

Friday, July 24, 2009

## About Power7

There was an interesting article at Reuters about the upcoming Power7 CPU. I don't want to talk about the point, that his announcement was well times to our announcement about the 1.6 GHz. I don't want to talk about the announcement, that systems will be available in 2010 when you take into consideration, that the last rollout of Power6 took almost a year.

No, i want to hint you on an interesting fact in this article. But at first you should take into consideration that IBM tries to tell the world, that per core performance is the most important factor at all. Now look at the following paragraph the Reuters article: It also said POWER7 will be more efficient than the POWER6, which was launched in May 2007, capable of two to three times performance while using the same amount of energy. Two to three times the performance ... sounds nice. Well ... it doesn't sound that nice, when you do some math on the numbers. Let's assume the Power6 is a 2-core system. Let's further assume, that Power7 is a 4 core architecture. Then Power7 delivers the same performance per core than Power6 ... just more cores. Nice, but nothing earthshattering. But it gets even more interesting. Back in July 2008 Ashlee Vance wrote in the Register about Power7: IBM looks set to join the seriously multi-core set with the Power7 chip. Internal documents seen by The Register show Power7 with eight cores per processor and also some very, very large IBM boxes based on the chip. Okay, let's do the same math again: Let's assume the Power6 is 2-core, let's further assume that Power7 is a 8-core design. Obviously a single core of the Power7 has just half the performance of the Power6 core. Hey, that will be interesting talks for the IBM sales people explaining that they were incorrect with the assumption that speed per core is the only interesting metric.

Posted by Joerg Moellenkamp in English, The IT Business at 16:32

you should check the discussions on realwordtech.com.

power 7 will be 8 core, 4 threads per core 3-4ghz. a 32thread cpu - just like rock should have been (will be?!)

anyhow - I hope there will be competitive sparc product in 2010 and larry sticks to his promise to increase investment in the architecture.

venus might come to the rescue - fujitus projections for the coming fiscal year are very bullish - wonder if they think that they can cash in on the lack of rock...

Anonymous on Jul 24 2009, 17:57

I'm wating for a comment by you on the article on POWER in the latest iX.

We're still waiting for Rock...

Anonymous on Jul 24 2009, 18:10

Everybody (also at realworldtech) is missing the point completely : technology is irrelevant in this kind of economy!

- The risk Power7 will be late just like Power6 was : new architecture AND die-shrink in one step? (perhaps done in a new factory in NY? Yikes, recipe for disaster anyone?)

- Most importantly (and ultimately decisive) : customers on P5 will have to write off DDR1 memory to buy DDR2 memory with their P6, which has to be written off in one year to buy DDR3 with their P7 in 2010. At IBM prices! To probably end up paying more per core for software, as like to add insult to injury.

Are they dreaming or what???

Anonymous on Jul 27 2009, 22:18

There are even some hints that upgrading from p6 to p7 will be limited to p570/p590, at least a document issued by ibm doesn't talk about al other systems.

Anonymous on Jul 27 2009, 22:26

First of all you are assuming that the 2 core P6 chip uses the same amount of power as the 4 core P7. That can not be concluded from the first article. Still, if it happened to be true, the conclusion will be that each core will be up to 50% faster.

From the register article you assume that the 8 core P7 will use the same amount of energy as a 2 core P6. That assumption is clearly wrong. From the same article you were citing they in fact claim that the P7 will be twice as fast per core.

"Obviously a single core of the Power7 has just half the performance of the Power6 core"

A 4GHz P7 core is designed to do more work per clock cycle than the 5GHz P6. Your conclusion does not seem obvious at all.

Anonymous on Aug 11 2009, 22:45

## Blog Export: c0t0d0s0.org, <http://www.c0t0d0s0.org/>

At first i didn't talked about energy in my article. Wondering, why you try to talk about it in your comment. Just about basic math. I didn't expected to explain this.

Performance\_P6=1  
Performance\_P7=2 (sorry for being conservative)

Assumption 1:  
Cores\_P6=2  
Cores\_P7=4

$PerCorePerformance = Performance / Cores$

$PerCorePerformance_{p6} = 1/2 = 0.5$   
 $PerCorePerformance_{p7} = 2/4 = 0.5$

Equal, so nice, but not especially earthshattering. A hooray for better process technologies forcing more cores on the same room.

Assumption 2:  
Cores\_P6=2  
Cores\_P7=8  
 $PerCorePerformance_{p6} = 1/2 = 0.5$   
 $PerCorePerformance_{p7} = 2/8 = 0.25$

I have no problem with the reduced performance per core, because the industry goes this way. But: I remember so many comments and articles and presentations from IBM that PerCorePerformance is the most important factor at all. And now IBM divert to the same conclusion Sun made some years ago with the introduction of the UltraSPARC T1 line: It's not the single core that count's ... it's the throughput of the CPU.

And that's the whole point of the article.  
Anonymous on Aug 12 2009, 06:48

Your logic is wrong as I have already explained. You cherry pick from many rumours around P7 and combine it with wrongful assumptions. It is wishful thinking not reality.

You conclude that the core performance will be half but from the very same article from where your conclusion is derived it states: " Just on the gigaflop basis, that makes Power7 twice as fast per core as today's dual-core Power6 chips"

IT is OK to be biased but this is something way beyond that.

Btw. not one comment on Rock yet? I know it is not officially cancelled but rumour and speculation does not stop you in other respects. Why the silence?  
Anonymous on Aug 12 2009, 09:55

At first: Gigaflop/s isn't really an important benchmarks outside of the HPC business. GF is calculated (Frequency \* number of concurrently executed FP commands), not proofed by computation. And yes i know, that the core of the P7 yields 32 GF/s where as the core of the P6 yields up to 19-20 GF/s. As P7 was the answer of IBM to the DARPA project, i'm would wondering if they just doubled the number of FPUs. But in commercial interest this isn't important. Heck ... you aren't even allowed to use FP in commercial IT for many tasks.

Sorry, for determining the possible performance of the CPU i use the most actual information available, that looks as authoritative: "It also said POWER7 will be more efficient than the POWER6, which was launched in May 2007, capable of two to three times performance while using the same amount of energy." This was said by IBM as stated by Reuters, which i consider as a reputable news agency.

The assumption, that the two to three performance hike is "per proc" and not per core can be substantiated based on a press release of IBM: [http://www-03.ibm.com/systems/resources/systems\\_power\\_news\\_20090721\\_annnc.pdf](http://www-03.ibm.com/systems/resources/systems_power_news_20090721_annnc.pdf)  
"POWER7 processors will come in four, six and eight-core versions, and offer two to three times the performance of POWER6 using the same amount of energy."

Thus the math still holds true.  
Anonymous on Aug 12 2009, 11:13

The only thing that is bordering on certain is that the P7 will be 2-3 times more power efficient than the P6. The halving of performance per core is only in your fantasy and not backed anywhere else. Your math is trivial but your logic and assumptions are wrong. IBM will not balk out of the profitable high end market and do it the SUN way, namely release a CPU that competes with Dell/x86.  
Anonymous on Aug 12 2009, 12:00

1. It's a document of IBM stating "POWER7 processors will come in four, six and eight-core versions, and offer two to three times the performance of POWER6 using the same amount of energy." (page 2, last paragraph of the mentioned document [http://www-03.ibm.com/systems/resources/systems\\_power\\_news\\_20090721\\_annnc.pdf](http://www-03.ibm.com/systems/resources/systems_power_news_20090721_annnc.pdf))

2. Even when you suggest, that there will be version using more energy thus delivering more performance, there is a point substantiate a different view: The statement of direction regarding the update of 570/590 states: "The upgrade is planned as a simple replacement of the processor books and two system controllers with new POWER7 components, within the existing system frame." (page 4, last paragraph) As there were no mentioned changes to cooling or energy grid of the system, it is a reasonable assumption that the new processors will be in the same thermal ballpark than the P6. The heat dissipation is correlated to the energy consumption. So it looks like a reasonable assumption that we will see a P7 system to be 2 to 3 times as performant as a P6 system.

## Blog Export: c0t0d0s0.org, <http://www.c0t0d0s0.org/>

Now comes the core count into play in conjunction with the relative performance of the system.  $2/8 = 0.25$ . Even when we are not progressive, the single core has a relative performance of 0.375 instead of 0.5 and is thus still slower than the single core of a P6 but i tend to calculate with the lower numbers even with Sun processors as long i'm not answering to mabushi using peaks.

I do not see a problem with it: Intel and AMD have high clocked singlecores and lower clocked quad cores. Sun and Fujitsu has Sparc64 and UltraSPARC T. IBM would have it's p6 and it's p7 line. But IBM can't tell everywhere that per core performance is the most important as they have now a multithreaded proc with less performance per core as well. And this will be an interesting job IBM sales reps.

As i wrote you before: It's basic math. Of course it's based on assumptions as there is no reasonable benchmark available for the p7. But my assumption aren't based on wishful thinking, they are based on some not totally unreasonable considerations based on the information available.

Anonymous on Aug 12 2009, 13:09

I'm not even close to an expert on chips, but if they are going to build the fastest supercomputer based on this chip (Power7) then it ain't too slow. But the IBM salesmen aren't going to be using performance as a selling point so much as the whole package. Sun can't figure out which why is up when it comes to virtualization. Let's see, today it's LDoms, yesterday it was Containers / Zones and the day before that Domains. IBM did virtual in 1967 (yes, they were too stupid to know just what they had, but with a little help from VMWare they seem to be about to figure it out) better than Sun did it in 2007 or will in 2017. Until Sun starts acting like one company with all divisions working together as a team they are going to continue to lose market share and leading edge tech. Where is anything from Sun to compete with Live Partition Mobility or Live Application Mobility? It doesn't exist. Sun may come out with the fastest chip the world has ever seen, though I seriously doubt it, but it will NOT matter unless they can get their whole act together and focus on the whole solution.

Rots a Ruck!

Anonymous on Aug 25 2009, 20:55

At first: That's a common misunderstanding to think "When you build a supercomputer from it, it must be fast". No, you need only enough of it.

Nevertheless i'm really interested in the HotChips presentation and hope to get some good data about Power7 out of it.

Sorry, but IBM was able to understand something, that many of the outside proponents of their virtualisation technology doesn't understand. There is no universal virtualisation technology that fit all needs. This is the reason, why IBM has several virtualisation technologies in it's portfolio. There are LPARs, WPARs, all the \*PAR in the portfolio of IBM.

And it's essentially the same with Sun. We have fault-isolating domaining with domains on Uni-XSBs. We have a more granular domaining with Quad-XSBs. We have LDOMS on the T-class developed in a way, that complements optimally the hardware. Those methods of virtualisation were developed to have separate operating systems, in the case you need or want it.

On the other side you have a method to virtualize the operating system with almost no overhead. It's called containers.

It's not yesterday Containers, the day before that Domains, today LDOMS. It's today Containers, today Domains, today LDOMS, today Resource Management. At the end it's about choice.

Two things at the end:

- Do you really think, we stop the development with the current warm migration feature in the LDOMS?
- WPARs are a can full of problems. Just read their own redbooks describing compatibility modes for application migration. As far as i see it from IBMs own documentation, it's possible to maneuver yourself into a situation, where you can migrate to another server, but not back. It's nice on whitepapers but the practical side is questionable. The "without noticable interruption" is a joke when you migrate something larger than a vi over the network. I wrote about it some time ago in <http://joerg.moellenkamp.org/archives/3421-Some-thoughts-about-IBMs-Workload-Partition..html>

And by the way: From my point of view virtualisation is vastly overhyped at least in conjunction with UNIX. Consolidating workloads can be done with Unix without the help of hypervisors. Virtualisation provides administrative domains.

Anonymous on Aug 25 2009, 21:43