electric power to retain its data.

Both types of memory are based on the storage of an electrical charge in the memory cell, representing a value of 0 or 1.

Volatile memory

DRAM (Dynamic Random Access Memory) and SRAM (Static Random Access Memory) are the two main forms of volatile memory. DRAM is widely used in PC's, laptops etc and the electrical charges in the memory cells need to be regularly refreshed, e.g. every 64 milliseconds for a common DRAM chip.

Even though SRAM does not need to be refreshed like DRAM, it's still considered volatile memory, as the stored data will be lost when power is switched off. SRAM is faster than DRAM but also more expensive because SRAM can't be manufactured in the same densities as DRAM. SRAM is used in automotive electronics, medical devices and industrial applications as well as in a computer's Central Processing Unit (CPU) cache memory to store frequently used data to enable faster operating speeds.

Non-volatile memory

Memory types that retain their stored data, when the electrical power is switched off are categorized as non-volatile memory. One example is Read-only Memory (ROM), which is typically used in situations where stored data only needs to be read from the memory device and never written to it. ROM's were invented more than sixty years ago and had one main drawback, which was that they could only be programmed once.

Through a number of iterations between the 1950's and mid 1980's, ROM evolved into the EEPROM, or Electrically Erasable Programmable ROM. As the name suggests, EEPROM's can be programmed and erased multiple times. For instance, a Wi-Fi router will have an embedded ROM chip that contains the factory-programmed firmware that controls and monitors the device. Once in a while, e.g. once a year, a firmware update may become available that can be downloaded by the user and is automatically programmed (written) into the EEPROM.

This cycle of Programming and Erasing is what's known as a P/E cycle, or write cycle, and is an important metric in memory chips.

Modern day EEPROM's can handle between 100,000 and 1 million write cycles but they are very low-density devices.

From the early EEPROM's, the semiconductor industry went on to develop Flash memory which is in heavy demand for data storage applications in mobile phones, laptops, wearables, data centers etc.

Hard Disk Drives (HDD) are another means of permanent data storage, i.e. non-volatile. However, unlike DRAM, SRAM and Flash memory, which comprise of integrated semiconducting circuits (IC's), HDD's are based on magnetic polarization, i.e. data is written to the hard disk by applying magnetic charges to the disk. HDD's are widely used in PCs, laptops, servers, data centers etc, but are increasingly replaced by Flash memory due to better speed, power consumption and heat characteristics.

Solid State Drives gaining ground

Another name for Flash memory is Solid State Drive (SSD). Because of the fast data access, SSD data storage allows for very quick boot-up times for PC's. All data required to start up Windows, for instance, is directly available for use whereas a PC using a Hard Disk Drive, will need to load the operating system and all programs from the Hard Disk Drive upon start-up, which may take several minutes. As a consequence, PC manufacturers are increasingly using SSD instead of HDD's, especially since SSD price points, while still higher per bit than HDD's, are becoming increasingly attractive for mainstream use and are starting to approach HDD price points.

Additional benefits of SSD include lower power consumption, a smaller physical size, lower weight and mechanical shock resistance, which makes SSD storage ideal for use in portable devices.

From memory stick to data center

Flash memory/SSD was initially used in memory sticks/cards and USB flash drives (Figure 6). However, in the last 15 years, as the technology evolved, the benefits became more pronounced and storage capacities increased. Consequently, the use of Flash memory expanded into applications such as PC's, iPods, digital cameras, mobile phones, synthesizers, video games, robotics, and medical electronics. Nowadays even data centers, with very high data storage capacity requirements, are increasingly moving away from hard disk storage in favor of SSD.

FIGURE 6: FLASH-BASED PRODUCTS: MEMORY STICK, FLASH CARD AND A MICRO SD CARD

Source: SanDisk, TMT Analytics

Basic types of Flash memory are NAND Flash and NOR Flash

NAND and NOR Flash are the two main types of Flash memory in use today. Apart from manufacturing differences, a key distinction between NAND Flash and NOR Flash is the way in which the individual memory cells are interconnected. Cells in NOR Flash are connected in parallel, which allows cells to be programmed and read individually. Individual cells in NAND Flash, on the other hand, are connected in series, which lowers energy consumption and saves space.

Consequently, one of the key benefits of NAND Flash compared to NOR Flash is that NAND Flash can be approximately 40% smaller than NOR Flash, which makes it better suited for memory storage in mobile applications, such as tablets, mobile phones etc where chip size is a very important factor (Volume production of single level NAND Flash is below 20nm). Additionally, because of NAND Flash's serial connections, it is substantially faster to erase and write to than NOR Flash.

NOR Flash more reliable

On the other hand, NOR Flash is more reliable when it comes to data retention as NAND Flash suffers much more from so-called bit flipping issues where the value of an individual memory cell is reversed (or reported by the system as having been reversed). Bit flipping can be caused by slow changes in the cell's voltage levels, interference from a neighboring cell as well as outside interference. These reliability issues become more pronounced with smaller cell sizes, as we will elaborate on below. Because of its higher reliability, NOR Flash is commonly used to store the more critical information, such as program code. It is available in capacities up to 256MB.

NAND Flash is the workhorse of high capacity data storage

Because of its advantages with regards to price, speed, size, power consumption and storage capacity, NAND Flash can be considered today's workhorse when it comes to high capacity data storage applications. In addition to the mobile applications and data center storage mentioned earlier, NAND Flash is also increasingly replacing hard disk drives in PC's and laptops. While SSD storage is still more expensive than HDD storage, the price gap is narrowing resulting in increasing affordability beyond just high-end laptops and PC's.

Flash memory is approaching its physical limits

In the manufacturing process of computer chips, lithography is a key process step. While explanation of the entire lithography process is beyond the scope of this report, it is important to describe several key aspects of the technology in order to better understand the hurdles facing Flash memory.

FIGURE 7: A STYLIZED PHOTOMASK, THE LITHOGRAPHY PROCESS AND AN ACTUAL PHOTOMASK

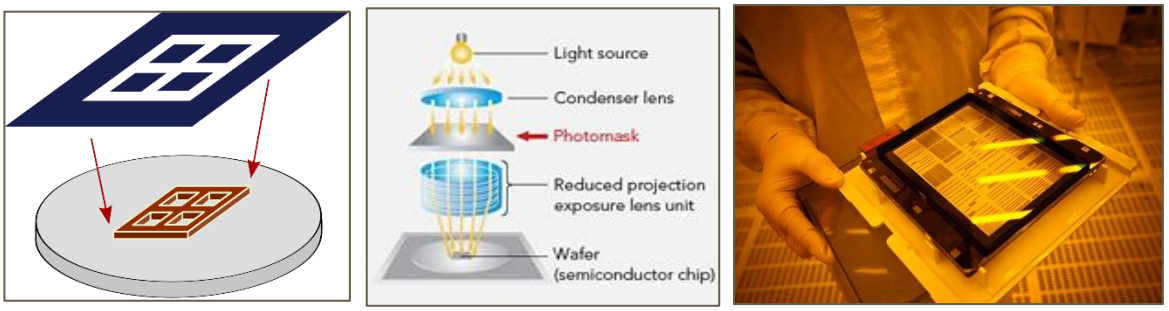

Source: Wikipedia, Hoya Corp., Wired, TMT Analytics

In a nutshell, the lithography process (Figure 7) enables the patterning of the electronic circuitry on the different layers of a chip. In the first step of the process, laser light is flashed onto the chip surface through a so-called photomask that contains all the circuitry of the particular chip layer that is being manufactured. The exposed areas are etched away and filled with a conductive material, such as aluminum or copper. The process is repeated many times to construct and connect the various layers of a chip.

Volume production resolutions

The linewidth of the individual circuits of a computer chip is measured in nanometers (nm). One nanometer is equal to one millionth of a millimeter, 100,000 times smaller than the width of a human hair.

Today's leading edge fabs (semiconductor manufacturing facilities) can manufacture DRAM chips with a linewidth, or resolution just under 30nm. Single level, or 2D, NAND Flash is currently produced at 16nm by companies such as Micron and Samsung.

NAND Flash is approaching physical scaling limits

Facilitated by fundamental physics research, the technological progress made in lithography and adjacent technologies over the years have allowed chip manufacturers to produce semiconductors at ever-higher resolutions (narrower linewidths and smaller cells), resulting in ever-higher storage capacity. The benefits of smaller cells include lower power consumption, the ability to put more functionality on the same chip surface (or the same functionality on a smaller surface), lower prices and faster speeds.

However, the main problem with ever-smaller geometries is cell-to-cell interference as the spacing between two cells reaches critical levels. The cell walls become so thin that electric charges leak to the adjacent cell, potentially corrupting the data stored in that cell. For NAND Flash specifically, this causes a further reduction of the reliability of the memory cells.

Furthermore, as resolutions become smaller and smaller, patterning of the chip circuitry in the lithography process becomes increasingly difficult, and thus more expensive as technical solutions and optical tricks need to be developed to overcome these difficulties.

In addition, due to higher resolutions and smaller chip sizes, the number of P/E cycles for NAND Flash has fallen to approximately 100,000 from nearly 1,000,000 for EEPROM many years ago. Increasingly smaller chip circuitry still needs to endure relatively high voltages, limiting the durability of NAND Flash. In other words, the endurance drops with higher resolution. This is particularly problematic for high-end applications that require many P/E cycles over the lifespan of the chip.

> *These physical scaling limits, where memory cells start to interfere with each other, have become particularly problematic for NAND Flash at resolutions below 20nm.*

Building skyscrapers to further improve cell density

One solution the semiconductor industry has applied to the scaling limits of NAND Flash is to go vertical, i.e. stacking memory cells on top of each other to build so-called 3D NAND Flash memory. Even though the memory cells themselves cannot get much smaller anymore, stacking them creates higher densities and thus lower costs per unit of storage (bit).

Additionally, as memory cells in 3D NAND Flash are physically closer together, similar to apartments in a skyscraper as opposed to freestanding houses in a suburb, interconnections

can be shorter, allowing for lower energy consumption and higher speeds. Furthermore, because resolutions in stacked memory cells don't need to be as high as in regular (2D or planar) NAND Flash cells, 3D NAND Flash can maintain the reliability levels achieved in the past at these lower resolutions, i.e. roughly between 40nm and 50nm for 3D NAND Flash versus ~16nm in 2D NAND Flash cells.

3D NAND Flash has been in commercial production since 2013 when Samsung introduced Vertical NAND, or VNAND (Figure 8). The chip architecture is different from planar NAND Flash and VNAND was initially produced with 24 layers of memory cells providing 128GB of storage. Samsung's current VNAND chips feature 48 layers providing 0.5TB of storage while the company anticipates to be able to offer storage capacity up to 1TB in 2017 by adding layers and further improving the resolution of the chip.

FIGURE 8: BIRD'S EYE VIEW OF VNAND STRUCTURE

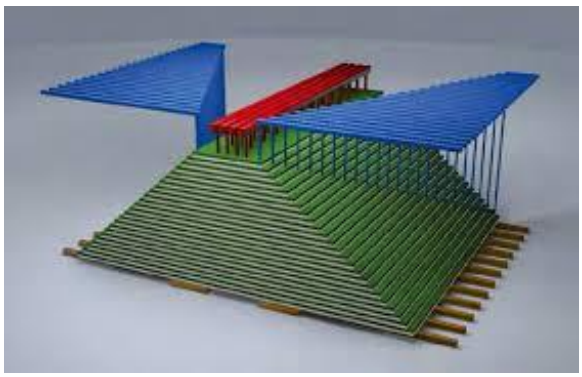

Source: Samsung, TMT Analytics

In 2015 the market for 3D NAND Flash amounted to US$ 4.5BN, a large part of that attributable to Samsung. With players like Micron, SanDisk (recently acquired by Western Digital), SK Hynix and Toshiba going into commercial production in 2016, the 3D NAND Flash market is expected to nearly triple, according to Webfeet research.

Building higher structures has its limits too

While the number of stacks that can be supported in 3D NAND Flash can be quite high, theoretically enabling many years of runway before 3D NAND Flash will reach its limits, practical problems, like proper failure analysis inside the internal layers of a 3D structure, limit the practical number of layers in 3D NAND Flash. Current consensus puts the maximum number of layers in 3D NAND Flash at 96 and potentially 128.

However, there are more ways to increase cell densities in a memory chip.

Putting more data in the same cell is another way to increase density

Traditionally, memory cells had two possible states, 1 and 0, and therefore could contain 1 bit of data. These cells are categorized as Single Level Cell (SLC). The stored charge in Multi-Level Cells (MLC) can be interpreted as a variety of values and can store 2 bits of data in a single cell. MLC Flash is the most common type of Flash used today, especially in consumer electronics, such as mobile phones, camera's, tablets etc.

SLC's have lower power consumption and therefore a longer lifespan compared to MLC, i.e. the number of P/E cycles is up to 10 times higher in an SLC (~100,000 cycles versus ~10,000 for a MLC). However, MLC's are cheaper to manufacture per unit of storage, hence its popularity for consumer applications. Because of higher reliability and faster speeds, SLC

Flash is Industrial/Enterprise grade and can be found in higher-end storage applications, including data center storage.

Triple Level Cells (TLC) can store 3 bits of information and are therefore the cheapest to manufacture per unit of storage. However, with a relatively limited number of P/E cycles (between 1,000 and 5,000) most TLC Flash is best suited to lower-end consumer applications, such as USB sticks, media players etc, not yet high-end storage.

MLC + 3D NAND Flash = very high cell densities

Combining the higher storage capacity of MLC with the skyscraper approach of 3D NAND Flash has resulted in very high capacity SSD storage devices. For instance, Micron manufactures MLC's in 3D, which provides 256Gb of storage. The company's TLC in 3D even provides 384Gb of storage capacity.

FIGURE 9: SEMICONDUCTOR MANUFACTURING ROADMAP

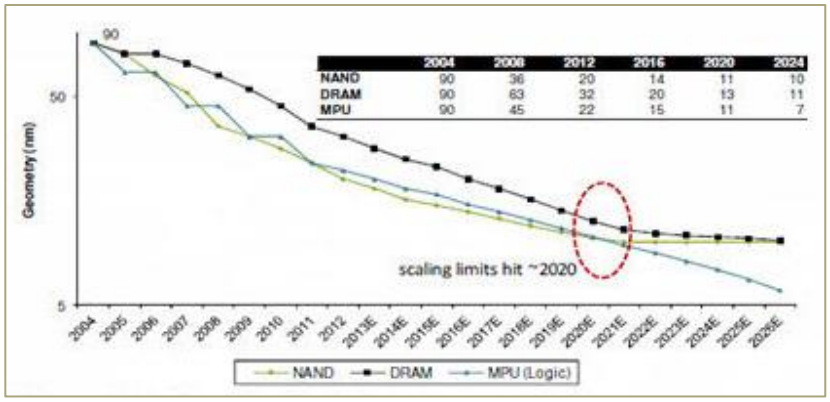

Source: ITRS, SK Hynix, TMT Analytics

In light of the physical scaling limits of 2D NAND Flash we discussed earlier, solutions such as MLC and 3D NAND have extended the life of NAND flash by another 5 to 7 years, i.e. roughly into the 2021-2023 time frame.

> *However, with resolution improvements expected to flat line around 10nm for 2D NAND Flash (see Figure 9) and a limit to the number of layers that can be stacked in 3D NAND Flash, it is expected that cost effective density improvements in NAND Flash will eventually run out of steam.*

Several possible technologies for the post-NAND era

Various technologies that could potentially replace NAND Flash in different applications have emerged in recent years. Specifically, a lot of R&D effort has been dedicated to five technologies that appear to be in contention for the post-NAND era: Magneto resistive RAM, Ferroelectric RAM, spin transfer torque RAM (STT-RAM), phase-change memory (PCM) and Resistive RAM (ReRAM). PCM and ReRAM currently seem best placed to fill the shoes of NAND Flash in certain application areas.

Outsiders Magneto resistive RAM (MRAM) and Ferroelectric RAM (FRAM)

MRAM uses a magnetic polarization to change the electric resistance of a memory cell to represent a value of 0 or 1. One of the main drawbacks of MRAM is that it is hard to reduce cell sizes, and thus to improve cell densities in a single memory chip, as the electric current required to change the magnetic state of the cell remains unchanged, regardless of how small the cell size becomes. If the circuitry of a chip shrinks and the electric current running through these thinner wires remains unchanged, it will result in faster degrading of these wires and thus lower endurance of the memory chip as a whole, which limits the application areas for MRAM.

FRAM uses a ferroelectric layer to achieve a non-volatile state. It uses relatively low power and has a high write speed. However, it is relatively expensive and limited in storage capacity. Furthermore, due to the method in which an FRAM cell is being read, the durability of FRAM cells is lower than for other technologies.

Even though MRAM chips are being sold for niche applications today, both MRAM and FRAM seem to be trailing where R&D for future replacement of NAND Flash in high densities and high volumes is concerned.

Spin transfer torque RAM (STT-RAM)

STT-RAM is a form of MRAM that uses the change of direction of spinning electrons to change the magnetic orientation, and thus the value, of a memory cell. STT-RAM allows for lower power consumption, and thus better scalability, compared to MRAM's. Additionally, STT-RAM has a far superior number of P/E cycles compared to other non-volatile RAM technologies as a result of its specific switching technology, which doesn't degrade the switching materials. Furthermore, data retention can be tuned to be anywhere between 10 and 20 years, making STT-RAM a lead candidate for long-term storage requirements.

Quite some R&D effort has been invested into STT-RAM in the last ten years by companies such as Hitachi, Qualcomm, IBM and Samsung. The latter two are working to commercialize STT-RAM within the next several years.

Good prospects for Phase Change Memory (PCM) as well

The cell switching principal of PCM is based on the various states the glass-like switching material chalcogenide can assume under the influence of heat. This can be either crystalline (low electric resistance) or amorphous when heated (high electric resistance), representing a 1 or 0 respectively.

PCM has clear advantages compared to NAND Flash in the areas of speed and endurance (>10 million P/E cycles). More importantly, PCM holds the promise of being much more scalable to higher resolutions than NAND Flash.

FIGURE 10: PHASE CHANGING IN PCM

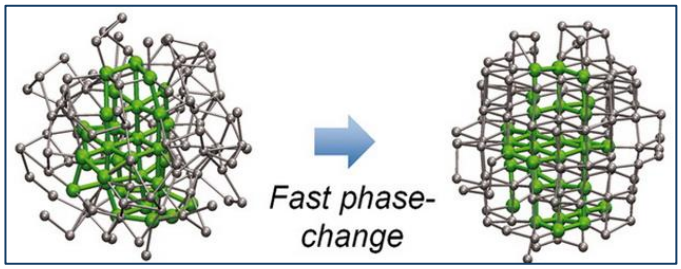

Source: Columbia University, TMT Analytics

Intel/Micron collaboration to launch 3D XPoint late 2016

Early endeavors into PCM by Micron were terminated due to the emergence of 3D NAND and MLC technologies. However, Intel and Micron have teamed up to develop 3D XPoint, which was introduced in 2015 with production anticipated to start ramping later this year. While expert opinions vary as to whether 3D XPoint is actually a form of PCM given Intel's and Micron's secrecy around the manufacturing process, the technology holds a lot of promise, including the speed, endurance and scalability advantages mentioned above. However, these advantages are likely to come at a price premium compared to NAND Flash, at least in the near to medium term.

FIGURE 11: 3D XPOINT ARCHITECTURE

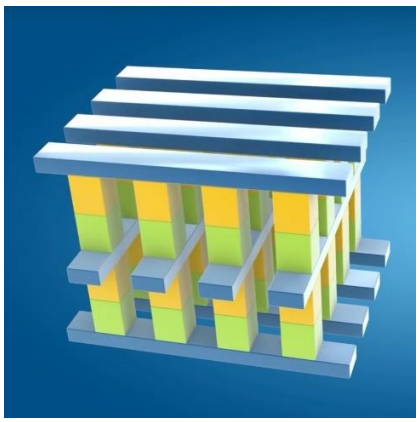

Source: Intel, TMT Analytics

IBM demonstrated lower cost PCM version

A solution to the relatively high cost of PCM compared to NAND Flash was offered by IBM earlier in 2016 when the company unveiled a TLC version of a PCM chip, i.e. a chip with all the advantages of PCM, but with triple the storage capacity offered by a single level PCM cell. While IBM's technology is highly promising, the company will need a manufacturing partner to bring TLC PCM into volume production. Given the upcoming commercial launch of Intel and Micron's 3D XPoint later this year, logical candidates for IBM would be SK Hynix, Samsung and Toshiba. However, even if IBM were able to team up with one of these companies, we would expect volume production of TLC PCM chips to be at least 18 to 24 months away.

> *All in all, though, we believe PCM is one of the two leading candidates to gradually replace NAND Flash in due course, the other one being Resistive Random-Access Memory (ReRAM).*

Addressable markets growing fast

The Storage Class Memory (SCM) market, i.e. cloud storage and data centers, is growing at double digits globally. Moreover, the substitution of Hard Disks by NAND Flash in the cloud and data center market has created an additional driver for SSD growth in this space.

Hard Disk Drives in decline, but sales of NAND Flash showing healthy growth

As illustrated in Figure 12 the market for Hard Disk Drives has been in decline for a number of years. Unit shipments keep falling, in part due to larger storage capacities per unit, but mostly because NAND Flash is increasingly replacing HDD's, both in consumer and enterprise applications such as data centers.

FIGURE 12: GLOBAL QUARTERLY HARD DISK DRIVE SHIPMENTS (IN M)

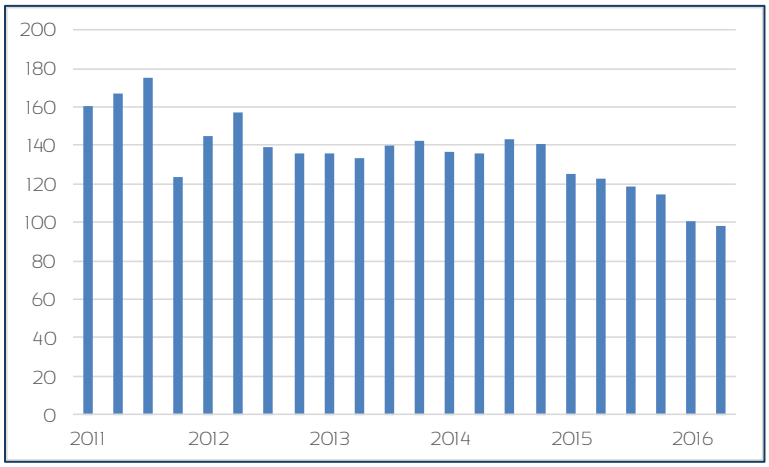

Source: Statista, TMT Analytics

In terms of size, NAND Flash accounts for approximately one third of the overall memory/storage market (excluding HDD) as illustrated in Figure 13. In absolute terms, NAND Flash revenues amounted to US$ 26.8BN in 2015, while growth from 2015 through 2019 is expected to amount to almost 9% CAGR.

With annual NAND Flash bit growth in datacenter applications in excess of 50% in the past few years, this segment currently accounts for about 25% of overall annual NAND Flash consumption. On the back of an apparently insatiable demand for storage capacity, this segment is expected to continue to grow its share of bit consumption in the foreseeable future (Figure 13).

FIGURE 13: MEMORY MARKET IN CY15 (US$83.4BN)

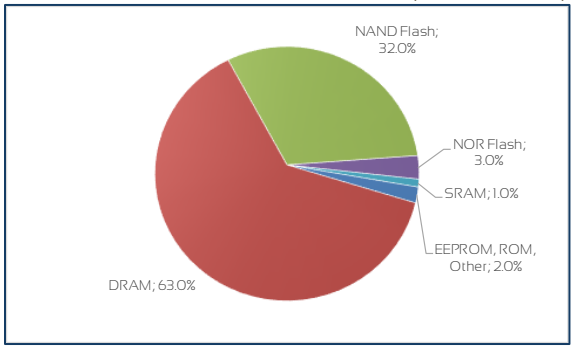

NAND FLASH MEMORY MARKET SIZE (US$BN)

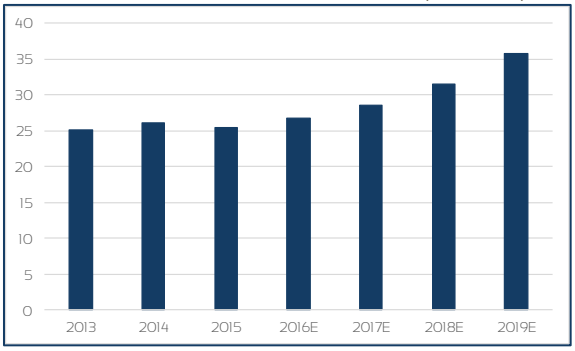

Source: IC Insights, Statista, TMT Analytics

Storage Class Memory is 4DS' sweet spot

This strong demand for storage capacity is driven by the fast growth of online video content, e.g. YouTube and the increasing use of video in Social Media. Additionally, Big Data analytics, the growth of the Internet of Things and generally strong demand for colocation and off-premise managed services are contributing to increased spending on cloud and data center storage capacity.

> *Given the characteristics of 4DS' Interface Switching ReRAM technology, including high endurance, low energy consumption and hence lower operating temperatures as well as the potential to manufacture its memory chips in high cell densities, the most logical addressable market for 4DS' technology is Storage Class Memory, i.e. the data center and cloud storage market, especially given the strong underlying fundamentals of this market.*

This doesn't rule out mobile use of the technology, i.e. in mobile phones, tablets and laptops, but demand drivers clearly favour the cloud and data center storage markets that are not bound by population growth. Rather this segment benefits from high, secular growth in demand for storage capacity.

The NAND Flash market is fairly consolidated

Historically, Samsung has always been one of the key players in NAND Flash memory (Figure 14). However, the company has been consistently losing market share over the past twelve years but seems to be relatively stable now with a share between 30% and 35% since 2014. The SanDisk/Toshiba collaboration has taken the main chunk of market share from Samsung, growing from around 30% market share ten years ago to a range of 35% to 40% in recent years.

FIGURE 14: NAND FLASH MEMORY MARKET SHARE (1Q16)

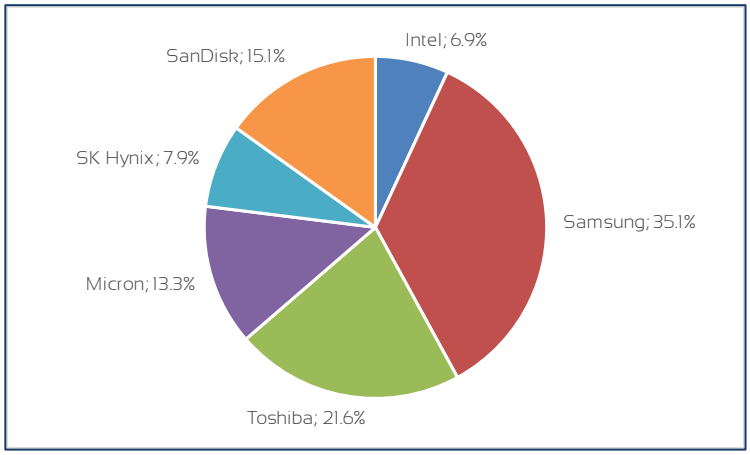

Source: Statista, TMT Analytics

While SK Hynix and Micron have had a fairly stable NAND Flash market share in recent years, Intel is the new kid on the block, growing its share from zero ten years ago to around 7% currently.

Appendices

Board of Directors

James Dorrian (Chairman): Mr. Dorrian is a former partner at Crosspoint Venture Partners, a Silicon Valley based early stage venture capital firm. He has served as both CEO and Board Member of several Silicon Valley companies and has in depth experience in M&A and IPOs. Prior to this, Mr. Dorrian was the Founder and CEO of Arbor Software and has held management roles with a number of multinational IT companies. He is a founding member of the OLAP standards council, an industry consortium for On-Line Analytical Processing. Mr. Dorrian received a Bachelor of Arts from Indiana University in Economics and Communications.

Dr. Guido Arnout (Executive Director and Chief Executive Officer): Dr. Arnout has helped guide multiple Silicon Valley companies through commercialization or sale. He was the founding President & CEO of PowerEscape, which introduced the first tools for the development of low-power software executing on multicore devices. He was also founding President & CEO of CoWare, which pioneered system level design tools for hardware-software co-design and the time-based licensing business model. Dr. Arnout co-founded the Open SystemC Initiative (OSCI), an industry consortium to standardize a language for system level design, and as its President submitted the SystemC language to IEEE. He served as VP of Engineering and later senior VP of marketing of CrossCheck Technology. He co-founded and later became VP of Engineering of Silvar-Lisco, the first commercial EDA (electronic design automation) company. Dr. Arnout received his PhD in electrical engineering from the University of Leuven in Belgium.

David McAuliffe (Non-Executive Director): McAuliffe is an experienced board director and entrepreneur who has had over twenty years' experience, mostly in the international biotechnology field. During that time, he was involved in numerous capital raisings and inlicensing of technologies. He is a founder of several companies in Australia, France and the United Kingdom, many of which have become public companies. Mr. McAuliffe has an Honours degree in Law and a Bachelor of Pharmacy degree.

Howard Digby (Non-Executive Director): Mr. Digby started his career at IBM and has spent over 25 years managing technology related businesses across the Asia Pacific region, including 12 years being based in Hong Kong. Before returning to Perth, he was with The Economist Group as regional managing director. Prior to this he held senior management roles at Adobe and Gartner where his clients included major semiconductor players including Samsung, Hynix and TSMC. Mr. Digby is a Non-Executive Director of Estrella Resources (ASX: ESR) and is currently an advisor to geospatial imagery company Spookfish (ASX: SFI). Mr. Digby has a Bachelor of Engineering (Mech, Hons) from The University of WA.

FIGURE 15: ANNUAL DIRECTOR REMUNERATION

Guido Arnout (exclusive of superannuation)	USD 185,000
James Dorrian	AUD 40,000
David McAuliffe	AUD 30,000
Howard Digby	AUD 30,000

Source: Company, TMT Analytics

FIGURE 16: 4DS CAPITAL STRUCTURE (IN M SHARES)

Ordinary shares	545.8
Escrowed shares	113.4
Performance shares	67.6
Options	106.2
Fully diluted	833.0

Source: Company, TMT Analytics

GENERAL ADVICE WARNING, DISCLAIMER & DISCLOSURES

The information contained herein ("Content") has been prepared and issued by TMT Analytics Pty Ltd ABN 17 611 989 774 ("TMT Analytics"), an Authorised Representative (no: 1242594) of Bellmont Securities ABN 47 119 852 890 AFSL 331625. All intellectual property relating to the Content vests with TMT Analytics unless otherwise noted.

DISCLAIMER

The Content is provided on an as is basis, without warranty (express or implied). Whilst the Content has been prepared with all reasonable care from sources we believe to be reliable, no responsibility or liability shall be accepted by TMT Analytics for any errors or omissions or misstatements howsoever caused. Any opinions, forecasts or recommendations reflect our judgment and assumptions at the date of publication and may change without notice. TMT Analytics will not accept any responsibility for updating any advice, views, opinions or recommendations contained in this document.

No guarantees or warranties regarding accuracy, completeness or fitness for purpose are provided by TMT Analytics, and under no circumstances will any of TMT Analytics, its officers, representatives, associates or agents be liable for any loss or damage, whether direct, incidental or consequential, caused by reliance on or use of the Content.

GENERAL ADVICE WARNING

The Content has been prepared for general information purposes only and is not (and cannot be construed or relied upon as) personal advice nor as an offer to buy/sell/subscribe to any of the financial products mentioned herein. No investment objectives, financial circumstances or needs of any individual have been taken into consideration in the preparation of the Content.

Financial products are complex, entail risk of loss, may rise and fall, and are impacted by a range of market and economic factors, and you should always obtain professional advice to ensure trading or investing in such products is suitable for your circumstances, and ensure you obtain, read and understand any applicable offer document.

DISCLOSURES

TMT Analytics has been commissioned to prepare the Content. From time to time, TMT Analytics' representatives or associates may hold interests, transact or hold directorships in, or perform paid services for, companies mentioned herein. TMT Analytics and its associates, officers, directors and employees, may, from time to time hold securities in the companies referred to herein and may trade in those securities as principal, and in a manner which may be contrary to recommendations mentioned in this document.

TMT Analytics may receive fees from a company referred to in this document, for research services and other financial services or advice we may provide to that company. The analyst has received assistance from the company in preparing this document. The company has provided the analyst with communication with senior management and information on the company and industry. As part of due diligence, the analyst has independently and critically reviewed the assistance and information provided by the company to form the opinions expressed in the report. Diligent care has been taken by the analyst to maintain an honest and fair objectivity in writing this report and making the recommendation. Where TMT Analytics has been commissioned to prepare Content and receives fees for its preparation, please note that NO part of the fee, compensation or employee remuneration paid will either directly or indirectly impact the Content provided.

RECOMMENDATIONS

TMT Analytics' issues a BUY recommendation in case of an expected total shareholder return (TSR, share price appreciation plus dividend yield) in excess of 25% within the next twelve months, an ACCUMULATE recommendation in case of an expected TSR between 5% and 25%, a HOLD recommendation in case of an expected TSR between -5% and +5% within the next twelve months and a SELL recommendation in case of an expected total return lower than -5% within the next twelve months.