
Subject: Custom vs. commodity networks

Posted by [Patrick R. Haspel](#) on Mon, 24 May 2004 14:31:56 GMT

[View Forum Message](#) <> [Reply to Message](#)

Dear colleagues,

I would like to hear your thoughts about the suitability of these 2 classes for the FutureDAQ networks.

We could start with advantages/disadvantages of the two classes and afterwards rank them with respect to the use in this project.

Tell me your thoughts and concerns...

Thanks for all participation,

Patrick

Subject: Re: Custom vs. commodity networks

Posted by [Alexander Mann](#) on Tue, 01 Jun 2004 08:30:53 GMT

[View Forum Message](#) <> [Reply to Message](#)

To start the discussion, I have just collected some arguments pro and against the two network classes:

commodity networks:

(Ethernet, ATM, SCI, HIPPI, ...) everything one can simply buy.

- + No major hardware development needed. Chips are widely used and well tested.
- + Low price because of high quantities.
- + Components are available within short time. No expensive stock-keeping for spare parts.
- + Switches or interface cards come fully tested from the manufacturer. We do not have to worry about making them work or about repairing them.
- + Configuration by protocol software, so changes are easy to apply.

- Our problem might not perfectly fit in the existing network technologies, so we have to oversize the network system.
- High bitrate / bandwidth components still are quite expensive.

custom networks:

- + We can optimize the network implementation to our specific needs.
- + High usable bitrate, due to low protocol overhead.
- + Additional features (like time and clock distribution) can be implemented.
- All hard- and software design work has to be done by ourselves.
- If chips are not tested well, design changes might become very expensive.
- If the network requirements change, the custom implementation might become obsolete and a new implementation might be necessary.

Some other interesting questions are:

How shall the network behave in error conditions (defective computing / routing node or link). Should the network find a way to work around the error, or should there be a supervisor which determines the entire network routing?

How is the performance affected by these two implementations?

What error detection or correction is needed?

Can our network requirements be mapped on existing network hardware or do we need something completely different?

Regards,

Alexander Mann

Subject: Re: Custom vs. commodity networks

Posted by [Krzysztof Korcyl](#) on Tue, 08 Jun 2004 10:02:06 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello,

I would like add yet another "pro" for the commodity networks - mainly a possibility to profit from continuous evolution and upgrade of existing standards - examples being: Ethernet (100Mbps, 1Gbps, 10Gbps, 40Gbps....), ATM (155 Mbps, 622 Mbps) and perhaps others.

And a comments on "cons" for the commodity network.

The required oversize of the network due to not perfect match to our requirements is not a serious limitation. Only in the point-to-point connection one can use fully the available network bandwidth. In a distributed networks one has to install additional bandwidth (and buffering) to accomodate for asynchronous communications. As the networks should be undersubscribed then the additional bandwidth from the commodity standards may actually be cheaper (as the components are cheaper) than from dedicated networks.

And the second comment: high price for the high rate components are subject to continuous and ever-faster drop than we used to see over previous years.

best regards,

Krzysztof.

Subject: Re: Custom vs. commodity networks

Posted by [Patrick R. Haspel](#) on Thu, 10 Jun 2004 14:07:14 GMT

[View Forum Message](#) <> [Reply to Message](#)

I can pretty much agree to most of the mentioned points.

We could roughly sum up all of them by stating, that if there is a cost reasonable commodity

network that fulfills all requirements, it makes no sense to reinvent the wheel by designing a custom network.

But if we have special requirements unlikely implemented in commodity hardware we need to consider a custom solution.

So the key issues needed to be figured out are the communication requirements.

Just some words about the cost of a custom network, just to rebut the common view. Cost effectiveness tremendously depends on the needed quantity:

Mask cost for 180nm CMOS ASIC are about 300kUSD, cost per 200mm wafer are about 2kUSD. A sophisticated NIC design takes about 32qmm. With a pessimistic yield of 75% you will get about 750 dies out of a wafer. The part cost would then be:

@ 750 chips: 400 USD per chip

@ 1000 chips: 304 USD per chip

@ 5000 chips: 62 USD per chip

@ 10000 chips: 32 USD per chip

@100000 chips: 5 USD per chip

Compared to costs of a sophisticated FPGA of between 1,5 and 2,5kUSD one must take quantities into account for any ASIC/FPGA implementation. Of course there are applications where flexibility is key and therefore FPGAs are the only solution.

Best regards,
Patrick

{some markup added by moderator}

Subject: Re: Custom vs. commodity networks

Posted by [Walter F.J. Müller](#) on Tue, 15 Jun 2004 08:05:57 GMT

[View Forum Message](#) <> [Reply to Message](#)

I'd like to inject two new arguments into the discussion:

The main usage of networking components will be to interconnect programmable logic bases processing units. This means that the link termination is, at least partially, in programmable logic resources. So the complexity and cost of the link termination is a significant factor, and as a consequence, one would like to have very lean and light-weight protocols.

The shopping list of commodity networks given earlier (Ethernet, ATM, SCI, HIPPI, ...) contains the traditional products, developed for LAN and SAN applications. A network designed for short and medium distance chip-to-chip communication might much better match our requirement profile. PCI Express Advance Switching is positioned in exactly this market and now fully specified. I wonder whether somebody had a serious look at PCIe-AS so far.

Subject: Re: Custom vs. commodity networks

Posted by [Patrick R. Haspel](#) on Tue, 15 Jun 2004 09:10:33 GMT

[View Forum Message](#) <> [Reply to Message](#)

Dear Walter,

reading your args, some questions popped up:

Regarding 1:

Is the final solution PLD based?

Do we make a difference between the interconnections of the different levels?

Regarding 2:

What is the average/maximum distance between the processing units?

Subject: Re: Custom vs. commodity networks

Posted by [Walter F.J. Müller](#) on Tue, 15 Jun 2004 09:29:36 GMT

[View Forum Message](#) <> [Reply to Message](#)

Is the final solution PLD based?

The current working assumption is, that the first layer(s) are based on PLD's and that the final trigger decision is done in a CPU (which might be in PC).

What is the average/maximum distance between the processing units?

In general, links connecting units on the detector as well as in the counting house will be 2-10 m, while links between detector and counting house will be significantly longer, say order 50m. Which ones are just point-to-point and or require a true network is one of the basic architecture questions.

Subject: Re: Custom vs. commodity networks --> PCIe-AS

Posted by [Walter F.J. Müller](#) on Mon, 21 Jun 2004 08:50:45 GMT

[View Forum Message](#) <> [Reply to Message](#)

As stated in the previous posting, I wonder whether somebody had a serious look at PCIeexpress Advanced Switching. Some material is under

<http://www.asi-sig.org/education/usage>