

Re: RC Transmission Lines (Wafer-Scale)

Source: <http://sci.tech-archive.net/Archive/sci.electronics.design/2007-10/msg04141.html>

- *From:* Robert Baer <robertbaer@xxxxxxxxxxxxx>
 - *Date:* Sun, 21 Oct 2007 21:01:48 -0700
-

Phil Hobbs wrote:

Robert Baer wrote:

Guy Macon wrote:

Content-Transfer-Encoding: 8Bit

Guy Macon <<http://www.guymacon.com/>> wrote:

I have been doing a web searches on this, and found these:

[_The Future of Wires_](#)

<http://www.cs.utah.edu/classes/cs6943/papers/horowitz99future.pdf>

[_The Wire_](#)

<http://www.vlsi.uwaterloo.ca/~manis/ece730/lecture2.pdf>

Parasitic Extraction and Performance Estimation from
Physical Structure

<http://lsiwww.epfl.ch/LSI2001/teaching/webcourse/ch04/ch04.html#4.2>

<http://lsiwww.epfl.ch/LSI2001/teaching/webcourse/toc.html>

I also found this in an abstract (the actual paper isn't online):

A 3Gb/s/wire Global On-Chip Bus with Near
Velocity-of-Light Latency

"We successfully show the practical feasibility of a purely electrical global on-chip communication link with near velocity-of-light delay. The implemented high-speed link comprises a 5mm long, fully shielded, repeaterless, on-chip global bus reaching 3Gb/s/wire in a standard 0.18 μ m CMOS process. Transmission-line-style interconnects are achieved by routing signal wires in

Re: RC Transmission Lines (Wafer-Scale)

the thicker top metal M6 layer and utilizing a metal M4 ground return plane to realize near velocity-of-light data transmission. The nominal wire delay is measured to 52.8ps corresponding to 32% of the velocity of light in vacuum.

<http://www.google.com/search?hl=en&q=%22A+3Gb+s+wire+Global+On-Chip+Bus+with+>

I sure as all heck do not call 32% " of C(vac) as being anyway "near velocity-of-light delay".

Lines with repeaters run at about $c/10$.

You can get more bandwidth on the fatwire levels (high in the stack), basically because the capacitance per unit length is the same, and the resistance goes down. Of course, the number of wires available goes down too, and you have to get through all the lower wiring levels to get up to the fatwires. This leads to wireability problems.

In something like a highly multicore processor, you need lots and lots of fast wires. There's a current DARPA program called TELL (for terabit electrical links at low power, or something like that). TELL is all about finding out how far electrical links can be pushed (and therefore at what point you have to go to optical links or accept reduced performance). Since wiring capacitance tends to be independent of size scaling, it's really hard to get below 2 pF/cm, so to save power, links need to use very low voltage, at which point you get into nasty problems with noise, drift, offset voltages, and crosstalk.

This gets especially difficult starting with the 32 nm node (iirc again), because the threshold voltages of the FETS are hard to control...it turns out that you have to worry about statistical variations in the number of dopant atoms in the FET channel. A 30-nm cube of silicon, doped to 10^{20} per cc, contains 2700 dopant atoms. If a chip has a billion transistors, you'll have lots of 6-sigma outliers, which will be off by ± 300 atoms, or 11% of nominal, which causes a nasty threshold voltage shift. Smaller devices get worse fast.

The wire guys think they can overcome these problems, and maybe they can...but my money says that we'll see on-chip optical signalling in the next 10 years. (I also have to try keeping my management convinced.) ;)

Cheers,

Phil Hobbs

Forgive me for a rather nasty question.

What is this phoney push for more junk on a piece of silicon to support what used to be relatively simple applications?

What was wrong with KISS?

In a personal computer, one does not need 2^{10} core CPUs, or even dual core; any CPU speed over 1Ghz is

Re: RC Transmission Lines (Wafer-Scale)

wasted, and for 99+% uses Win98Se is more than good enough.

Now, if one gets into graphics (read: games, design PCBs or other complex artwork), then more speed becomes useful and Win2K becomes a better choice.

Oh, you say, we "need" dual (or quad) core for graphics.

What the hell is that large graphics chip on the fancy video card for? Boat anchor?

In fact, what good was the MMX instruction set for, since the sound card already supported those functions.

On a cell phone guess what – its purpose is to send and receive calls, period.

Want to do something else like portable music – players have been around for over 10 years that do that; they just get smaller and store more.

Etc etc and etc (courtesy of Yul Brynner in the King and I).

.