A Look at Memory and Storage Trends

Tom Coughlin President Coughlin Associates, Inc. www.tomcoughlin.com

© 2018 Coughlin Associates

Outlook

- Drivers for Memory and Storage
- DRAM and NAND Trends
- China Memory Investments and the Trade War
- Emerging Memory Trends
- Other Storage Trends and Conclusions

Drivers for Storage and Memory

- Increasing storage demands—IDC 163 Zetabytes of data created by 2025 (16 ZB in 2016)
- New sources for unstructured data from media and entertainment, IoT, medicine, geo-science and big data
- Growth in local storage, storage at the edge (or the fog) and storage in large data centers (the cloud)
- There is a need for fast memory and storage to support processing and accessing this data and cheap storage to keep it for the long term



HDD versus SSD Raw Capacity Price Comparison

- This slide is from the WD MAMR HDD announcement in 2017
- The introduction MAMR (or HAMR) HDDs will allow continued cost advantages for projected HDDs versus SSD raw storage costs



Digital Storage and Memory Tiering

2018 Emerging Memories Report, Coughlin Associates, www.tomcoughlin.com

Flash Memory Trends and Projections

-ASC

© 2018 Coughlin Associates

Flash Memory

- Flash memory is increasing in storage capacity (density) and decreasing in \$/GB pricing but still more expensive than HDDs
- Flash memory is winning more applications as its price (\$/GB) drops
- Development of NVMe and NVMe-oF has enabled better access to the performance capabilities of flash memory
- In many data center applications, flash memory is now the primary storage
- Flash Memory can also handle more rugged environments, making this a favored storage media for remote location—such as for edge storage

3D Flash

- 64 layer shipments common today
- Announcements up to 96 layers—1 Tb per die (volume by 2019)
- Technology projections of hundreds of layers (up to 500?)
- Announced quad-level cells for higher density

NAND Flash Expectations

- Flash Memory has moved from primarily planar to planar + 3D flash
- New 3D flash fab has reached parity with planar production in 2018, easing supply constraints
- This has resulted in a drop in flash prices in 2018
- Projections that 3D flash could go out many generations—
- The price reductions for 96 layer and higher will be less than going to 64 layer, because of slower process speeds

NAND Flash Technology Roadmap Update

NAND Shipment Trends

Looking at NAND from the long run THE EB SHIPPED CAGR IS SLOWING, BUT STEADY GROWTH¹ 1,000,000 50 33% CAGR 45 42% CAGR - NAND Sales 40 SB 100,000 Petabytes Shipped 71% CAGR 35 Revenue 10,000 30 25 188% CAGR 1,000 20 Indusry 15 100 10 abytes shipped on a 10 traight line = Same CAB eadate Market Research 1 Over the last 13 years

- NAND revenues are generally increasing (with a particularly large bump between 2016 and 2017). The chart also shows the growth of NAND flash shipping capacity. The chart suggests that the growth of flash memory shipped capacity is slowing and may now be on the order of 33% cumulative annual growth rate with total shipped capacity at about 170-180 Exabytés
- This is from a Seagate slide from the last Flash Memory Summit in August 2017

NAND Flash Memory and DRAM Projections

- Projections of NAND and DRAM shipped capacity in Petabytes (Coughlin Associates, Emerging Memory Report)
- Prices will follow boom and bust trends following fab construction trends
- Big revenue hit from oversupply likely in 2019-2020 time frame

NVMe Flash versus Enterprise Performance HDDs

- Enterprise high performance HDDs (10k and 15k RPM HDDs) have benefited from the higher cost of flash memory in 2017 (prices 2X more than in 2016)
- In 2018 with flash memory prices declining again, the competition will get more intense for enterprise performance HDDs
- In performance applications the market is moving to flash as primary storage
- Even so the decline in high performance HDDs will be more gradual than it could be because of the large number of systems out there using these drives replacements and extensions of current infrastructure will continue for a few more years
- However, by the early 2020's we expect a significant decline in high performance enterprise HDDs, while enterprise capacity hard disk drives will experience significant growth

Solid State Memory Investments (With a Focus on China)

Memory is Driving Semiconductor Equipment and Fab Spending

- According to SEMI, the semiconductor industry is nearing a third consecutive year of record equipment spending with projected growth of 14 percent (YOY) in 2018 and 9 percent in 2019, this would extend the streak to a historic fourth consecutive growth year.
- Only once before, in the mid-1990's has the industry had four consecutive years of capital equipment spending growth.
- The highest growth in capital spending is in Korea and China. Samsung dominates semiconductor capital spending and China is surging ahead of other markets, as shown in the next slide. This chart from June 2018 includes new and refurbished equipment.
- Samsung made a significant investment in 2017 and according to a recent report (DigiTimes, 7/13/18), is increasing their 2019 equipment investment to about \$9 B (40% higher than in 2018)
- SK Hynix is also increasing its Korean capital spending.

2017 End of Year Equipment Forecast, SEMI

 Note that Korean investment in 2017 and 2018 exceeds that of China

Equipment Spending by Region

Chinese Semiconductor Fab Spending (Semi+)

- China's equipment spending is forecast to increase 65 percent in 2018 and 57 percent in 2019.
- 58 percent of investments in China in 2018 and 56 percent in 2019 are by companies with headquarters in other regions such as Intel, SK Hynix, TSMC, Samsung, and GLOBALFOUNDRIES.
- Samsung will expand its capex for NAND flash memory to about \$9 B in 2019, up from about \$6.4 B in 2018, spending mainly on 3D NAND chip output in Pyeongtaek, South Korea and Xian, China (DigiTimes, 7/13/18)
- In 2018, Chinese-owned companies are expected to invest about US \$5.8 billion, while non-Chinese will invest US \$6.7 billion.
- Many new companies such as Yangtze River Storage Technology, Fujian Jin Hua, Hua Li, and Hefei Chang Xin Memory are investing heavily in the region. These companies are expected to double their equipment investments again in 2019.

Chinese Semiconductor Fab Spending (Semi)

- Construction spending (not capex) will reach all-time highs with China construction spending taking the lead at US\$6 billion in 2017 and US\$6.6 billion in 2018, establishing another record: no region has ever spent more than US\$6 billion in a single year for construction.
- In March 2018 SEMI said that 3D NAND will lead overall product sector spending, growing 3 percent each in 2018 and 2019, to US\$16 billion and US\$17 billion, respectively. DRAM will see robust growth of 26 percent in 2018, to US\$14 billion, but is expected to decline 14 percent to US\$12 billion in 2019. Foundries will increase equipment spending by 2 percent to US\$17 billion in 2018 and by 26 percent to US\$22 billion in 2019, primarily to support 7nm investments and ramp of new capacity.
- At the 2018 SEMICON, this last week, I heard some talk that capital equipment projections for 2019 are softening, except probably in China (and Korea)
- In the next slide is a SEMI chart from March 2018 on semiconductor fab construction, showing the growth in Chinese based fabs in 2017.

Number of Volume Fabs Starting Construction

(All probabilities, including discretes)

Source: World Fab Forecast report, Feb. 28, 2018, SEMI

China's Push for IC Production

- In 2017, China consumed \$138 billion, or 38%, of the world's chips, according to IC Insights. IC production in China reached \$18.5 billion in 2017, equating to 13.3% of the world's production, according to the firm.
- In total, the nation is currently building 19 new fabs, with 10 of those projects being 300mm plants, according to SEMI analysts Dan Tracy and Clark Tseng. The numbers include both domestic and multinational chipmakers.
- Note the commitments of up to \$150 B in memory fabs in China comes in stages and will be spread out over many years
- "We expect wafer fab equipment investment in China being up in 2018 by about \$2 billion compared to 2017," said Arthur Sherman, vice president of marketing and business development at Applied Materials

Chinese IC Demand and Production

Worldwide Semiconductor Capital Spending Trends (2000-2022F)

© 2018 Coughlin Associates

Non-Chinese Company China Fabs

- SK Hynix has the biggest memory fab in China, a DRAM operation in Wuxi that is expected to sell \$8.4 B in memories in 2022
- Samsung has expanded its 3D NAND production in China which is estimated to reach sales of \$6.1 B in 2022
- SMIC is expected to reach \$4.6 B of sales in 2022
- SMC is expected to hit \$1.8 B of sales in 2022

China's Memory and IC Market

2017 Bank	Company	2012 IC Sales (SM)	2013 IC Sales (SM)	2014 IC Sales (SM)	2015 IC Sales (SM)	2016 IC Sales (\$M)	2017 IC Sales (\$M)		Products	2022 IC Sales (SM_Ecst)
Kank	Company SK Hunix*	2 360	3 200	4 040	4 100	3 680	6.480	DRAM		8 400
	SK Hynix	2,500	0,200	270	2 370	2 725	3,800	30	NAND Elach	6,100
2	Samsung	1 5 4 2	1 062	1.070	2,070	2 014	3,000	30	Foundar	0,100
3	SMIC	1,542	1,902	1,970	2,230	2,914	3,100	Foundry		4,600
4	Intel*	2,380	2,650	2,710	1,830	520	1,500	3D NAND Flash		2,700
5	Hua Hong Semi	571	585	665	650	721	807	Foundry		1,100
6	TSMC*	470	510	550	590	635	700	Foundry		1,700
7	Shanghai Huali	70	210	255	305	460	550	Foundry		1,200
8	XMC	160	150	165	175	205	250	Foundry/3D NAND		D 1,800
9	CR Micro	179	165	180	190	205	220	Foundry/Std ICs		300
10	Diodes-BCD	139	155	170	180	200	220	Foundry/Std ICs		350
11	ASMC	135	117	130	119	120	150	Foundry		250
_	HeJian**	215	0	0	0	0	0	Foundry		0
	Others	590	575	590	605	625	690		- /	5,100
_	Total Chinese IC Production	\$8,811	\$10,279	\$11,695	\$13,350	\$13,010	\$18,467		- (\$33,600
_	WW IC Market (\$B)	\$259.3	\$271.9	\$291.6	\$286.9	\$297.7	\$363.6		_	\$466.8
_	Chinese Companies' Share of WW IC Market	3.40%	3.78%	4.01%	4.65%	4.37%	5.08%		_ /	7.20%

Will China Succeed in Memory?, Semiconductor Engineering, Feb. 19, 2018

*China fab production only. **Merged with UMC beginning in 2013.

Source: IC Insights' Strategic Reviews database,

CCID, CSIA, PwC

Made in China 2025

- China has made chip development a key element of its Made in China 2025 plan.
- The Chinese government wants local chips to make up at least 40% of china's semiconductor needs by the middle of the next decade
- These efforts have been hindered by the blocking of a series of chip-related purchases in the US:
 - A Chinese deal for the U.S. semiconductor testing company Xcerra Corp was shot down by a U.S. national security panel in February, while the \$1.3 billion acquisition of the U.S. chipmaker Lattice Semiconductor Corp was blocked last year.
- It has also been difficult for China to attract top end talent and overcome some technical hurdles in making memory technology
- It may be difficult for China to make advanced memory products for use outside of China unless it creates or acquires IP currently owned by the major NAND and DRAM manufacturers

History of Chinese Memory Fab Investment

- China's most notable efforts in memory started in 2006, with the emergence of Wuhan Xinxin Semiconductor Manufacturing Corp. (XMC). Wuhan-based XMC is a NOR flash foundry vendor. In addition, XMC has been developing 3D NAND as part of an alliance with Spansion, now part of Cypress.
- Also in 2006 SK Hynix build a DRAM fab in China
- In 2016, Tsinghua Unigroup acquired a majority stake in XMC. Then, XMC was moved under a new group called Yangtze River Storage Technology (YRST). There seem to be some confusion between whether this group is called YRST or YMTC (Yangtze Memory Technology Company)
- Tsinghua Unigroup and its memory unit, YRST, have announced three major memory projects in China in recent times. Here is the latest activity:
- In December 2016, YRST unveiled a \$24 billion 3D NAND memory project in Wuhan, the capital of the Hubei province in central China. YRST hopes to build three fabs, which will each produce 100,000 wpm.
- In February 2017, Tsinghua Unigroup announced a \$30 billion memory project in Nanjing, the capital of China's eastern Jiangsu province. The goal is to make DRAMs first, then 3D NAND.
- In January of 2018, the group announced another 3D NAND project in Chengdu, the capital of southwestern China's Sichuan province. The goal is to build three fabs for a total investment of over \$30 billion over the next 10 years.

History of Chinese Memory Fab Investment (2)

- China's Tsinghua Unigroup, a state-run electronics giant, has three big memory projects at a total cost of \$84 billion. The goal is to build nine fabs, but only one fab is actually under construction today. It is focusing on 3D NAND, where it is sampling a 32-layer device with a 64-layer technology in R&D.
- Two other Chinese companies, Jinhua Integrated Circuit Co. (JHICC) and Innotron, are separately expected to ramp up 22nm DRAMs in new 300mm fabs. (Innotron is sometimes referred to as Rui-Li.)

Summary of Chinese Fab Investments

- China is currently building 19 new fabs, with 10 of those projects being 300 mm wafer plants
- A sizable number of international fab companies have fabs in China, including TSMC, UMC and soon, Global Foundries
- Intel and Samsung product 3D NAND in China, although the total output from those fabs is a tiny fraction of overall Chinese demand
- China's Tsinghua Unigroup announced 3 big memory projects costing \$84 B to build nine fabs but only one fab is currently under construction
- According to Semico, China's domestic memory vendors are projected to increase from about zero today to more than 300,000 wafers per month by 2021 (this would be less than 10% of WW estimated memory capacity by 2021)
- Note that Samsung invested \$26 B in capex in 2017, 2X normal level, probably to counter efforts by the Chinese to ramp up their own production. Capex spending on memory by Samsung is estimated to increase again in 2019 (40% higher than in 2018)

China Owned Memory Fabs (2018E-2021E)

	NAND				DRAM					
	2018E	2019E	2020E	2021E		2018E	2019E	2020E	2021E	
YMTC (XMC, Wuhan)					Innotron Memory					
Capex + Fab	~\$6.3 B 32 layer			64 layer	Capex + Fab	~\$7.2 B 22 nm, DDR3/4			17 nm	
Wafers/mo	5,000	10,000		100,00 0	Wafers/mo		20,000		125,000	
YMTC (Chengdu)	••••				JHICC					
Capex + Fab	\$30 B over 10 years 64 layer				Capex + Fab	~\$5.7B 22 nm, DDR3/4			\$2.3B	
Wafers/mo	,				Wafers/mo			20,000		
					YMTC (Tsinghua Unigroup, Nanjing)					
					Capex + Fab	~\$6.4B 22nm, DDR3/4				
					Wafers/mo					
					Xi'an UnilC (Qimonda/Inspur Group)	Shipped 4 GB DDR4 Chips, 30+ nm?				

Impact of USA & China Trade War

- The Trump administration set tariffs on \$34 billion in Chinese goods, effective July 6, 2018. These tariffs target electronic components in addition to steel and aluminum.
- Some of the impacts include:
 - The trade war may be an element in Micron losing a patent suite in China to UMC
 - China has held up Qualcomm's acquisition of NXP Semiconductors
 - HDDs made in china will be subject to tariffs but Seagate and WD plan to move their production of HDDs outside of China
 - Trump has also said that he would block Chinese investments in US company, which could dry up a source of investment capital

China's consumption, of Micron's DRAM and NAND flash bit output for 2018

- Fuzhou Intermediate People's Court of China preliminary injunction against Micron (brought by UMC), enjoining the company from selling related items in China.
- This will have very negative implications for Chinese companies like Alibaba, Baidu and Tencent, especially since DRAM is in short supply versus demand—need to buy from other suppliers
- With current DRAM shortage this will likely have little negative impact on Micron
- Chart at right shows Chinese company's dependence upon non-Chinese suppliers

What Goes Up, Will Go Down

- NAND and generally, DRAM, demand has been steadily increasing but supply follows a boom/bust cycle
- This is because of the lag time to adding capacity
- This has caused historical price cycles for memory
- Over supply in 2019 and 2020 will cause significant \$/GB declines for NAND and DRAM

Objective Analysis Data, 2018

Emerging Storage and Memory Trends and Conclusions

Emerging Non-Volatile Memories

- There is intense effort to commercialize several non-volatile memories that could replace current volatile memories, such as DRAM and SRAM
- These technologies can be applied to stand along memory chips as well as in embedded memory
- This could reduce energy expenditure in battery and low power devices and also create more efficient data centers
- These NV memories will enable both IoT devices as well as data centers at the edge or in the cloud
- The memory technologies under consideration include magnetic random access memory (MRAM), resistive RAM (RRAM or ReRAM), phase change RAM (PRAM) and ferroelectric RAM (FRAM or FeRAM)

Memory Technology Comparison

Technology	FeRAM	MRAM	ReRAM	РСМ	DRAM	NAND Flash	
Nonvolatile	Yes	Yes	Yes	Yes	No	Yes	
Endurance	1012	1012	106	10 ⁸	10 ¹⁵	10 ³	
Write Time	100ns	~10ns	~50ns	~75ns	10ns	10µs	
Read Time	70ns	10ns	10ns	20ns	10ns	25µs	
Power Consumption	Low	Medium/Low	Low	Medium	Very High	Very High	
Cell Size (f ²)	15-20	6-12	6-12	1-4	6-10	4	
Cost (\$/Gb)	\$10/Gb	\$30-70/Gb	Currently High	\$0.16/Gb	\$0.6/Gb	\$0.03/Gb	
Companies	Ramtron, Fujitsu	Everspin,GF, Samsung, others	WD, Sony, Adesto, others	Intel, Micron	Many	Many	

How Persistent Memory will Success, Tom Coughlin and Jim Handy, 2018 SNIA Persistent Memory Summit, January 2018 ³⁵ © 2018 Coughlin Associates

MRAM and PRAM

- MRAM
 - Everspin shipped over 70 M MRAM Chips. Company has partnership with Global Foundries, who is building 300 mm wafers and targeting embedded memory applications
 - Samsung--plans to ship STT MRAM product samples by 2018.
 - Seagate was showing an Everspin MRAM boot SSD at the 2017 FMS
- PRAM
 - Intel Optane NVMe products shipped in 2017.
 - Micron planning to introduce DIMM-based 3D XPoint product
 - Intel introduced their Optane DIMM products in June 2018

MRAM Product Generations

© 2018 Coughlin Associates

Latency Reduction Requirements

Latency (uS)

NVMe Architectures and DIMMS

Flash in the Memory Channel

- The memory channel is where you plug DIMM DRAM boards into a device
- Memory channel performance is higher than PCIe bus
- Several start-ups and established companies (e.g. HPE's The Machine) working to put non-volatile storage on the memory bus

Big and Affordable Memory

High Performance Storage

Direct Load/Store Access

Native Persister

128, 256, 512GB

(intel)

XEON

DDR4 Pin Compatible

Hardware Encryption

High Reliability

NOW SHIPPING SAMPLES AD DEVELOPER ENGAGEMENT

BLUEPRIN

Intel Optane DIMMS

- Samples in 2018
- Plan production starting in 2019
- This was harder than earlier NVMe SSDs
- Speeds up databases, increases number of containers possible, faster boots
- With Fabric can do data replication between CPUs
- Open API community

GET MORE WITH HIGH PERFORMANCE STORAGE

alts have been estimated bas ware and workloads used in are, operations and function mance of that product when

VS. COMPARABLE SERVER SYSTEM WITH DRAM AND NAND NVME DRIVES WHEN USING APACHE* CASSANDRA-4.0

vistems, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance if for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, suse the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the primation go to www.intel.com/benchmarks.

OPTANE DC 🖇

PERSISTENT MEMORY OVER FABRICS (PMOF) For data replication with direct load/store access

PM Needs New Software

- Move away from "Storage vs. Memory" approach
 - Store at the byte level, not blocks
 - Avoid the storage stack
 - Avoid things like flash translation

Memory-centric Computing

For many emerging challenges, memory capacity, memory access latency and memory bandwidth are more constrained than compute resources

- Memory Disaggregation Remove memory from behind the processor
- Memory Pooling & Sharing Enable efficient use of memory. Address new class of problems with large memory footprint
- Heterogeneous Compute
 Enable multi-vendor heterogeneous
 compute (e.g. ML accelerators)

CPU-centric architecture

Data-centric architecture

43

19

Upcoming PM Report

- Examines the PM Ecosystem
 - Technologies
 - Companies
 - Markets
 - Support requirements
- Forecasts PM consumption
 - Embedded PM
 - Discrete PM
- Objective Analysis/Coughlin Associates (www.tomcoughlin.com)
 - Due date: July 2018

Digital Storage Capacity Projections

- The growth and processing of data will lead to the use of many types of digital storage
- SSDs will dominate for high performance storage and higher total revenue
- HDDs will be high capacity and used for colder storage
- Magnetic tape will be used by some organizations for the lowest cost (currently <1 cent/GB)

Conclusions

- Digital content is exploding, driving the need for memory and storage to support rapid processing of that data as well as inexpensive storage for long term retention
- Flash memory will continue for many generations, driven by 3D technology and demand for fast storage
- China is moving to make more internal solid state memory to meet its internal demand
- It remains to be seem what the impact of the USA/China trade war will have on memory companies and customers.
- It also remains to be seen how well China will execute on it's plans with such strong international competitors
- Emerging memory technologies will feed the memory-centric processing future world
- Evolving storage technologies don't seem to go away, but their uses change.

References

- Digital Storage Technology Newsletter, Coughlin Associates, https://tomcoughlin.com/product/digital-storagetechnology-newsletter/
- 2017 FMS Keynote Presentations by Intel, Micron, Samsung, Seagate, Western Digital
- Touch Rate: A metric for analyzing storage system performance, Steven Heltzer and Tom Coughlin, 2015, <u>http://www.tomcoughlin.com/techpapers.htm</u>

Questions?