



# Low Power processors in HEP

Michele Michelotto – INFN Padova

# Intel Xeon v5 compare (V1, V2, V3) 32 bit

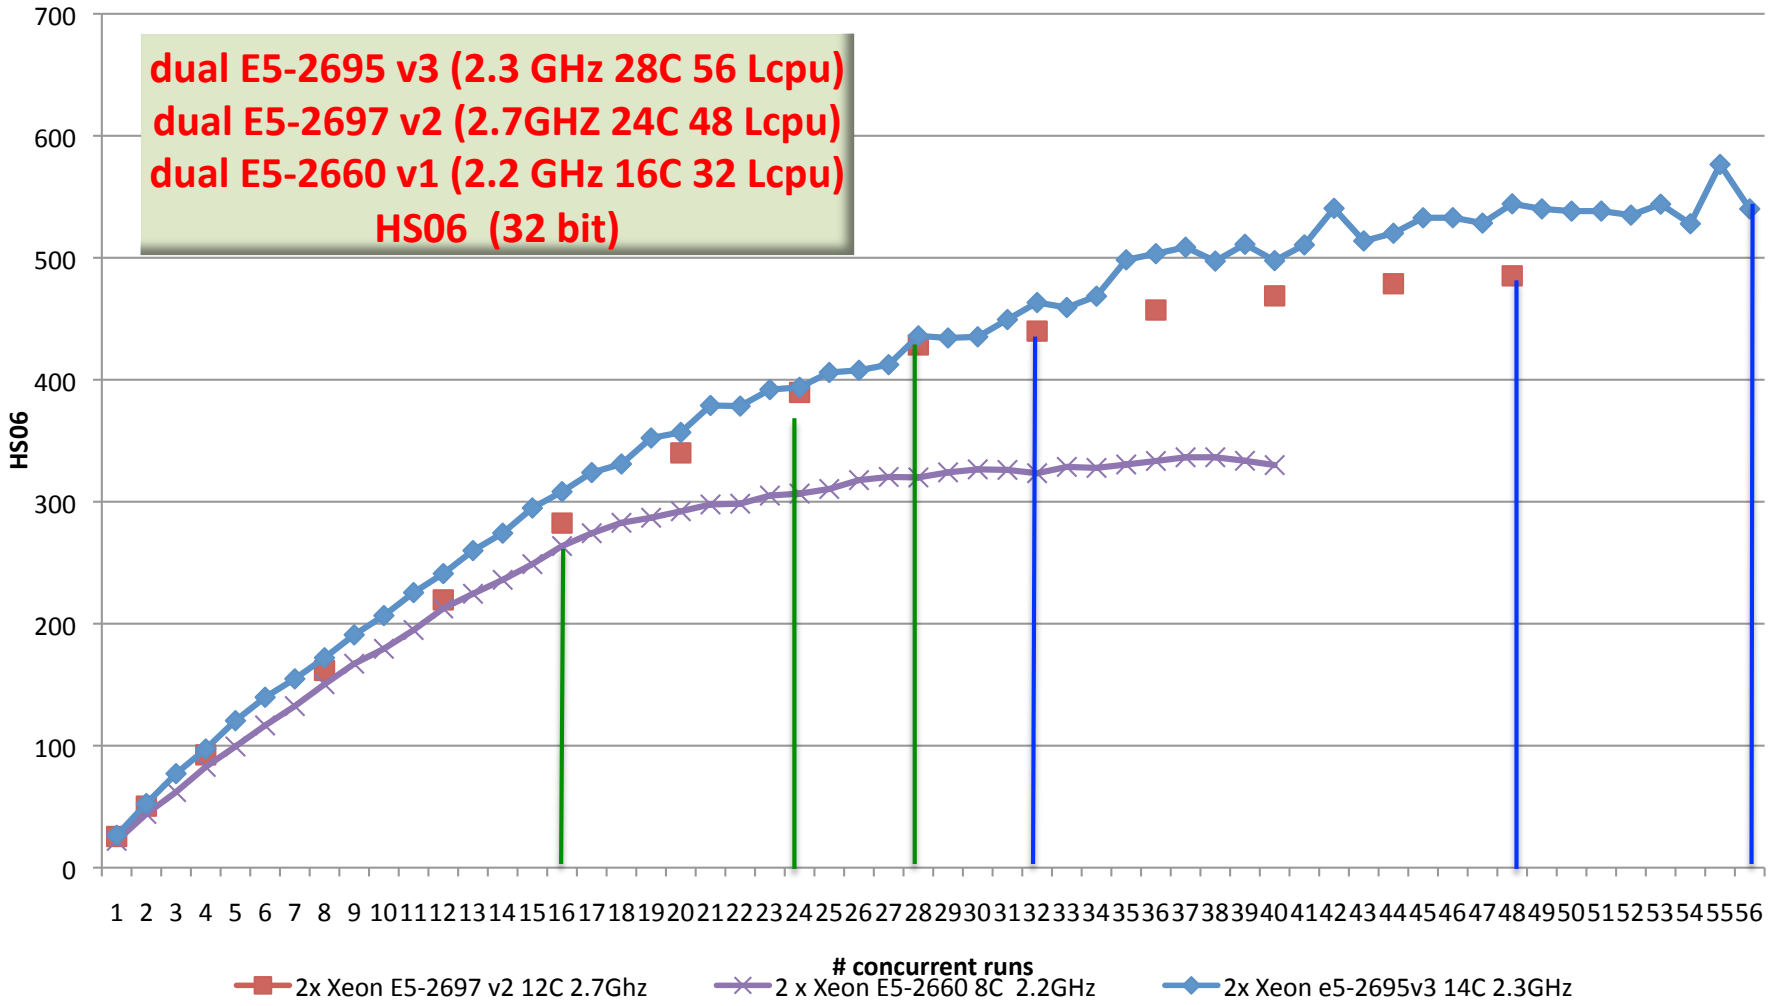

# Intel Xeon v5 compare (V1, V2, V3) clock normalized (HS06/GHz)

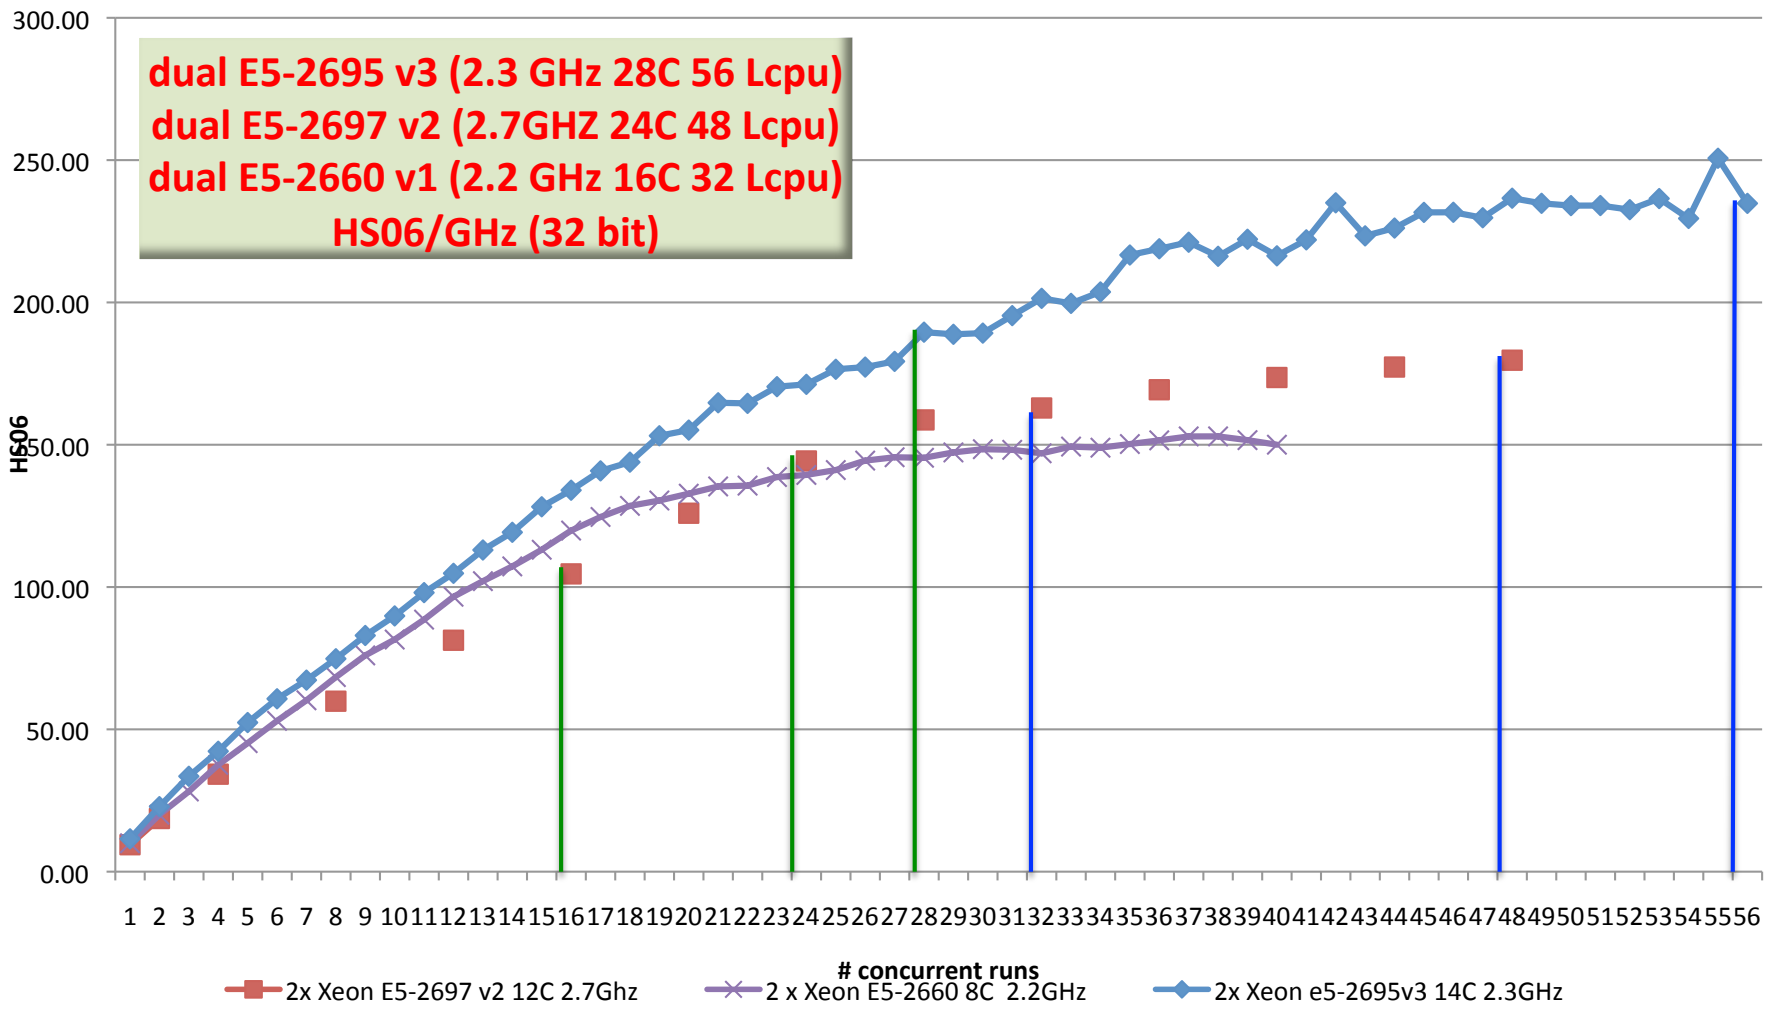

# How many HS06 I get from a new job when N-1 core are loaded?

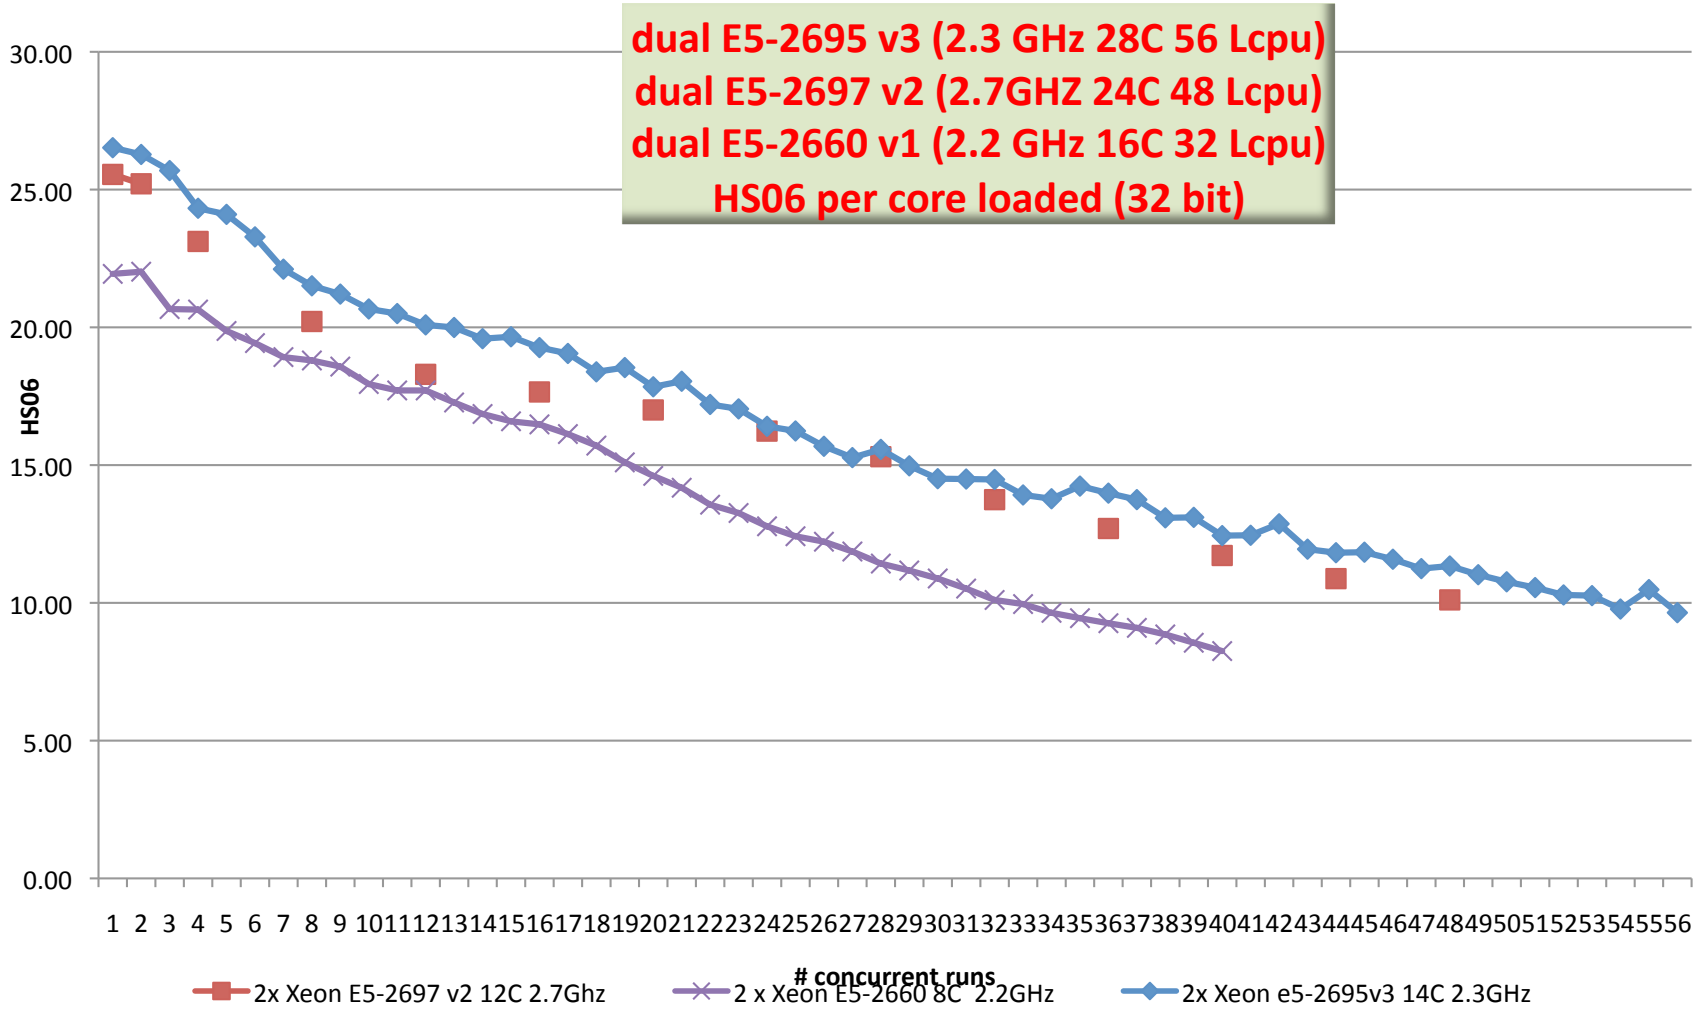

# HS06 vs HS06 compiled at 64 bit

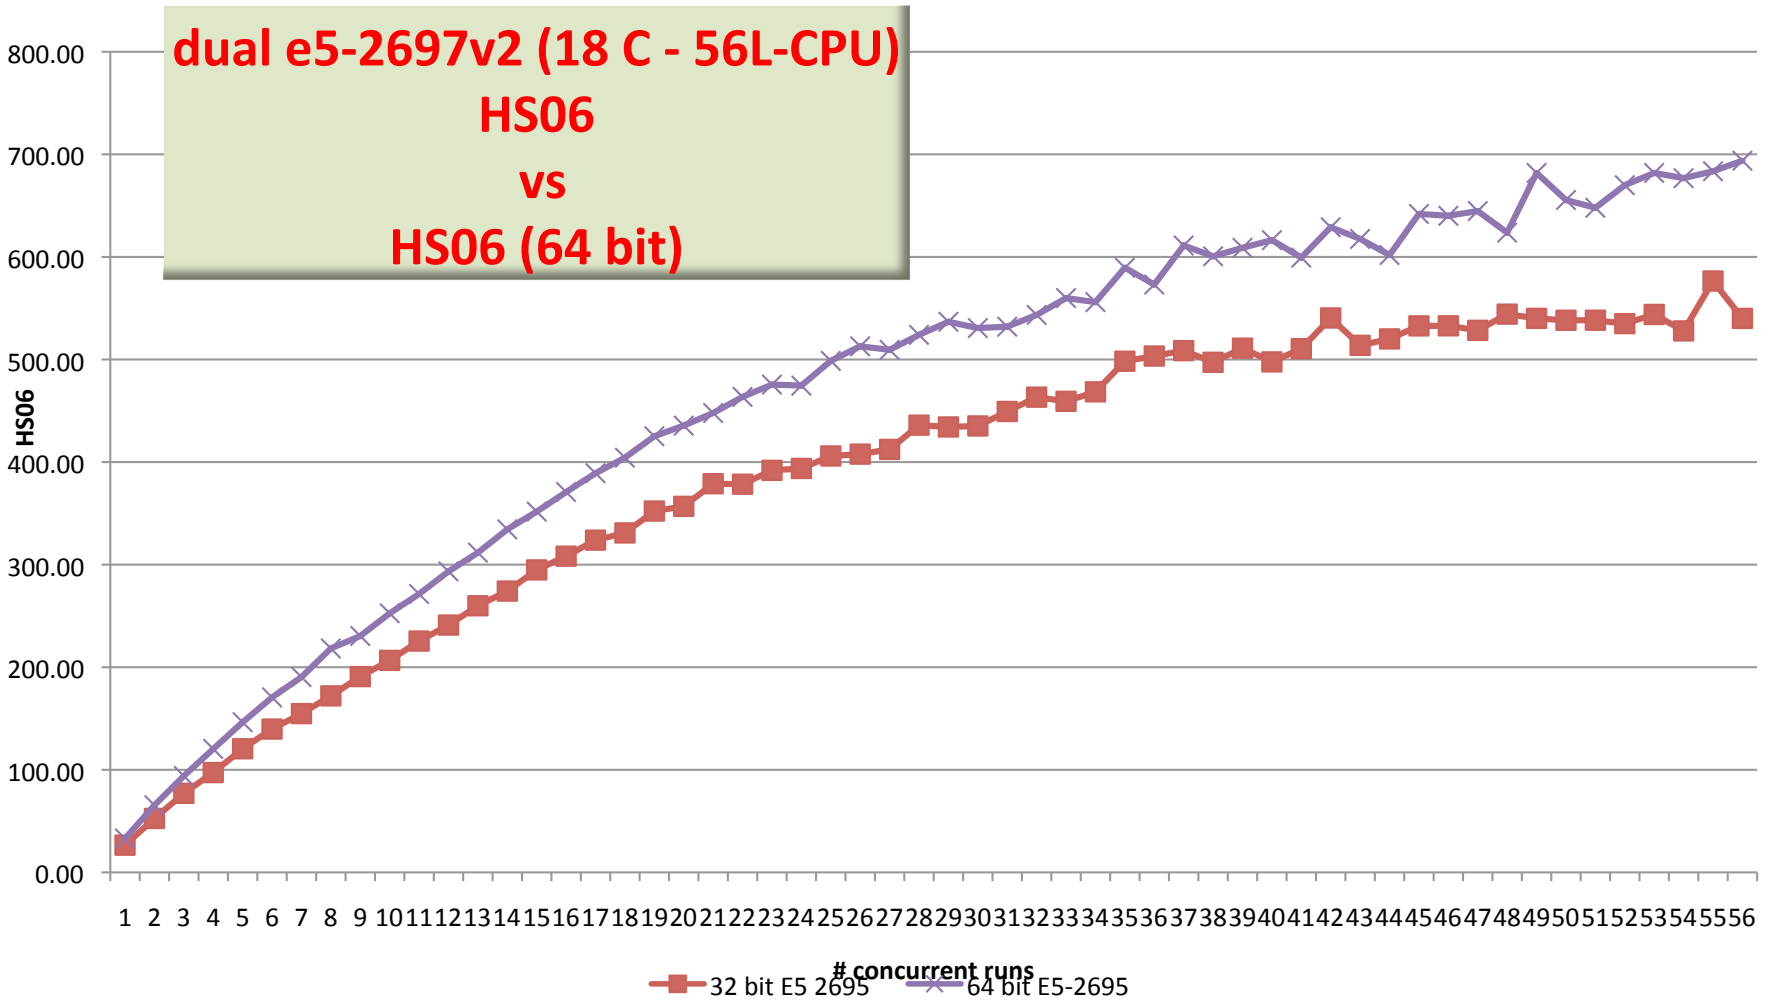

# CPU requested LHC on all Tier1

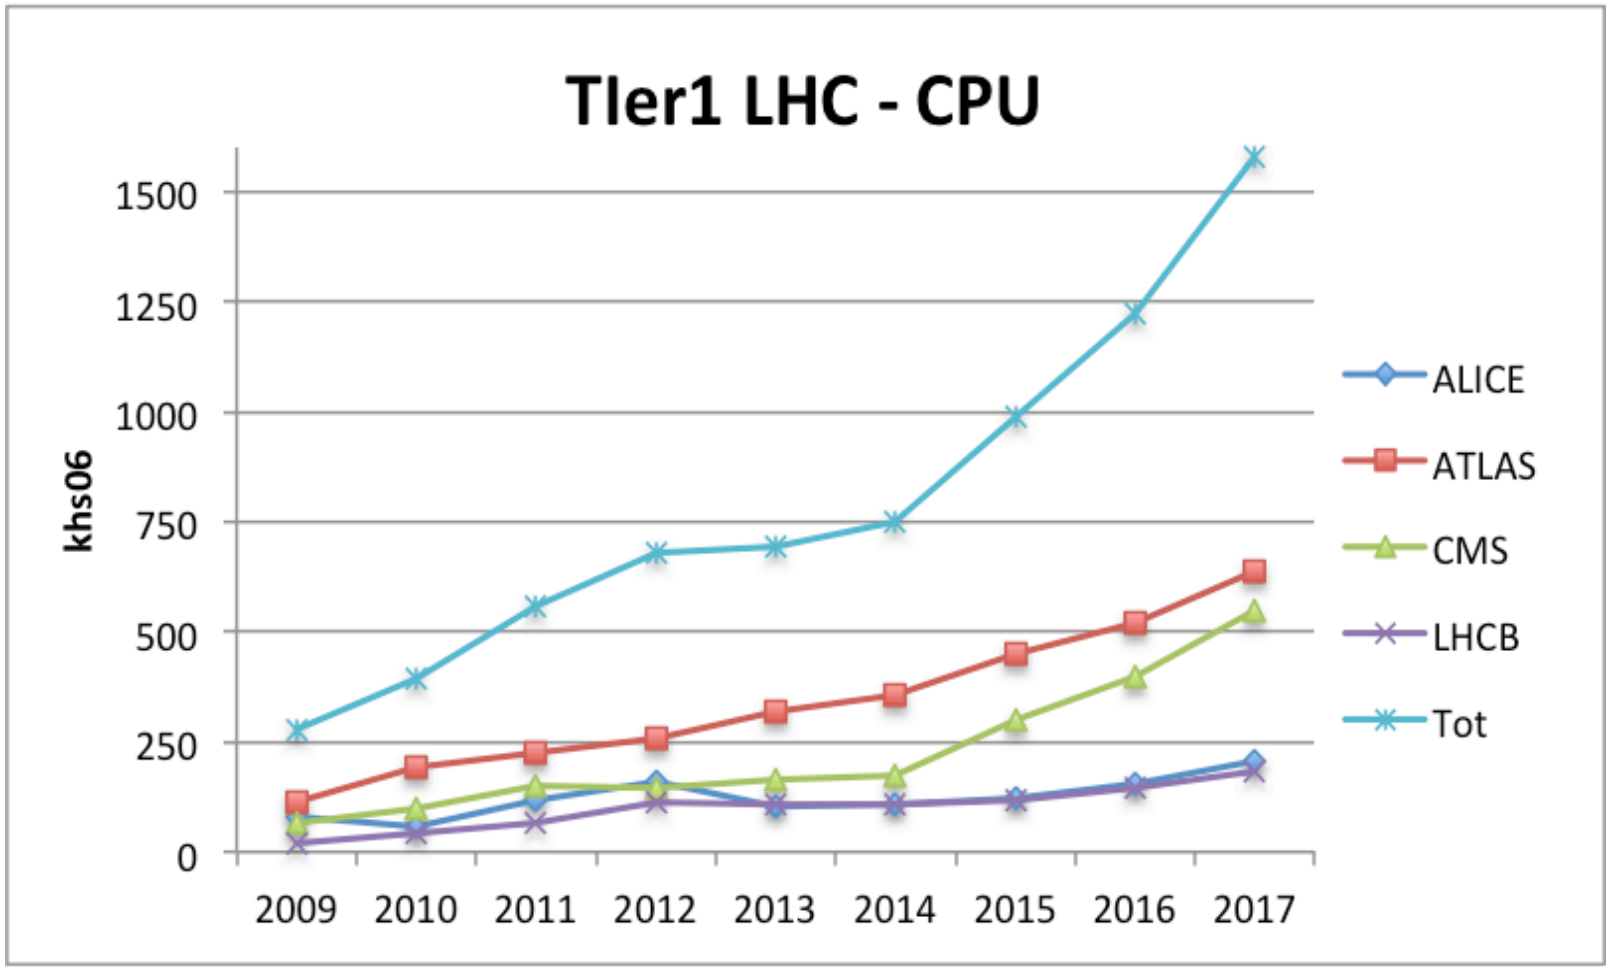

# LHC data forecast

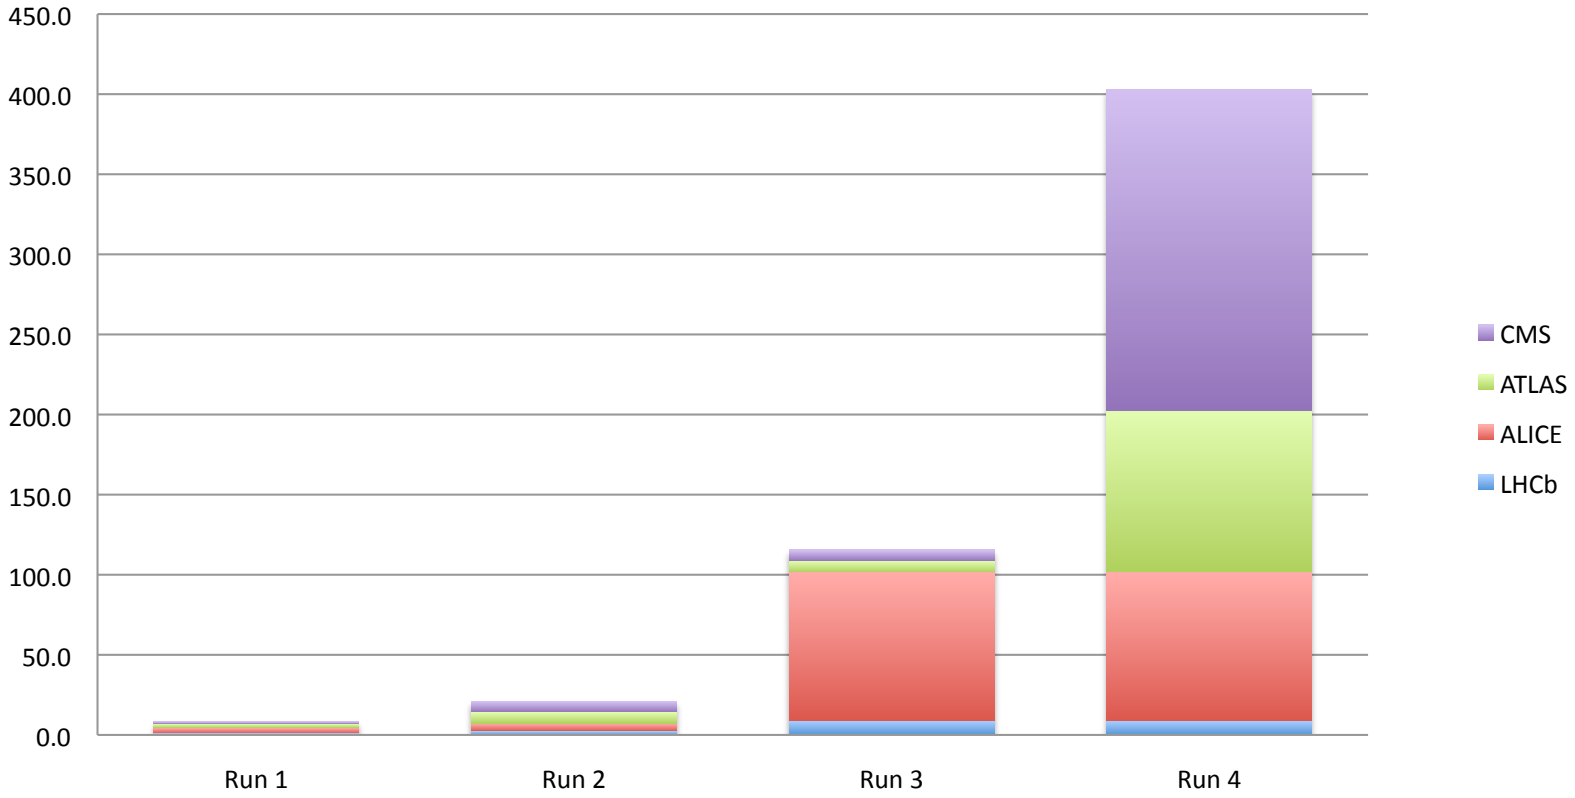


| | years | Luminosity | Integr. Lum. | Pileup |
|---|---|---|---|---|
| run2 | 2015-2018 | 1.5 $10^{34}$ $cm^{-2}s^{-1}$ | 150 $fb^{-1}$ | 40 |
| Long shutdown 2 | 2019-2020 | | | |
| run3 | 2021-2023 | 2.5 $10^{34}$ $cm^{-2}s^{-1}$ | 300 $fb^{-1}$ | 60 |
| long shutdown 2 | 2024-2025 | | | |
| run4 | 2026-2029 | 5.0 $10^{34}$ $cm^{-2}s^{-1}$ | 1000 $fb^{-1}$ | 140-200 |

# Historical Trend (data from M.Alef)

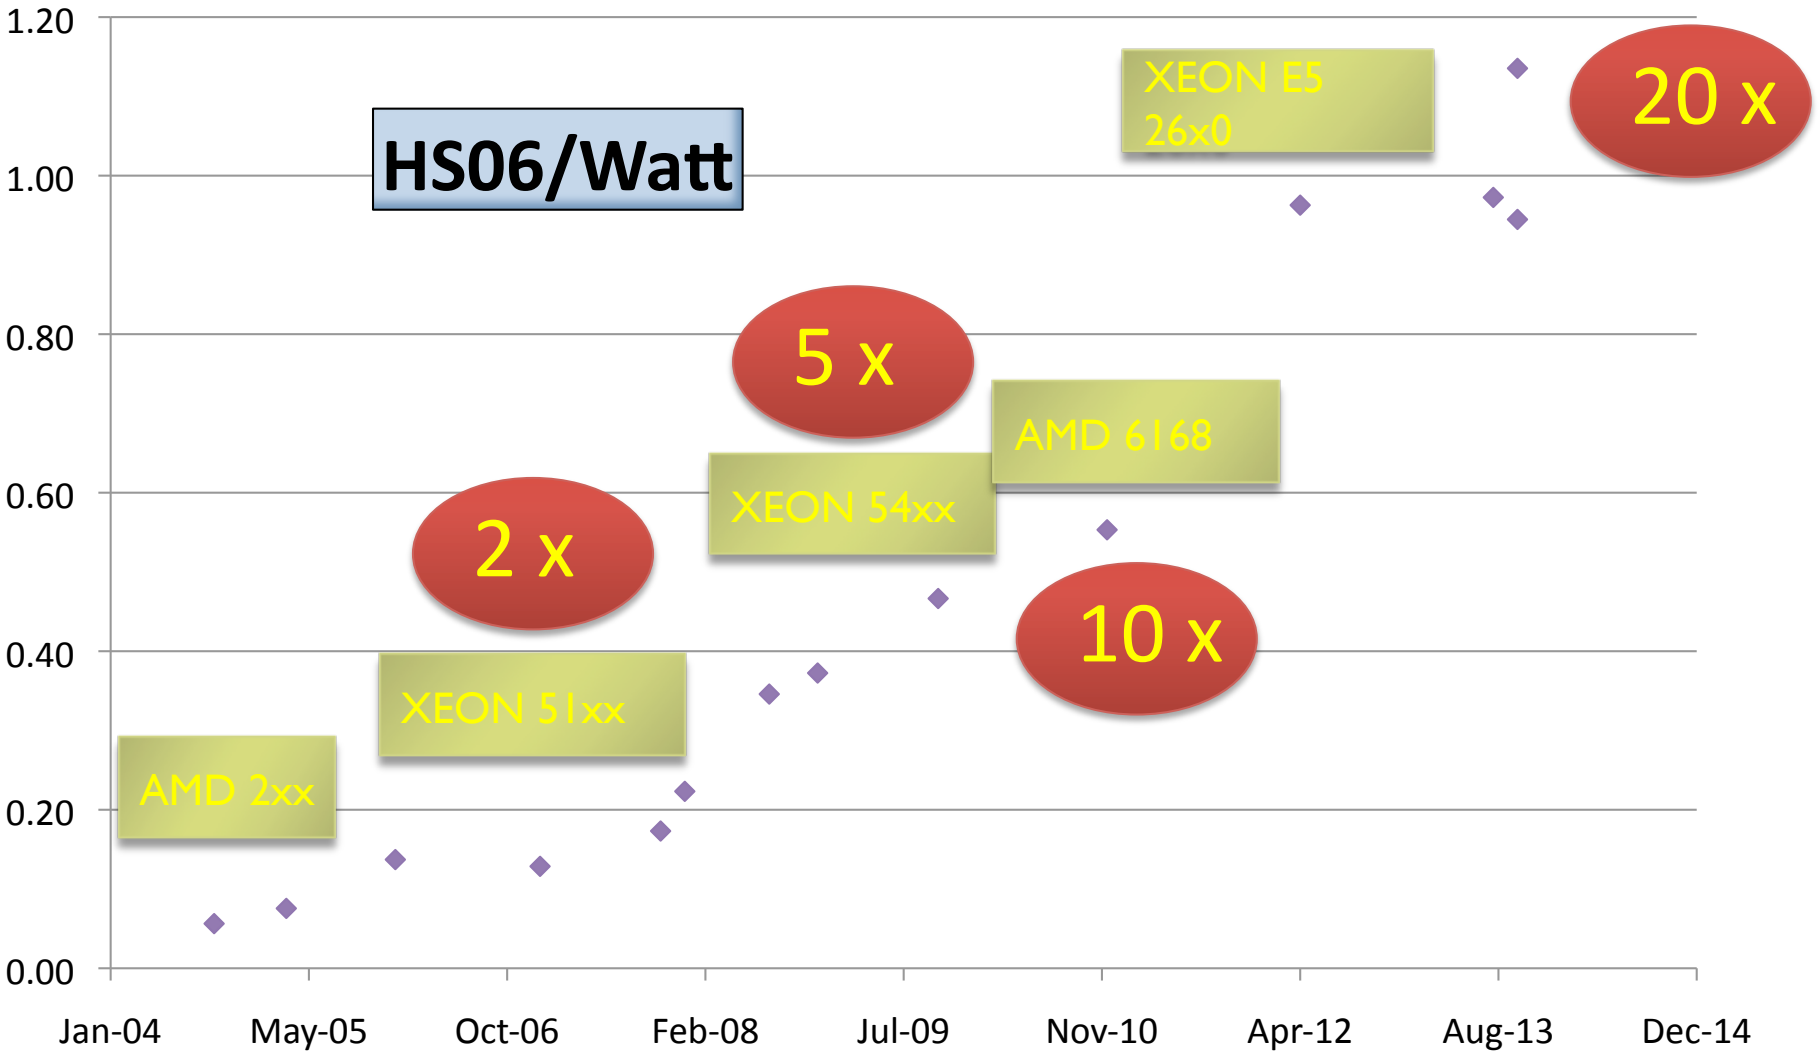

# Courtesy of Peter Elmer, Princeton Univ.

| Type | Cores | Power | Events/min /core | Events/min /Watt |
|---|---|---|---|---|
| Exynos4412 Prime @ 1.704GHz | 4 | 4W? | 1.14 | 1.14 |
| Xeon L5520 @ 2.27GHz | 2x4 | 120W? | 3.50 | 0.23 |
| Xeon E5-2630L @ 2.0GHz | 2x6 | 190W? | 3.33 | 0.21 |

# What drives the market?

- Videogames
  - High level GPUs
- Smartphones / Tablet
  - Low consumption CPUs, low cost (50$ each vs 1000$ for a Xeon)

- In both cases, expect many cores CPUs, O(1000) if you include GPU cores

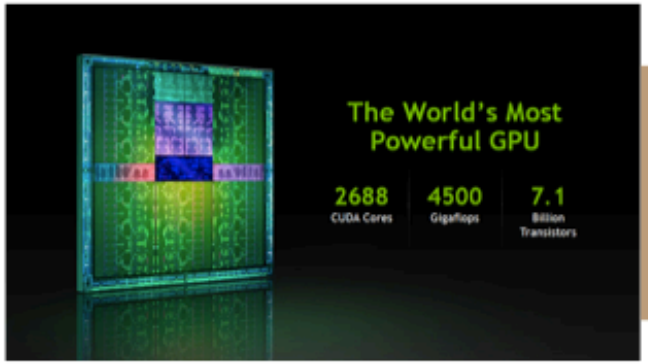


**Nvidia GTX Titan**
4500 Gflops
250 W
(~1000$)

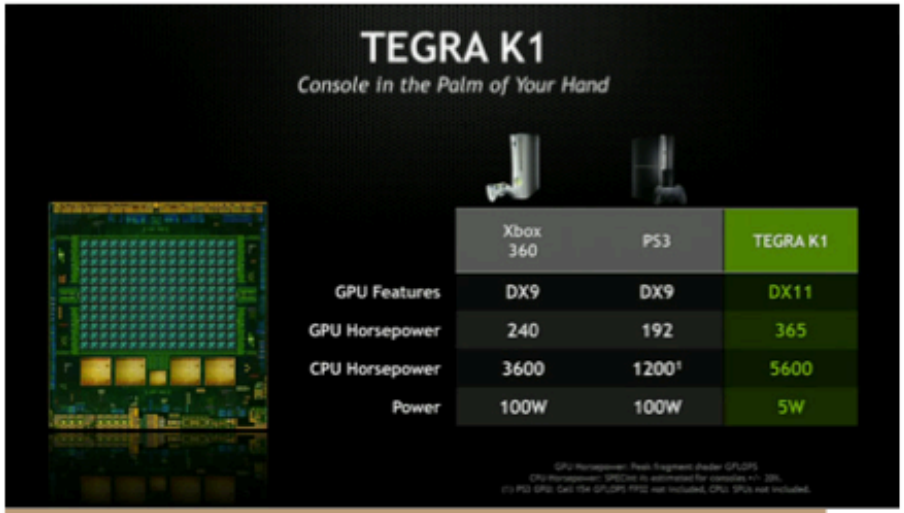


**Tegra K1**
365 Gflops
5 W (includes also 2 ARMv8 cores)
($50?)

**Xeon E5-4650L**
1000 Gflops
115W
2000$

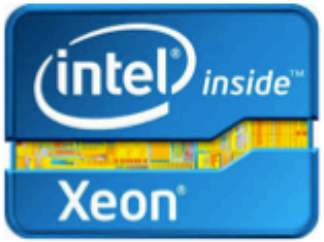

# Shipment

ARM: Smarphone, Tablet, Cars, Home Media
X86: Personal computer, notebook

Billions !!!

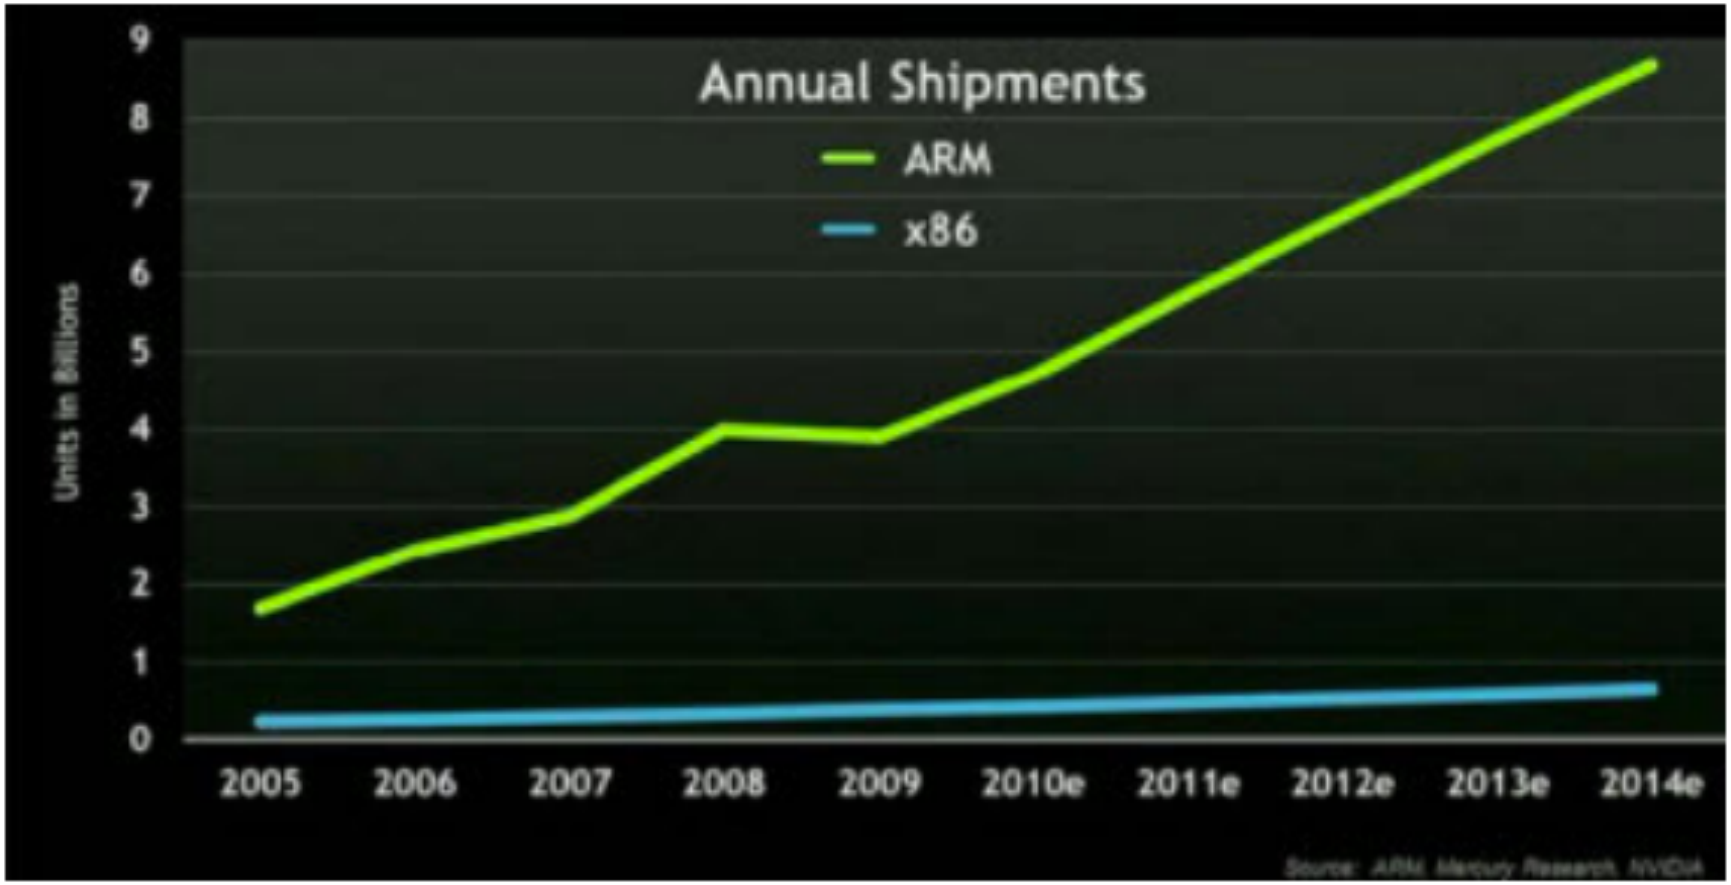

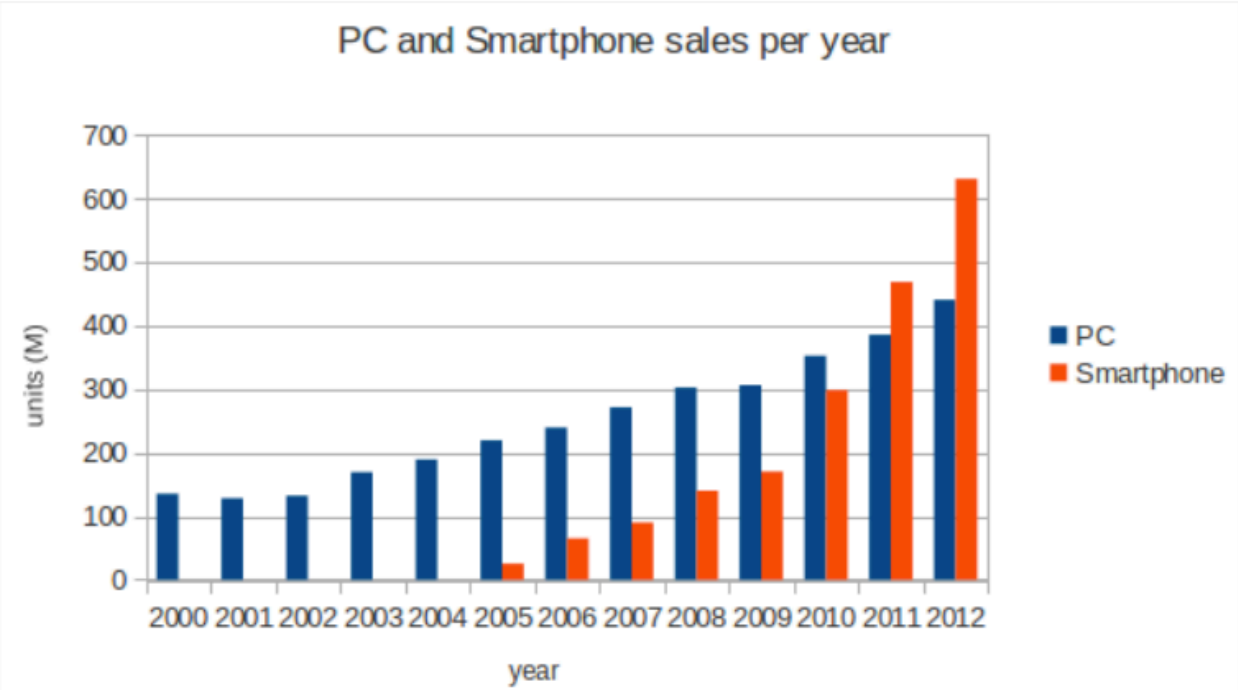


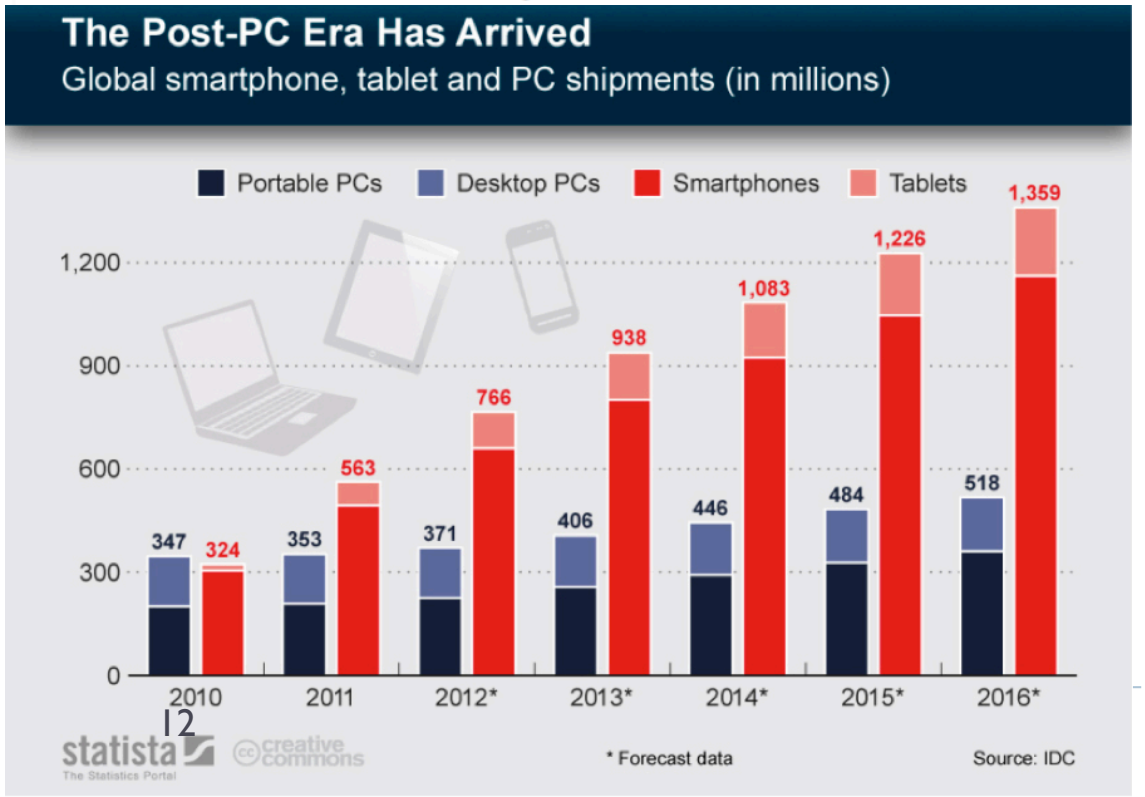


| | Units shipped 2012 → 2013 |
|---|---|
| Simple Phones | 750 → 900 |
| Smartphones | 722 → 1000 |
| Tablets | 128 → 210 |
| Notebooks | 202 → 200 |
| Desktops | 148 → 142 |
| Server | 10 → 9.5 |
| HPC | 0.1 → 0.1 |

HEP!

# ARM introduction

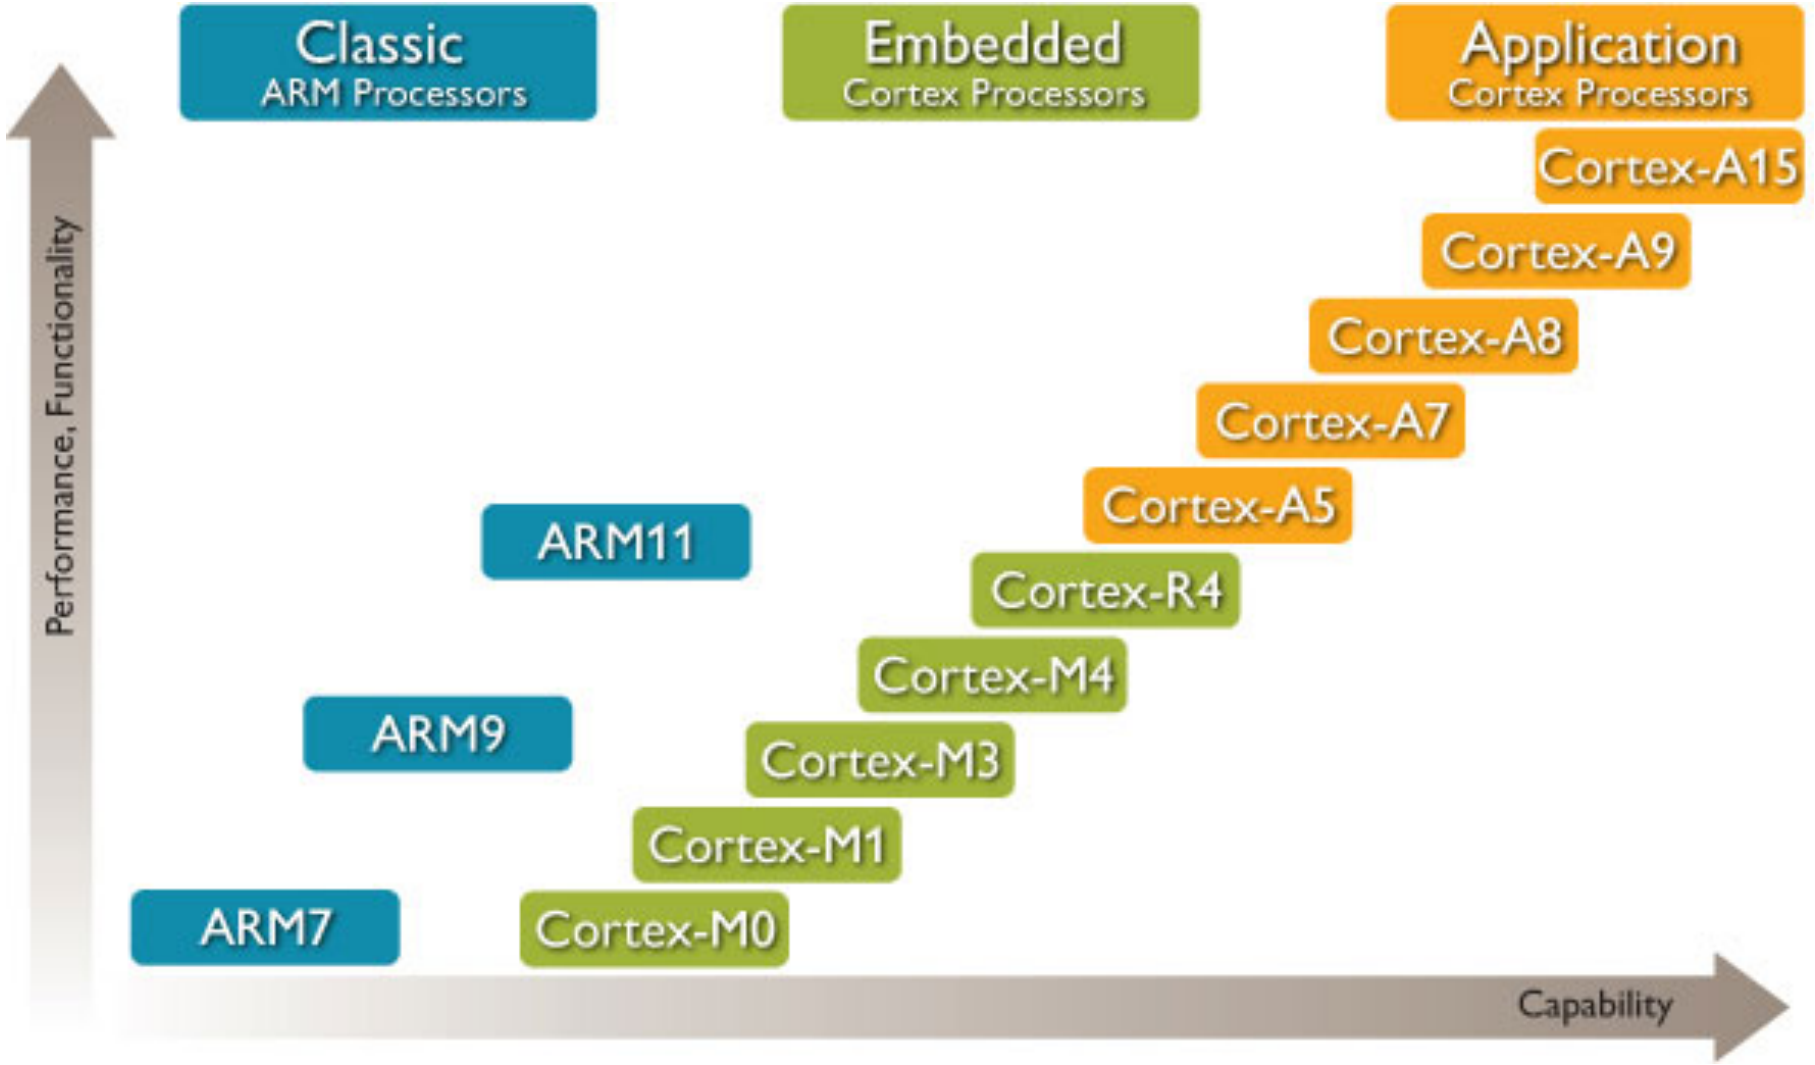

# The ARM Cortex A Series

| MODEL | Architecture | Process | out of order | Core | Cache L1 | Cache L2 | Dmips/MHz |
|---|---|---|---|---|---|---|---|
| | | nm | | min-max | | | |
| A5 | Arm v7 32 bit | n.d. | no | 1-4 | 4-64 | n.d. | 1.57 |
| A7 | Arm v7 32 bit | 40-28 | no | 1-8 | 8-64 | up to 1 MB | 1.9 |
| A8 | Arm v7 32 bit | 45-65 | no | 1 | 32+32 | 256-512 KB | 2 |
| A9 | Arm v7 32 bit | 28-65 | yes | 1-4 | 32+32 | up to 1 MB | 2.5 |
| A12 | Arm v7 32 bit | n.d. | yes | 1-4 | 32+32 | up to 8 MB | 3 |
| **A15** | **Arm v7 32 bit** | **28-32** | **yes** | **2-8** | **32+32** | **up to 8 MB** | **3.5 - 4.0** |
| A17 | Arm v7 32 bit | 28 | yes | 4 | 32+32 | up to 8 MB | 4 |
| A53 | Arm v8 64 bit | 28 | no | 1-4 | 8-64 + 8-64 | 128 KB - 2MB | n.d. |
| **A57** | **Arm v8 64 bit** | **28-20** | **yes** | **1-4** | **48+32** | **512 KB - 2MB** | **n.d.** |

# Odroid XU

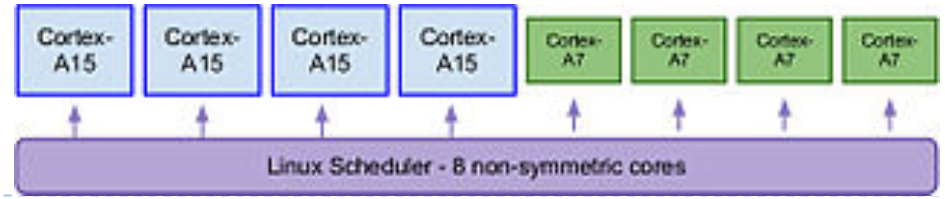


- Exynos5 Octa Cortex A15 1.6 Quadcore and A7 Quadcore CPUs
- 2GByte LPDDR3 RAM
- $169
- Fedora 18, armV7, gcc4.8, ODROID kernel

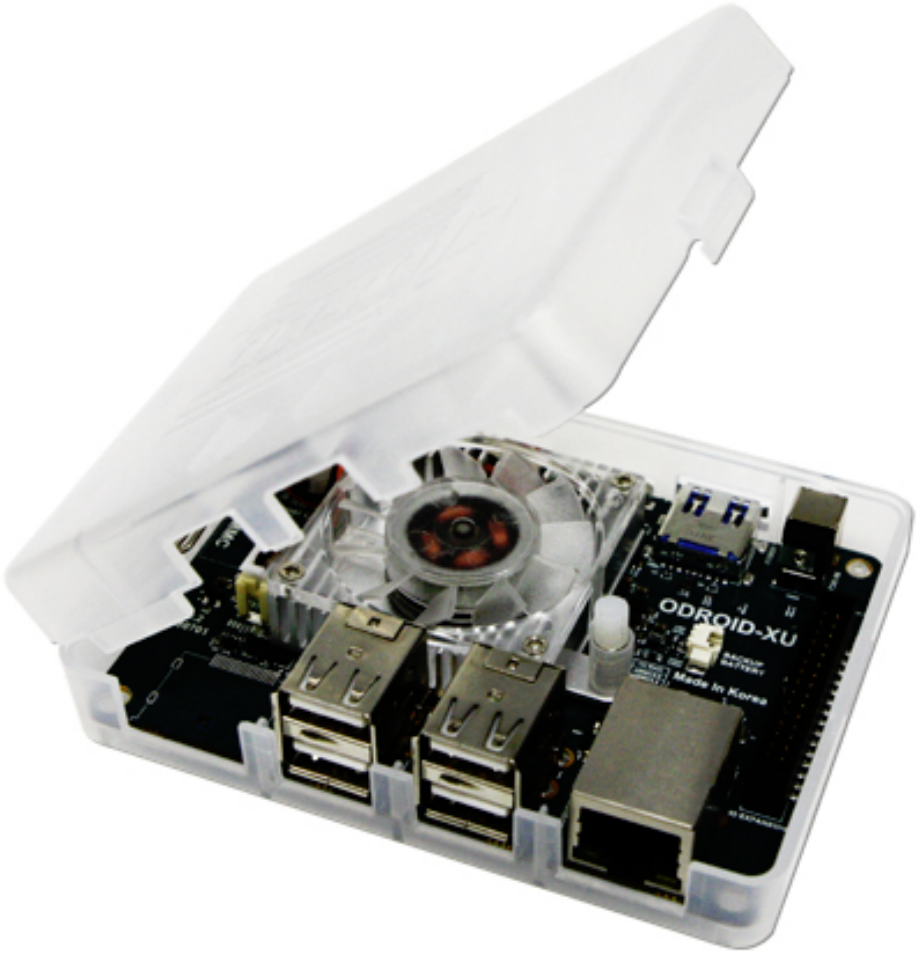

# The big LITTLE architecture

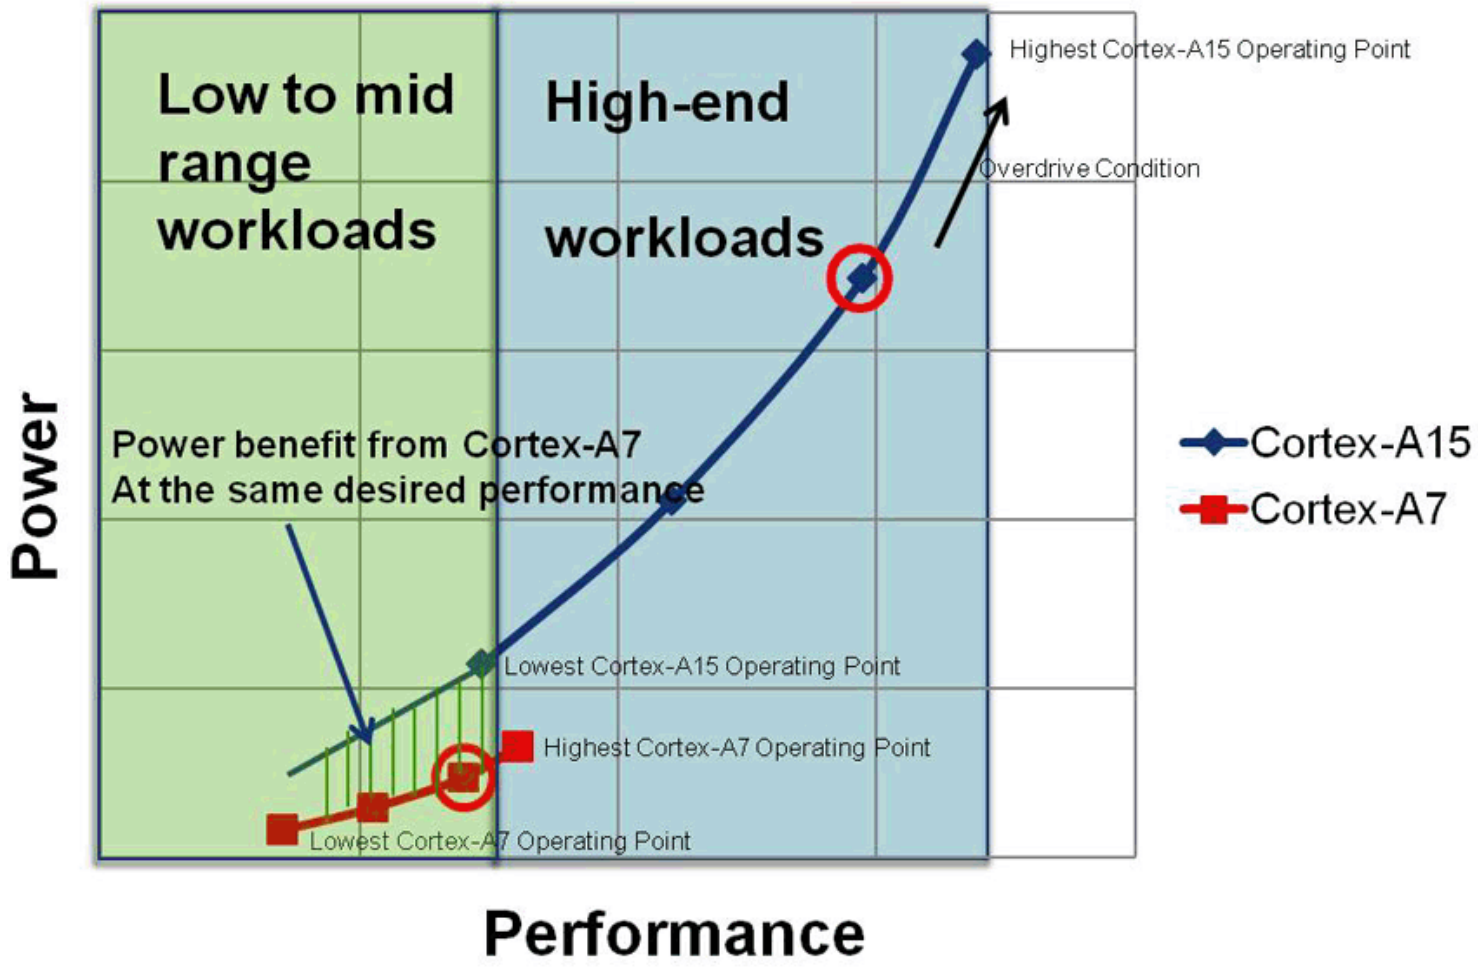

# Jetson Tegra K1

- Nvidia Tegra K1 SOC
  - Kepler GPU with 192 CUDA cores (unused in my tests)
  - 4+1 quad core A15 cpu
  - 2200 MHz
- 2GB x 16
- Linux Ubuntu/Linaro 4.8.2-19

# Nvidia Tegra K1 Building Blocks

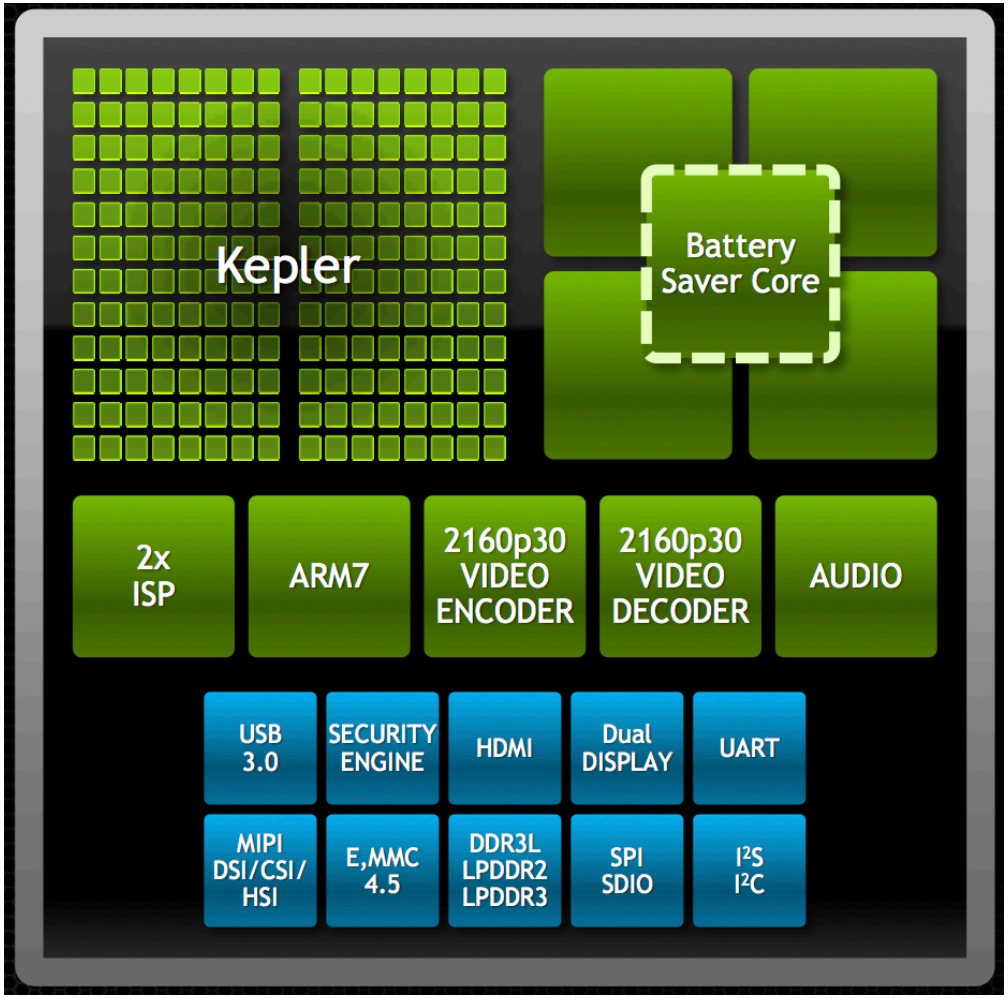

# Intel Atom Server

| Processor | Power | Frequency (Boost[10]) | Cores | Memory Channels | DIMMs per Channel | Memory Type | Memory Speed | Max. Memory Capacity | Max. PCIe Lanes | PCIe* Controllers | Ports |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Intel® Atom™ Processor C2750 | 20W | 2.4 GHz (2.6 GHz) | 8 | 2 | 2 | ECC DDR3/L | 1600 MHz | 64 GB | 16 | 4 PCIe* 2.0 | 4 x 2.5 GbE 2 x SATA 3 4 x SATA 2 4 x USB 2 |
| Intel® Atom™ Processor C2730 | 12W | 1.7 GHz (2.0 GHz) | 8 | 2 | 2 | ECC DDR3/L | 1600 MHz | 32 GB | 8 | 2 PCIe 2.0 | 2 x 2.5 GbE 2 x SATA 3 4 x USB 2 |
| Intel® Atom™ Processor C2550 | 14W | 2.4 GHz (2.6 GHz) | 4 | 2 | 2 | ECC DDR3/L | 1600 MHz | 64 GB | 16 | 4 PCIe 2.0 | 4 x 2.5 GbE 2 x SATA 3 4 x SATA 2 4 x USB 2 |
| Intel® Atom™ Processor C2530 | 9W | 1.7 GHz (2.0 GHz) | 4 | 2 | 2 | ECC DDR3/L | 1333 MHz | 32 GB | 8 | 2 PCIe 2.0 | 2 x 2.5 GbE 2 x SATA 3 4 x USB 2 |
| Intel® Atom™ Processor C2350 | 6W | 1.7 GHz (2.0 GHz) | 2 | 1 | 2 | ECC DDR3/L | 1333 MHz | 16 GB | 4 | 1 PCIe 2.0 | 4 x 2.5 GbE 2 x SATA 3 4 x USB 2 |

- Not a SOC but a low power architecture
- X86-64

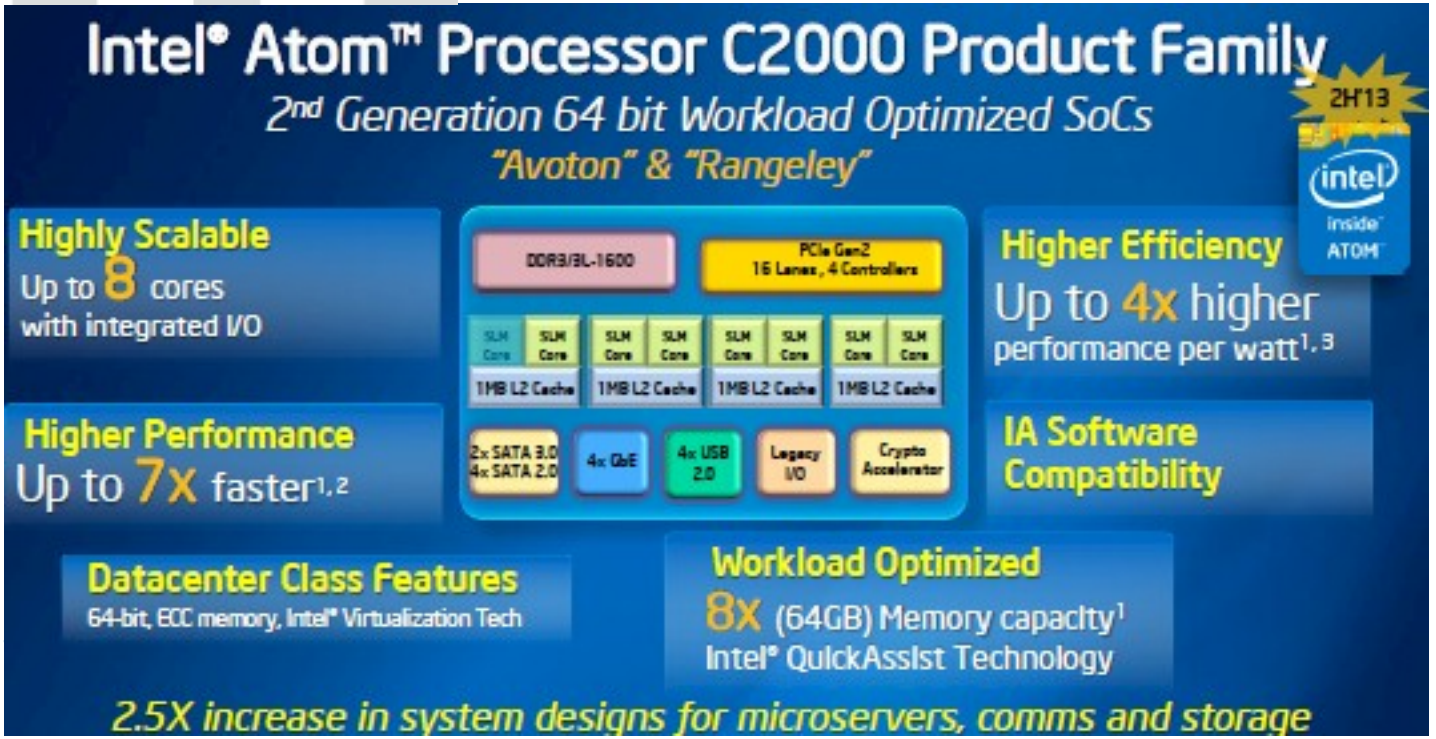

# Server based on Intel server

- Atom server in 2U rack mount chassis
- Each card has 8core Avoton computer node and up to 32GB
- 92 host is this prototype

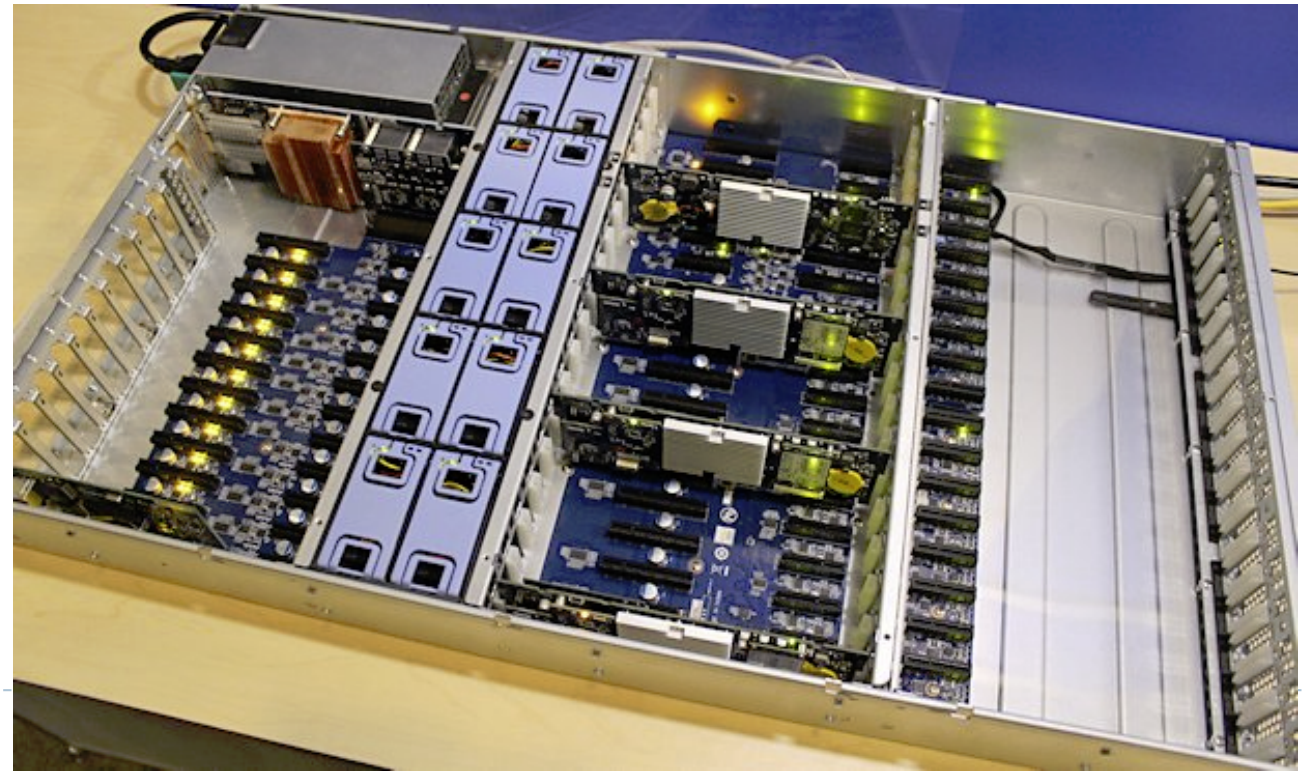

# ARM server for HTC


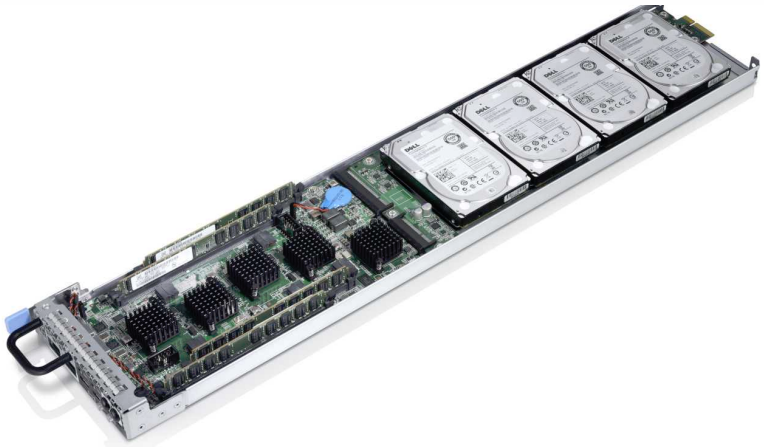

# HS06 on Exynos5, TegraK1 and Atom C2750

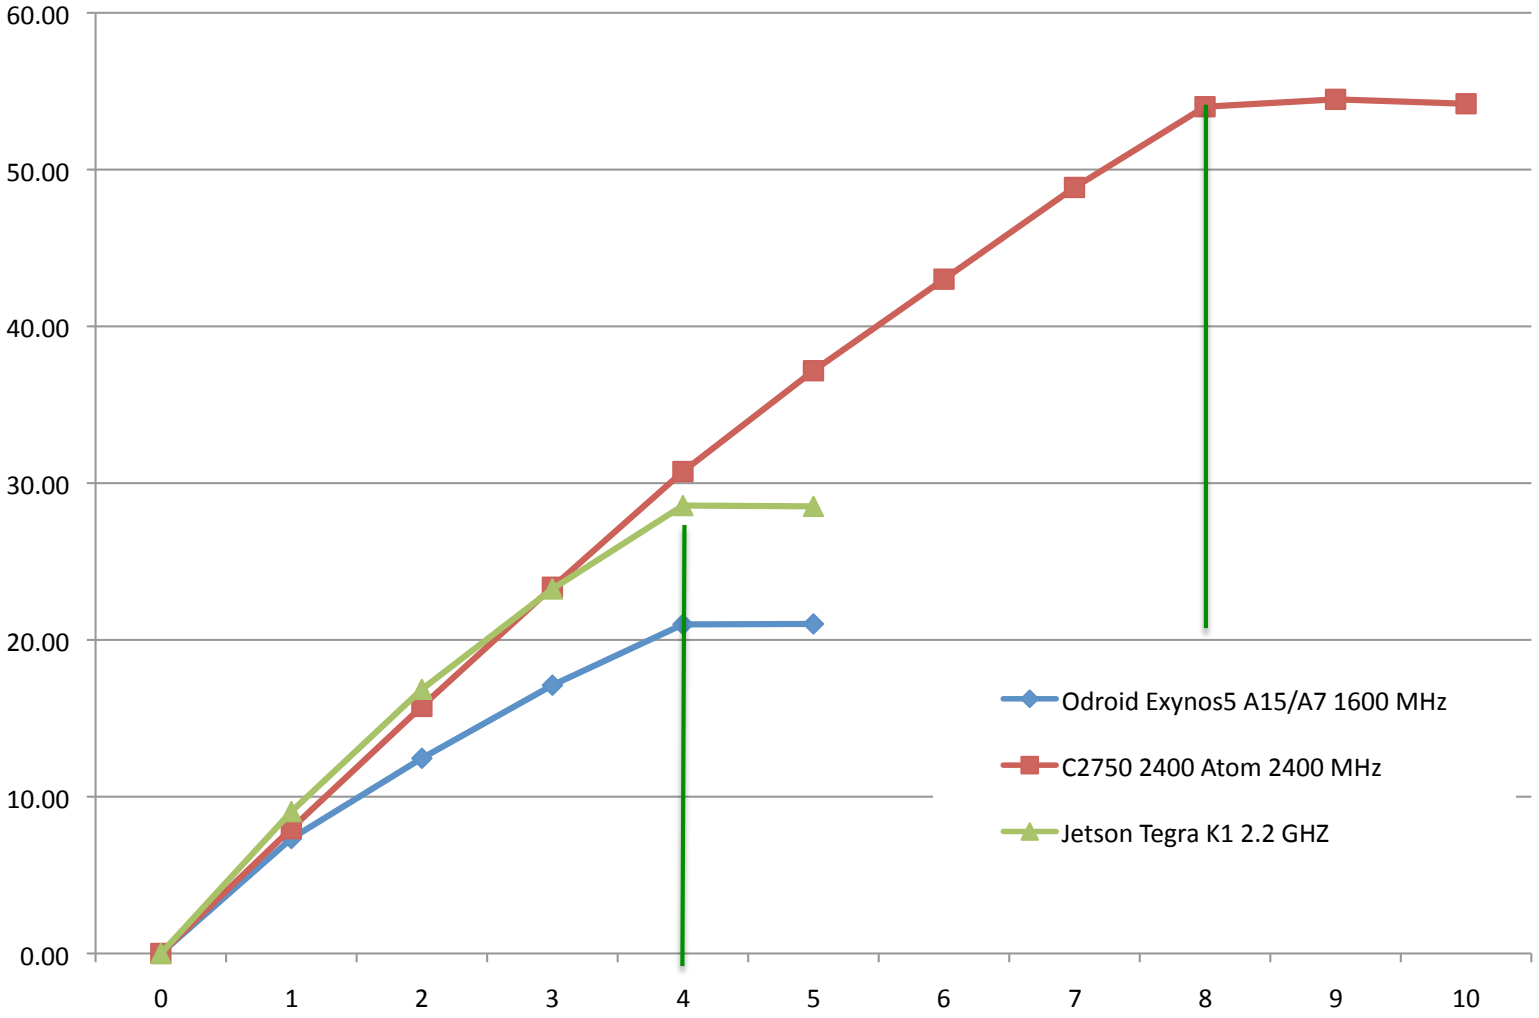

# HS06 on Exynos5, TegraK1 and Atom C2750 – Per core loaded

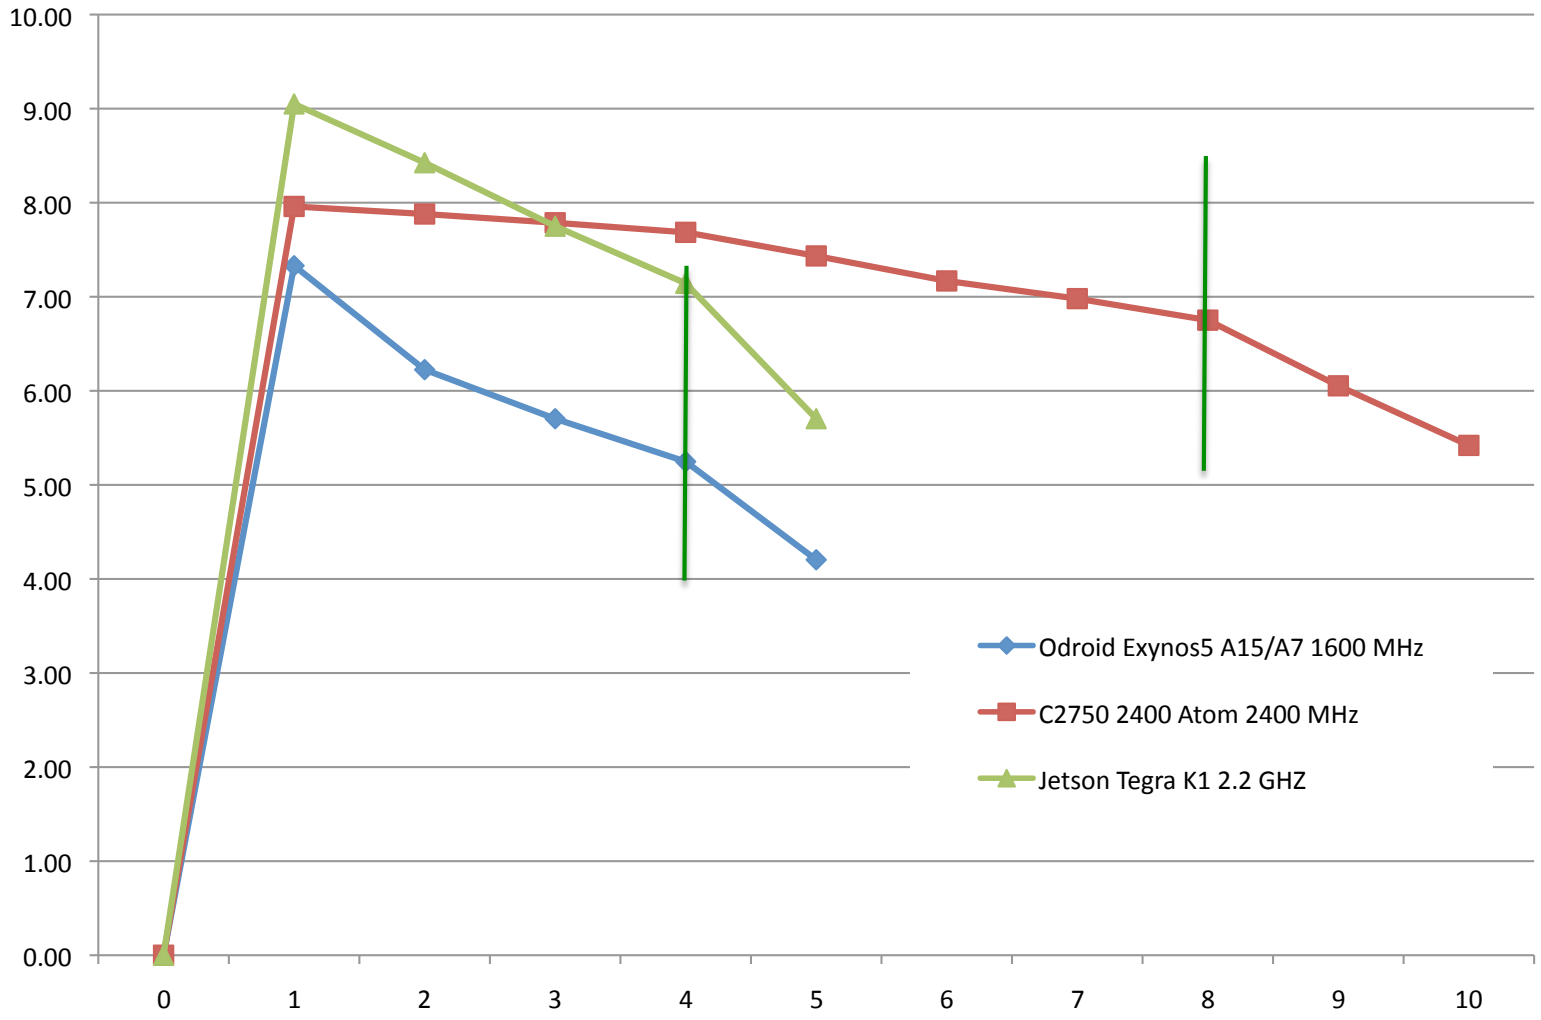

# HS06 on Exynos5, TegraK1 and Atom C2750 – Per core loaded and per clock

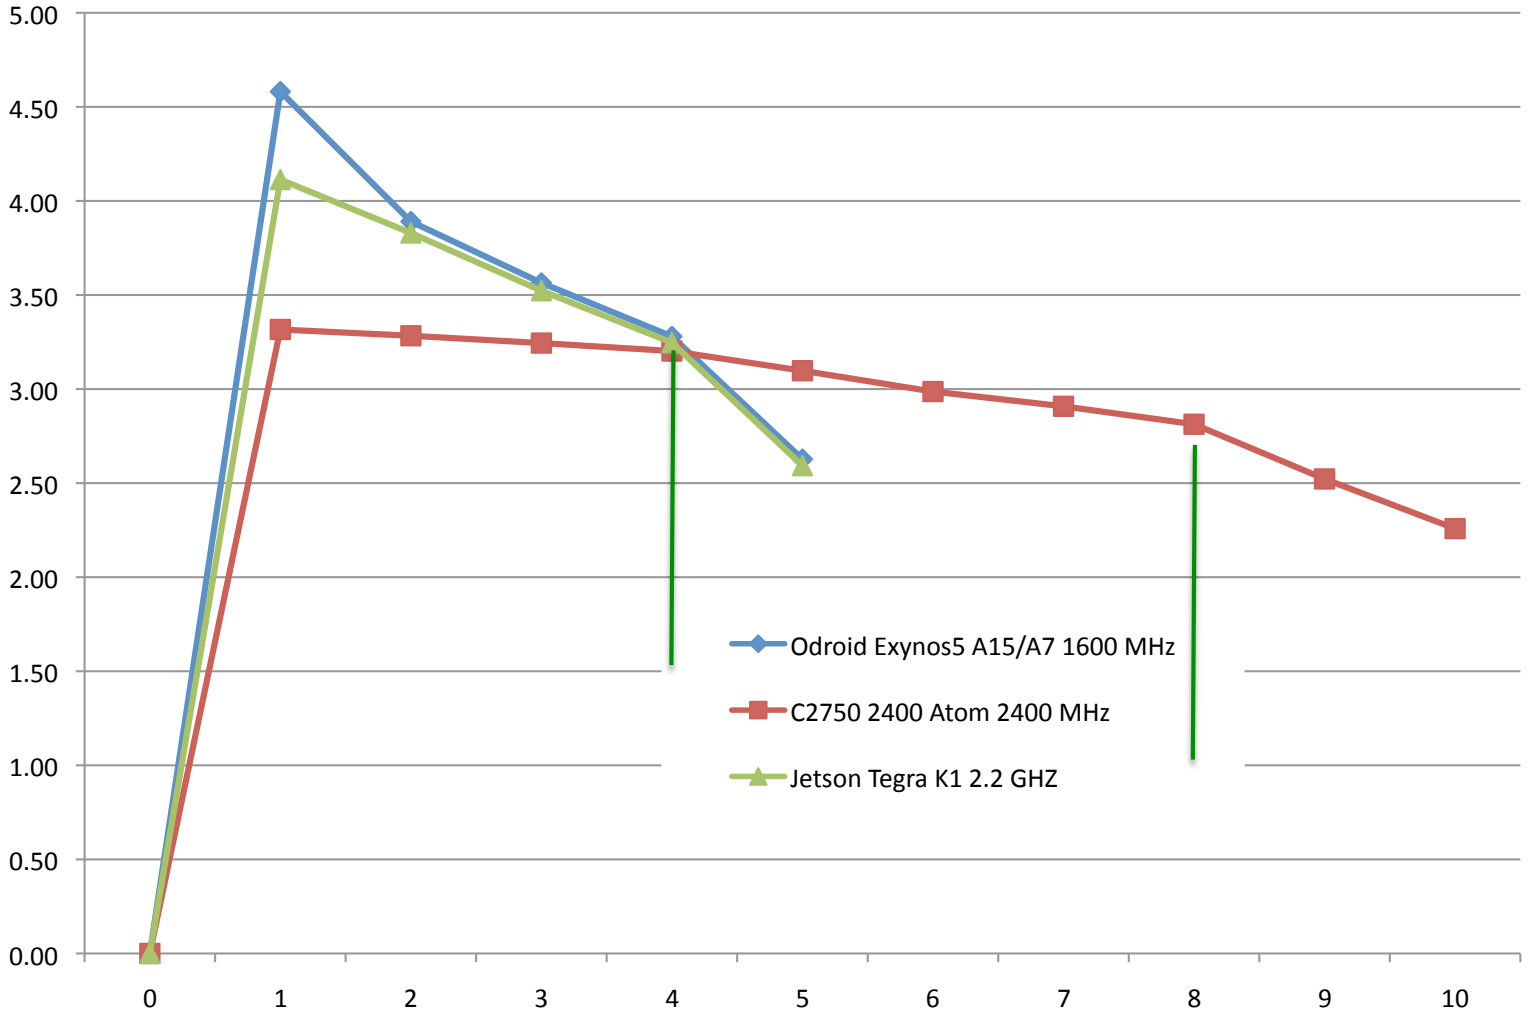

# Xeon D Broadwell SoC

| New Intel Xeon D-1500 Series Processors (to be updated) | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| Launch | Intel Xeon D Model | Cache (MB) | # of Cores | # of Threads | Processor Base (GHz) | Max Turbo (GHz) | TDP | RCP | Cores * Base GHz | Threads * Base GHz |
| Q1 2015 | D-1540 | 12 | 8 | 16 | 2 | 2.6 | 45 | $581 | 16 | 32 |
| | D-1520 | 6 | 4 | 8 | 2.2 | 2.6 | 45 | $199 | 8.8 | 17.6 |
| Q4 2015 | D-1518 | 6 | 4 | 8 | 2.2 | 2.2 | 35 | $234 | 8.8 | 17.6 |
| | D-1521 | 6 | 4 | 8 | 2.4 | 2.7 | 45 | | 9.6 | 19.2 |
| | D-1527 | 6 | 4 | 8 | 2.2 | 2.7 | 35 | $259 | 8.8 | 17.6 |
| | D-1528 | 9 | 6 | 12 | 1.9 | 2.5 | 35 | $389 | 11.4 | 22.8 |
| | D-1531 | 9 | 6 | 12 | 2.2 | 2.7 | 45 | | 13.2 | 26.4 |
| | D-1537 | 12 | 8 | 16 | 1.7 | 2.3 | 35 | $571 | 13.6 | 27.2 |
| | D-1541 | 12 | 8 | 16 | 2.1 | 2.7 | 45 | $581 | 16.8 | 33.6 |
| | D-1548 | 12 | 8 | 16 | 2 | 2.6 | 45 | $675 | 16 | 32 |
| Q1 2016 (Feb) | D-1557 | 18 | 12 | 24 | 1.5 | | 45 | | 18 | 36 |
| | D-1567 | 18 | 12 | 24 | 2.3 | | 65 | | 27.6 | 55.2 |
| | D-1571 | 24 | 16 | 32 | 1.3 | | 45 | $1,222 | 20.8 | 41.6 |
| | D-1577 | 24 | 16 | 32 | 1.3 | | 45 | | 20.8 | 41.6 |
| | D-1581 | | | | | | | | 0 | 0 |
| | D-1587 | 24 | 16 | 32 | 1.7 | | 65 | | 27.2 | 54.4 |
| © 2016 ServeTheHome.com | | | | | | | | | | |

# Xeon D vs Avoton

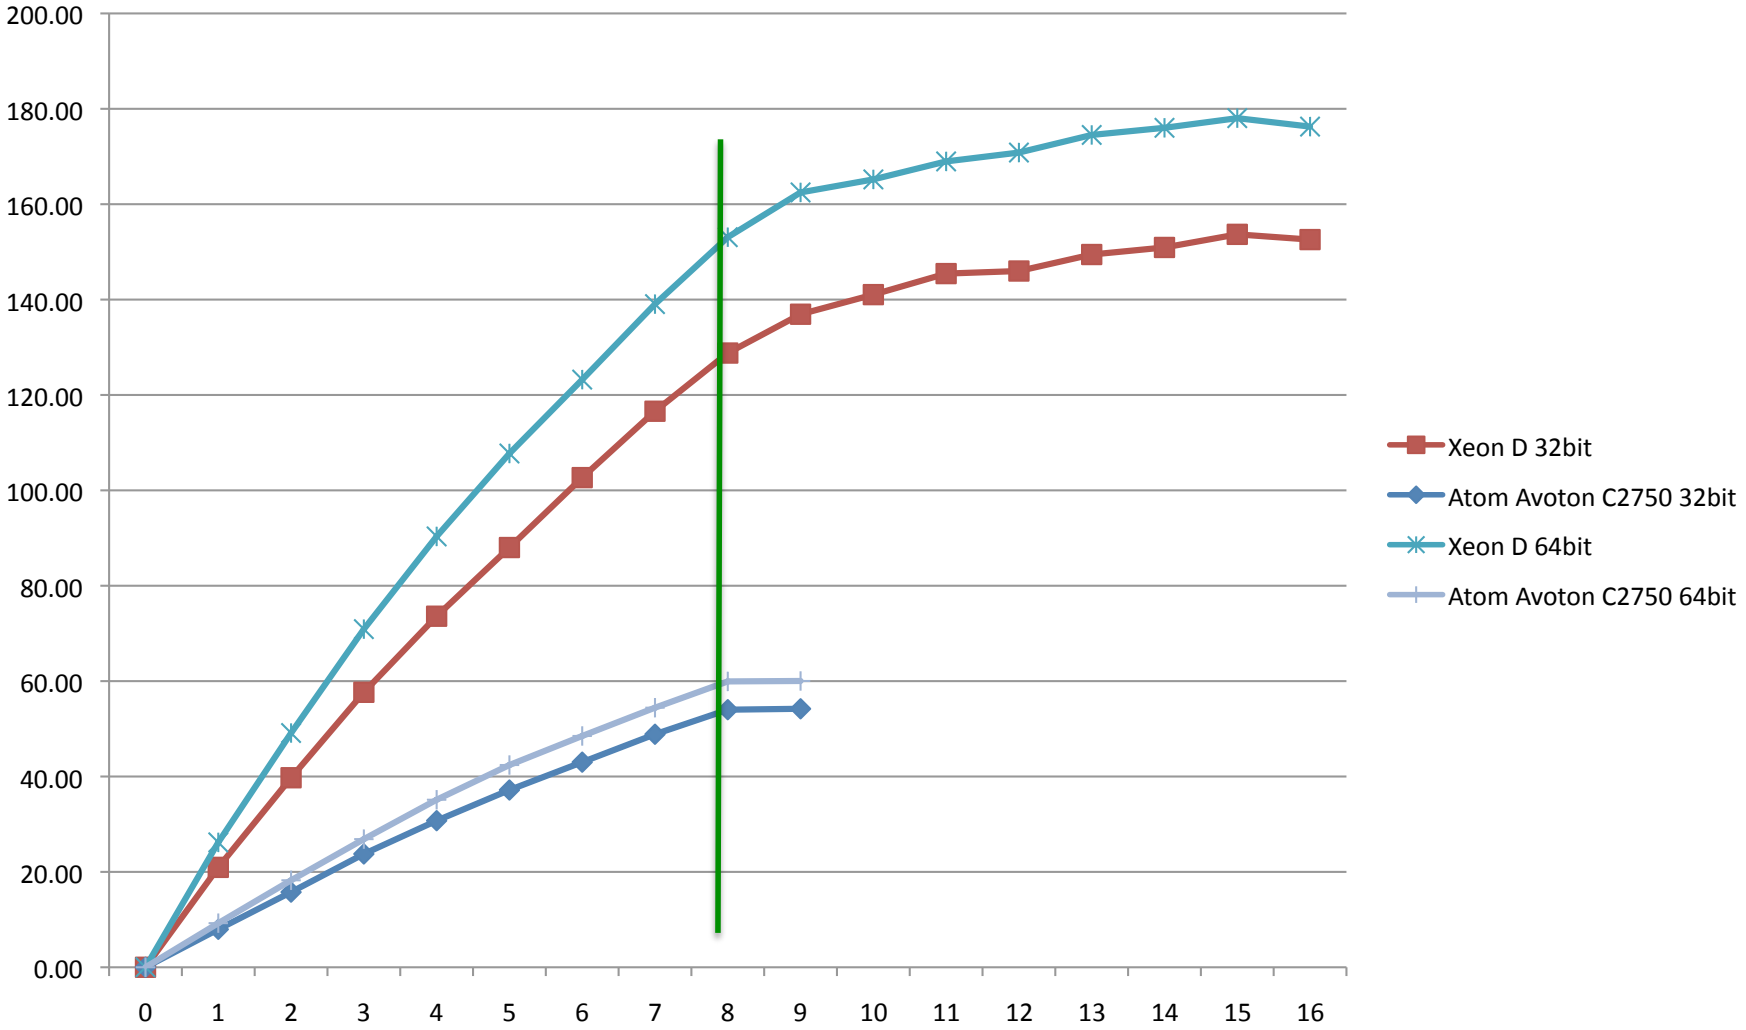

# Xeon D vs Avoton per core

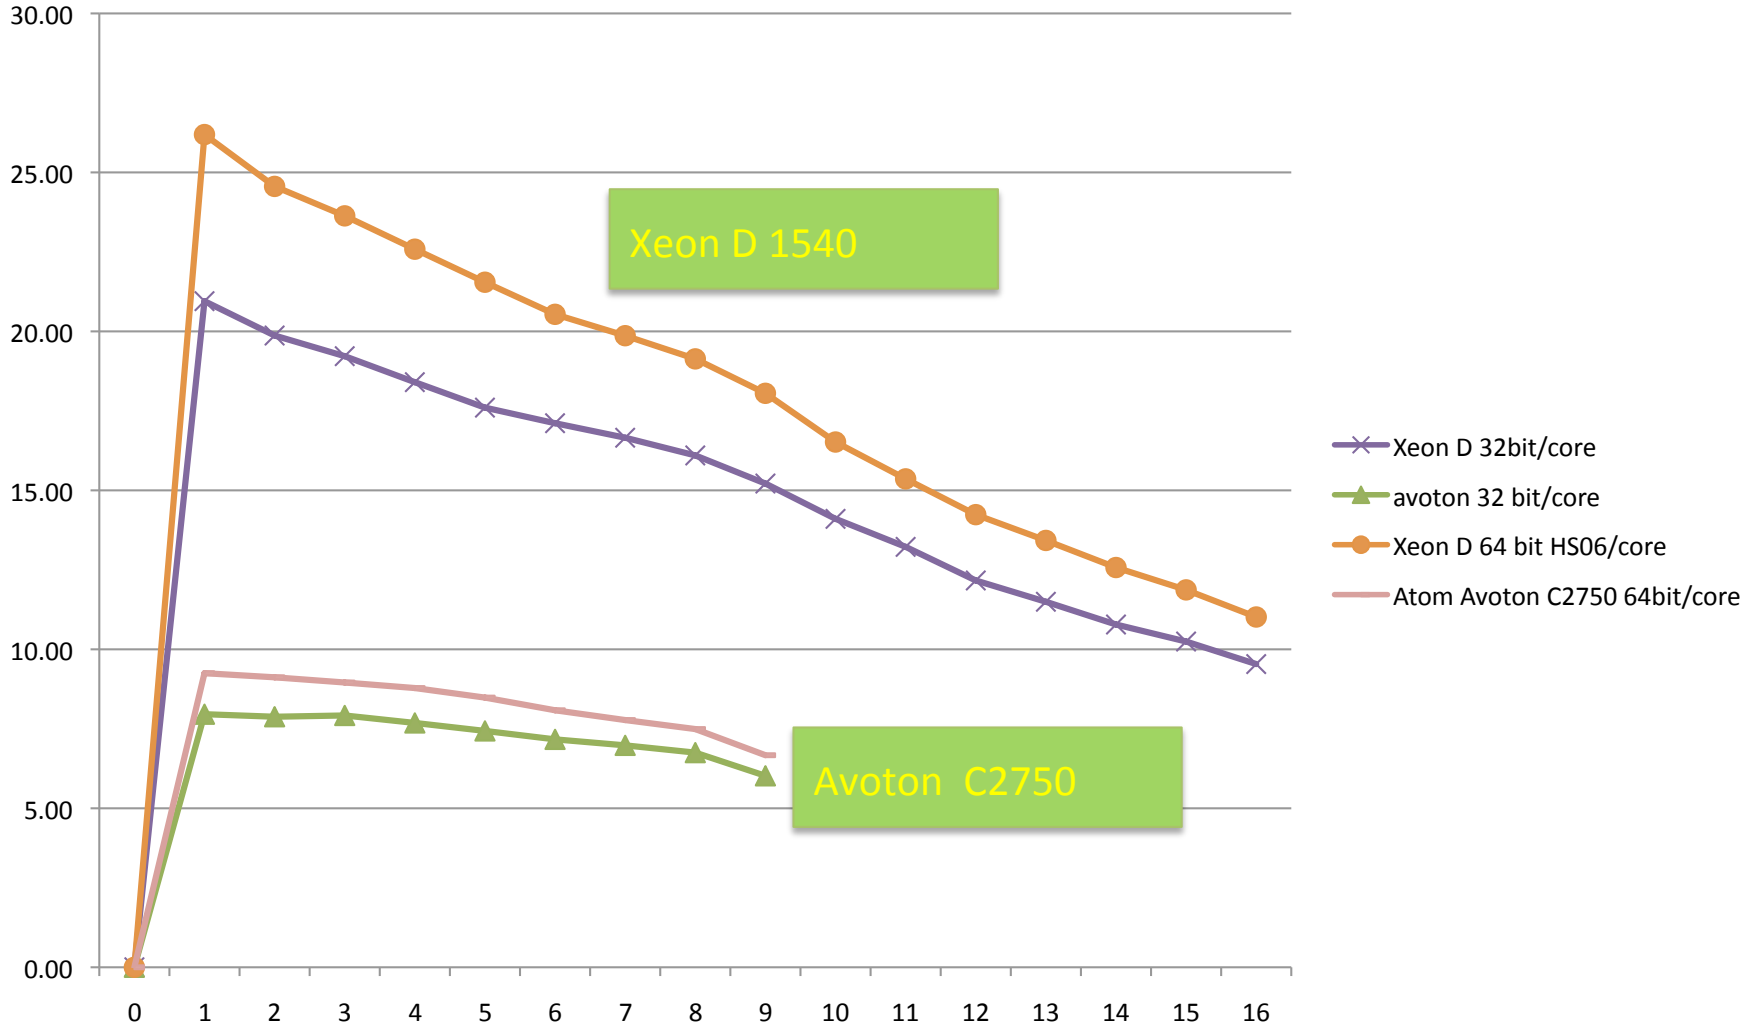

# Performance per core

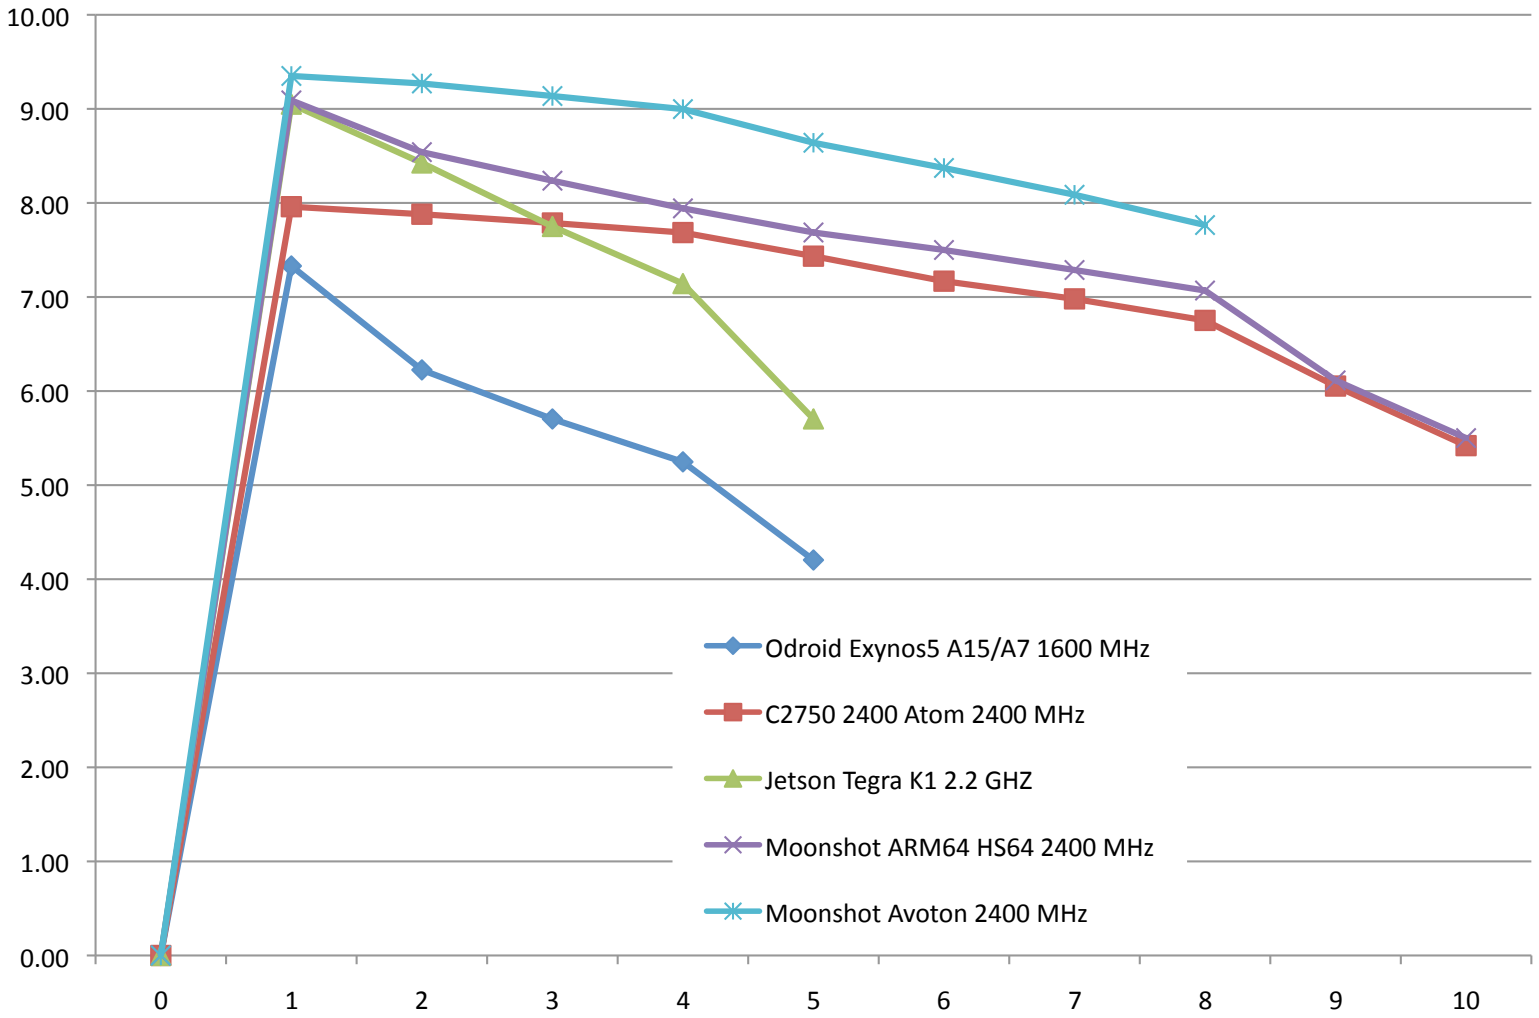

# Measurements of power consumption

- Measurements of voltage, amperage and power consumption
- The power logger
- Measurements setup


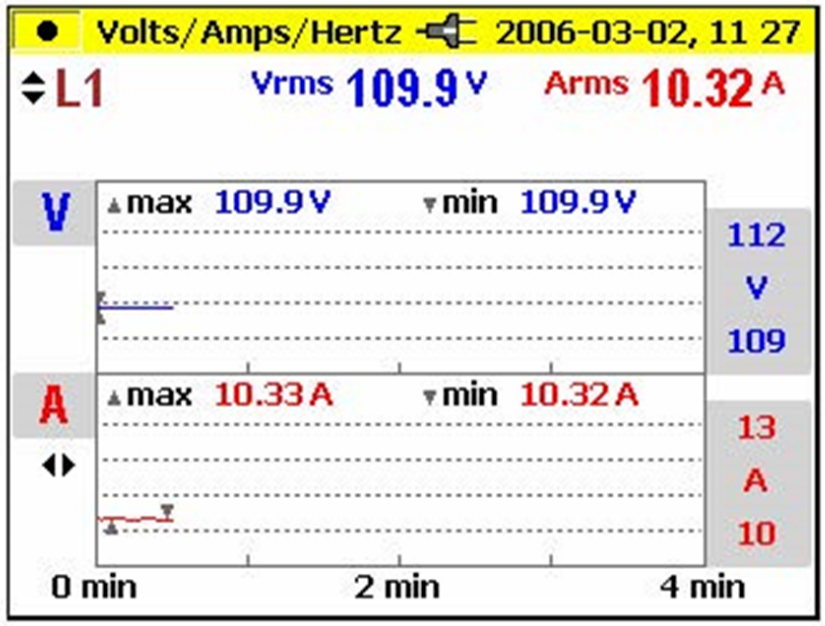

# Power logger sw

# Power consumption (Watt) on Intel Xeon E5 2660

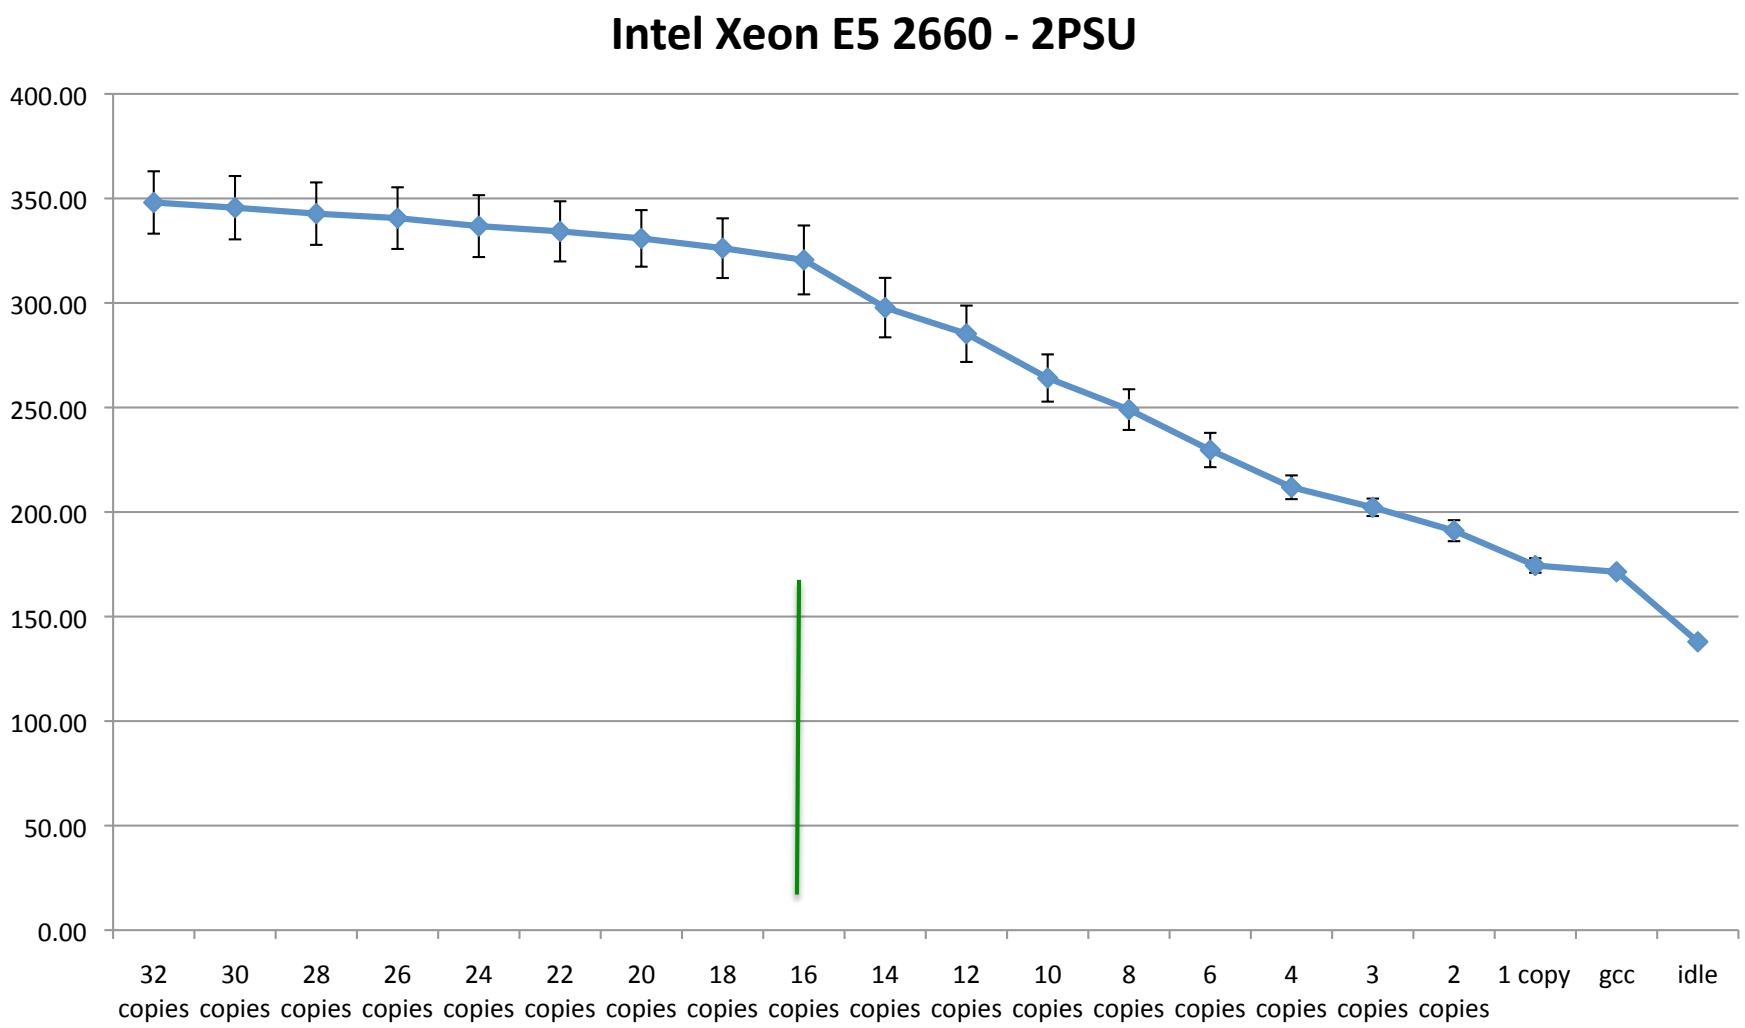

# Efficiency HS06/Watt

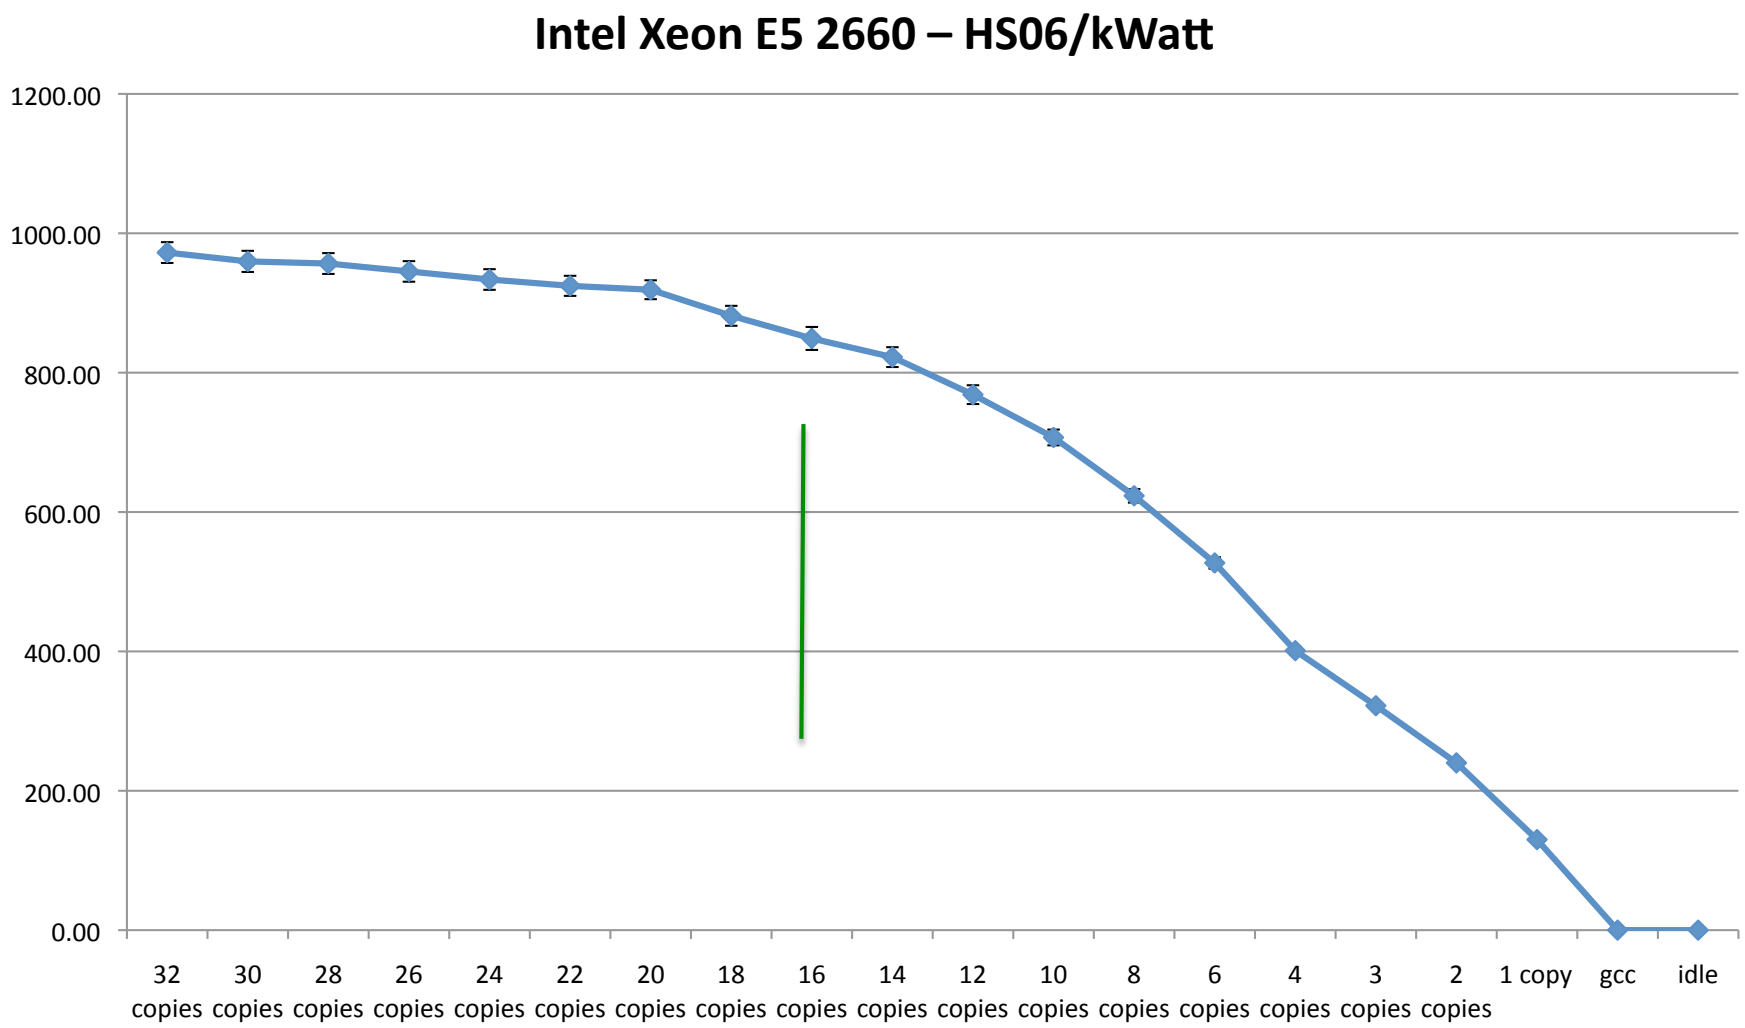

# Power consumed on Intel ATOM C2750 vs XEON E5

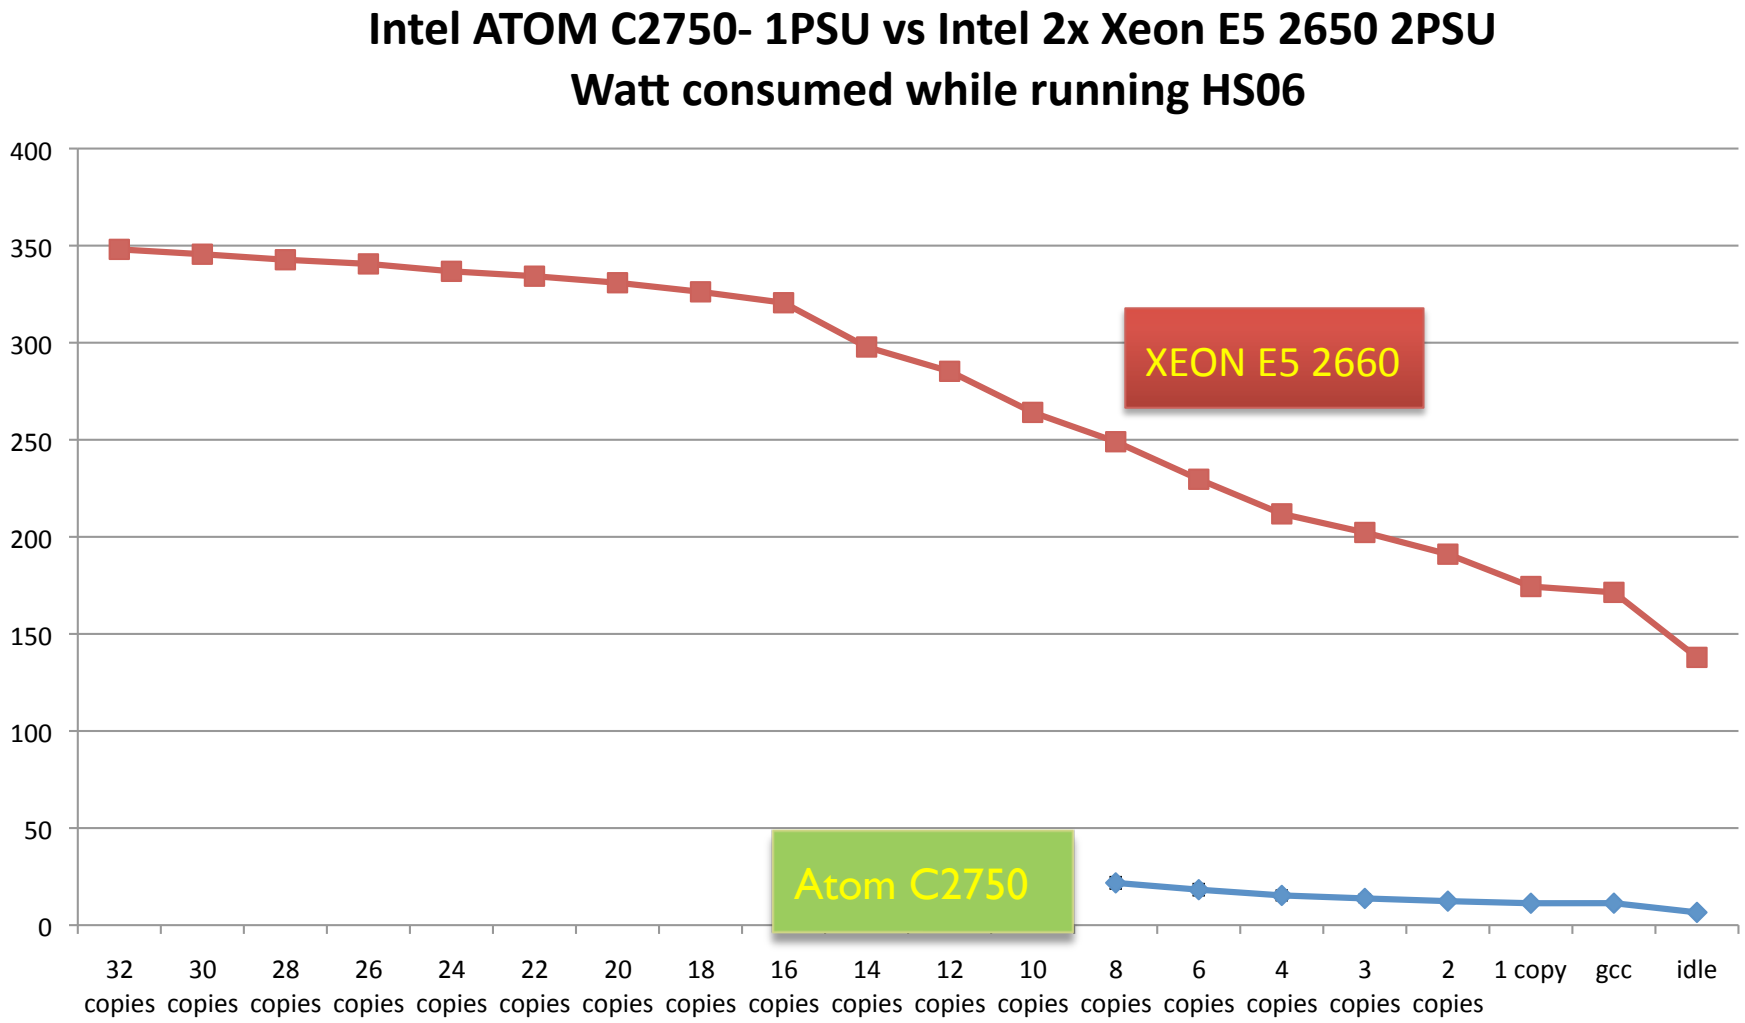

# Power Efficiency HS06/Watt
# Intel ATOM C2750 vs XEON E5

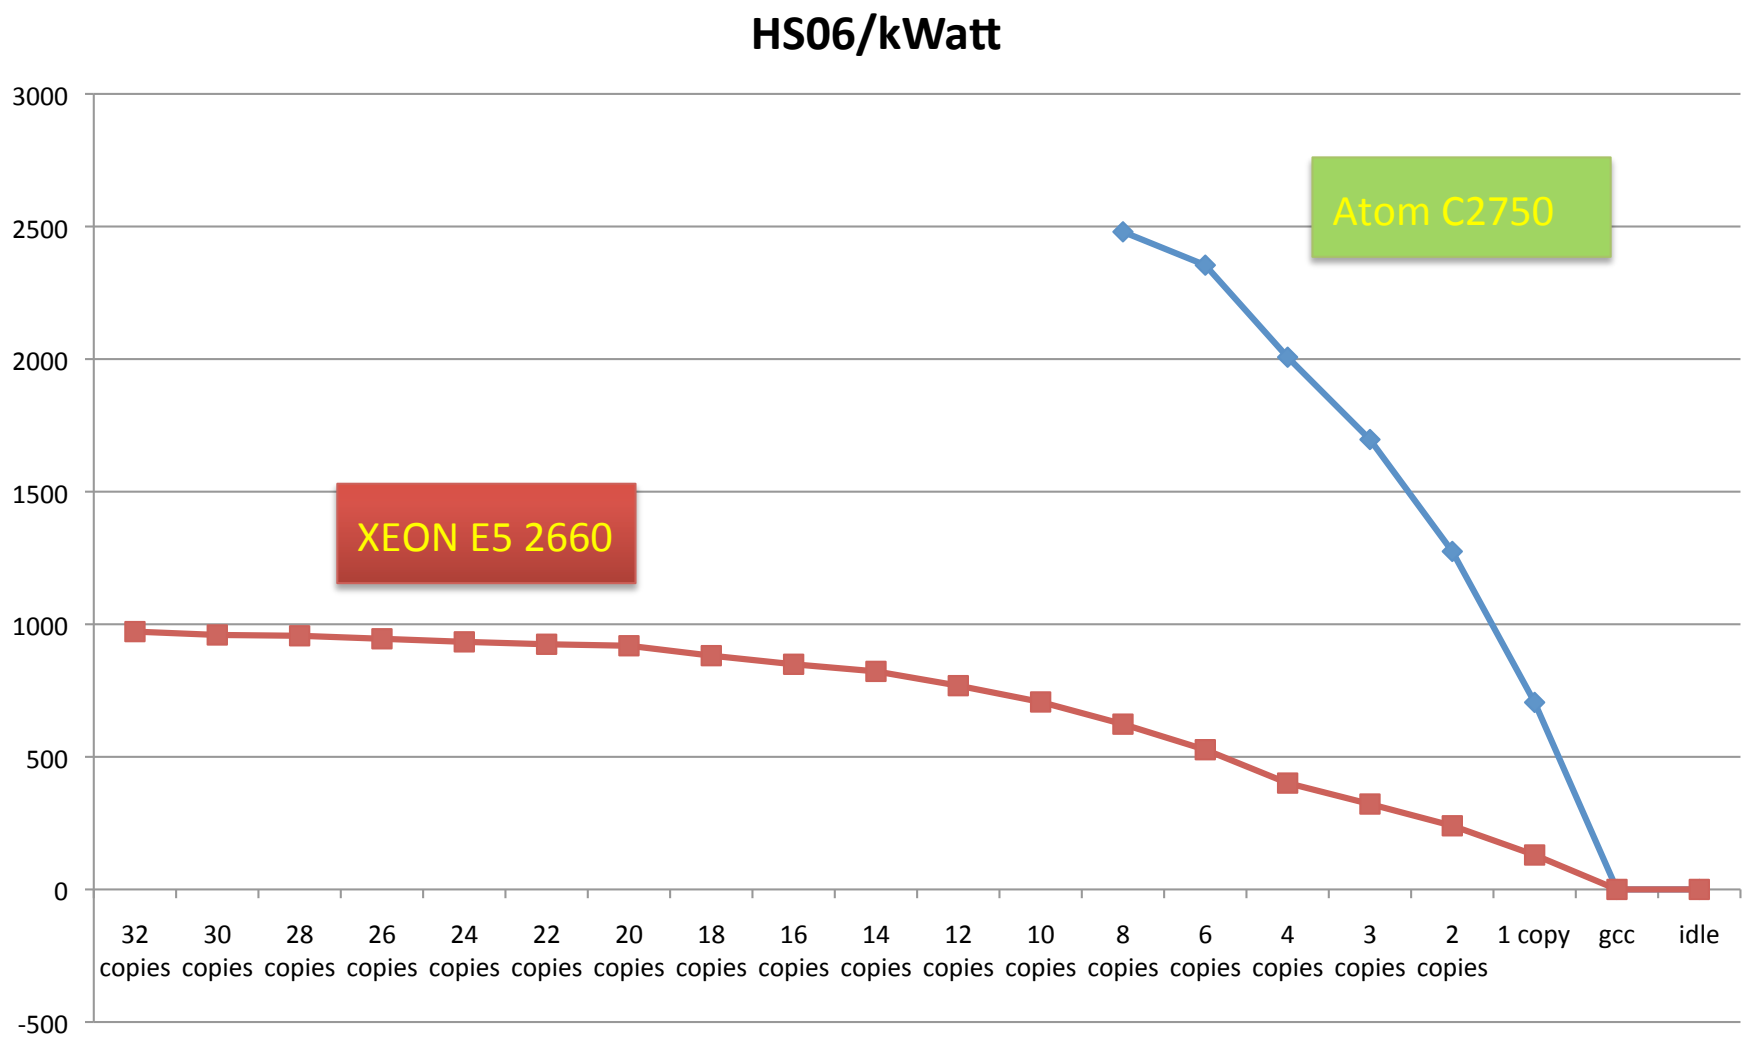

# HS06/Watt on ARM vs Intel

- I measured also the Tegra and Odroid dev kit
  - Efficiency less than Intel Atom but with cheap PSU, 20W consumed while 5W rumored
  - Waiting for next generation of ARM at 64bit

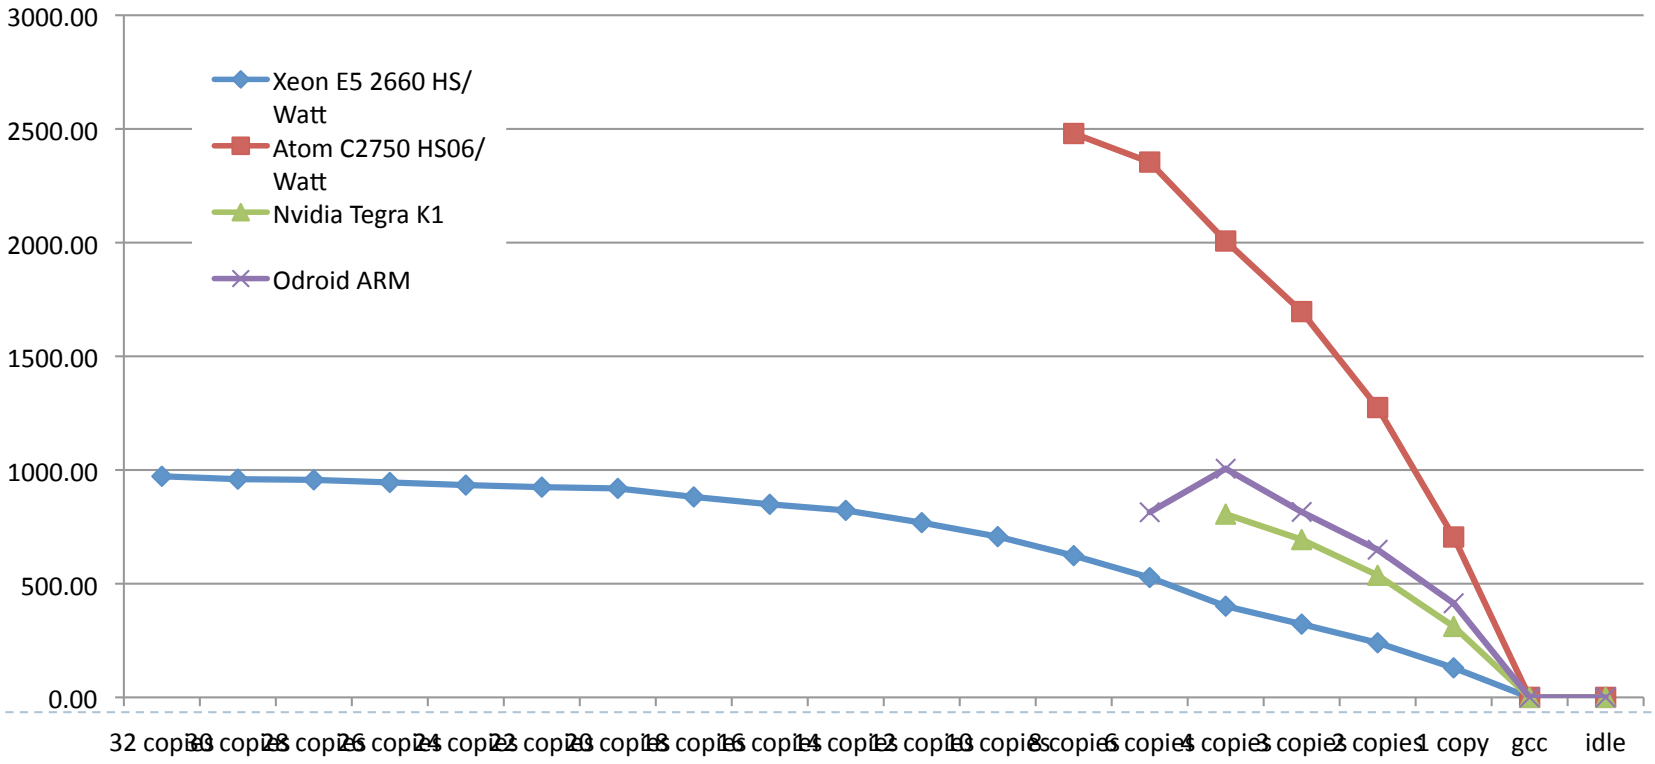

# AMD A1100 – A57 ARM

## “SEATTLE” SOC OVERVIEW

**Power Efficient Cores**
- Up to Eight ARM Cortex-A57 cores
- Up to 4MB shared L2 cache total

**Cache Coherent Network**
- Full cache coherency
- 8MB L3 cache
- SMMU: I/O address mapping and protection

**High Performance, Flexible Memory**
- Two 64-bit DDR3/4 channels with ECC
- Two DIMMs/channel up to 1866Mhz
- SODIMM, UDIMM, RDIMM support
- Up to 128GB per CPU

**Highly Integrated I/O**
- 8x SATA 3 (6Gb/s) ports
- Two 10GBASE-KR Ethernet ports
- 8 lanes PCI-Express® Gen 3, supports x8, x4, x2

**System Control Processor**
- TrustZone® technology for enhanced security
- Dedicated 1GbE system management port (RGMII)
- SPI, UART, I2C interfaces

**Cryptographic Coprocessor**
- Separate Cryptographic algorithm engine for offloading encryption, decryption, compression, decompression computations

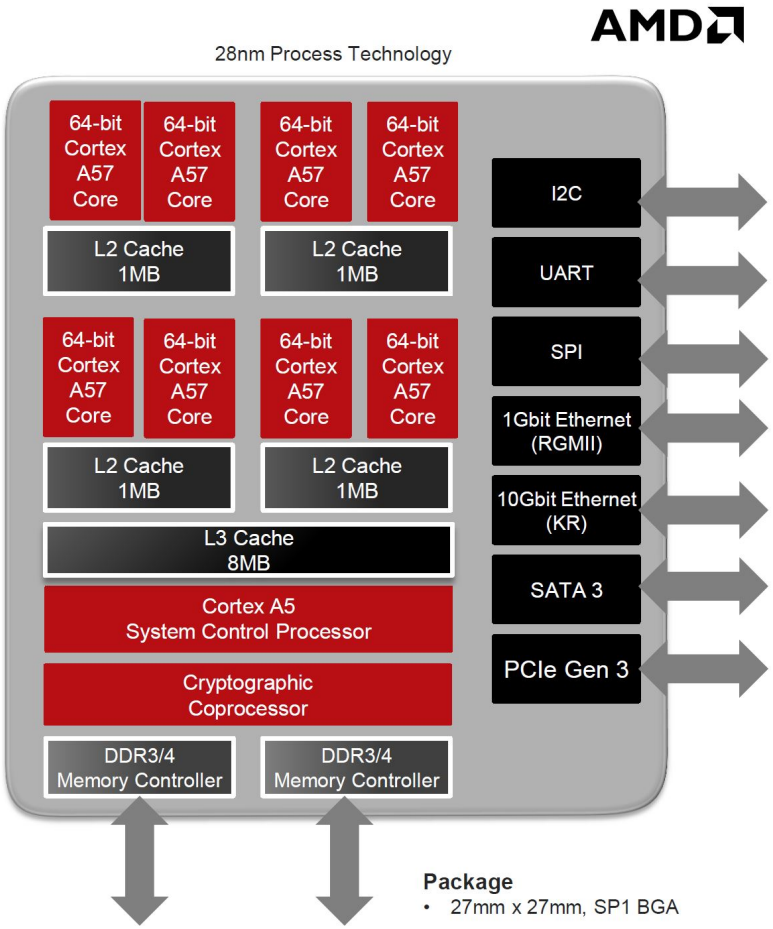

# HS06 – XeonD, Avoton, Amd A1100

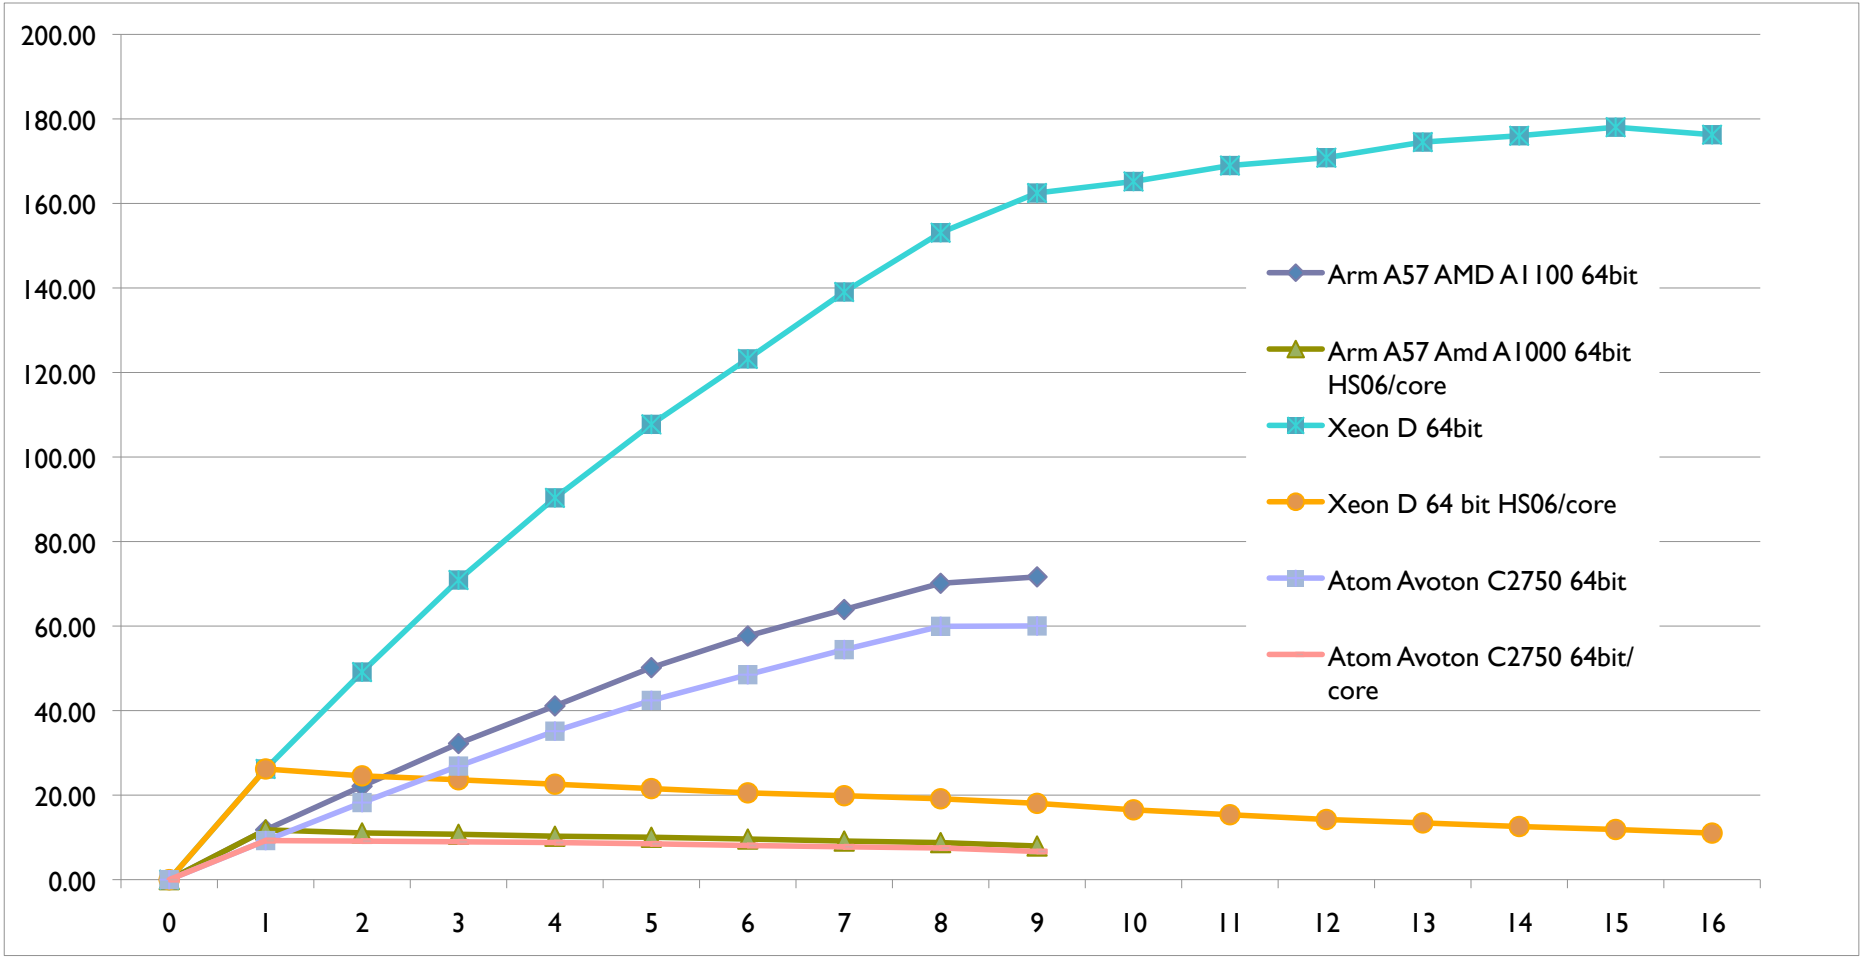

# Per core loaded

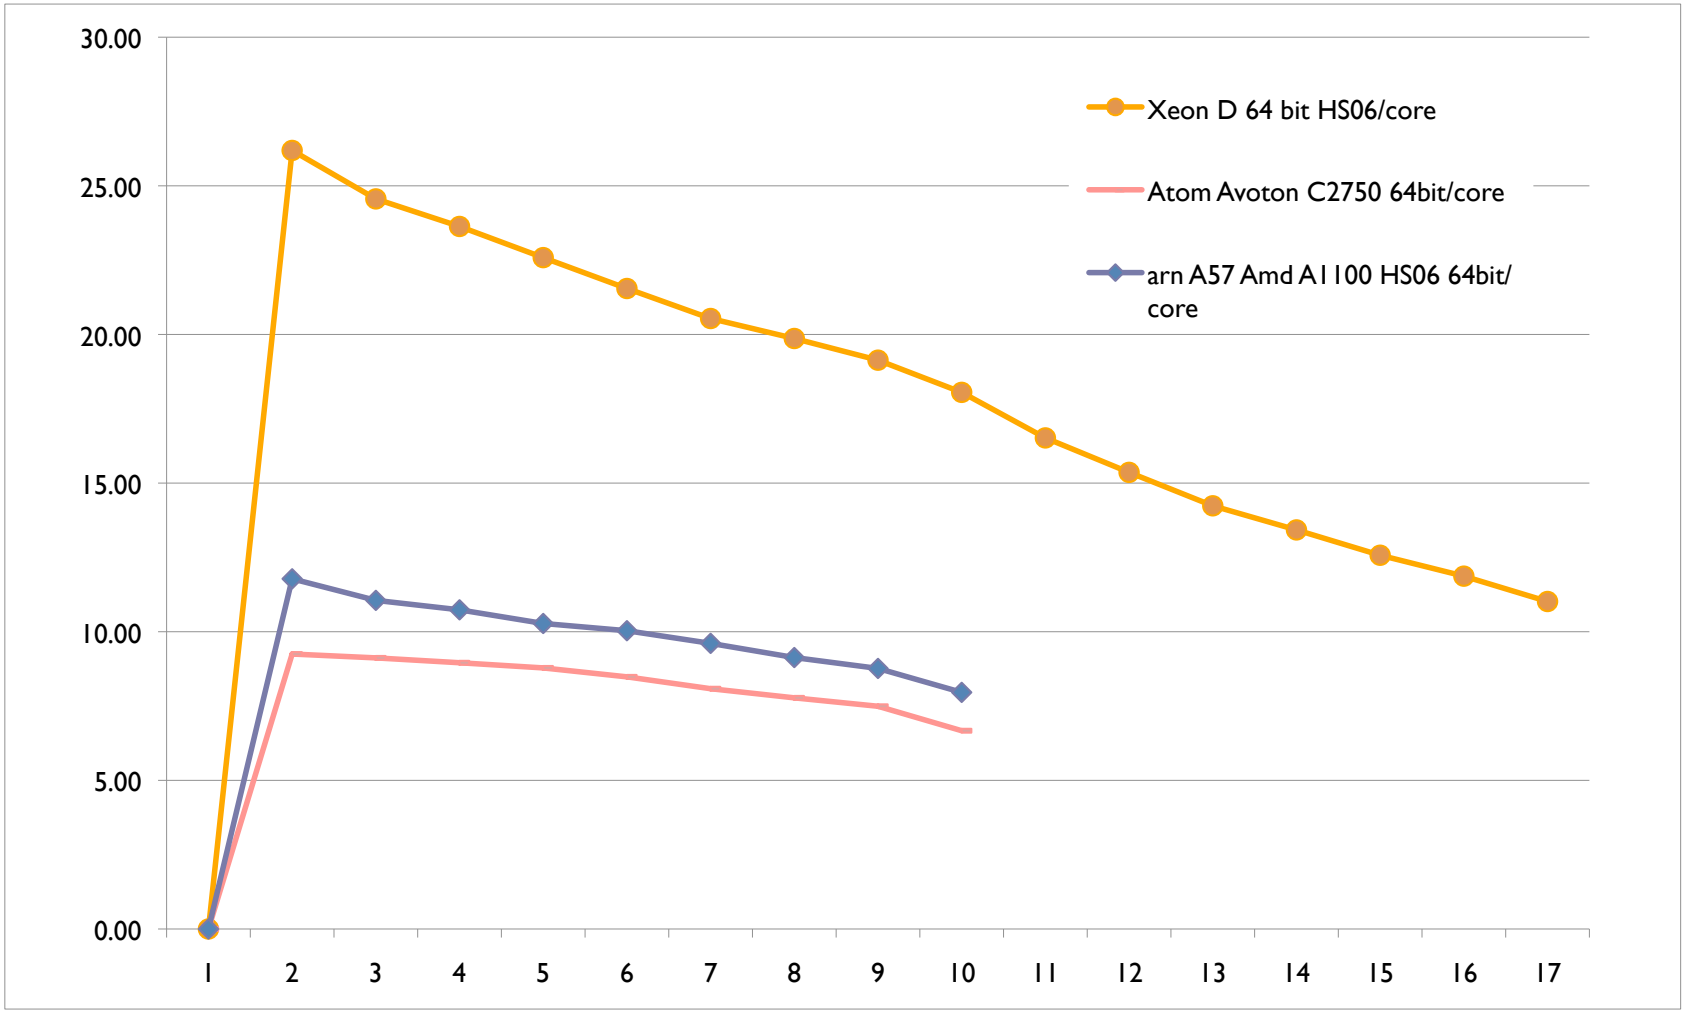

# HS06/Watt on A57 and Xeon D

- Xeon D perform as expected
- Amd A1100 – A57 has good performances but disappointing power consumption. Not yes production grade configuration. Power supply?

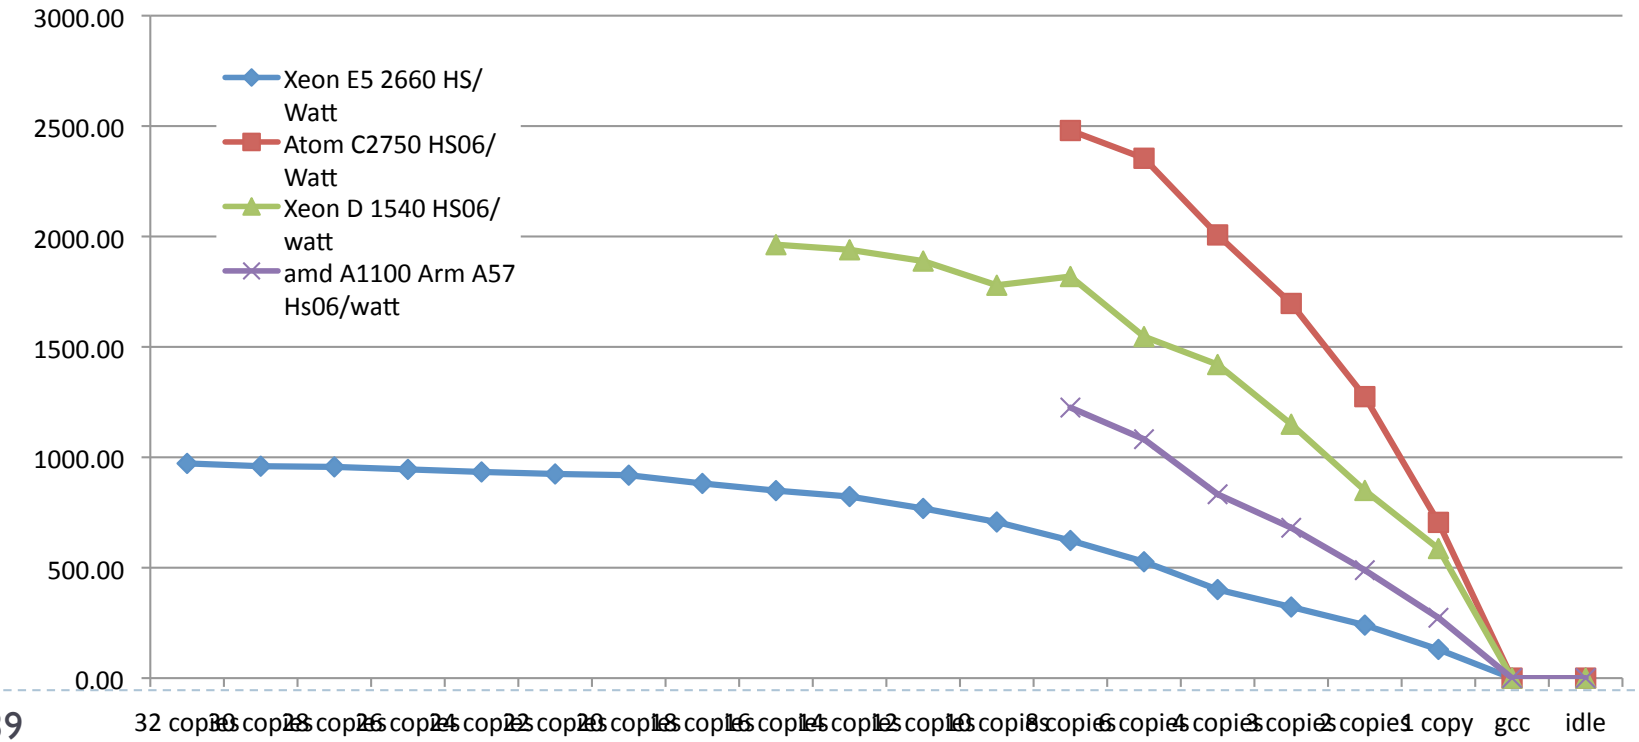

# Historical Trend with SoC

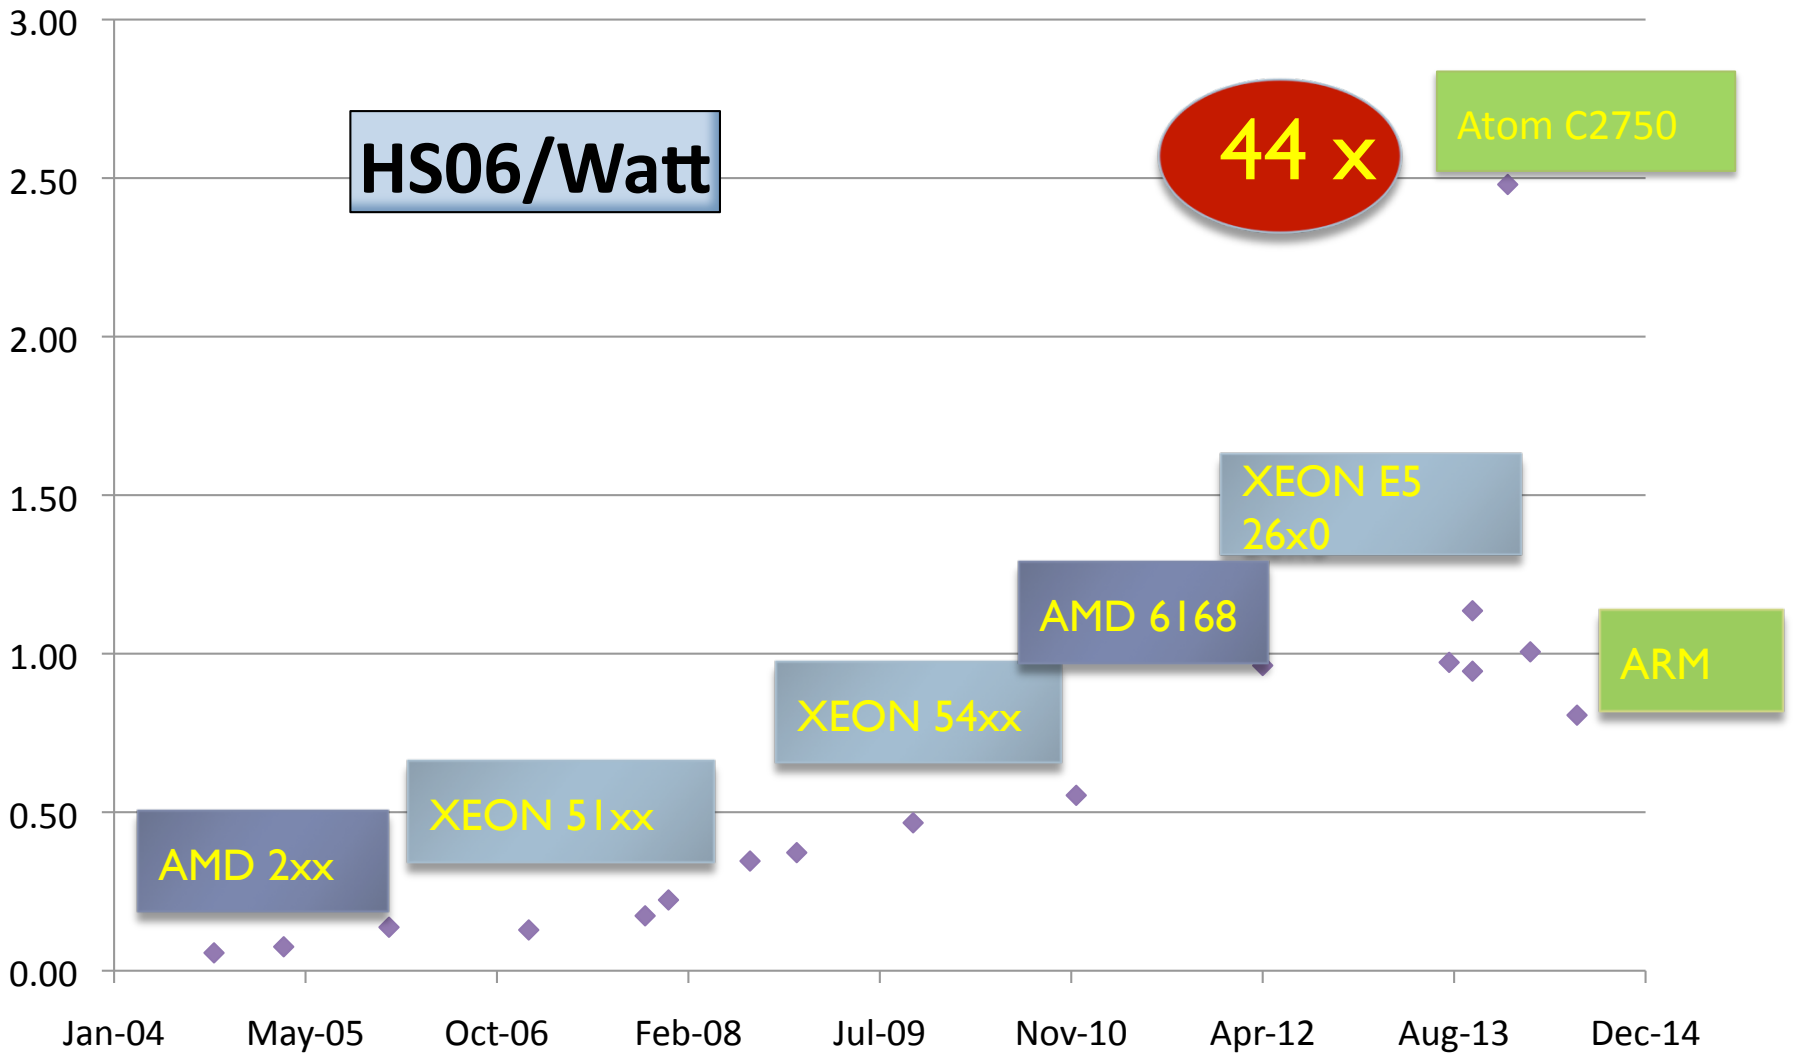

# HS06/Watt x86-64 vs LowPower

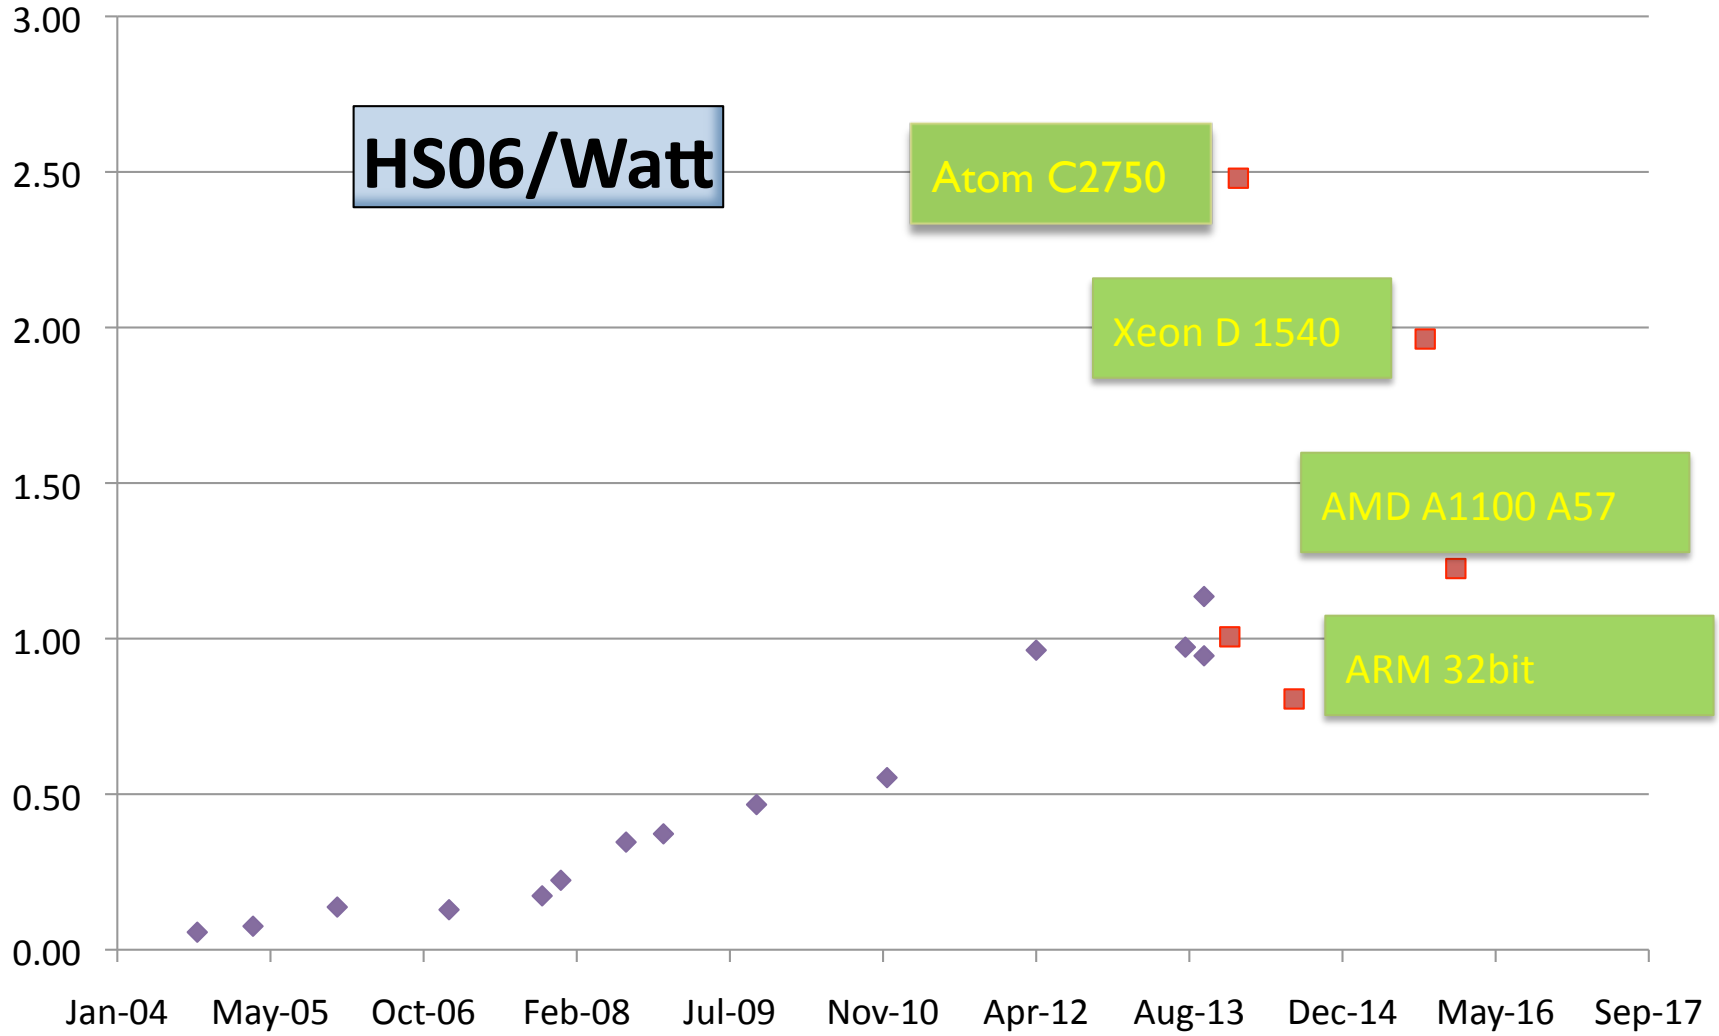

# Power consumption AMD A1100

## AMD OPTERON A1100 SKUs

| Model Number | OPN | TDP | Core Count | L2 Cache | L3 Cache | CPU Clock GHz | Max DDR3 Rate | Max DDR4 Rate | Temp Range (Tdie Max) | ECC |
|---|---|---|---|---|---|---|---|---|---|---|
| A1170 | OA1170AQD8NAD | 32W | 8 | 4M | 8MB | 2.0 | 1600 | 1866 | 0C – 80C | Yes |
| A1150 | OA1150AQD8NAD | 32W | 8 | 4M | 8MB | 1.7 | 1600 | 1866 | 0C – 80C | Yes |
| A1120 | OA1120ARD4NAD | 25W | 4 | 2M | 8MB | 1.7 | 1600 | 1866 | 0C – 80C | Yes |

All information provided for reference only and subject to change. Always refer to latest AMD's technical documentation for support details.

Base frequency with TDP = 15W; Base frequency is higher when TDP > 15W

9 | AMD OPTERON A1100 | JANUARY, 2016 | UNDER EMBARGO UNTIL JANUARY 14, 2016 @ 9:00 AM EASTERN U.S. TIME

**AMD Opteron A1100 vs. X2150**

| | CPU Core Configuration | CPU Frequency | SPECint_rate Estimate | SPECint per Core | Estimated TDP |
|---|---|---|---|---|---|
| AMD Opteron A1100 | 8 x ARM Cortex A57 | >= 2GHz | 80 | 10 | 25W |
| AMD Opteron X2150 | 4 x AMD Jaguar | 1.9GHz | 28.1 | 7 | 22W |

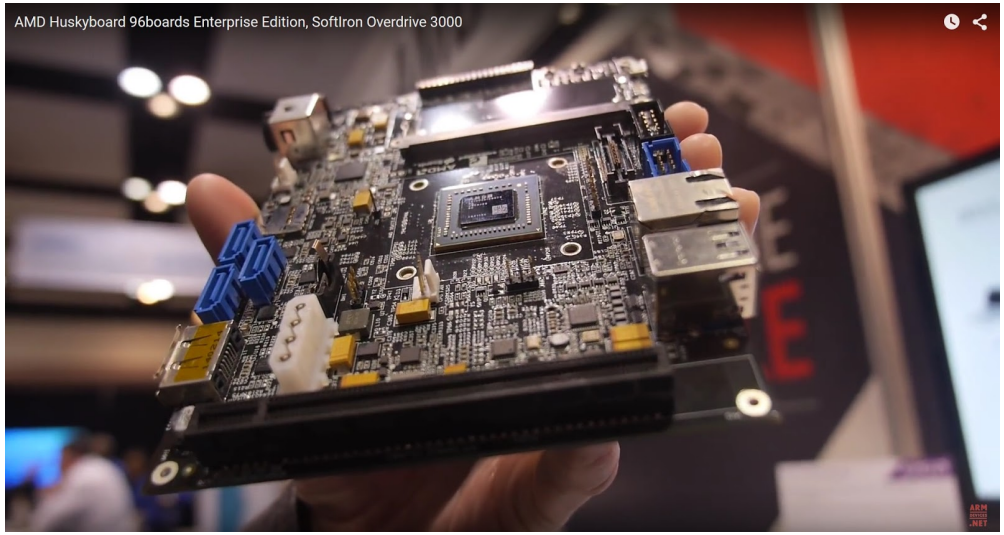

# Conclusion

- SoC processor can give at least a 2x increase in HS06/Watt in HS06 and probably in HEP code
- The number of core/processor is very small. So special high density (proprietary) chassis are needed
- ARM64 look interesting but more evaluation is necessary
- Intel is leading in term of HS06/Watt


TIME FOR QUESTIONS
Köszi
akk
iitos