Solid State Drive versus Hard Disk Drive

Background

- Data storage devices
  - Popular in our lives
  - Known as secondary storage devices
- 2 types of data storage devices
  - Hard Disk Drive (HDD)
  - Solid State Drive (SSD)

- HDD:
  - IBM 1956
  - Magnetic coating on metal platters
  - Data are read/written by movements of magnetic heads.

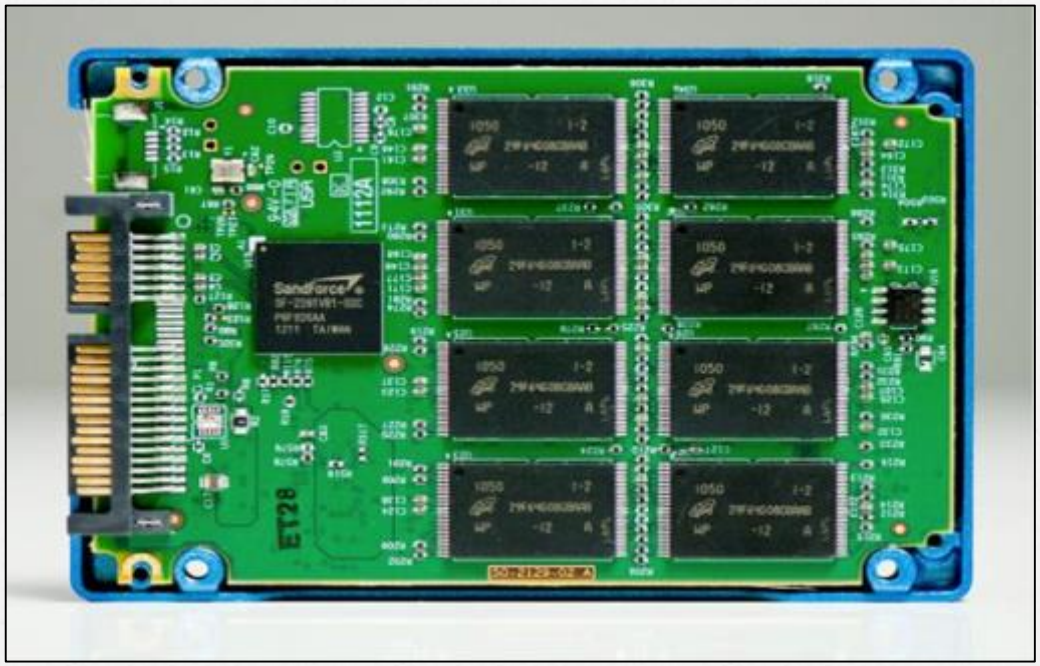

- SSD:
  - 1950s
  - Controller and NAND flash memory chips
  - Data are read/written by read/write functions

1. Comparison of NAND and NOR Flash Memory

	NAND flash memory	NOR Flash memory
Cost per bit	Low	High
Capacity	High	Low
Write Speed	High	Low
Read Speed	Medium	High
Code Execution	Hard	Easy
Active Power	Low	High
Stand by Power	Low	Medium
file storage use	Easy	some how hard

- **Cost-per bit**: NAND flash memory has low cost per bit compared to NOR.
- **Capacity**: NAND flash has high density and capacity compared to NOR.
- **File storage usage**: usage of file storage is easy in NAND device and hard in NOR, this may be the reason of widely used flash drive.
- **code of execution**: code of execution is easy in NOR and hard in NAND , that is why , in portable device its easily used to execute.
- **read speed**: read speed is faster in NOR flash that is why its being used in portable device.
- **Write speed:** write speed faster in NAND flash device.
- **Consumption of Power**: Power consumption in NOR is less and read is faster that

2. Why NAND flash memory is being used most widely?

NAND is most widely used flash memory because of few reasons:

- low cost per bit of storage
- NAND flash has high capacity and high density compared with NOR
- storage of file is easy

3. Why NOR flash memory used in Portable devices?

NOR Flash memory is being used in portable devices like mobile, smart phone, palmtop, tab, handy devices etc due to following reasons or you can say following are the main reason.

- **execution of code** directly in NOR flash memory is very easy and hard in NAND flash memory
- **Faster read**: read is faster in NOR
- **consumes less power**: consumption of power in NOR flash memory is less compared to NAND

on average if you observed carefully, read and write ration in any application is 80:20 approx. So in portable devices read operation is more that the write.

Solid State Drive(SSD)

- SSD is an PC storage device that uses Solid State memory to store information.
- SSD uses non volatile NAND Flash Memory , which enable[illegible] to retain data when the po[illegible] removed.

NAND Flash Memory

- NAND Flash Memory is the key component of SSD.
- It is a specific type of EEPROM chip.
- It has a grid of columns and rows with a cell that has two transistors at each intersection as Control gate and Floating gate transistor.
- The principle of operation is based on MOSFETs.

Types of NAND Flash

- Single Level Cell (SLC)
- Multi Level Cell (MLC)

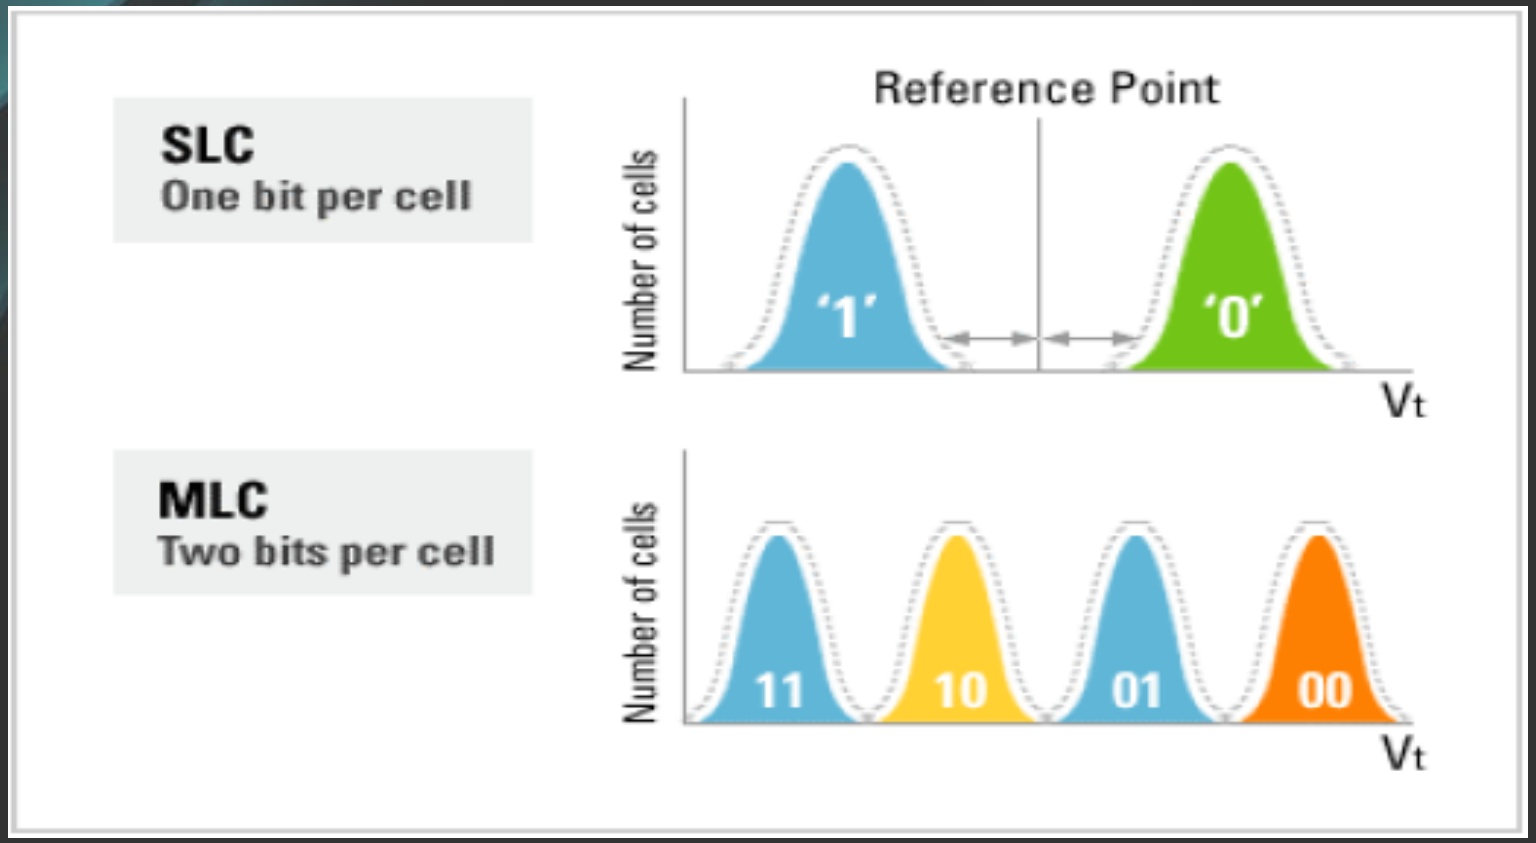

SLC vs. MLC

ITEM	SLC	MLC
Voltage	3.3V/1.8V	3.3V
Chip Size	0.12um	0.16um
Page Size/Block Size	2KB/128KB	512B/32KB or 2KB/256KB
Access Time(Max.)	25us	70us
Endurance	100K	10K
Cost per Bit	more	Comparatively less
Write Data Rate	8MB/s+	1.5MB/s

SSD Elements

SSD consists of basically three elements:

- SSD Controller.
- SSD Flash.
- SSD Interface.

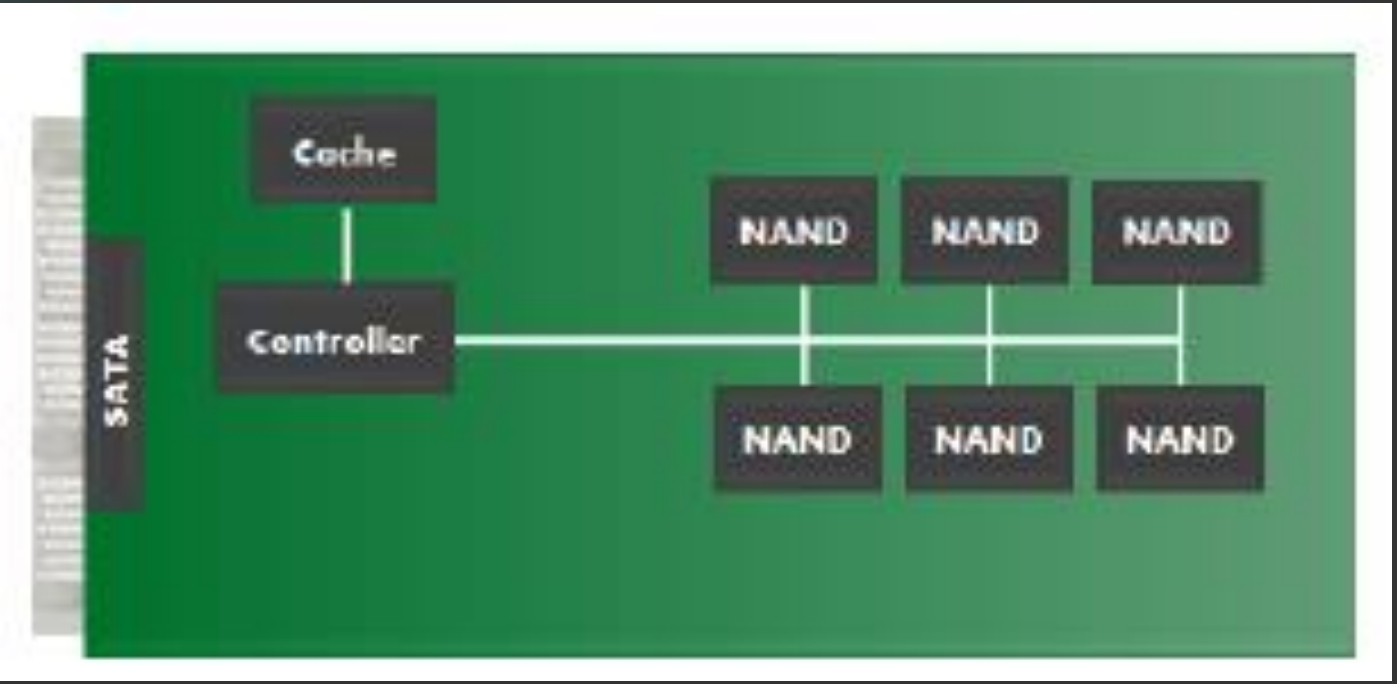

SSD Controller

- Flash controller includes the electronics that bridge the Flash memory components to the SSD input/output interfaces.
- The controller is an embedded processor that executes firmware-level software.

SSD Controller Block Diagram

- Processor
- ECC
- Flash Controller
- DRAM Controller
- I/O Interface
- Controller Memory
- Chip Configuration

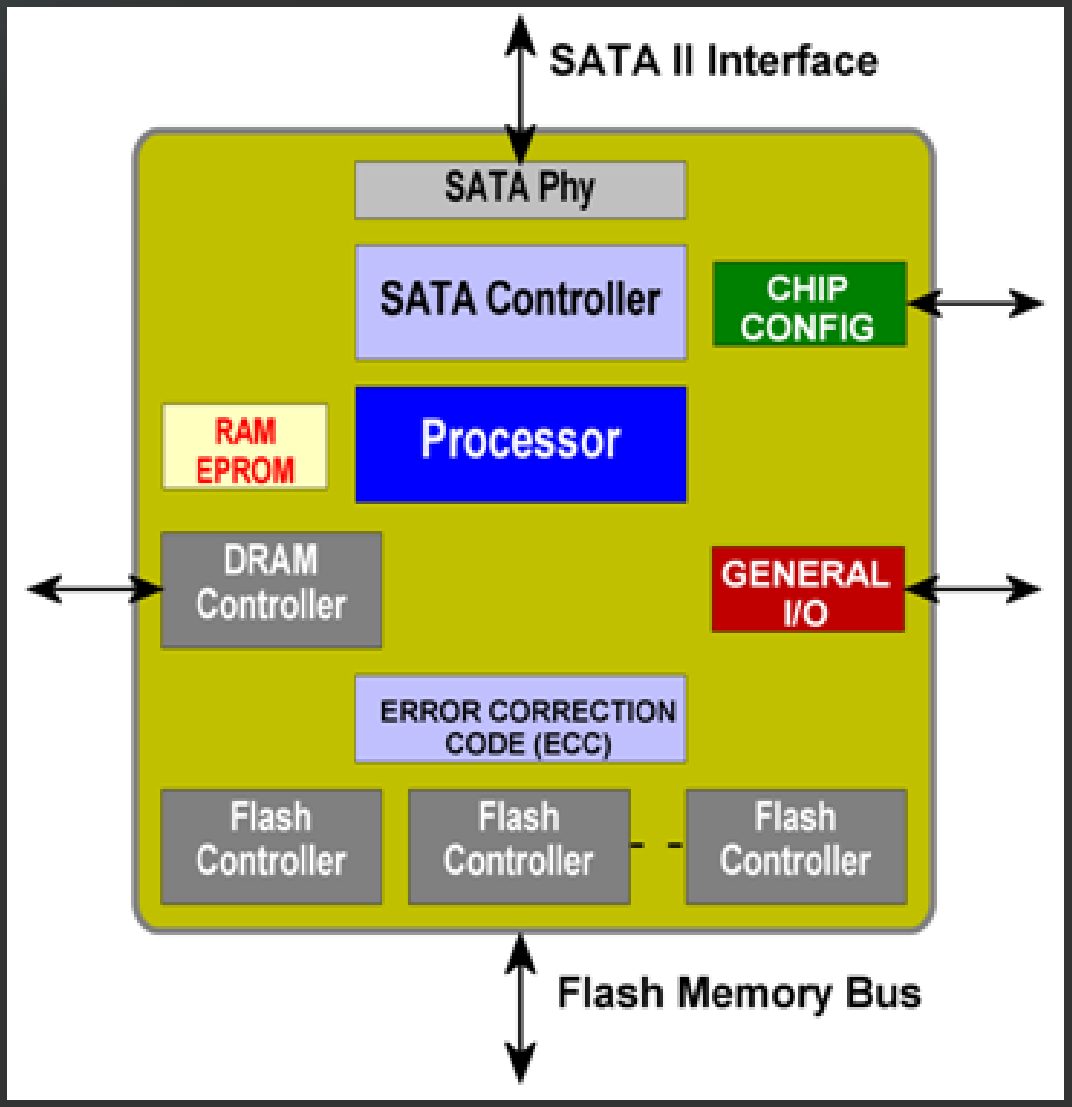

SSD Flash

- Solid State Flash memory is available using NAND technology.
- NAND Flash components come in densities from 1Gb (gigabit) to 64Gb per chip.
- NAND Flash components have structures called pages and blocks.
- There is an Error Correction Code (ECC) associated with each sector.
- NAND Flash has a limited rewrite endurance of about 1,000,000 times per block.

SSD Flash Block Diagram

- Data Interface
- I/O Controller
- Control Logic
- Address Register
- Data/Cache register
- Status register
- Row/Column Decode
- Flash Array

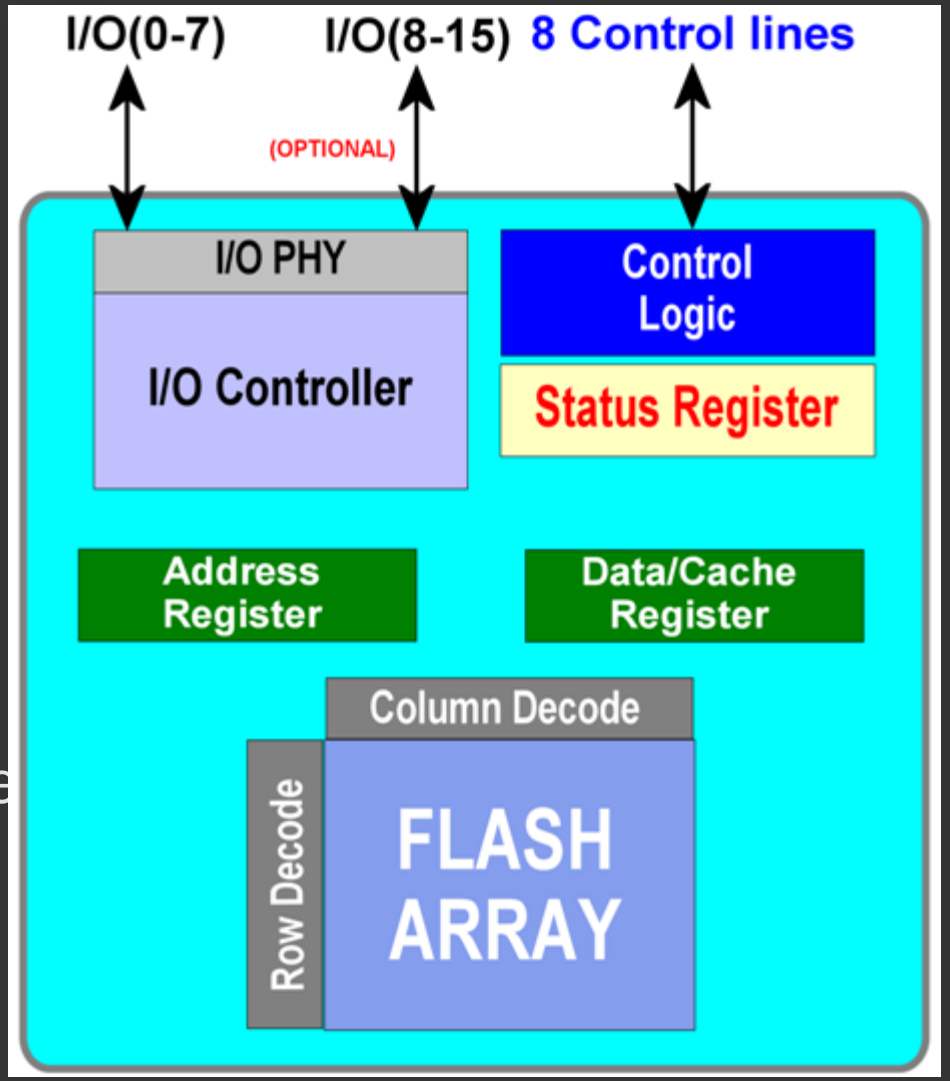

SSD Interface

- Interface is used to connect the SSD with the computer.
- Also, since SDDs are generally used in conjunction with magnetic disk drives, a common mass storage bus interface is used in most cases.
- SSD are available with a variety of system interfaces based primarily on the performance requirements:
  - Serial ATA
  - Serial attached SCSI
  - Fiber Channel
  - USB

SSD Storage Scheme

- Groups of NAND flash cells are organized into pages and these pages are organized into blocks.
- Read and write operations can be performed on pages, but erase operations can only be performed at the block level.
- This means that when rewriting a page, the entire block must be erased first.
- The SSD controller manages this process.

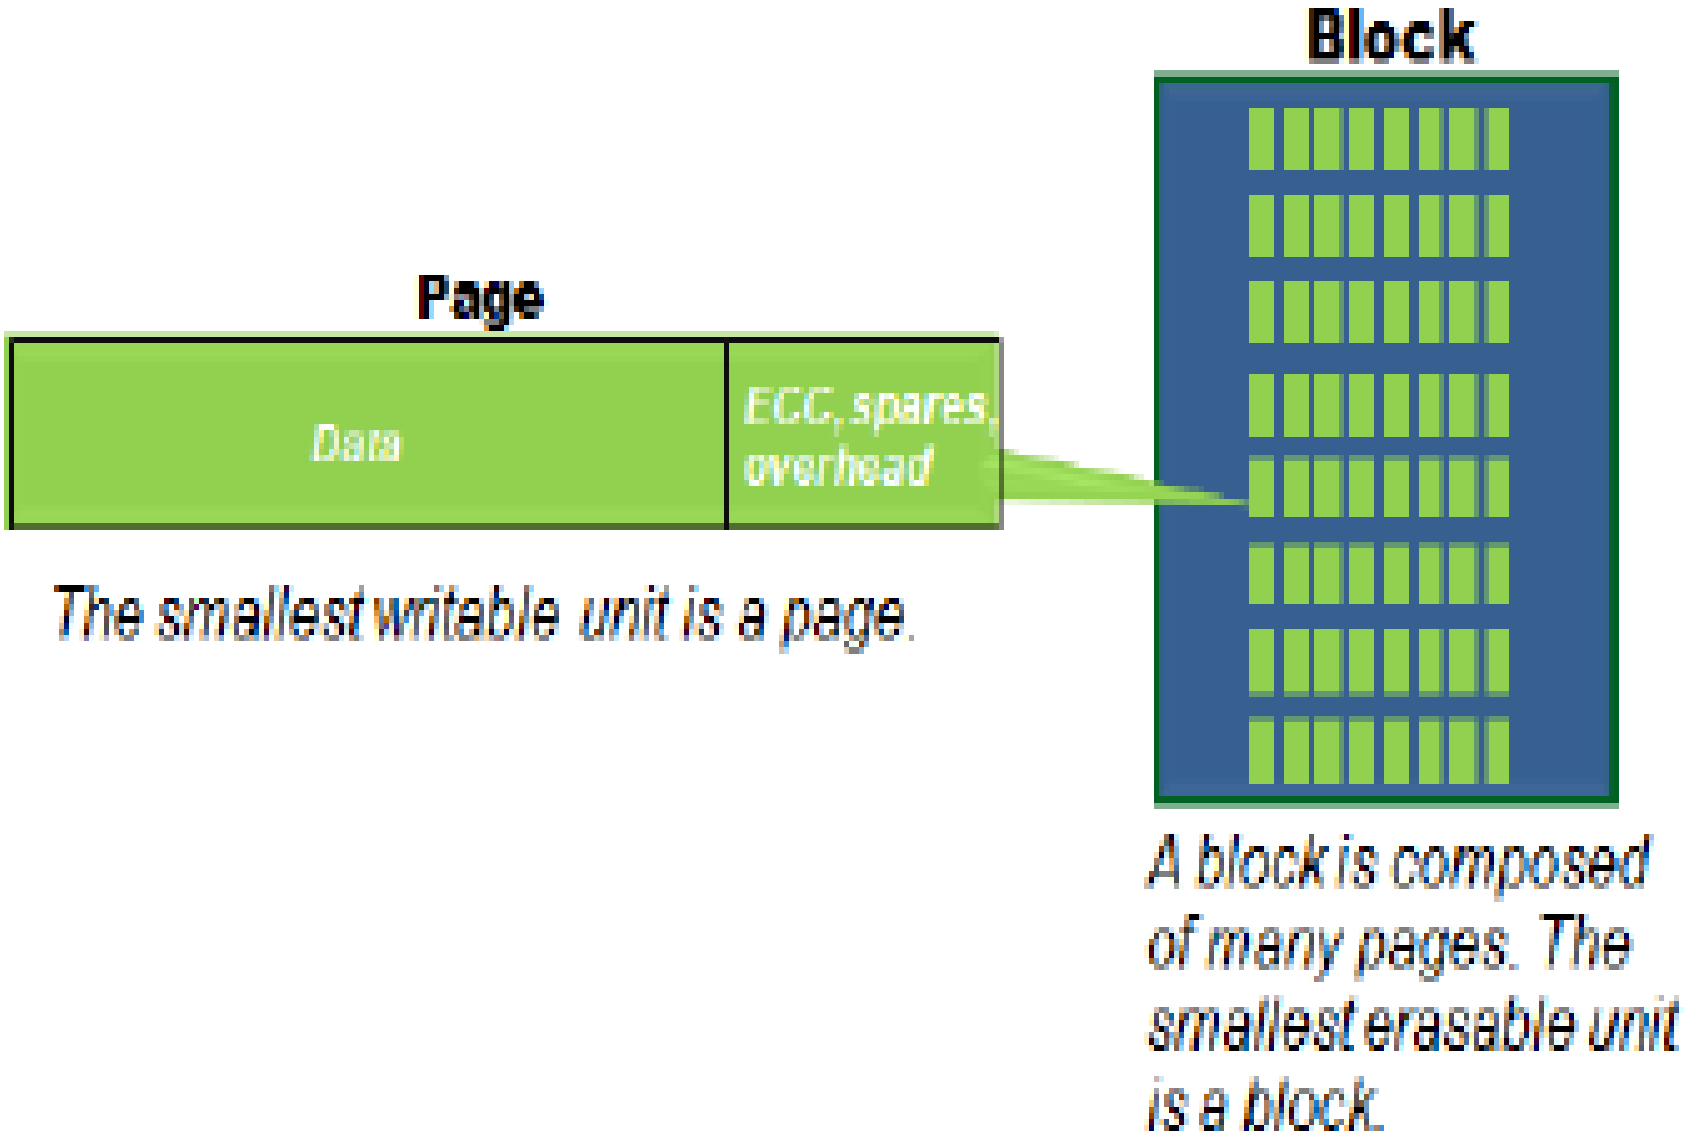
Block
Page
Data
ECC, spares, overhead
The smallest writable unit is a page.
A block is composed of many pages. The smallest erasable unit is a block.

SSD Around Us

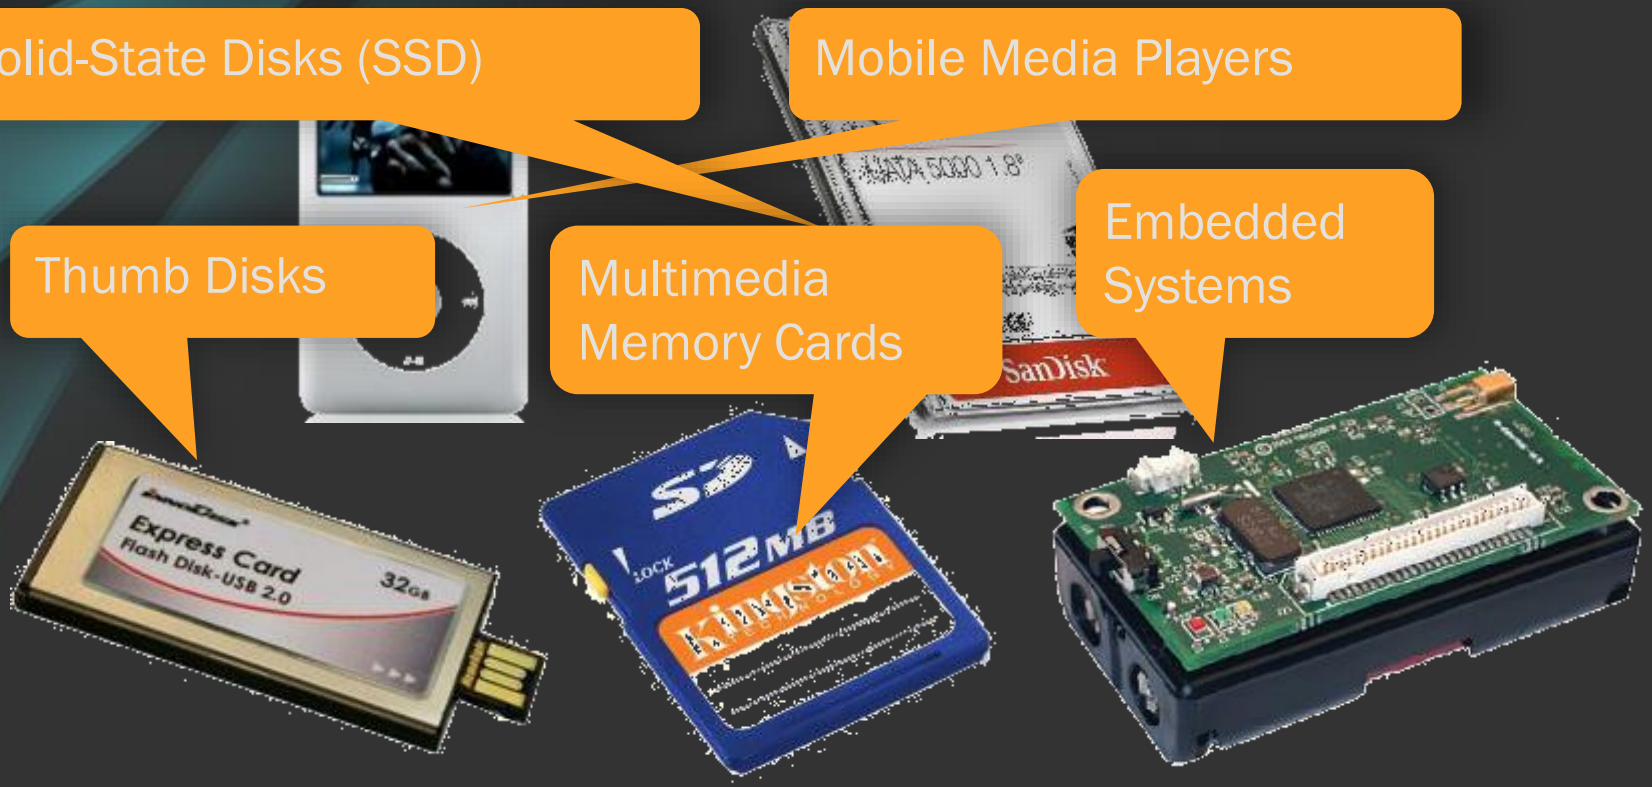

Intel x25 SSD- An Example

- Certified for 25 MB/s read speed and 70 MB/s write speed.
- MLS SSD standard.
- 10-channel memory controller.
- Each channel is “responsible” for two memory chips.
- http://www.insidehw.com/Reviews/Storage/Intel-X25-M-SSD.html

Intel x25 SSD

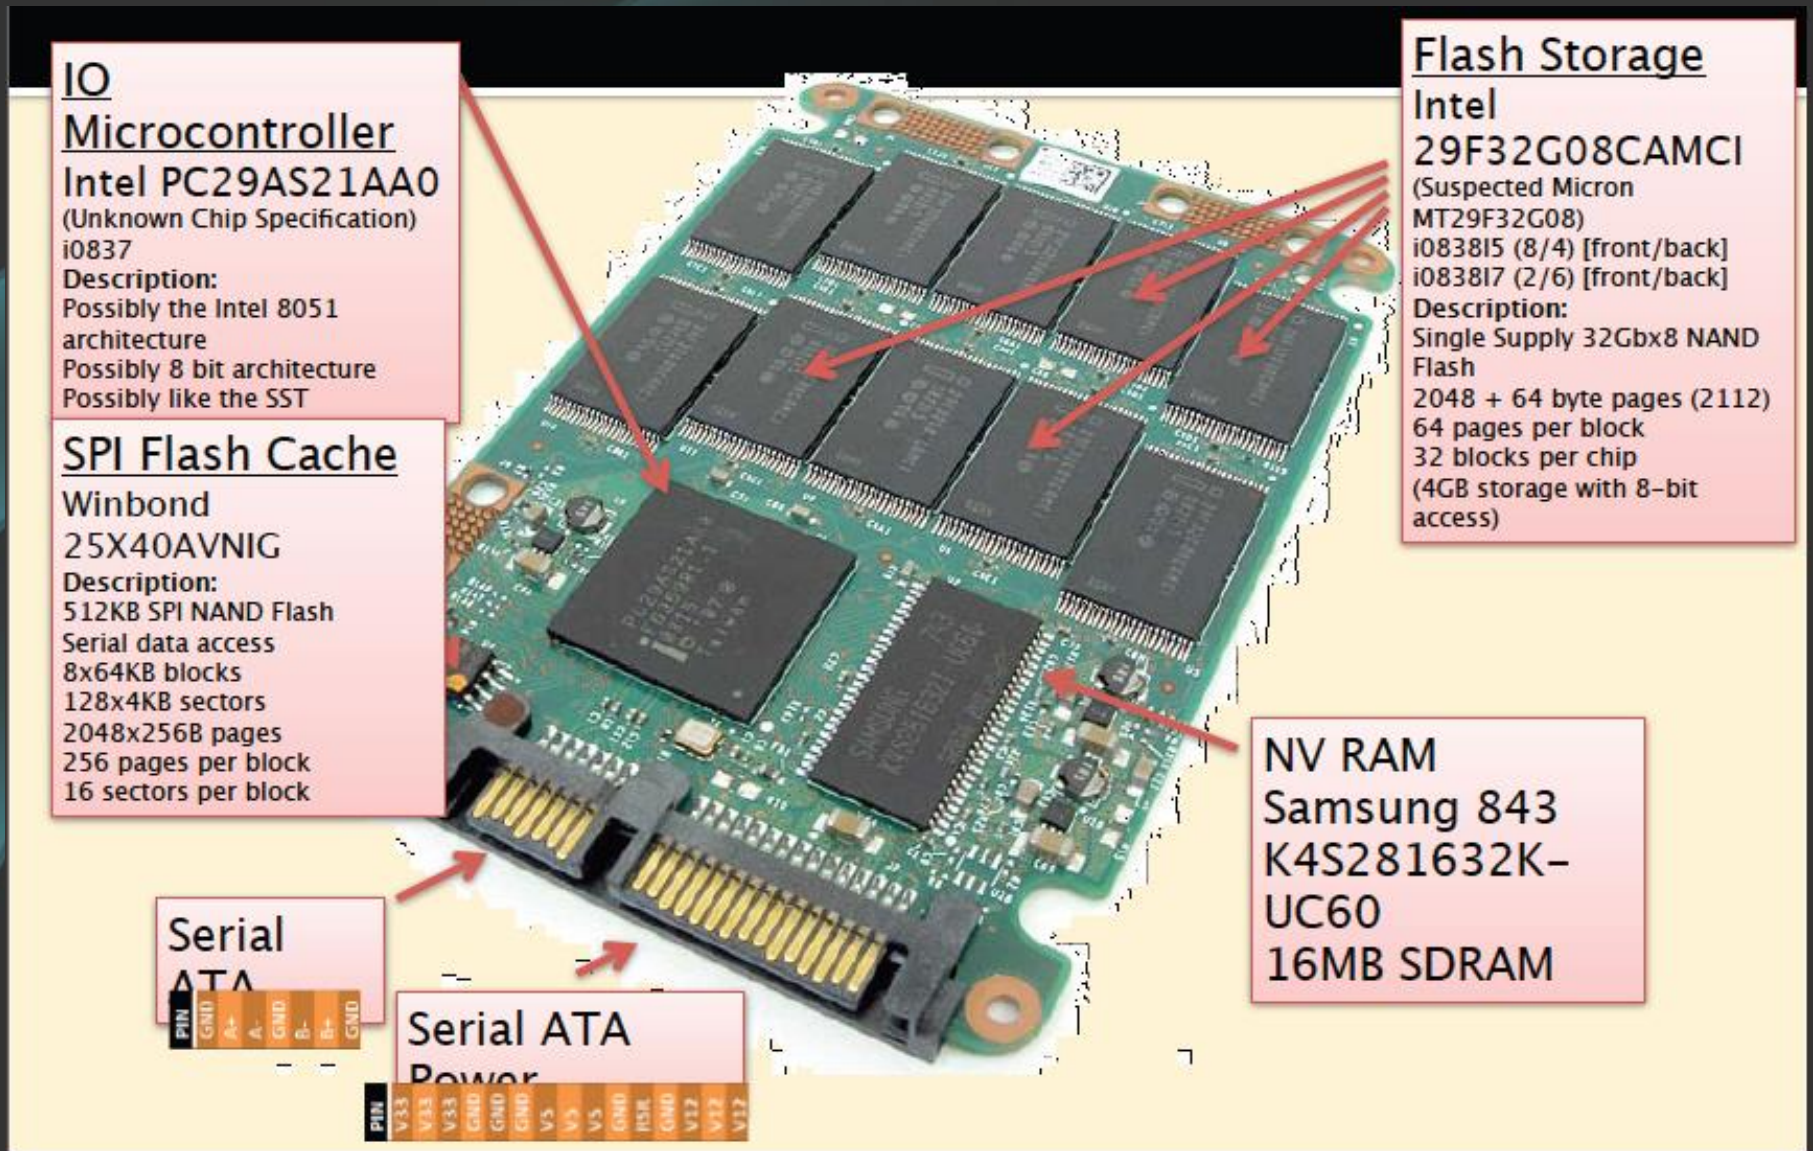

Differences

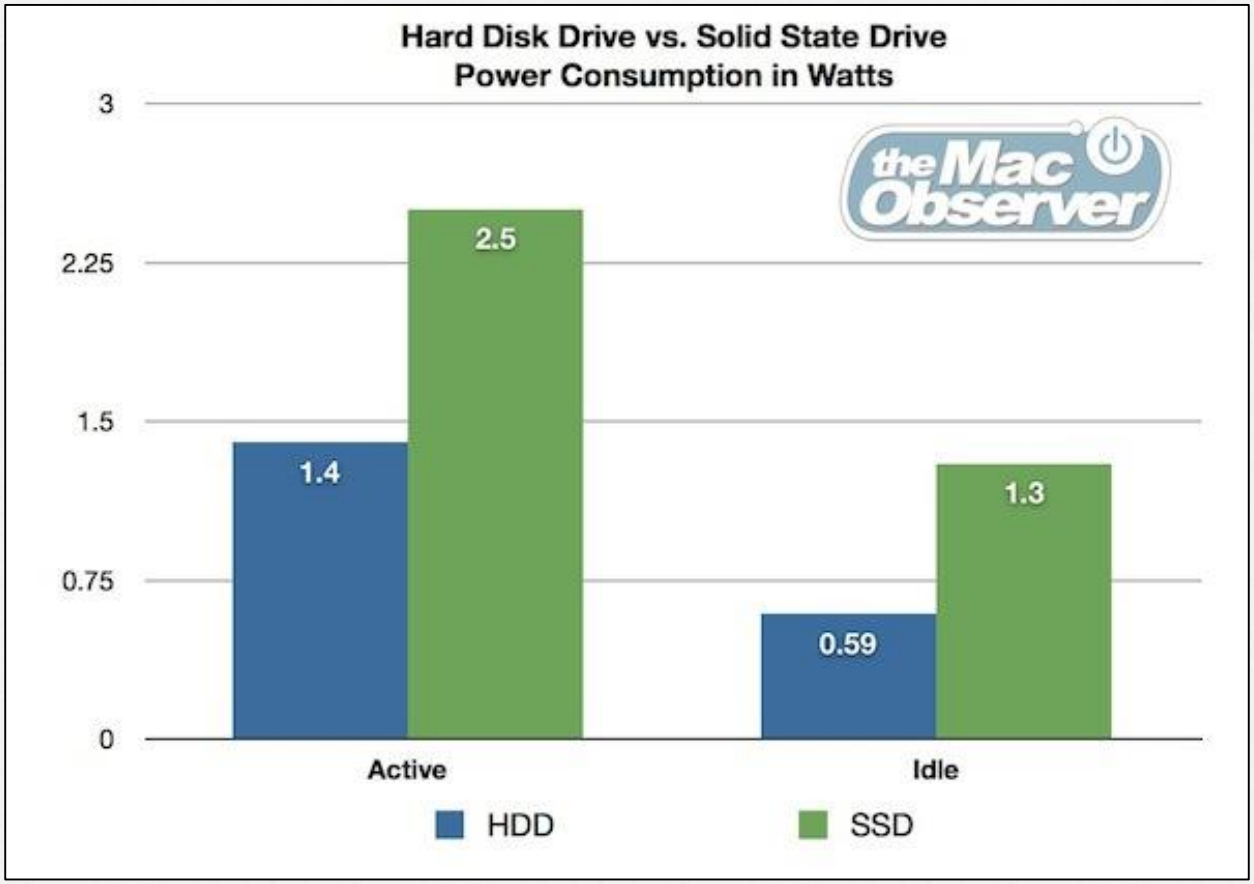

Differences

- F
- H

Make and model	Type	Capacity	Total cost	Price per GB
OCZ Vertex	SSD	128GB	£249.99 €276.10	£2.30 €2.54
Intel X25-M (160)	SSD	160GB	£344.99 €381.03	£2.15 €2.37
OCZ Agility	SSD	128GB	£229.99 €254.01	£1.80 €1.99
Intel X25-M (80)	SSD	80GB	£179.99 €198.79	£2.24 €2.47
WD VelociRaptor	HDD	300GB	£164.99 €182.23	£0.55 €0.61
WD RE4-GP	HDD	2TB	£215.99 €238.55	£0.11 €0.12

Primary Advantages

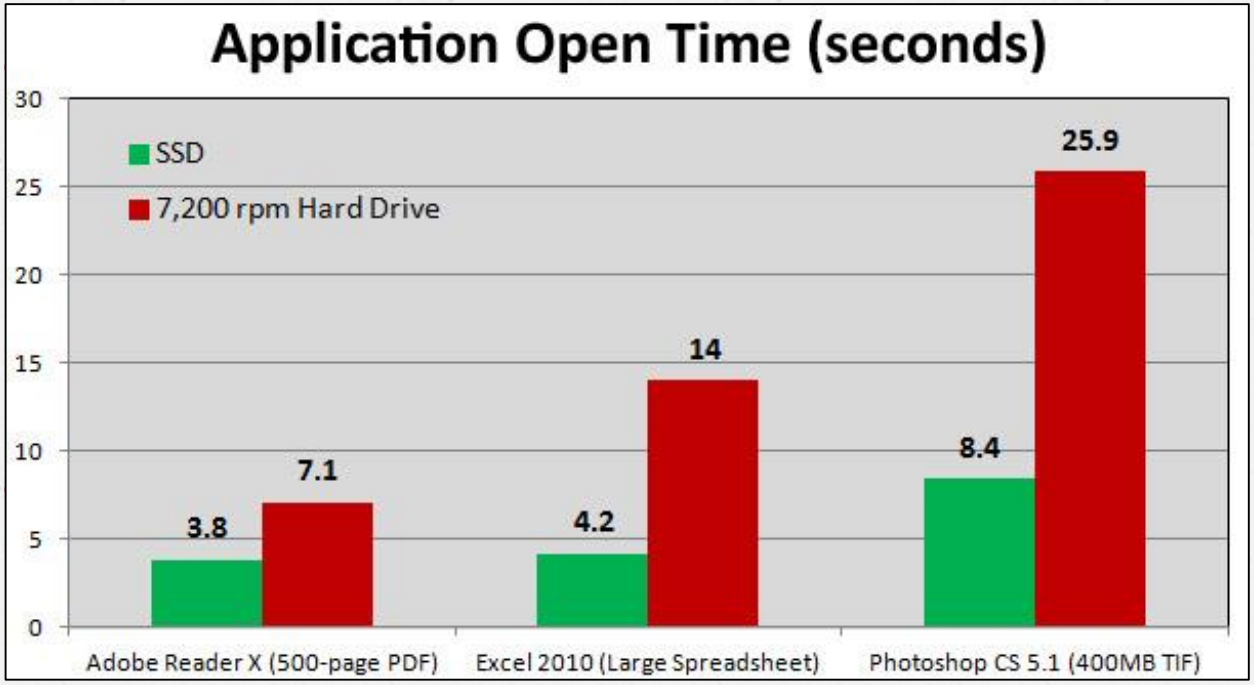

Primary Advantages

- Reasons why HDD lost in performance:
  - More time spent waiting for platters to access data
- Reasons why SSD lost in storage capacity:
  - Limitations in size of transistors and parts in SSD

Research Gap

- Little has shown:
  - Temperature fluctuation at different intervals
  - Statistical figures of boot up and read/write speed with a fair environment

Method

- The aim of the experiments were to obtain data to evaluate the performance of each of the drives.

- Three experiments:
  - Boot-up speed
  - Read & Write speed
  - Operating temperature

Materials

- 13-inch Apple MacBook Pro x2
- 256GB SSD
- 500GB HDD
- Stopwatch

Materials (Cont'd)

- **Blackmagic Disk Speed Test**
  - A software used to assess storage drives' read and write performance.
  - It obtains the drives' performance by comparing the time taken to write and read large data to and from the drive.

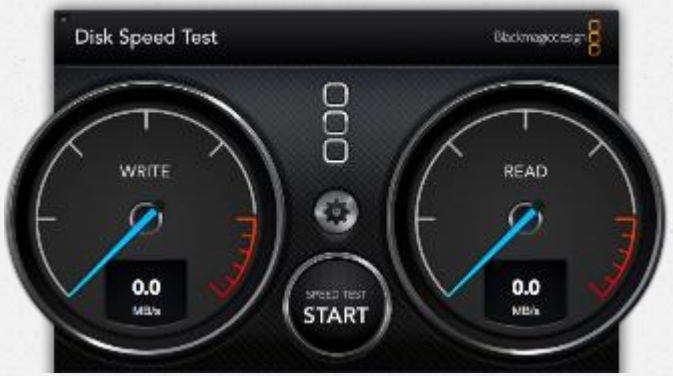

Materials (Cont'd)

- **iStat Menus**
  - A software that allow users to monitor detailed information of the computer in real time.
  - It makes use of the sensors built inside of the computer to obtain the hardware information.

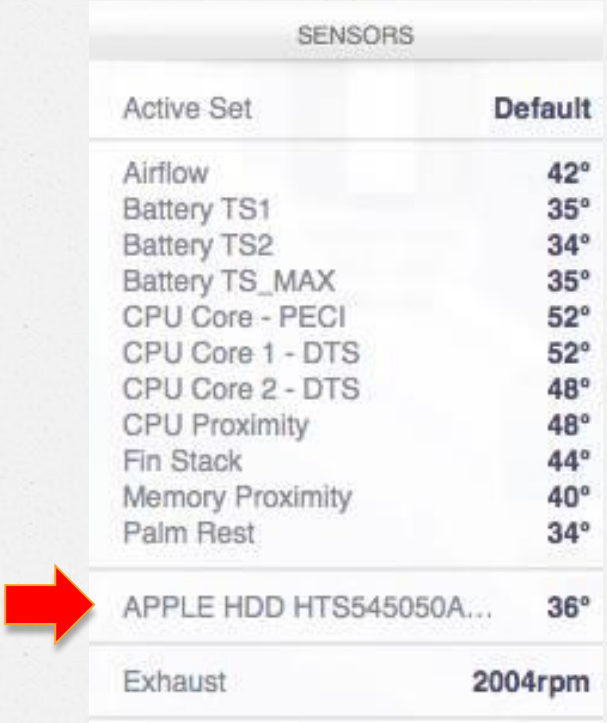

Procedures

1. Boot-up speed

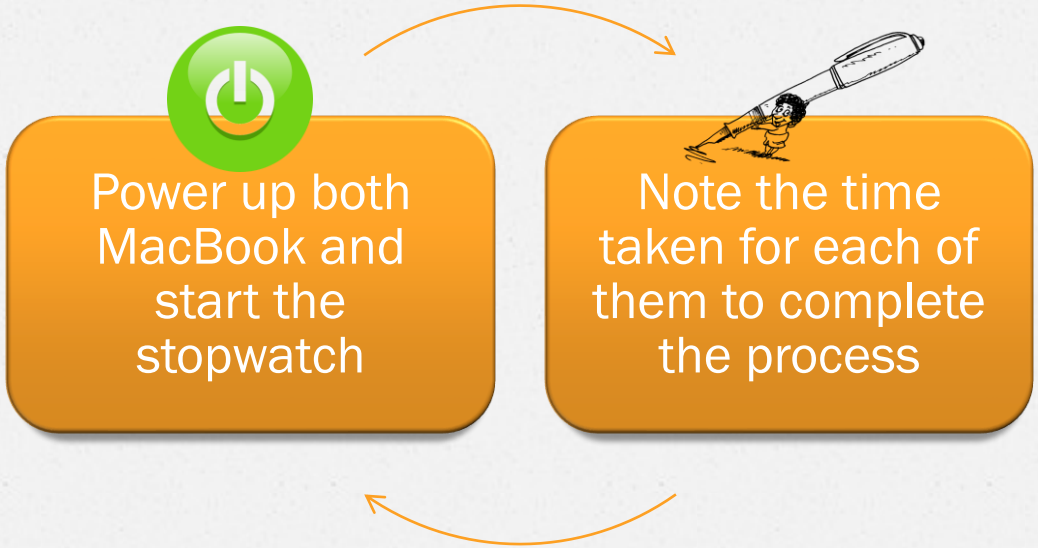

Finally, obtain the average from the results gathered.

Procedures (Cont'd)

2. Read & Write speed

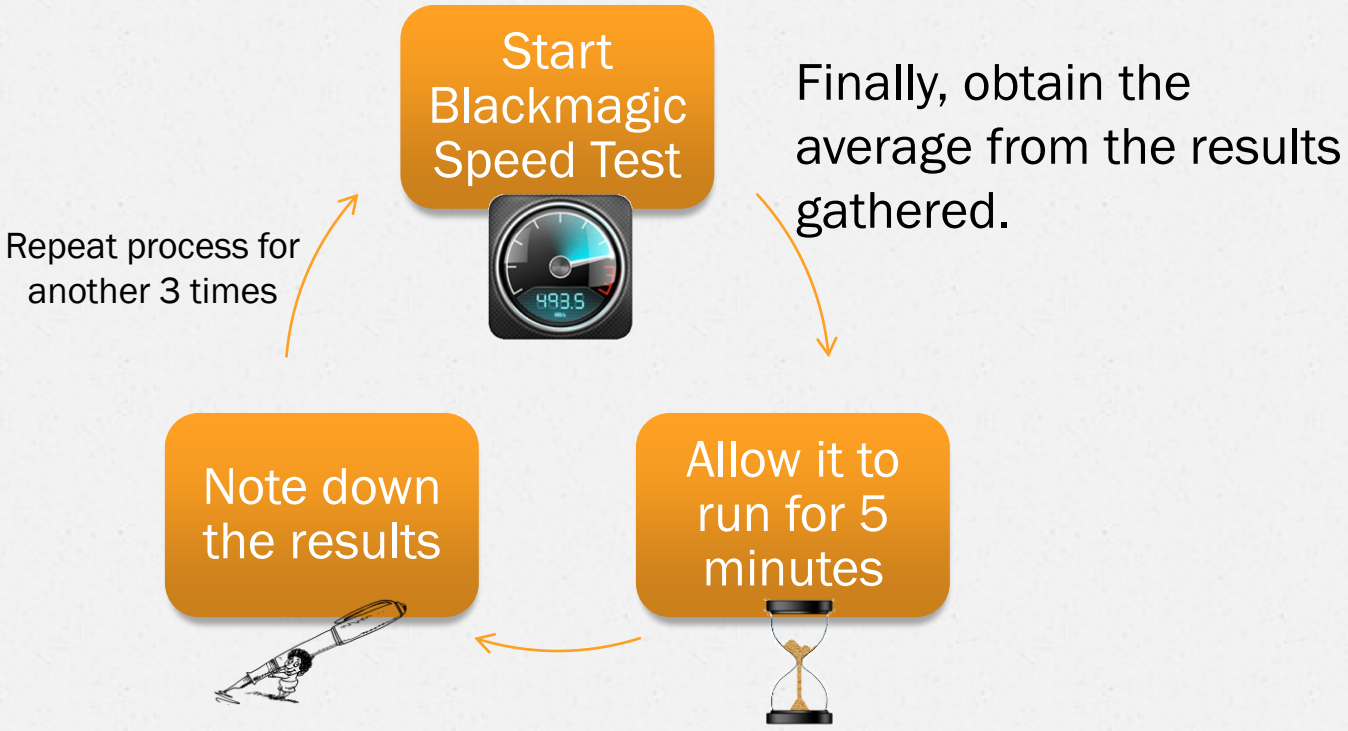

Procedures (Cont'd)

3. Operating temperature

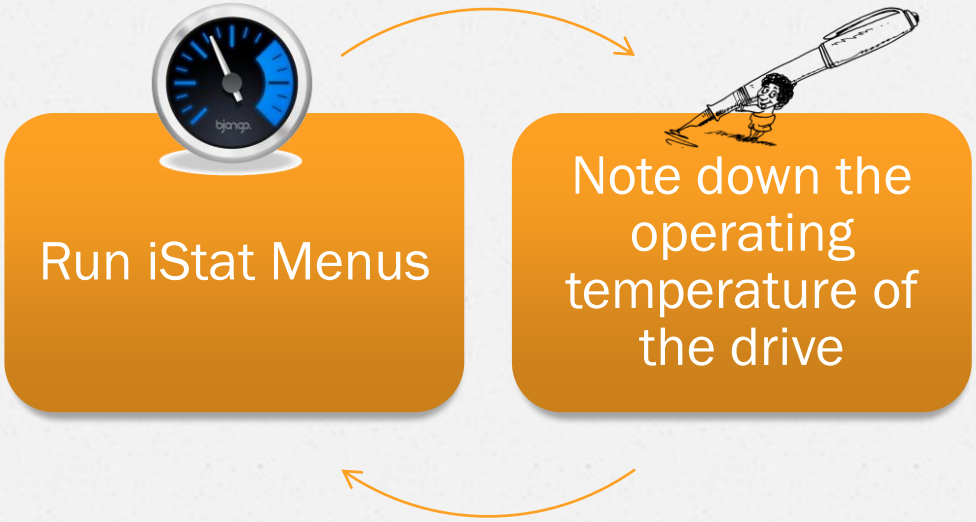

Repeat process for another 3 times every 15 minutes

Finally, obtain the average from all the results.

Results

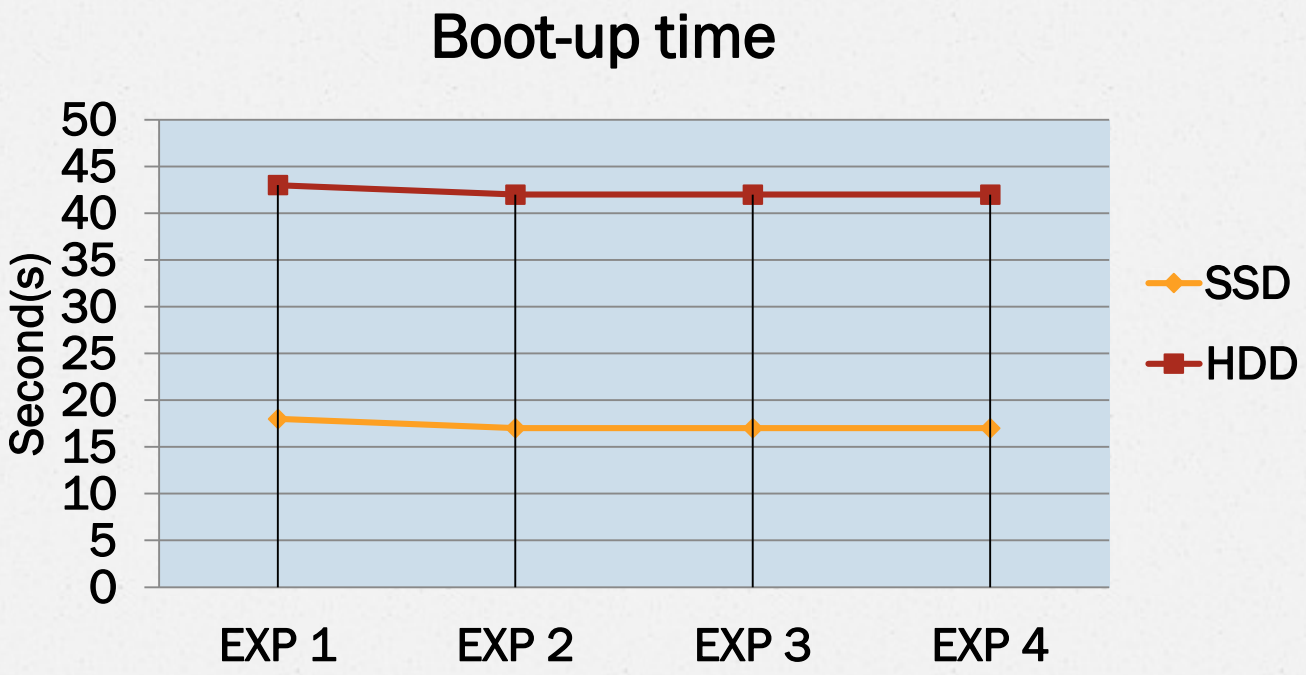
Boot-up time
Second(s)
50
45
40
35
30
25
20
15
10
5
0
EXP 1
EXP 2
EXP 3
EXP 4
SSD
HDD

Results (Cont'd)

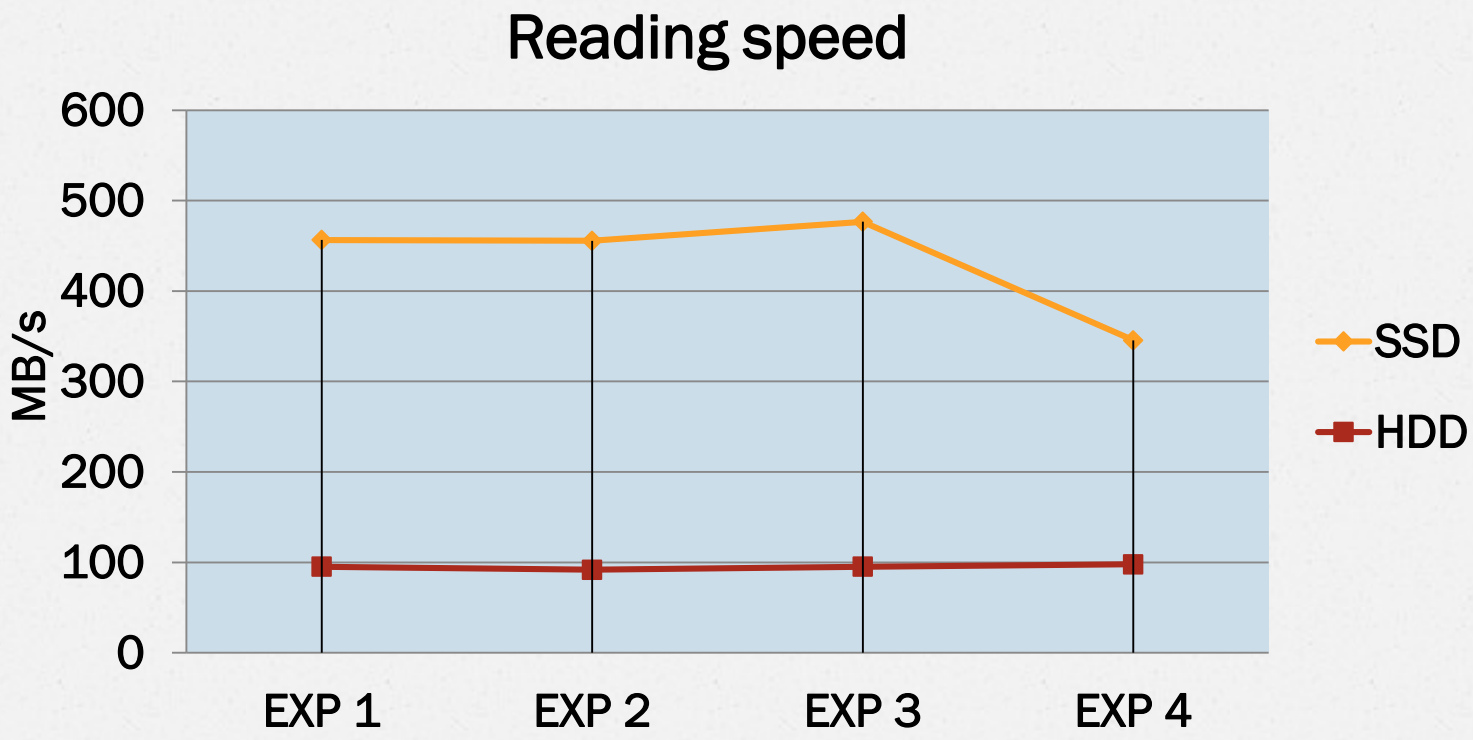

Results (Cont'd)

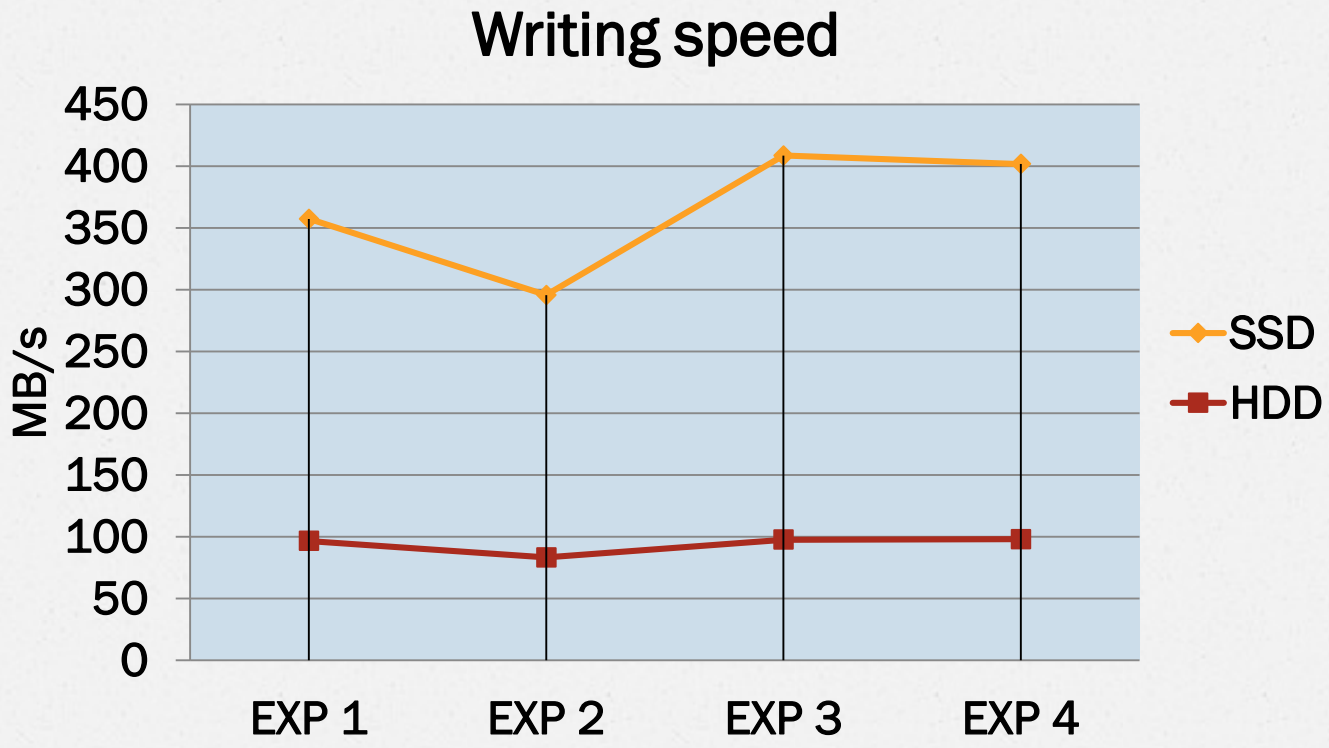

Results (Cont'd)

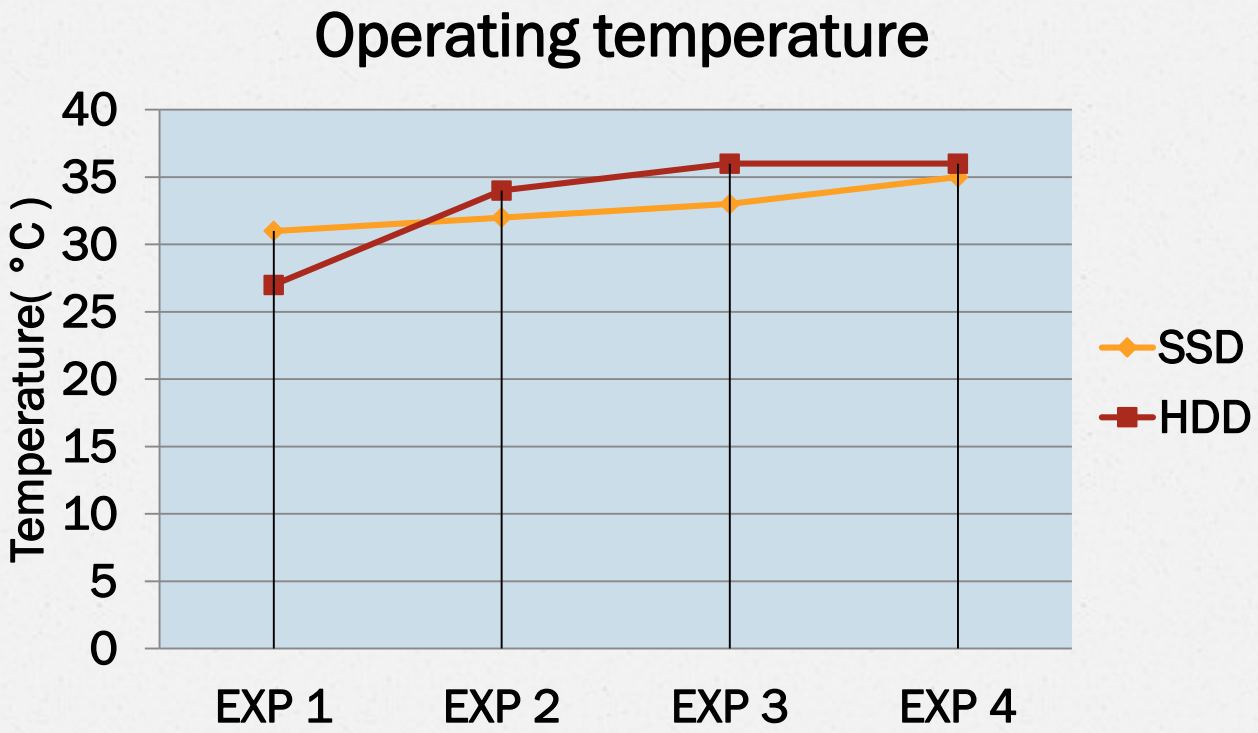

Final results

Result	SSD	HDD
Boot-up	17.2sec	42.2sec
Read	443MB/sec	95MB/sec
Write	366MB/sec	94MB/sec
Temperature	33°C	34°C

Why SSD is better than HDD?

Reason: The difference in technology.

HDD

SSD

Conclusion.....

1. Limitation
2. Recommendation
3. Application

Final results

Result	SSD	HDD
Boot-up	17.2sec	42.2sec
Read	443MB/sec	95MB/sec
Write	366MB/sec	94MB/sec
Temperature	33°C	34°C

New Technologies

- SATA Express

Com o mesmo conector, o usuário pode plugar dois HDs/SSDs SATA comuns ou usar um conector único SATAe que encaixa nas três entradas, acessando o conector de 4 pinos (localizado ao lado dos conectores SATA na imagem acima), passando a acessar as linhas PCI Express da placa-mãe. Usando o PCI Express 2.0, as velocidades máximas passam para cerca de 780 MB/s, chegando até 1580 MB/s com o PCIe 3.0 (contra cerca de 550 MB/s em SSDs SATA), mostrando o real benefício de usar o SATA Express.

New Technologies

- SATA Express

Why not SATA IV?

Doesn't support 1200 Mbps
Increases heating
Increases power consumption

New Technologies

- M.2

Assim como o SATA Express, o SATA M.2 é uma implementação híbrida, que pode usar tanto o PCI Express quanto o SATA, só que é ainda mais versátil. Para começar, não existe um tamanho específico para SSDs M.2 (antes conhecido como Next Generation Form Factor – NGFF), com larguras que variam de 16 mm até 110 mm (com o modelo de 22 mm sendo o mais comum), e comprimentos que variam de 30 mm até 110 mm, assim como as velocidades máximas de transferência.

New Technologies

- M.2

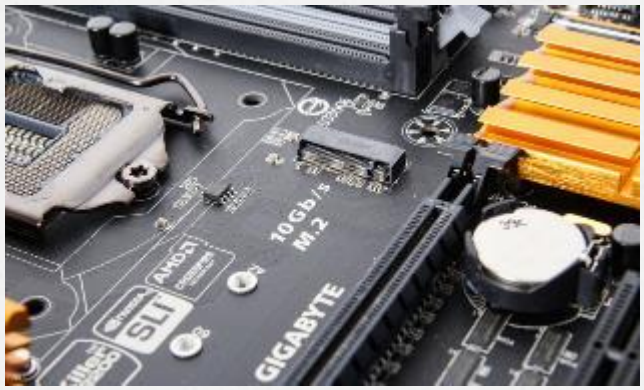

NVMe

NVMe (Non-Volatile Memory Express) é uma interface e unidade de comunicações que define um conjunto de comandos e recursos para SSDs com base em PCIe com a finalidade de aumentar o desempenho e a interoperabilidade em uma ampla variedade de sistemas empresariais e de clientes.

NVMe foi projetado para SSDs. Ele faz a comunicação entre a interface de armazenamento e a CPU do sistema utilizando soquetes PCIe de alta velocidade, independentemente do formato de armazenamento. As tarefas de Entrada/Saída realizadas utilizando drivers NVMe são iniciadas mais rapidamente, transferem mais dados e são concluídas com mais rapidez do que nos modelos de armazenamento mais antigos utilizando drivers mais antigos, como AHCI (Advanced Host Controller Interface), uma característica dos SSDs SATA. Como foi projetada especificamente para SSDs, NVMe está se tornando o novo padrão do setor para servidores no datacenter e em dispositivos de clientes como laptops, PCs e até consoles de jogos de última geração.

A tecnologia NVMe está disponível em uma série de formatos como o slot de cartão PCIe, M.2 e U.2. Embora existam SSDs que usam o slot PCIe, SATA e M.2 que são AHCI e não NVMe, U.2 é um formato que usa exclusivamente o protocolo NVMe.

NVMe


Unidades de comunicação
Usadas pelos sistemas operacionais para a comunicação de dados com os dispositivos de armazenamento
AHCI
Projetadas para Discos Rígidos com tecnologia Spinning Disk
1
Possui apenas 1 fila de comando
32
Pode enviar apenas 32 comandos por fila
Os comandos utilizam ciclos de alta utilização de CPU, ou High CPU
A latência é de 6 microssegundos
Deve se comunicar com o controlador SATA
IOPs até 100K
NVMe
Projetadas para SSDs com tecnologia Flash
64K
Possui 64K comandos por fila
Pode enviar 64K comandos por fila
Os comandos utilizam ciclos de baixa utilização de CPU, ou Low CPU
A latência é de 2,8 microssegundos
Comunica-se diretamente com a CPU do Sistema
IOPs acima de 1 mihão

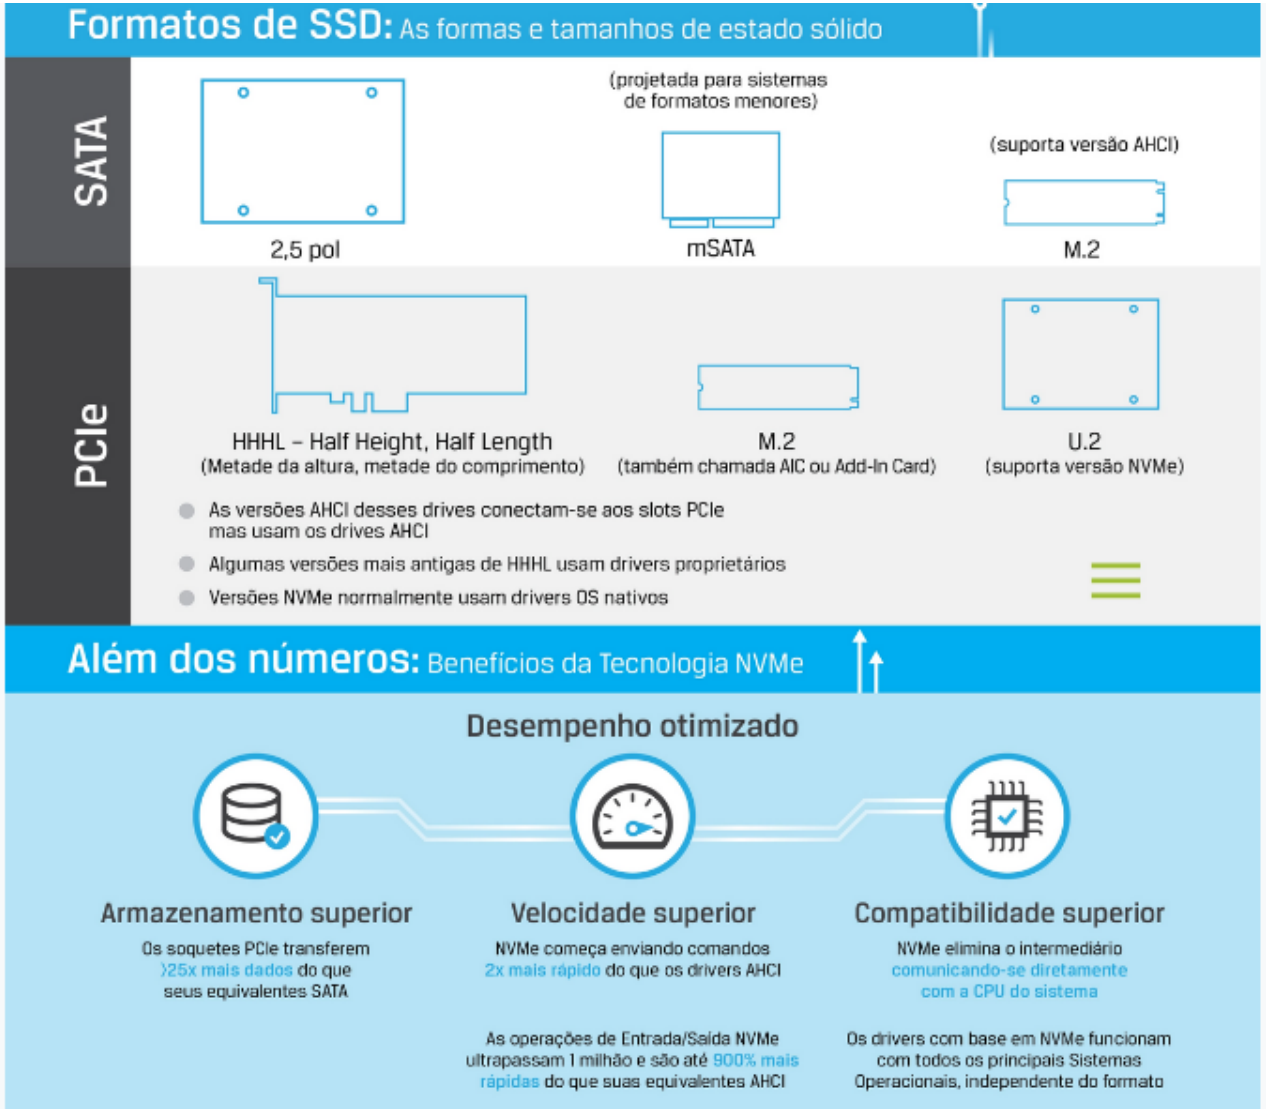
Formatos de SSD: As formas e tamanhos de estado sólido
SATA
(projetada para sistemas de formatos menores)
(suporta versão AHCI)
2,5 pol
mSATA
M.2
PCIe
HHHL – Half Height, Half Length
(Metade da altura, metade do comprimento)
M.2
(também chamada AIC ou Add-In Card)
U.2
(suporta versão NVMe)
As versões AHCI desses drives conectam-se aos slots PCIe mas usam os drives AHCI
Algumas versões mais antigas de HHHL usam drivers proprietários
Versões NVMe normalmente usam drivers OS nativos
Além dos números: Benefícios da Tecnologia NVMe
Desempenho otimizado
Armazenamento superior
Os soquetes PCIe transferem 25x mais dados do que seus equivalentes SATA
Velocidade superior
NVMe começa enviando comandos 2x mais rápido do que os drivers AHCI
As operações de Entrada/Saída NVMe ultrapassam 1 milhão e são até 900% mais rápidas do que suas equivalentes AHCI
Compatibilidade superior
NVMe elimina o intermediário comunicando-se diretamente com a CPU do sistema
Os drivers com base em NVMe funcionam com todos os principais Sistemas Operacionais, independente do formato

New Technologies

- Intel Optane

Uma das primeiras confusões a respeito dos Optane da Intel disponíveis atualmente está em associá-los com SSDs comuns. Na verdade, a Intel tem a intenção de lançar no futuro SSDs com grandes quantidades de espaço criados a partir da tecnologia 3DXpoint, que é a base dos produtos Optane.

O que existe é um tipo de acelerador de cache Optane. Na prática, as unidades que a Intel comercializa no momento, de 16 e 32 GB, servem como um cache de alta velocidade para o seu computador. O resultado é que elas fazem com que uma máquina que use apenas disco rígido magnético atinja velocidades equiparáveis a de um computador equipado com SSDs.

No resumo: um Intel Optane, atualmente, pode ser encontrado com 16 e 32 GB e se destina a fazer com que computadores que usam discos rígidos magnéticos, tenha desempenho similar ao de uma máquina com SSDs.

New Technologies

- Intel Optane