# Study of Microprocessor and AMD FX Series

Rochan Mehrotra, Sarthak Bhudhiraja
*Information technology*
*Dronacharya College of Engineering, Gurgaon, India*

*Abstract*— **As personal computers have become more prevalent and more applications have been designed for them, the end-user has seen the need for a faster, more capable system to keep up. The computing power and clock frequency of a single processor reached their peak a few years ago, it just isn't easy to create more powerful and/or faster processors than the current ones; so the major CPU manufacturers (Intel, AMD) switched strategy and went multicore. In order to increase clock speeds, the silicon transistors on the chip need to be able to switch faster. Although chip speed has increased exponentially over the years, that time is ending and manufacturers have shifted toward multicore**

Fig. 01. CPU transistor counts against dates of introduction

**processing that result in greater leakage, power consumption and heat output. You eventually reach a point where you cannot increase clock rates any further without requiring excessive amounts of power or using exotic cooling solutions.**

## I. INTRODUCTION

The trend of increasing a processor's speed to get a boost in performance is a way of the past. Multicore processors are the new direction manufacturers are focusing on. Using multiple cores on a single chip is advantageous in raw processing power, but nothing comes for free.

*a. Moore's Law*

One of the guiding principles of computer architecture is known as Moore's Law. In 1975 Gordon Moore stated that the number of transistors on a chip will roughly double in two year. The graph in Figure 1 plots many of the microprocessors against the number of transistors per chip.

If you were to look at processor speeds from the 1970's to 2009 and then again in 2010, one may think that the law has reached its limit or is nearing the limit. In the 1970's processor speeds ranged from 740 KHz to 8MHz; notice that the 740 is KHz, which is Kilo Hertz – while the 8 is MHz, which is Mega Hertz.

From 2000 – 2009 there has not really been much of a speed difference as the speeds range from 1.3 GHz to 2.8 GHz, which suggests that the speeds have barely doubled within a 10 year span. This is because we are looking at the speeds and not the number of transistors; in 2000 the number of transistors in the CPU numbered 37.5 million, while in 2009 the number went up to an outstanding 904 million; this is why it is more accurate to apply the law to transistors than to speed. On 13 April 2005, Gordon Moore stated in an interview that the law cannot be sustained indefinitely: "It can't continue forever. The nature of exponentials is that you push them out and eventually disaster happens". He also noted that transistors would eventually reach the limits of miniaturization at atomic levels. In terms of size of transistors you can see that we're approaching the size of atoms which is a fundamental barrier, but it'll be two or three generations before we get that far—but that's as far out as we've ever been able to see. We have another 10 to 20 years before we reach a fundamental limit. By then they'll be able to make bigger chips and have transistor budgets in the billions.

## II. NEED OF MULTICORE PROCESSOR

The computing power and clock frequency of a single processor reached their peak a few years ago, it just isn't easy to create more powerful and/or faster processors than the

current ones; so the major CPU manufacturers (Intel, AMD) switched strategy and went multicore. In order to increase clock speeds, the silicon transistors on the chip need to be able to switch faster. These higher speeds require higher input voltages and semiconductor manufacturing processes that result in greater leakage, both of which increase power consumption and heat output. You eventually reach a point where you cannot increase clock rates any further without requiring excessive amounts of power or using exotic cooling solutions.

The trend towards multiple cores is an engineering approach that helps the CPU designers avoids the power consumption problem that came with ever increasing frequency scaling. As CPU speeds rose into the 3-4 Ghz range the amount of electrical power required to go faster started to become prohibitive.

### III. MULTICORE CHALLENGES

Having multiple cores on a single chip gives rise to some problems and challenges. Power and temperature management are two concerns that can increase exponentially with the addition of multiple cores. Memory/cache coherence is another challenge, since all designs discussed above have distributed L1 and in some cases L2 caches which must be coordinated. And finally, using a multicore processor to its full potential is another issue. If programmers don't write applications that take advantage of multiple cores there is no gain, and in some cases there is a loss of performance. Application need to be written so that different parts can be run concurrently.

*a. Power and Temperature*

If two cores were placed on a single chip without any modification, the chip would, in theory, consume twice as much power and generate a large amount of heat. In the extreme case, if a processor overheats your computer may even combust. To account for this each design above runs the multiple cores at a lower frequency to reduce power consumption. To combat unnecessary power consumption many designs also incorporate a power control unit that has the authority to shut down unused cores or limit the amount of power. By powering off unused cores and using clock gating the amount of leakage in the chip is reduced

To lessen the heat generated by multiple cores on a single chip, the chip is architected so that the number of hot spots doesn't grow too large and the heat is spread out across the chip. The majority of the heat in the CELL processor is dissipated in the Power Processing Element and the rest is spread across the Synergistic Processing Elements. The CELL processor follows a common trend to build temperature monitoring into the system, with its one linear sensor and ten internal digital sensors.

*b . Cache Coherence*

Cache coherence is a concern in a multicore environment because of distributed L1 and L2 cache. Since each core has its own cache, the copy of the data in that cache may not always be the most up-to-date version. For example, imagine a dual-core processor where each core brought a block of memory into its private cache. One core writes a value to a specific location; when the second core attempts to read that value from its cache it won't have the updated copy unless its cache entry is invalidated and a cache miss occurs. This cache miss forces the second core's cache entry to be updated. If this coherence policy wasn't in place garbage data would be read and invalid results would be produced, possibly crashing the program or the entire computer. In general there are two schemes for cache coherence, a snooping protocol and a directory-based protocol. The snooping protocol only works with a bus-based system, and uses a number of states to determine whether or not it needs to update cache entries and if it has control over writing to the block. The directory-based protocol can be used on an arbitrary network and is, therefore, scalable to many processors or cores, in contrast to snooping which isn't scalable. In this scheme a directory is used that holds information about which memory locations are being shared in multiple caches and which are used exclusively by one core's cache. The directory knows when a block needs to be updated or invalidated.

AMD's Athlon 64 has to monitor cache coherence in both L1 and L2 caches. This is sped up using the Hyper Transport connection.

*c. Multithreading*

The last, and most important, issue is using multithreading or other parallel processing techniques to get the most performance out of the multicore processor. "With the possible exception of Java, there are no widely used commercial development languages with [multithreaded] extensions." Rebuilding applications to be multithreaded means a complete rework by programmers in most cases. Programmers have to write applications with subroutines able to be run in different cores, meaning that data dependencies will have to be resolved or accounted for (e.g. latency in communication or using a shared cache). Applications should be balanced. If one core is being used much more than another, the programmers is not taking full advantage of the multi-core system. Some companies have heard the call and designed new products with multicore capabilities; Microsoft and Apple's newest operating systems can run on up to 4 cores.

## IV. AMD FX-SERIES

A few weeks ago AMD announced the latest additions to their FX-series processor lineup: the FX-9590.The FX-9590's high clock speeds mean each individual processor has to go through a high level screening process to validate its ability to consistently hit its Turbo target of 5GHz and base clock of 4.7GHz. Not many Vishera cores will make the grade and these "failed" CPUs will eventually be rolled into the FX-9370's embrace and may even cascade down into other FX SKUs. Even the FX-9370 has a stringent binning process in an effort to weed out any unfit chips.

At its heart, the FX-9590 isn't all that much different from previous FX-series CPUs. Like the FX-8350, it uses AMD's 32nm Piledriver architecture, comes with 8MB of L2 cache and eight cores, supports DDR3 speeds of up to 1866MHz and can be used in conjunction with any supporting AM3+ motherboard. The 5GHz core frequency has been attained through the use of AMD's Turbo Core 3.0 which allows a base clock of



4.7GHz to hit higher levels when the right conditions present themselves. Unfortunately, that doesn't happen very often as our processor remained at 4.7GHz in most workloads. As one might expect, actually getting a lower clocked architecture to hit such high levels requires some heavy-duty muscle alongside stringent binning. In this case, a massive amount of voltage - 1.5V- has been applied and this has a secondary, nasty side effect: a substantial increase in heat production and power requirements. While the FX-8350's TDP of 125W was deemed inefficient, the FX-9590 brings things to a whole new level with an estimated thermal output of 200W-220W. That's an important number to remember when choosing a cooler since only the very best air-based solutions will be able to keep temperatures under control.

### a. *System Power Consumption*

For the CPU power consumption test, we use the standard testing system (with an NVIDIA GTX 670 installed) and wait until the system and discrete GPU are at idle speeds in order to log the idle power consumption. After this, WPrime 1024M is looped for 15 minutes while the power consumption is logged with a calibrated power meter to determine the peak watts. We have found that simply plugging in a power meter to a wall outlet or UPS will NOT give you accurate power consumption numbers due to slight changes in the input voltage. Thus we

use a Tripp-Lite 1800W line conditioner between the 120V outlet and the power meter.

With a TDP of about 220W, it the FX-9590's chart-topping performance shouldn't come as any surprise. Pushing the 32nm Piledriver architecture to 4.7GHz (or 5GHz in some situations) takes a fair amount of voltage and no small amount of excess PSU capacity. As a result, the FX-9590 requires about 70W more power than an FX-8350 while the current generation Haswell processors aren't even in the picture. Unfortunately this leads to AMD's latest processor having an extremely poor performance per watt ratio.

### b. *Temperature Testing*
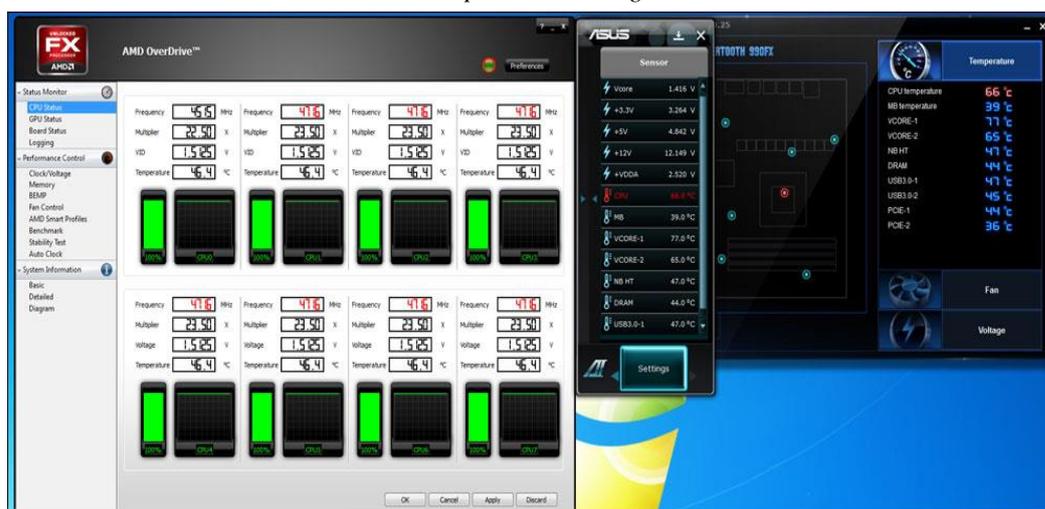
The FX-9590 is a hot running processor and we don't mean hot in any conventional meaning of the word either. This thing is like having a miniature nuclear reactor strapped to your motherboard. It claimed the chip idled at 19°C (ambient temperature was 23°) while load temperatures. Our FX-9590 began throttling some cores down to the 4.515GHz mark after 20 minutes or so of continual full-load testing. Another possibility is that AMD has set Turbo Core 3.0 to begin

**Fig. 02. AMD FX-9590 Power testing**

throttling downwards when core temperature hits that 65°C mark in an effort to cap thermals and power consumption.

The lack of accurate temperature logging software poses a large problem for anyone with one of these 220W TDP chips: they have no way of knowing how hot (or cool) their processor is running.

## V. CONCLUSION

Judging from the path this launch is taking, we can assume that AMD has refreshed their high-end FX series in an effort to drum up some attention. Steamroller-based processors are still some time away and this is an excellent opportunity to flex Vishera's considerable multi-threaded muscles.

REFERENCES

[1]  http://www.hardwarecanucks.com/forum/hardware-canucks-reviews/62166-amd-fx-9590-review-piledriver-5ghz.html

[2]  W. Knight, "Two Heads Are Better Than One", IEEE Review, September 2005

[3]  L. Peng et al, "Memory Performance and Scalability of Intel"s and AMD"s Dual-Core Processors: A Case Study", IEEE, 2007

[4]  D. Geer, "For Programmers, Multicore Chips Mean Multiple Challenges", Computer, September 2007

[5]  T. Holwerda, "Intel: Software Needs to Heed Moore"s Law",        http://www.osnews.com/story/17983/Intel-Software-Needs-to-Heed-Moores-Law/

[6]  *Multicore Processors* – A Necessity. By Bryan Schauer

[7]  http://en.wikipedia.org/wiki/Moore's_law#mediaviewer/File:Transistor_Count_and_Moore%27s_Law_-_2011.svg