BiG Research Sector Note: Semiconductors

In this report we analyze 6 players: Applied Materials, Lam Research, Micron, Intel, AMD and NVIDIA. Below is a scheme to show their place in the supply chain:

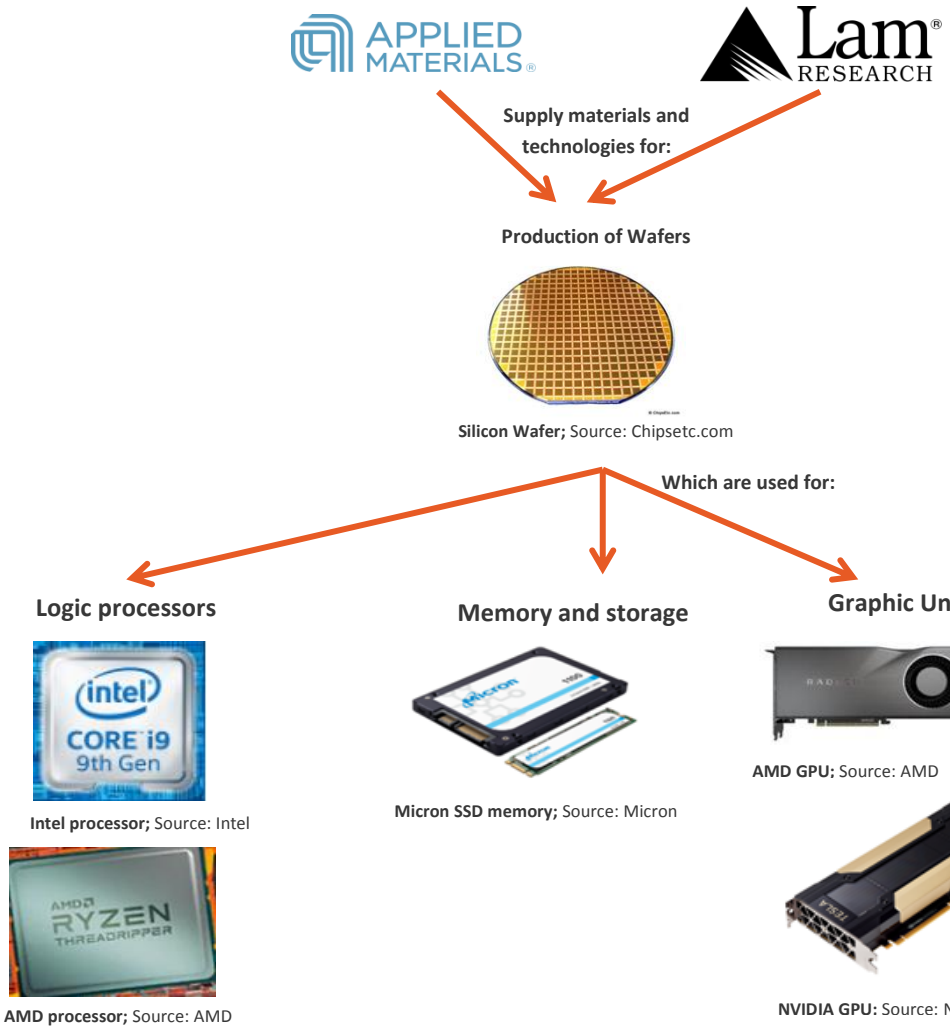

Relative Valuation and top picks

From the 6 players we studied and taking into account the dynamics of the sector we prefer Intel and Applied Materials, due to their exposure to the growing logic segment and low P/E ratios. Despite the lower P/E of 14.5x of Lam Research and Micron (7.6x) we expect the downfall in the memory segment to continue to affect the companies. AMD and NVIDIA despite the strong growth prospects in CPUs and GPUs respectively, quote at expensive multiples.

Name	Country	Market Cap (mn)	Currency	P/E 2019E	P/E 2020E	EV/EBITDA	YTD	Div. Yield	NetDebt/ EBITDA	Margin EBITDA
LAM RESEARCH CORP	UNITED STATES	31,066	USD	14.5	14.3	8.1	52.2%	2.1%	-0.4	30.1%
APPLIED MATERIALS INC	UNITED STATES	47,489	USD	17.1	13.8	10.0	54.9%	1.6%	0.4	27.7%
ADVANCED MICRO DEVICES	UNITED STATES	35,531	USD	51.5	32.8	46.8	78.0%		0.3	9.7%
INTEL CORP	UNITED STATES	229,894	USD	12.0	11.4	7.7	9.4%	2.4%	0.4	46.2%
MICRON TECHNOLOGY INC	UNITED STATES	52,088	USD	7.6	18.2	2.1	48.7%		-0.2	60.8%
NVIDIA CORP	UNITED STATES	104,334	USD	32.2	24.4	31.6	28.3%	0.4%	-1.6	30.5%
Average				22.5	19.1	17.7	45.26%	1.62%	-0.2	34.2%

Source: BiG Research

Analyst:
João Calado, CFA

Research:
research@big.pt

This document has been prepared exclusively for informative purposes, and is based on publicly available information, retrieved from sources deemed trustworthy. BiG does not assume any responsibility for the full correction of the information provided, and the information here provided should not be interpreted as an indicator that any results will be achieved. We emphasize that the projected results are susceptible to alterations due to changes in the assumptions that have served as basis to the information here provided. We forewarn that the previous performance of a security is not a guarantee of identical performance in the future. Changes in exchange rates of securities denominated in a currency different from that of the investor may lead to a negative impact on the value, price or return of such securities. BiG, or its employees, may have, at any moment, an investment, subject to alterations, in any security mentioned in this document. BiG may provide additional information, if so is requested. This document is not a sale proposal, nor a purchase solicitation for the subscription of any securities. BiG ensures independence in investment recommendations under the terms of points 7.23 and 7.24 of BiG's Conflict of Interest Management Policy. BiG continuously monitors the recommendations issued by analysts under the terms of the Code of Conduct and Personal Transactions Policy.

Current themes in the industry

Memory chips demand: Due to high inventories in the end of 2018 and supply problems of CPUs from Intel in the beginning of the year, memory chips demand fell leading to a correction in their prices. Most of players in the industry foresee further weakness in this segment for the rest of the year.

Cryptocurrency: The frenzy of 2017 around Bitcoin and other cryptocurrencies mining led to an exponential increase in demand for GPUs (graphic units) which were used in the mining process. This impacted favorably both AMD and NVIDIA sales, however the correction of Bitcoin price in 2018 ended the high growth of demand of GPUs for cryptocurrency mining.

Race for more advanced chips:

- **Moore's Law:** This was an observation made in 1965 by Gordon Moore, co-founder of Intel. It said that the number of transistors per square inch on ICs had doubled every year since it was invented and this trend would continue into the future. Currently it doubles at around each 18 months.

- **14nm vs 10nm vs 7nm:** These are node names but that actually do not identify the size of chips as it used to happen, but are mainly used as marketing tools to show the advancement of processors. The new 10nm and 7nm processes are based on more complex versions of the 16/14nm finFET transistors.

- **AMD vs Intel:** Despite appearing that AMD is more advanced in terms of technology since it advertises its future generation of chips as 7nm, the fact is that its performance will be very similar to Intel's 10nm. However is very likely that AMD will be able to launch commercially the 7nm chips before Intel's 10nm which may lead to a decrease in Intel's market share, something that didn't happen in the last 20 years.

Exascale supercomputers: Today's supercomputer can do up to a quadrillion calulations each second (1,000,000,000,000,000). The exascale supercomputer will be able to do a quintillion calculations each second (1,000,000,000,000,000,000). This will allow to do more precise and fast calculations that may unlock discoveries in medicine, climate, water usage, nuclear physics etc. This supercomuter will allow simulations while using AI to improve the technology like for example simulate air flow in an airplane and at the same time experiment through AI, different aerodynamic schemes to improve fuel efficiency in airplanes.

Analyst:
João Calado, CFA

Research:
research@big.pt

This document has been prepared exclusively for informative purposes, and is based on publicly available information, retrieved from sources deemed trustworthy. BiG does not assume any responsibility for the full correction of the information provided, and the information here provided should not be interpreted as an indicator that any results will be achieved. We emphasize that the projected results are susceptible to alterations due to changes in the assumptions that have served as basis to the information here provided. We forewarn that the previous performance of a security is not a guarantee of identical performance in the future. Changes in exchange rates of securities denominated in a currency different from that of the investor may lead to a negative impact on the value, price or return of such securities. BiG, or its employees, may have, at any moment, an investment, subject to alterations, in any security mentioned in this document. BiG may provide additional information, if so is requested. This document is not a sale proposal, nor a purchase solicitation for the subscription of any securities. BiG ensures independence in investment recommendations under the terms of points 7.23 and 7.24 of BiG's Conflict of Interest Management Policy. BiG continuously monitors the recommendations issued by analysts under the terms of the Code of Conduct and Personal Transactions Policy.

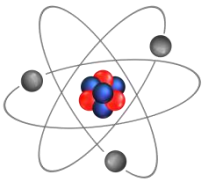

Atom; Source: pixabay

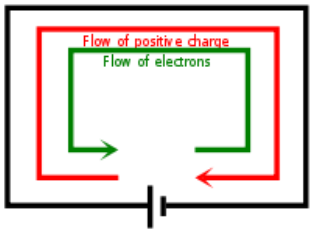

Flow of electrons; Source: Wikipedia

Basics of Semiconductor's Industry

What is an Atom? An atom is composed by a nucleus of protons and neutrons and orbiting the nucleus are the electrons. These electrons are responsible for electricity. The type of atom is defined by the number of protons in the nucleus which will give the number of the element in the periodic table. However they may have different number of neutrons or electrons, being in this case called isotopes.

How electricity works? Electrons are far lighter than the protons in the nucleus so can easily move and it is this movement that forms an electric current. The atoms' protons account for the positive charge of the nucleus and the electrons for the negative charge. When protons equal electrons then the atom its at its lowest possible energy level and net electric charge is zero. We can change the atom's charge by adding or removing electrons. More electrons than protons means the atom is negatively charged and vice versa. When charged the atom is called either positive ion or negative ion. Nature always tries to reach equilibrium so if a conductor loses electrons and becomes positively charged and then touched another conductor, electrons from the second are transferred to the first creating an electrical charge. Materials with electron mobility are called conductors (like copper). Materials with low electron mobility are called insulators (like plastic). The key to electricity is to generate a circuit of flowing electrons from a negative source to a positive source just like the two sides of a battery. Copper electrons can move randomly but for an electricity circuit there needs to be a flow, so we need a source of electrons that pushes them and a positive source that attracts them, which is called the voltage.

Voltage? The more voltage, the more current and the more power generated. this current is measured in Amperes. The current may be alternating or direct. The plugs at home provide alternating current (AC) so the electrodes do not flow in a continuos loop but instead alternate between moving backwards and forward. Laptops and mobile phones use direct current (DC) which will flow in a closed circuit. AC is more efficient to transport electricity over long distances. AC is converted to DC by using an inverter.

Short circuits: It is dangerous to connect directly two sides of a battery (called a short circuit) because the lack of something to absorbe the elctricity will cause the circuit to overheat, and this is why buildings have fuses that cut the current when it becomes to high. Fuses are very cheap and are composed with a more delicate cable that in case the current exceed a certain amount of amps, it gets destroyed interrupting the circuit without causing damage to the main cable. Resistors are used to slow the current in order to avoid overheating of the cable and damage to the circuit. The more devices connected the higher the current and that's why you have switches at home that interrupt the circuit when the current exceeds a certain level.

0s and 1s; Source: Pixabay

How computers work? Computers need to do 4 things: receive input, storage information, process that information and create an output. The input received from a keyboard or microphone or camera is transformed into a bunch of 0 and 1 which are registered through electrical signals being on (1) or off (0) in a certain circuit. Logic gates transform multiple signals into an different output, for example two signals of 1 may transform into a 0 while a 1 and a 0 may transform into a 1. Each 1 or 0 is called a bit. And 8 bits correspond to 1 byte. These logic gates can also do sums or multiplication or divisions in order to compute multiple results. Each operation done like this could be done by a human but as there are a lot of them in a chip and can work at the same time, more complex calculations can be done at a faster rythm. Electricity moves at the speed of light so if the circuit is smaller, the speed is faster. Modern circuits can perform billions of calculations per second.

Analyst:
João Calado, CFA

Research:
research@big.pt

This document has been prepared exclusively for informative purposes, and is based on publicly available information, retrieved from sources deemed trustworthy. BiG does not assume any responsibility for the full correction of the information provided, and the information here provided should not be interpreted as an indicator that any results will be achieved. We emphasize that the projected results are susceptible to alterations due to changes in the assumptions that have served as basis to the information here provided. We forewarn that the previous performance of a security is not a guarantee of identical performance in the future. Changes in exchange rates of securities denominated in a currency different from that of the investor may lead to a negative impact on the value, price or return of such securities. BiG, or its employees, may have, at any moment, an investment, subject to alterations, in any security mentioned in this document. BiG may provide additional information, if so is requested. This document is not a sale proposal, nor a purchase solicitation for the subscription of any securities. BiG ensures independence in investment recommendations under the terms of points 7.23 and 7.24 of BiG's Conflict of Interest Management Policy. BiG continuously monitors the recommendations issued by analysts under the terms of the Code of Conduct and Personal Transactions Policy.

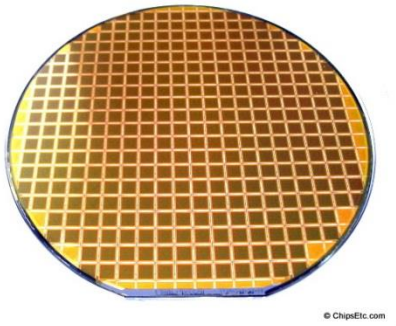

Silicon Wafer; Source: Chipsetc.com

BINARY CODE ALPHABET REFERENCE

1	A	00001	14	N	01110
2	B	00010	15	O	01111
3	C	00011	16	P	10000
4	D	00100	17	Q	10001
5	E	00101	18	R	10010
6	F	00110	19	S	10011
7	G	00111	20	T	10100
8	H	01000	21	U	10101
9	I	01001	22	V	10110
10	J	01010	23	W	10111
11	K	01011	24	X	11000
12	L	01100	25	Y	11001
13	M	01101	26	Z	11010

Binary code for alphabet; Source: Our Family Code

Intel Core i9 chip; Source: Intel

How a chip is produced? First you design circuits. Chips are made by sand and are called silicon wafers. These wafers take an oxide layer and then a photoresist chemical. Photolitographic technics transfer the circuit structures to the waffers like a slide projection. UV light transfers the circuits despicted in a mask to the waffer. Then they go to ion implantation where the electrical properties of transistors will be specified. Silicon is a semiconductor which means that its conductivity can chage by high precision placement of dopen atoms. Then using high temperatures, these atoms will acquire a fixed position in the atomic structure changing the silicon surface and creating this way a transistor. Transistor are the conections between multiple circuits. After copper is deployed to fill the trenches. Then excess copper is eliminated in order to isolate each interconnect. In two months the waffer is ready. Then the chips are cut off form the waffer and bond to the frame which is then sealed with a cover.

Binary system: This system uses the digits 0 and 1 to represent any number. While in the decimal system each position of the digit is multiplied by 10 (so 1, 10, 100, 1000) in binary system is multiplied by 2: so number one is 0001, number 2 is 0010, number 4 is 0100, number 10 is 1010. This example is using 4 wires (0000) however using 32 wires you can store numbers of over 4bn. So through this system you can represent any number you like. Text can be also represented as each letter would be a number and any word or paragraph would be a sequence of numbers, which can be stored as on or off electrical signals. Images are composed by tiny pixels in which each pixel has a colour that is represented by a number. Regarding sounds, any sound is basically a vibration, which may be represented in a graph and by numbers also.

CPU (Central Processing Unit): These are the chips that do all the computation needed in a PC. The main processor players are Intel and AMD. Chips grab information from RAM and process it to produce an output which then it is stored again in the PC. A CPU can have multiple cores, in which each core is performing an individual task (also known as command or thread). A core may be able to perform 2 threads at the same time which enables current chips to have for example 8 cores and 16 threads, which basically means that the chip will process 16 commands at the same time and consequently will be faster. The speed at which the processor runs is called the system clock and is measured in GHz, as an example a 3.1GHz processor has a clock that beats 3.1bn times per second. This means that in the case of 64 bits processor (the most used currently), it processes 64 bits in each clock time. Higher GHz speed also consumes more energy and generates more heat and while important for the overall chip speed, increasing it is not even as efficient as increasing the number of cores for example. The speed also depends on the cooling system as the chip is programmed to decrease speeds if the heating increases. There are 2 types of clocks, internal and external. Internal is the speed at what the chip processes the information it already has, the external is the speed at which the processor can fetch information from external components to then process and so has also a significant effect in the chip speed. The CPU may also have an integrated graphics card eliminating the need of having an external graphics card if there are no demanding graphic tasks such as high quality gaming or picture and video rendering edition. The cache refers to the memory integrated in the processor which is also called SRAM and its detailed below. CPUs used for servers are very different from CPUs for consumers and that is why both major players (Intel and AMD) have 2 distinct categories of CPUs.

Server CPUs: Servers CPUs usually run at more cores but at a lower clock speed in order to save energy, while consumer chips usually have less cores but run at higher clock speed. One socket server means that it is a server based on a single motherboard and a single CPU socket. With more sockets, servers have more cores, but at a higher price, energy consumption and costs. One socket servers became popular since 2016 as the number of cores in each CPU increased significantly.

Analyst:
João Calado, CFA

Research:
research@big.pt

This document has been prepared exclusively for informative purposes, and is based on publicly available information, retrieved from sources deemed trustworthy. BiG does not assume any responsibility for the full correction of the information provided, and the information here provided should not be interpreted as an indicator that any results will be achieved. We emphasize that the projected results are susceptible to alterations due to changes in the assumptions that have served as basis to the information here provided. We forewarn that the previous performance of a security is not a guarantee of identical performance in the future. Changes in exchange rates of securities denominated in a currency different from that of the investor may lead to a negative impact on the value, price or return of such securities. BiG, or its employees, may have, at any moment, an investment, subject to alterations, in any security mentioned in this document. BiG may provide additional information, if so is requested. This document is not a sale proposal, nor a purchase solicitation for the subscription of any securities. BiG ensures independence in investment recommendations under the terms of points 7.23 and 7.24 of BiG's Conflict of Interest Management Policy. BiG continuously monitors the recommendations issued by analysts under the terms of the Code of Conduct and Personal Transactions Policy.

Nvidia GPU; Source: Nvidia

GPU (Graphics Processing Unit): While a CPU is built to be able to solve complex problems, a GPU is more oriented to solve many repetitive simple problems like for example geometry calculations that enable graphics in gaming to be built. That is why this component is so important for the gaming industry and for the bitcoin industry given that both need many simple calculations to be done at the same time. CPUs may have between one or 10 cores while GPUs usually have hundreds of cores that are less powerful but are perfect for simple repetitive tasks. The information processed will then be stored in Video RAM also known as GDDR (explained below). For the GPU to work at full capacity, it has to have an efficient cooling system which usually determines different prices for the whole GPU depending on the cooler supplier.

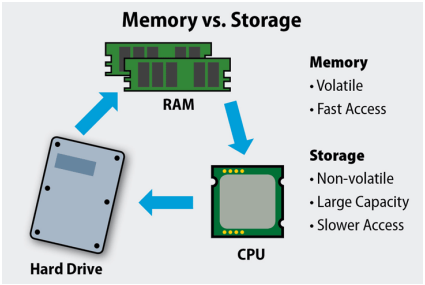

Memory vs Storage; Source: Enterprise Storage Forum

Memory vs storage in computers: Memory as commonly known as RAM (random access memory) is volatile which means if you turn off the PC it disappears. Storage (Hard drive) is non volatile, so it will remain stored in the computer until something else overrides that data. Nowadays Magnetic Core Memory is used. This works by using a lot of magnetic cores which by receiving an electric pulse it magnetizes in a certain direction, and stays that way even if the electric pulse stops. If you pass an electric pulse in another direction, the magnetic direction of the core changes and consequently you changed the core from 0 to 1 or vice versa. For a program to work it needs to be stored in hard drive and then loaded in RAM for which CPU will access the data. If there is not enough RAM then the RAM memory can´t store the whole information there and has to send and receive part of the information from the hard drive which is slower. That is why a computer with more RAM performs faster than a computer with less RAM.

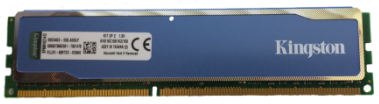

RAM; Source: Pixabay

RAM: Memory is stored in motherboards in diferent modules called DIMMs (Dual Inline Memory Module). Synchronous Dynamic RAM contains capacitors which have constantly to be refreshed by electricity in order to store information and that is why when you turn off the computer, the RAM memory is erased. The SDRAM works in sync with the system clock. When we speak about 32 or 64 bit data path, it means the amount of bits of data that are transferred in 1 clock cycle. The more bits per cycle, the faster the computer. DIMMS transfer 64 bits per clock cycle. If a SDRAM operates at 100 MHz and has an 8 byte wide bus it means that its speed is 100MHz x 8 bytes = 800MB/s. The invention of DDR (Double data rate) enabled that twice of the information to be sent per clock. DDR RAMS includes both the clock speed and total bandwith in its name, so a DDR-333 PC-2700 means 333MHz of clock speed and 2700 is the MB/s (333MHz x 8 bytes = 2700MB/s). The last generation is DDR4 which has 288 pins has more speed with a bandwith around 4266MHz so in total 34100MB/s. In servers which cannot go down for any reason, RAM modules will have 9 memory chips instead of 8 because the additional chip will perform Error Correcting Code (ECC).

NVDIMMS: These Non-Volatile Dual In-line Memory Modules use volatile memory during normal operation and dump the contents into non-volatile memory if the power fails, using an on-board backup source.

TSV RDIMMS: These Through Silicon Via-based chip packaging technology allows to a more efficient connection of DRAM chip dies which replace wire bonding. This allows higher die stacks possible and so more memory, like for example the latest TSV DDR4 DRAM with 128GB RDIMM modules.

GDDR RAM; Source: HotHardware.com

GDDR: While CPU uses its specialized RAM (DDR), graphic uses GDDR. On graphics, the last technology is GDDR6 (Graphics Double Data Rate). This type of RAM uses less energy and so produces less heat, however has a higher bandwith. The GDDR6 is capable of operating at 1.35V and offers an increased per pin bandwith of up to 16Gbits/s, reaching a speed of 768GB/s.

Analyst:
João Calado, CFA

Research:
research@big.pt

This document has been prepared exclusively for informative purposes, and is based on publicly available information, retrieved from sources deemed trustworthy. BiG does not assume any responsibility for the full correction of the information provided, and the information here provided should not be interpreted as an indicator that any results will be achieved. We emphasize that the projected results are susceptible to alterations due to changes in the assumptions that have served as basis to the information here provided. We forewarn that the previous performance of a security is not a guarantee of identical performance in the future. Changes in exchange rates of securities denominated in a currency different from that of the investor may lead to a negative impact on the value, price or return of such securities. BiG, or its employees, may have, at any moment, an investment, subject to alterations, in any security mentioned in this document. BiG may provide additional information, if so is requested. This document is not a sale proposal, nor a purchase solicitation for the subscription of any securities. BiG ensures independence in investment recommendations under the terms of points 7.23 and 7.24 of BiG's Conflict of Interest Management Policy. BiG continuously monitors the recommendations issued by analysts under the terms of the Code of Conduct and Personal Transactions Policy.

SRAM; Source: RS Components

SRAM: Static RAM is what is used in the CPU cache. It is faster than RAM but more expensive also. Its function is to store copies of data from RAM which is waiting to be used by the CPU. The CPU always checks first it the information it needs is in cache, and if not it gets it from RAM. That is why a computer with more cache will be faster. There are 3 levels of cache. The number one runs at the same speed as the processor and its located in the processor itself so its the fastest cache on the PC. The level 2 cache is on the processor and catches data not caught by the level 1. The same goes for level 3 vs level 2 cache. So the processor gets first data from level 1 and if its not there, goes to level 2 and level 3 consequently. Level 1 and level 2 are dedicated to each CPU core while level 3 is shared by the different cores.

MRAM: Magneto-Resistive Random Access Memory enables a symmetrical speed of reading vs writing and at the same speed as SRAMs. It is nonvolatile, as content is not lost on power down and writes to the memory do not damage it. However this technology has been studied from 1990s and some research companies gave up on it and started to deploy resources in other type of technologies. The main difficulties is to be able to reduce the size of it given that magnets have to keep a certain distance to not influentiate each other and the cost per bit it is still higher than Flash, DRAM or even SRAM.

ReRAM: Resistive Random Access Memory has not been commercially launched yet. ReRAM is non volatile so it does not need power to keep its data. This would enable it to replace flash and increase the speed of processing of the PC. It works by changing the resistance across a dielectric solid state material, which will this way record data (0s and 1s).

Intel Optane; Source: Intel

3D XPoint: This is a non-volatile memory developed jointly by Intel (Optane) and Micron Technology (QuantX) and already commercially available and its price is lower than DRAM but higher than flash. It is a non-volatile memory and its one of the fastest SSDs available, being fast enough to be used as non-volatile RAM. With support of the new chip Cascade Lake from Intel, it may also be used as cache memory. This memory has no transistors, so does not deteriorate significantly with writing cycles and uses less space. Its a grid of parallel and perpendicular wires and between the wires there is a material that changes state which allows for information to be stored as 0s or 1s. Curently it is still only used as an extra to improve speed of program loading, instead of a full replacement of RAM and Flash.

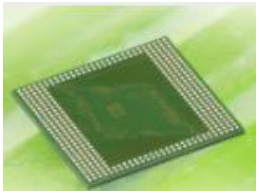

LPDDR4; Source: Micron

LPDRAM: This Low Power DRAM allows the use of RAM in smaller devices and which consumes less energy. the last version, LPDDR4 can have between 4Gb and 32Gb with either 2x or 4x 16bits channels and voltages between 0.6V and 1.1V. The clock frequencies reach up to 2133MHz.

Flash memory; Source: Pixabay

Flash memory (NAND): Flash memory is not volatile and its called flash due to the need of erasing every byte in each section in order to rewrite the information there. These are composed by chips with a lot of circuits just like processor chips. This is why after writing a lot of information in the memory, when you need to rewrite it will take some time for it to be processed. NAND flash has to be accessed sequentially while NOR Flash may be randomly accessed do its faster and usually used for BIOS, boot codes and configuration setting constantly updated. Each time the memory is erased is called a Erase cycles and with each cycle the gate oxides begin to deteriorate. Usually each NAND Flash can take around 1,000 to 10,000 erase/write cycles while NOR Flash takes between 10,000 and 100,000 erase/write cycles. NAND Flash is the cheapest memory, even cheaper than RAM.

Analyst:
João Calado, CFA

Research:
research@big.pt

This document has been prepared exclusively for informative purposes, and is based on publicly available information, retrieved from sources deemed trustworthy. BiG does not assume any responsibility for the full correction of the information provided, and the information here provided should not be interpreted as an indicator that any results will be achieved. We emphasize that the projected results are susceptible to alterations due to changes in the assumptions that have served as basis to the information here provided. We forewarn that the previous performance of a security is not a guarantee of identical performance in the future. Changes in exchange rates of securities denominated in a currency different from that of the investor may lead to a negative impact on the value, price or return of such securities. BiG, or its employees, may have, at any moment, an investment, subject to alterations, in any security mentioned in this document. BiG may provide additional information, if so is requested. This document is not a sale proposal, nor a purchase solicitation for the subscription of any securities. BiG ensures independence in investment recommendations under the terms of points 7.23 and 7.24 of BiG's Conflict of Interest Management Policy. BiG continuously monitors the recommendations issued by analysts under the terms of the Code of Conduct and Personal Transactions Policy.

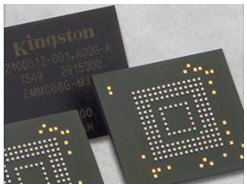

eMMC; Source: Kingston

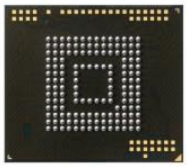

UFS; Source: Samsung

Compact Flash; Source: Kingston

SSD; Source: Intel

Managed NAND: Usually is a combination between raw NAND and a hardware which does bad block management, ECC and wear leveling. Main technologies are eMMC, eSD, CompactFlash, SSDs.

eMMC: The embedded MultiMediaCard is mostly used in mobile devices or small PCs, and the storage is soldered directly in the motherboard. The speed will be lower than SSDs and as there are less memory gates, the volume of operations at the same time will be lower.

Universal Flash Storage: This kind is also used in smartphones and the 2.1 version was used for Samsung Galaxy S10. The last version UFS 3.0. is very fast and consumes less energy, and consequently is replacing the eMMC. It is capable of reaching 23.2Gbps in two lanes (each with 11.6Gbps each), which is another advantage vs eMMC which only has one lane.

eSD: The embedded Secure Digital cards are also used for mobile devices.

Compact Flash: was created by SanDisk in 1994 and uses flash memory technology to store data on a very portable device, now used mainly for saving photos in digital cameras.

SSD: Solid State Drivers use less power than HDD because they are solid state chips which also increases its speed. They are similar to flash memory. NAND Flash memory has many cells in which bits turn on or off according to electric charge which enables the storage of information. The number of bits in each cell may be only 1 (Single Level Cell or SLC Flash), 2 (Multi Level Cell or MLC), 3 (Triple Level Cell or TLC) or the latest technology of QLC (Quad Level Cell with 4 bits per cell).

SLC SSDs: These are the fastest and most accurate NAND Flash memory in both reading and writing, and has the longest lifespan and charge cycles (between 90,000 and 100,000). However its also the most expensive and often has smaller capacities. Usually used in servers that require heavy reading or writing cycles.

MLC SSDs: These are less accurate and have shorter cycles (10,000), however its production cost is lower. It is usually used by the everyday consumer or PCs for gamers.

TLC SSDs: These not only have shorter cycles (3,000 to 5,000) but are also slower so only suitable to consumer and usually for netbooks or tablets.

QLC SSDs: These can store more data and are cheaper, but run at a slower pace and the cycles are lower also.

NVMe: This is an interface protocol built specially for SSDs, which works with PCIe. NVMe began in the enterprise market but is entering the consumer market. Intel SSD 600p series has 1TB of memory and has a throughput of 8Gbps in each direction so a total of 16 Gbps vs conventional SATA SSDs of 6Gbps. While the latest SATA protocol allowed for 1 queue of commands and had a limit of 32 commands in that queue, M.2 NVMe allows 64k queues and 64k commands per queue.

HDD: In an Hard Disk Drive a magnetic disc spins and lets data be saved and read.

eMCP: This refers to embedded Multi Chip Package and is basically a combination of different chips for memory and processing into a compact module in order to decrease its size to enable its use in IoT.

Analyst:
João Calado, CFA

Research:
research@big.pt

This document has been prepared exclusively for informative purposes, and is based on publicly available information, retrieved from sources deemed trustworthy. BiG does not assume any responsibility for the full correction of the information provided, and the information here provided should not be interpreted as an indicator that any results will be achieved. We emphasize that the projected results are susceptible to alterations due to changes in the assumptions that have served as basis to the information here provided. We forewarn that the previous performance of a security is not a guarantee of identical performance in the future. Changes in exchange rates of securities denominated in a currency different from that of the investor may lead to a negative impact on the value, price or return of such securities. BiG, or its employees, may have, at any moment, an investment, subject to alterations, in any security mentioned in this document. BiG may provide additional information, if so is requested. This document is not a sale proposal, nor a purchase solicitation for the subscription of any securities. BiG ensures independence in investment recommendations under the terms of points 7.23 and 7.24 of BiG's Conflict of Interest Management Policy. BiG continuously monitors the recommendations issued by analysts under the terms of the Code of Conduct and Personal Transactions Policy.

DISCLOSURES

- Banco de Investimento Global, S.A. is an institution registered on and regulated by the Bank of Portugal and by the Portuguese Securities Market Commission, the two main entities that regulate financial activities in Portugal.

- BiG has a Code of Conduct, applicable to all its employees that carry out activities as financial analysts, with the aim to ensure the continuation of the accuracy, competence and excellence that characterize its institutional image. This document is available for external consultation, if required.

- The analysts in BiG's Research Team do not, and will not, receive any kind of compensation in the scope of the regular carrying out of its recommendations, which reflect strictly personal opinions.

- There isn't a predefined coverage policy in regards to the selection of stocks that are subject to investment recommendations.

- Clarification of the qualitative terms implied in the recommendations:
  - Buy, expected absolute return above 15%;
  - Accumulate, expected absolute return between +5% and +15%;
  - Keep/Neutral, expected absolute return between -5% and +5%;
  - Reduce, expected absolute return between -5% and -15%;
  - Sell, expected absolute return below -15%;

The investment framework aforementioned is merely indicative and not globally strict.

- Unless otherwise specified, the price-targets of the investment recommendations issued by BiG's Research Team are valid for 12 months.

- The update of the investment recommendations models and respective price-targets will occur, usually, in a period of 6 to 12 months.

- For more information please see our Disclaimer document online in this link, https://big.pt/InformacaoMercados/AnalisesBig/Outros , or please contact research@big.pt.

DISCLAIMER

This document has been prepared exclusively for informative purposes, and is based on publicly available information, retrieved from sources deemed trustworthy. BiG does not assume any responsibility for the full correction of the information provided, and the information here provided should not be interpreted as an indicator that any results will be achieved. We emphasize that the projected results are susceptible to alterations due to changes in the assumptions that have served as basis to the information here provided. We forewarn that the previous performance of a security is not a guarantee of identical performance in the future. Changes in exchange rates of securities denominated in a currency different from that of the investor may lead to a negative impact on the value, price or return of such securities. BiG may provide additional information, if so is requested. This document is not a sale proposal, nor a purchase solicitation for the subscription of any securities.

BiG ensures independence in investment recommendations under the terms of points 7.23 and 7.24 of BiG's Conflict of Interest Management Policy. BiG continuously monitors the recommendations issued by analysts under the terms of the Code of Conduct and Personal Transactions Policy.

Analyst:
João Calado, CFA

Research:
research@big.pt

This document has been prepared exclusively for informative purposes, and is based on publicly available information, retrieved from sources deemed trustworthy. BiG does not assume any responsibility for the full correction of the information provided, and the information here provided should not be interpreted as an indicator that any results will be achieved. We emphasize that the projected results are susceptible to alterations due to changes in the assumptions that have served as basis to the information here provided. We forewarn that the previous performance of a security is not a guarantee of identical performance in the future. Changes in exchange rates of securities denominated in a currency different from that of the investor may lead to a negative impact on the value, price or return of such securities. BiG, or its employees, may have, at any moment, an investment, subject to alterations, in any security mentioned in this document. BiG may provide additional information, if so is requested. This document is not a sale proposal, nor a purchase solicitation for the subscription of any securities. BiG ensures independence in investment recommendations under the terms of points 7.23 and 7.24 of BiG's Conflict of Interest Management Policy. BiG continuously monitors the recommendations issued by analysts under the terms of the Code of Conduct and Personal Transactions Policy.