

Re: How to develop a random number generation device

needs
gigaflops.
It's
all about
reliability.

Until you can come up with
some sort of justification,
however vague, as
to why you think one cpu
per process is more reliable
than context
switches, this whole
discussion is useless.

You define yourself by the ideas you refuse
to consider. So I suppose
you'll still be running Windows 20 years
from now.

I run windows (on desktops) and Linux (on a desktop, a
laptop, and a
bunch of servers, and on a fairly high-reliability automation
system I
am working on), and I'd use something else if I needed an
OS in my
embedded systems. If something better came along, I'd use
that –
whatever is the right tool for the job.

The relevant saying is "keep an open mind, but not so open
that your
brains fall out". I'm happy to accept that doing things in
hardware is
often more reliable than doing things in software (I work
with small
embedded systems – I know when reliability is important,
and I know
about achieving it in practical systems). But what I am not
willing to
accept is claims that you alone understand the way to make
all computers
reliable,

I have made no such claims.

Re: How to develop a random number generation device

You have repeatedly said that current OS's (software OS's running on one or a few cores) is inherently unreliable, while your idea of a massively multi-core cpu running a task per core would be totally reliable. As far as I can see, you are the only person who believes this. If I've misunderstood (either about your claims, or if you can show that others share the idea), please correct me.

using a hardware design that is obviously (to me, anyway)

impractical,

Can't help what's obvious to you

and you offer no justification beyond repeating claims that

"hardware is always more reliable than software",

Isn't it?

No it isn't. At best, you can compare apples and oranges and note that a ram chip is more reliable than windows, despite the former having more transistors than the later has lines of source code.

We agree that typical hardware design processes are more geared to producing reliable and well-tested designs than common software design processes. But that does not translate into a generalisation that a given task can be performed more reliably in hardware than software.

and therefore you can

practically guarantee that the future of computing will be dominated by single task per core processors.

I can't guarantee it. My ideas are necessarily simplistic, and would

Perhaps "guarantee" was a bit strong – but you stated confidently that your 1024-core one-core-per-task devices were "gonna happen".

Re: How to develop a random number generation device

That's probably true. Sun will soon be shipping 8-core, multithread processors. Looks like the number of cores per chip is at least doubling every year, now that clock speeds are no longer the holy grail.

So, in 5 years, with $8 * 2^5 = 256$ cores, running maybe 1K threads, why context switch?

get more complex in a real system. Like, for example, my multicore chip would probably have a core+GPU or three optimized for graphics, and maybe some crypto or compression/decompression gadgets. There's no point sacrificing performance to intellectual purity.

This is beginning to sound a lot more like a practical system – devices exist today with several specialised cores, particularly in the embedded market. Arguably graphics cards fall into this category, as do high-end network cards with offload engines. But that's a far cry from your cpu-per-thread idea, and it is done for performance reasons – **not** reliability.

Well, the world is ready for OS reliability.

But the trend towards multiple cores, running multiple threads each, is a steamroller. So far, it's been along the Microsoft "big OS" model, but when we get to scores of processors running hundreds of threads, wouldn't a different OS design start to make sense? The IBM Cell is certainly another direction.

Forget windows – it's a bad example of an OS, and it's an extreme example of unreliable software. There is no "Microsoft big OS" model – they just have a bad implementation of a normal monolithic kernel OS.

Re: How to develop a random number generation device

There are uses for computers based on running large numbers of threads in parallel – the Sun Niagara processors can handle 64 threads in hardware (running on 8 cores). But these do not use a core (or even a virtual core) per thread – the cores have context switches as threads and processes come and go, or sleep and resume. Clearly you will get better *performance* when you can minimise context switching – but no one would plan for a system where context switching did not happen. There is nothing to suggest that the system could be made more reliable by avoiding context switches, except in the sense of reliably being able to complete tasks at the required speed – it's a performance issue.

I believe I have been open minded – I've tried to point out the problems with your ideas, and why I think it is impractical to design such chips,

Sorry, I missed that part. Why is it, or more significantly, why *will it* be impractical to design a chip that will contain, or act like it contains, a couple hundred CPU cores, all interfaced to a central cache?

Perhaps I didn't explain it well, or perhaps you didn't read these posts – it's hard to follow everything on s.e.d.

The problem with so many cores accessing a shared cache is that you have huge contention for the cache resources. RAM cells get bigger, slower and more complex the more ports they have – it's rare to get more than dual-ported RAM blocks. So if you have 1000 cores all trying to access the same cache, you're going to have huge latencies. You also need complex multiplexing hierarchies for your cross-switches – as each cpu needs to access the cache, you basically require a 1000:1 multiplexer. Assuming your cache has multiple banks and access to some IO or other buses, you'd need something like a 1000:10 cross-switch. That would be really horrible to implement – you'd need to find a compromise between vast switching circuits and multiple levels introducing delays and bottlenecks.

Here's a brief view of the Niagara II – your device would face similar challenges, but greatly multiplied:

<http://www.theinquirer.net/?article=42256>

If each core has an L1 cache to relieve some of the pressure (without it, the system would crawl), you then have a very nasty problem of tracking cache coherency. Current cache coherency strategies do not scale well – they are a big problem on multicore systems.

With existing multiprocessor systems, it is the cache and memory

Re: How to develop a random number generation device

interconnection systems that are the big problem. If you look at high-end motherboards with 8 or 16 sockets, the cross-bar switches that keep memory coherent and provide fast access for all the cores cost more than the processors themselves. Building it all in one device does not make it significantly easier (although it saves on some buffers).

There are alternative ways to connect up large numbers of cores – a NUMA arrangement with cores passing memory requests between each other would almost certainly be easier. But you would have very significant latencies and bottlenecks, a very large number of inter-core buses, and you'd still have trouble with the L1 cache coherence.

With a new OS, and certain significant restraints on the software, you could perhaps avoid many of the L1 cache coherence problems. In particular, being even more restrictive about memory segments would allow you to assume that L1 data is private, and thus always coherent. For example, if all memory came from either a read-only source for code, or was private to the task using it, then you'd have coherency. You'd need a system for read and write locks for memory areas, with a central controller responsible for dishing out these locks and broadcasting cache invalidations when these changed, but it might work.

However, you've lost out on a range of requirements here. First off, your cores are now far from simple, and the glue logic is immense. Thus you have lost all hope of making the device cheap and reliable. Secondly, you've still got significant latencies for all memory access, slowing down the throughput of any given core, crippling your maximum thread speed. The bottlenecks don't matter so much in the grand view of the device – the total bandwidth to the cpus should still be more than if it were a normal multi-core device.

Exactly! Except there is no context switch overhead.

Thirdly, you've lost

compatibility with all existing software – it won't run most programs, as they rely on being able to have shared data access.

Exactly! We can't run .NET forever.

Re: How to develop a random number generation device

and why they would be impractical for general purpose computing even if they were made.

Why? Because Windows, and other "big" OS's like Linux, don't support it?

Yes, that's about it. To be more precise, it will be impractical for general purpose computing because it won't run common general purpose programs.

Circular reasoning. Why aren't we still running 1401 code?

Even with the required major changes to the software and compilation tools, and without the cache restrictions mentioned earlier, it would run common programs painfully slowly.

If the cache throughput is the limit, you get the same amount of computing no matter how many CPUs are running. CPUs can also have a little bit of local instruction cache, since code does not have to be kept globally coherent.

I've repeatedly asked for justification for your claims, and received none of relevance. I am more than willing to discuss these ideas more if you can justify them – but until then, I'll continue to view massively multi-core chips as useful for some specialised tasks but inappropriate for general purpose (and desktop in particular) computing.

It's generally accepted that a microkernel-based OS will be more reliable than a macrokernel system, because of its simplicity, but the microkernel needs too many context switches to be efficient.

A microkernel *may* be more reliable because of its modular design – each part is relatively simple and communicates through limited,

Re: How to develop a random number generation device

Re: How to develop a random number generation device

controlled ports. That's far from saying it always **will** be more reliable. Much of the theoretical reliability gains of a microkernel do not actually help in practice. For example, the ability of low-level services to be restarted if they crash is useless when the service in question is essential to the system. Thus there are no reliability benefits from putting your memory management, task management, virtual file system, or interrupt system outside the true kernel – if one of these services dies, you're bugged whether it kills the kernel or not. A similar situation is found in Linux – because X is separate from the kernel, it can die and restart independently of the OS itself. But to the desktop user, their system has died – they don't know or care if the OS itself survived.

Most of the benefits of a microkernel can actually be achieved in a monolithic kernel – you keep your services carefully modularised, developed and tested as separate units with clear and clean interfaces. It's a good development paradigm – it does not matter in practice if the key services are directly linked with the kernel or not, since they are all essential to the working of the OS. About the only way a microkernel improves reliability is by enforcing this model – you are not able to cheat.

What **does** make sense is keeping as many device drivers as possible out of the kernel itself. Non-essential services should not be in the kernel.

<http://en.wikipedia.org/wiki/Microkernel>

So, let's get rid of the context switches by running each process in its own real or virtual (ie, multithreaded) CPU. Then nobody can crash the kernel. A little hardware protection for DMA operations makes even device drivers safe.

You underestimate the power of software bugs – you'll **always** be able to crash the kernel!

No. Not if it's small and correct, and it's absolutely protected by the hardware, and it runs on a CPU that runs nothing else. I've written RTOS's that never crashed.

There's nothing wrong with dreaming, quite the opposite. But you have to be able to see when it is nothing but a dream.

Re: How to develop a random number generation device

Do I have to give all that money back? Roughly \$200 million so far.

John

.