

A Comparative Study of 64 Bit Microprocessors Apple, Intel and AMD

Surakshitha S

Information Science and Engineering
JSSATE
Bangalore, India

Shifa A

Information Science and Engineering
JSSATE
Bangalore, India

Ameen Rekha P M

Information Science and Engineering
JSSATE
Bangalore, India

Abstract- In this paper, we draw a comparative study of microprocessor, using the three major processors; Apple, Intel and AMD. Although the philosophy of their micro architecture is the same, they differ in their approaches to implementation. We discuss architecture issues, limitations, architectures and their features. We live in a world of constant competition dating to the historical words of the Charles Darwin-‘Survival of the fittest’. As time passed by, applications needed more processing power and this lead to an explosive era of research on the architecture of microprocessors. We present a technical & comparative study of these smart microprocessors. We study and compare two microprocessor families that have been at the core of the world’s most popular microprocessors of today – 64 bit microprocessor & Apple microprocessor.

Key words: Thread count, supported instruction set, individual power consumptions, built-in GPUs supported, instruction sets and the Clock frequency.

I. INTRODUCTION

The Microprocessor is a multipurpose, clock driven, register based, programmable electronic device that accepts digital data or binary data as input, processes it according to instructions stored in its memory, and provides results as output. Microprocessors contain both combinational logic and sequential digital logic. Microprocessors operate on numbers and symbols represented in the binary numeral system.

Here we compare 64 bit microprocessors i.e., A **64-bit computing processor** uses data path widths, integer size, and memory address widths of 64 bits (eight octets). Also, 64-bit CPU and ALU architectures are those that are based on registers, address buses, or data buses of that size. From the software perspective, 64-bit computing means the use of code with 64-bit virtual memory addresses. The term 64-Bit describes a generation of computers in which 64-bit processors are the norm. 64 bits is a word size that defines certain classes of computer architecture, buses, memory and CPUs, and by extension the software that runs on them. Without further qualification, a 64-bit computer architecture Generally has integer and addressing registers that are 64 bits wide, allowing direct support for 64-bit data types and addresses. However, a CPU might have external data buses or address buses with different sizes from the

registers, even larger. The term may also refer to the size of low-level data types, such as 64-bit floating-point numbers.

Most CPUs are designed so that the contents of a single integer register can store the address of any data in the computer's virtual considered an appropriate size to work with for another important reason: 4.29 billion integers are enough to assign unique references to most entities in applications like databases.

Some supercomputer architectures of the 1970s and 1980s used registers up to 64 bits wide. In the mid-1980s, Intel i860 development began culminating in a 1989 release. However, 32 bits remained the norm until the early 1990s, when the continual reductions in the cost of memory led to installations with quantities of RAM(random access memory) approaching 4 GB, and the use of virtual memory spaces exceeding the 4 GB ceiling became desirable for handling certain types of problems. In response, MIPS and DEC developed 64-bit microprocessor architectures, initially for high –end workstation and server machines . By the mid-1990s, the HAL Computer Systems, Sun Microsystems, IBM(international business machines), Silicon Graphics and Hewlett Packard had developed 64 –bit architectures for their workstation and server systems. A notable exception to this trend were mainframes from IBM, which then used 32-bit data and 31-bit address sizes; the IBM mainframes did not include 64-bit processors until 2000. During the 1990s, several low-cost 64-bit microprocessors were used in consumer electronics and embedded applications. Notably, the Nintendo 64 and the PlayStation 2 had 64-bit microprocessors before their introduction in personal computers. High-end printers and network equipment, as well as industrial computers, also used 64-bit microprocessors, such as the Quantum Effect Devices R5000. 64-bit computing started to drift down to the personal computer desktop from 2003 onwards, when some models in Apple Macintosh lines switched to power PC 970 processors termed G5 by Apple) and AMD released its first 64 x 86-64 bit processors.

II. 64 BIT COMPUTING

In architecture, 64 bit computer computing is the use of data path processors that have widths, integer size and memory addresses of 64 bits with eight octets wide. Also, 64-bit CPU and ALU architectures are those that are based

on registers, address buses, or data buses of that size. From the software perspective, 64-bit computing means the use of code with 64-bit virtual memory addresses.

The term 64-bit is a descriptor given to a generation of computers in which 64-bit processors are the norm. 64 bits is a word size that defines certain classes of computer architecture, buses, memory and CPUs, and by extension the software that runs on them. 64 bit CPUs have existed in supercomputers since 1970s and RISC-based workstations and servers the early 1990s, notably the DEC Alpha, Sun Ultra SPARC, Fujitsu SPARC64 and IBM PowerPC-AS. In 2003 they were introduced to the mainstream personal computer arena in the form of the x86-64 and 64-bit PowerPC processor architectures and later in 2012 even in processors that were before mainly considered only as part of embedded systems with the introduction of the Arch64 processor architectures in ARMv8.

Without further Qualification a 64-bit computer architecture generally has integrand addressing registers that are 64 bits wide, allowing direct support for 64-bit data types and addresses. However, a CPU might have external data buses or address buses with different sizes from the registers, even larger (the 32-bit Pentium had a 64-bit data bus, for instance). The term may also refer to the size of low-level data types, such as 64-bit floating-point numbers.

III. ARCHITECTURE

A. AMD

8086-64 also known as x64, x86_64 and AMD64 is the 64-bit version of the x86 instruction set. The AMD Athlon 64-bit chip is shown in fig.1. It supports vastly larger amounts theoretically, 2^{64} bytes or 16 exbi bytes of virtual memory and physical memory than is possible on its 32-bit predecessors, allowing programs to store larger amounts of data in memory. x86-64 also provides 64-bit general-purpose registers and numerous other enhancements. It is fully backward compatible with the 16-bit and 32-bit x86 code. Because the full x86 16-bit and 32-bit instruction sets remain implemented in hardware without any intervening emulation, existing x86 executables run with no compatibility or performance penalties, whereas existing applications that are recoded to take advantage of new features of the processor design may achieve performance improvements.

The primary defining characteristic of AMD64 is the availability of 64-bit general-purpose processor registers, 64-bit integer arithmetic and logical operations, and 64-bit virtual addresses. The designers took the opportunity to make other improvements as well. Some of the most significant changes are described below.

64-bit Integer capability: All general-purpose registers are expanded from 32 bits to 64 bits, and all arithmetic and logical operations, memory-to-register and register-to-memory operations, etc., can now operate directly on 64-bit integers. Push and pops on to the stack uses default to 8-byte strides, and pointers are 8 bytes wide.

Additional registers: In addition to increasing the size of the general-purpose registers, the number of named general-purpose registers is increased from eight (i.e. eax, ebx, ecx, edx, ebp, esp, esi, edi) in x86 to 16 (i.e. rax, rbx, rcx, rdx, rbp, rsp, rsi, rdi, r8, r9, r10, r11, r12, r13, r14, r15). It is therefore possible to keep more local variables in registers rather than on the stack, and to let registers hold frequently accessed constants; arguments for small and fast subroutines may also be passed in registers to a greater extent.

AMD64 still has fewer registers than many common RISC instruction sets (which typically have 32 registers) or VLIW-like machines such as the IA-64 (which has 128 registers). However, an AMD64 implementation may have far more internal registers than the number of architectural registers exposed by the instruction set.

Additional XMM (SSE) registers: Similarly, the number of 128-bit XMM registers used for Streaming SIMD instructions is also increased from 8 to 16.

Larger virtual address space: The AMD64 architecture defines a 64-bit virtual address format, of which the low-order



Fig 1: The AMD Athlon 64-bit chip

48 bits are used in current implementations. This allows up to 256 TB (2^{48} bytes) of virtual address space. The architecture definition allows this limit to be raised in future implementations to the full 64 bit extending the virtual address space to 16 EB exabytes (2^{64} bytes). This is compared to just 4 GB (2^{32} bytes) for the x86. This means that very large files can be operated on by mapping the entire file into the process' address space which is often much faster than working with file read/write calls, rather than having to map regions of the file into and out of the address space.

Larger physical address space: The original implementation of the AMD64 architecture implemented 40-bit physical addresses and so could address up to 1 TB (2^{40} bytes) of RAM. Current implementations of the AMD64 architecture (starting from AMD 10h micro architecture) extend this to 48-bit physical addresses and therefore can address up to 256 TB of RAM. Fig. 2 shows picture of Athlon 64 and Pentium IV. The architecture permits extending this to 52 bits in the future (limited by

the page table entry format); this would allow addressing of up to 4 PB of RAM. For comparison, 32-bit x86 processors are limited to 64 GB of RAM in Physical Address Extension (PAE) mode, or 4 GB of RAM without PAE mode.

Larger physical address space in legacy mode: When operating in legacy mode the AMD64 architecture supports Physical Address Extension (PAE) mode, as do most current x86 processors, but AMD64 extends PAE from 36 bits to an architectural limit of 52 bits of physical address. Any implementation therefore allows the same physical address limit as under long mode.

Instruction pointer relative data access:

Instructions can now reference data relative to the instruction pointer (RIP register). This makes position independent code, as is often used in shared libraries and code loaded at run time, more efficient.

B. INTEL

Intel architecture chips have obviously undergone

many changes over the past 40+ years. A list of currently available Devices is available here. Early chips were given technical part numbers, such as 8086, 80386, or 80486. This led to the commonly used shorthand of 86 architecture, in reference to the last two digits of each chip's part number. Beginning in 1993, the "86" naming Convention gave way to more memorable product names such as Intel Pentium processor, Intel Celeron processor, Intel Core processor, and Intel Atom processor. Fig.3 shows the block diagram of Intel Itanium 2 (64 bit) Processor.

Although every branch of the broad Intel architecture (or x86) family tree retains the same basic features and functionality as the earlier chips, and retains backward compatibility with them, each new generation also adds its own unique features to the mix. For example, Intel Pentium processor added multimedia extensions that accelerated audio and video processing. Extended temperature Intel Pentium processor with MMX technology (multimedia technology) is with more streaming-media capabilities known as Intel Streaming SIMD (Single instruction Multiple data) Extensions and Intel Streaming SIMD Extensions 2. Floating-point units (FPUs) Went from optional upgrade to standard feature of Intel architecture processors, and today encryption/decryption extensions, power-management features, and multilevel caches are now found on most Intel architecture processors. Data paths have widened from 8 bits to 32 bits, 64 bits, and even 128 bits and more.

Operating frequencies have jumped from a few megahertz to 2 GHz (two billion cycles per second) and beyond.

C. APPLE

The **Apple A7** is a package on package (PoP) 64-bit system-on-a-chip (SoC system on a chip) designed by Apple. Its first appearance was in the iPhone 5S, which

was introduced on September 10, 2013. The chip would also be used in the iPad Mini 2 and iPad Mini 3. Apple states that it is up to twice as fast and has up to twice the graphics power compared to its predecessor the Apple A6. The A7 features an Apple-designed 1.3-1.4GHz 64-bit ARMv8-A dual-core CPU, called Cyclone, and an integrated Power VR G6430 GPU (graphics processing unit) in a four cluster configuration. It now has 31 general purpose registers that are each 64-bits wide and 32 floating-point/NEON registers that are each 128-bits wide. The **Apple A8** is a package on package (PoP) 64-bit system-on-a-chip (SoC) designed by Apple and manufactured by.




Fig 2: Picture of an Athlon 64 and Pentium IV

TSMC. Its first appearance was in the iPhone 6 and iPhone 6 Plus. A year later it would drive the iPad Mini 4. Apple states that it has 25% more CPU performance and 50% more graphics performance while drawing only 50% of the power compared to its predecessor, the Apple A7.

The A8 features an Apple-designed 1.4 GHz 64-bit ARMv8-A dual-core CPU. The A8 is manufactured on a 20 nm process by TSMC, which replaced Samsung as the manufacturer of Apple's mobile device processors. It contains 2 billion transistors. Despite that being double the number of transistors compared to the A7, its physical size has been reduced by 13% to 89 mm² (consistent with a shrink only, not known to be a new micro architecture). Apple A8X.

The **Apple A8X** is a 64-bit system on a chip (SoC) designed by Apple, introduced at the launch of the iPad Air 2. It is a high performance variant of the Apple A8. Apple states that it has 40% more CPU performance and 2.5 times the graphics performance of its predecessor, the Apple A7.

Unlike the A8, this SoC uses a triple-core CPU, a new octa-core GPU, dual channel memory and slightly higher 1.5 GHz CPU clock rate. It uses an integrated octa-core

Power VR GX6850 graphics processing unit (GPU) running at 450 MHz and a dual-channel memory subsystem.¹ It's manufactured by TSMC on their 20 nm fabrication process, and consists of 3 billion transistors.

The **Apple A9** is a 64-bit ARM based system on a chip (SoC) designed by Apple Inc. Fig. 4 shows the Apple A9 chip .It first appeared in the iPhone 6S and 6S Plus. Apple states that it has 70% more CPU performance and 90% more graphics performance compared to its predecessor, the Apple A8.

The **Apple A9X** is a 64-bit system on a chip (SoC) designed by Apple Inc. It first appeared in the iPad Pro. It purportedly offers double the memory bandwidth and double the storage performance of its predecessor, the Apple A8X.




Fig.3: Block diagram of Intel Itanium 2 (64 bit) Processor

IV. FEATURES

A. AMD

There are four variants: Athlon 64, Athlon 64 FX, Mobile Athlon 64 and the dual-core Athlon 64 X2. All Athlon 64s also support the NX bit, a security feature named "Enhanced Virus Protection" .And as implementations of the AMD64 architecture, all Athlon 64 variants are able to run 16 bit, 32 bit x86, and AMD64 code, through two different modes the processor can run in: Legacy mode and long mode. Legacy mode runs 16-bit and 32-bit programs natively, and long mode runs 64-bit programs natively, but also allows for 32-bit programs running inside a 64-bit operating system. All Athlon 64 processors feature 128 Kilobytes of level 1 cache, and at least 512 KB of level 2 caches.

On-die memory controller: The Athlon 64 features an on-die memory controller, a feature seen on only the Trans Meta Crusoe. It means the controller runs at the same clock rate as the CPU itself, it also means the electrical signals have a shorter physical distance to travel compared to the old north bridge interfaces .The result is a significant reduction in response time for access requests to main memory.

Memory and HT Northbridge buses: As the memory controller is integrated onto the CPU die, there is no FSB(front side bus) for the system memory to base its speed upon. Instead, system memory speed is obtained by using the following formula (using the ceiling function):

$$\frac{\text{CPU speed}}{\lceil \frac{\text{CPU multiplier}}{\text{DRAM divider}} \rceil} = \text{DRAM speed}$$

So, the memory is always running at a set fraction of the CPU speed, with the divisor being a whole number .A second bus, the north bridge, connected the CPU to the chipset and device



Fig. 4: The Apple A9 chip

attachment bus. This was implemented using a new high-performance standard, Hyper Transport. It was also useful in building multi-processor systems without additional glue chips.

Translation Look aside Buffers : Translation Look aside Buffers have also been enlarged (40 4k/2M/4M entries in L1 cache, 512 4k entries), with reduced latencies and improved branch prediction. This and other architectural enhancements, especially as regards SSE(System Security Engineering) implementation, improve instruction per cycle (IPC) performance over the previous Athlon XP generation. To make this easier for consumers to understand, AMD has chosen to market the Athlon 64 using a PR (Performance Rating) system, where the numbers roughly map to Pentium 4 performance equivalents, rather than actual clock speed.

Cool ' n 'Quiet : Athlon 64 also features CPU speed throttling technology branded *Cool ' n ' Quiet*, a feature similar to Intel's *Speed Step* that can throttle the processor's clock speed back to facilitate lower power consumption and heat production.^[51] When the user is running undemanding applications and the load on the processor is light, the processor's clock speed and voltage are reduced. This in turn reduces its peak power to as low as 32 W or 22W The Athlon 64 also has an Integrated Heat Spreader (IHS) which prevents the CPU die from accidentally being damaged when mounting and amounting

heat sinks. With prior AMD CPUs a CPU shim could be used by people worried about damaging the die.

NX bit: The No Execute bit (NX bit) supported by Windows XP Service Pack 2 and future versions of Windows, Linux 2.6.8 and higher and FreeBSD 5 for improved protection from malicious buffer overflow security threats. Hardware-set permission levels make it much more difficult for malicious code to take control of the system. It is intended to make 64-bit computing a more secure environment.

Semiconductor Technology: The Athlon 64 CPUs have been produced with 130 nm and 90 nm SOI process technologies. The feature of an amalgam of strained silicon and 'squeezed silicon', co-developed with IBM.

B. INTEL

Intel 64 architecture delivers 64-bit computing in embedded designs when combined with supporting software. Intel 64 architecture improves performance by allowing systems to address more than 4 GB of both virtual and physical memory.

Intel® 64 provides support for:

Some features of the 64-bit Intel architecture are

- 64-bit flat virtual address space
- 64-bit pointers
- 64-bit wide general purpose registers
- 64-bit integer support
- Allows addressing of up to 2^{64} bytes of memory or 16 EB (Exabyte) of RAM.

C. APPLE

True 64-bit microprocessor. 64-bit integer operations, 64-bit floating-point operations, 64-bit registers, 64-bit virtual address space.

High-performance microprocessor

150 peak MIPS (machine instructions per second) at 150MHz, 50 peak MFLOP/s at 150MHz, 96 SPECint92 at 150MHz, Two-way set associative caches

High level of integration

64-bit integer CPU, 64-bit floating-point unit, 16KB instruction cache; 16KB data cache - Flexible MMU with large TLB.

Low-power operation

3.3V or 5V power supply options, 20mW/MHZ typical internal power dissipation (2.0W @ 100MHz, 3.3V), Standby mode reduces internal power to 90mA to 400mW

•Fully software compatible with R4000 RISC (Reduced Instruction Set Computers) Processor Family

•Available in R4000PC/R4000PC pin-compatible 179-pin PGA or 208-pin MQUAD (Mega Quad(4)).

•50MHz, 67MHz, 75MHz input frequencies with mode bit dependent output clock frequencies, On-chip clock double for 150MHz pipeline

64GB physical address space

•Processor family for a wide variety of applications

Desktop workstations and PCs, Desk side or departmental servers, High-performance embedded applications (e.g. Color printers, multi-media and internetworking.), Notebooks.

V. LIMITATIONS

The AMD64 architecture as of 2011 allows 52 bits for physical memory and 48 bits for virtual memory. These limits allow memory sizes of 4 PB (4×10^{24} bytes) and 256 TB (256×10^{12} bytes), respectively. A PC cannot contain 4 petabytes of memory (due to the physical size of the memory chips, if nothing else) but AMD envisioned large servers, shared memory clusters, and other uses of physical address space that might approach this in the foreseeable future, and the 52-bit physical address provides ample room for expansion while not incurring the cost of implementing 64-bit physical addresses. Similarly, the 48-bit virtual address space was designed to provide more than 65,000 times the 32-bit limit of 4 GB (4×10^9 bytes), allowing room for later expansion without incurring the overhead of translating full 64-bit addresses.

Intel32 and Intel64 both have the same static data memory space i.e., about 2 GB each. Note that the limit on static and stack data is the same in both 32-bit and 64-bit variants. This is due to the format of the Windows Portable Executable (PE) file type, which is used to describe EXEs and DLLs (Data link layer) as laid out by the linker. It has 32-bit fields for image section offsets and lengths and was not extended for 64-bit variants of Windows. As on 32-bit Windows, static data and stack share the same first 2GB of address space.

If you exceed the limit on static code and data you may see the following error when the application is linked.

The other disadvantage of going to 64-bit is that your code size and DRAM fetches all inflate somewhat. Good code practices can keep this down, but compare any 32-bit program to a 64-bit version, and you'll find that the 64-bit flavor is slightly larger. That makes power efficiency more difficult, though Apple can still extract net gains here due to more efficient use of processor characteristics.

VI. CASE STUDY APPLICATIONS

Apple microprocessor chips are used in modern day mobile phones A7 in iPhone 5S, A8 in iPhone 6, A9 in iPhone 6S and iPhone 6S plus. These phones are fast processors and offer a variety of services when it comes to App – based technology, few can be mentioned like

Photoshop applications , or even using your Google Drive in your phone. One important feature is it helps in carrying out different jobs in different applications at the same time, simultaneously.

Intel of course, has been a processor favorite since a longtime. The latest updates come in the most recently released core processor chip i7. It comes a new platform architecture , higher performance multiprocessor systems with QPI(Quick path Interconnect) , improved cache latency response time with the smart L3 Cache , CPU boost performance through Intel boost facility , optimized performances through hyper threading and improved virtualization performance . Fig. 5 and Fig.6 shows Intel’s approach – 12 cores with 6 shared cores and AMD’S approach -12 cores with 12 dedicated threads respectively.

VII. CONCLUSION

We finally conclude Comparing all the three major Microprocessors – Apple, Intel and AMD families using this table alongside where we have compared various important parameters that contribute to the performance of a microprocessor in the current trend .The three micro-processor families described here really compete for the top spot in modern high performance microprocessors. The latest Intel, Apple and AMD microprocessor seem to be very close in terms of performance. Table 1 shows restricted briefing of differences between the considered 32 bit processors using restricted number of parameters of interest. In the table DDR stands for double data rate 3(memory).




Fig 5: Intel’ s approach – 12 cores with 6 shared cores




Fig 6 : AMD ‘s approach – 12 cores with 12 dedicated threads

Table 1:Comparative Study of Microprocessors

S/N	FEATURES	AMD (OPTERON 6300)	INTEL (ITANIUM 2)	APPLE (A7)
1.	Supported instruction set	MMX, SSE, SSE2, SSE3, SSE4	SSE	RISC (A64, T32, A32)
2.	FSB (Front Signal Bus)	1000MHz	773MHz	
3.	Clock Frequency	2.4	1.66	1.32
4.	Thread count	16	2	2
5.	Built in Graphics	No	Yes	Yes (g6340)
6.	Number of CPU cores	16	8	2
7.	XOP Instructions	Included	Excluded	Excluded
8.	Hardware channels	DDR3	DDR3	DDR3
9.	On-chip L2 + L3 cache	16+16MB	22 MB	5 MB
10.	Power consumption	115W	115W	4W

VIII. REFERENCES

- [1] Hadi , "Tablet ward: Apple ipad Vs blackberry playbook vs Samsung galaxy Tab", Setember 28,2010
- [2] Hewlett-Packard Development Company. AMD Opteron and Intel Xeon x86 for CPUs. 21st Euromicro International Conference on Parallel, Distributed, and processors in industry-standard servers. Technology 2012.
- [3] Kamanashis B. & Ashraful Md. I., "Hardware Virtualization Support in Intel, AMD And IBM Power Processors. (IJCSIS) International Journal of Computer Science and Information Security Vol. 4, No. 1 & 2, 2009
- [4] Weng, Ning, and Tilman Wolf. "Pipelining vs. multiprocessors choosing the right network processor system topology." Proc. of Advanced Networking and Communications Hardware Workshop (ANCHOR 2004) 2004.
- [5] Jing Fu and Hagsand, Olof, " Designing and Evaluating Network Processor Applications", IEEE Workshop on High Performance Switching and Routing: Hong Kong, PEOPLES R CHINA, MAY 12-14, 2005, 2005, 142-146
- [6] Crowley, Patrick "Characterizing processor architectures for programmable network interfaces.", ACM International Conference on Supercomputing 25th Anniversary Volume. 2014.
- [7] Danijela Jakimovska —Performance Estimation of Novel 32-bit and 64-bit RISC based Network Processor Cores, (JSAT), May Edition, 2011.
- [8] Erika Brown, E. C. (2008, 04 23). Apple buys Chip Designer. Retrieved 05 04, 2008, from Forbes: http://www.forbes.com/technology/2008/04/23/apple-buys-pasemi-tech-ebiz-cz_eb_0422apple.html
- [9] P .Crowley "Optimal network processor topologies for efficient packet process processing", IEEE GLOBECOM 2005.
- [10] Jing Fu and Olof Hagsand "Designing and evaluating network processor applications", IEEE 2005