Minutes of the Internal Review 24.07.2001

on FEM/RU

Chairman: Ph. Gavillet
Secretary: S. Schmeling

I. AGENDA

- Morning Session
  - FEM/RU Scope, Functionality - Assessment criteria (J. Harvey)
  - FPGA-based implementation
    - FPGA-based Readout Unit (J. Toledo)
    - The embedded MCU card on the RU (A. Guirao)
    - Overview and status of the project (H. Muller)
  - Network Processors-based implementation
    - Introduction to Network Processors (B. Jost)
    - IBM Network Processor, SW Development environment and LHCb Software (N. Neufeld)
    - Implementation, Applicability and Summary (B. Jost)

- Afternoon Session
  - General Discussion - questions to both approaches for clarification
  - Closed session of the referees
    - Impressions / preliminary conclusions from the review
    - Procedure to prepare the review report
  - Closeout - referees communicate their main conclusions

II. INTRODUCTION (PH. GAVILLET)

Overview of the agenda. Questions should be restricted preferably to the Q&A session.

III. FEM/RU SCOPE, FUNCTIONALITY, ASSESSMENT CRITERIA (JOHN HARVEY)

- Presentation of requirements (estimations from J. Christiansen's web site)
  - as FEM for data merging
    - data rate after mux: 40 kHz (eventually 100 kHz)
    - 4:1 multiplexing for VELO and IT with 25/27 units
    - minimal buffering
    - compatibility to RU required
  - as RU also for data merging
    - 203 links for 79 units estimated
    - rate should not exceed 50 MB/s
    - presently calculated cost for 110 units
    - in case of enlarging the operational trigger rate, one could envisage increasing the number of RU's to moderate the increased load
    - large fragments ~250 bytes/input
    - significant buffering
    - formatting and traffic shaping required
    - flexible destination determination (partitioning support, CPU failures, …)
  - data merging in Level 1
    - rate of Level 0 ~1.1 MHz

        - latency restriction ~10µs
        - small fragments ~30 bytes/fragment
        - medium buffering 10-100 events
- assessment criteria (proposal)
    - functionality (including environment)
    - performance
    - cost (~1 MCHF envisaged for 110 units)
    - acceptance test
    - support for integration, commissioning and consolidation
    - flexibility
    - maintenance and support
    - generality (use for more than one application in LHCb)
    - commonality (usability in more than LHCb)
- input to review
    - design process
    - design features
    - results of simulation
    - results of prototyping
    - implementation details
    - potential future program
    - cost estimates

Q&A:

- H. Dijkstra: Is traffic shaping a requirement?
    - J.P. Dufey: It may be for load balancing.
- Ph. Charpentier: Early availability seems to be a requirement to have a small system for test purposes already early enough.

IV. FPGA-BASED READOUT-UNIT

A. FPGA-BASED IMPLEMENTATION (H. MULLER)

- Traffic shaping was not foreseen so far.
- Applications:
    - SLink Multiplexer
        - 16:1 at 40 (80) kHz
        - 2.5 (5) MB/link
        - 2-10 events buffered (2000 available for XON/XOFF)
    - DAQ
        - 4:1 mux 40 (80) kHz @ 1 KB (250 bytes)
        - up to 1000 (4000) sub-events buffered
        - throttle latency < 10 ms
    - VELO Level 1
        - quad 1 MHz sub-event building for 50 bytes clusters
        - output 200 MB/s to shared memory via PCI-NIC
        - up to 5000 sub-event buffers (5 ms)
        - custom scheduling via Tagnet
    - Remote control and configuration
        - PCI, I2C, JTAG
- Sub-event building
    - MUX: FPGA push protocol
    - DAQ: NIC pull
    - VELO: 2x FPGA push
    - minimal overhead 4x 32bit word

    - transport transparent to RU
- FPGA evolution
    - now manufacturer independent
    - ASIC core available
    - VHDL verilog
- Environment
    - LEP crates: IEEE 960
    - 13 slots used (double width)
    - 1 kW/crate
    - crates from LEP -> free

Q&A:

- J. Christiansen: ASIC cores restrict to a certain manufacturer.
    - H. Muller: Libraries have to be changed but it is possible.
    - F. Formenti: VHDL core independent.
    - M. Schulz: Transition from ORCA to ALTERA. Porting took ~1 week.
- M. Schulz: Fastbus-crates are for free? Quantities?
    - H. Muller: in this order no problems to be expected
    - Ph. Gavillet: Power supplies have to be revised.

B. FPGA-BASED READOUT UNIT (J. TOLEDO)

- 4 SLink inputs to FIFOs, PCI output
- 2 FIFOs read by 1 FPGA
- memory implemented as circular buffers as memory is expensive and a fix assignment is not possible
- SLink output for Level 1
- PCI output for DAQ towards intelligent NIC that pulls via FPGA from memory
    - DMA
    - memory mapped into FPGAs
- PCI output for Level 1
    - more performance needed
    - push protocol
- Connectors (all standard)
    - FB power
    - PCI for test/debugging
    - 4 SLink cards
    - PCI 64bit PMC for NIC or SLink card
    - MCU PMC slot
    - Tagnet connectors
    - reset input
    - throttle output
- main components
    - 4 input FIFOs
    - 4 FPGA
    - PCI output
- performance (estimated for 50 MHz operation)
    - FEM: up to 100 kHz possible
    - DAQ: 40 kHz possible at 3 KB/event
        - restricted because of Gbit NIC
    - DAQ: 80 kHz possible at 1.5 KB/event
        - restricted because of Gbit NIC
    - VELO: 1.3 MHz possible at 200 B/event
    - conclusion

        - The requirements are met.
        - DCS connections implemented on-board
- HDL
    - Design in VHDL
    - Simulation also possible
    - Automatic code generation
    - Vendor specific at a late stage before bitstream generation
    - 10 min to generate code for FPGA after definition in HDL
    - Pros:
        - tools are standard
        - behavioral description possible
        - code independent from specific FPGA/ASIC
        - simple syntax
        - used and supported at CERN
    - Status and Conclusions
        - HARDWARE features
            - prototype finished
            - 5 PCBs produced
            - SLink-SLink under test
            - PCI loading for FPGAs tested
            - DPM access from MCU tested
        - FEM
            - VHDL code programmed
        - DAQ
            - RU-NIC protocol has to be agreed
        - Mezzanines to be integrated on-board

Q&A:

- J. Christiansen: All is based on VHDL?
    - J. Toledo: Complete simulation is possible. Performance data from analysis not from simulation as the simulation depends on event sizes.
    - J. Christiansen: Why not statistical distribution?
    - J. Toledo: Can be done. Functional simulation done. Buffer sizes critical.
    - H. Muller: Throttle crucial. Simulation has to be done for the whole system as all buffers may influence the system performance.
    - J. Toledo: Maybe more FIFOs needed to meet performance requirements.
    - J. Christiansen: Full model available now?
    - J. Toledo: To be put together in libraries, but can be done.
    - J. Christiansen: PCI model available?
    - H. Muller: No.
    - M. Schulz: Simulation anyway problematic. VHDL is for synthesis.
    - J. Toledo: Model of RU board available, but not PCI and NIC.
- Ph. Gavillet: When is the first performance estimate available?
    - J. Toledo: Few weeks.
- M. Schulz: Several performance figures given in brackets. What does that mean?
    - J. Toledo: INTEL restricts that. Waiting for information if MCU may be overfloated. PMC connector with 66 MHz PCI needed. Model would have to be bought.
- J. Christiansen: Assumptions about NIC: Does it exist as PMC card?
    - H. Muller: This was a design agreement that a card like that would exist. Gbit cards do not exist anymore. SCI PMC exist. Currently PMC card does not exist.
    - B. Jost: Gbit PMC probably exist. Availability Gbit PMC with all features desired has to be studied. Possibly this has to be designed. Any rate here is outside the mainstream of these products.

    - J. Christiansen: Test needed.
    - B. Jost: For testing one could use a standard card.
    - M. Schulz: 1U boxes with standard PCI available.
    - H. Muller: Additional PCI connector also meant to plug PC-standard PCI connector flat on board to use standard PC PCI cards. Decision for PMC/PCI was taken years ago. Possibility to put SLink inputs directly to FPGAs with modified core.
- T. Nakada: What has to be done for rate increase to 100 kHz?
    - J. Toledo: Nothing has to be done. A decrease of rate would simplify the board. Board not over designed for VELO, but overperformand for other applications.
- B. Jost: How big is the input FIFO?
    - J. Toledo: > 500 bytes
- B. Jost: What if routing fails?
    - J. Toledo: Error message may be sent. Plenty of free space in FPGAs to program.
- J. Christiansen: Not many gates needed for routing. No event processing.
    - J. Toledo: Only for pipelining.
- St. Wotton: Current use of FPGAs?
    - A. Guirao: 5%
- J. Christiansen: Error handling implemented? May need complex state machines.
    - J. Toledo: Not yet. If too complicated for FPGAs, may be done in MCU software.
    - H. Muller: Error handling not yet implemented because of lack of definition. Now just error history. Inconsistent data error detected. All links always carry data, even empty events. Sequence always +1 in event number.

C. THE EMBEDDED MCU CARD ON THE RU (A. GUIRAO)

- PCI host needed, implemented in standard PMC, called MCU.
    - PCI configuration agent
    - FPGA bitstream loading
    - Bus arbitration
    - process FPGA and NIC interrupts
- other functions
    - DCS connection via Ethernet
    - JTAG master for RU JTAG chain through PMC connector
    - I2C controller for RU clock generation through PMC connector
    - RU reset generation
- Every card following VITA-32 PMC standard, providing several signals through PMC can be used
- Connectors: FDD, HDD, Keyboard, Mouse, PCI
- OS: Linux; standard drivers
- Tec specs:
    - Cyrix 486 compatible core 120 MHz
    - 64MB SDRAM
    - PCI 32bit 33 MHz
    - ISA
    - Fast Ethernet
    - AT legacy
    - USB
    - I2C bus master
    - software JTAG
    - PMC IEEE P-1386.1
- Summary
    - low-cost solution
    - complete PC-on-a-board
- Status

- prototype available
- remote booting to be implemented
- software development during this year

Q&A:

- C. Gaspar: Available in industry now?
  - A. Guirao: Not standard for JTAG for I2C. Otherwise standard. So one could buy it.
  - H. Muller: Can be made custom.
- J. Christiansen: Do all chips have boundary scan capability for JTAG? Can you test the board?
  - A. Guirao: Big chips yes. Not tried yet.
  - H. Muller: Connector JTAG also on board as for pixel application, so scan could be done from their control unit.
  - H. Muller: FPGAs can be tested.
  - J. Toledo: Indirect testing possible, but not necessarily via boundary scans.

D. OVERVIEW AND STATUS OF THE PROJECT (H. MULLER)

- Revised (Jan 2001) RU available as prototype, 5 modules built
  - specs see above
- mezzanine used to be independent from technology choices
- remote and local diagnosis
  - possibility to use PCI tracer
- exerciser station (under construction)
  - Level 1 emulator
  - LHCb sub-event builder
- cost
  - low quantity
    - 4770 CHF for RU
    - SLink ~400 CHF
    - MCU ~600 CHF
  - high quantity 100+
    - estimated at 63%
    - 4 link MUX 4.97 kCHF
    - 16 link MUX 5.45 kCHF
    - DAQ/VELO RU 4.67 kCHF + NIC
- work plan
  - until October
    - finalize FEM application
    - RU exerciser
    - Slink receiver cards test
    - FPGA register definition
    - test of halogen free RU vs. FR4
  - long-term (1 year)
    - migrate from 9U into PCI
    - parallel optical links for input on single mezzanine, RJ45 output
    - everything in single FPGA with external RAM
    - reuse STF sub-event building
    - 60 links input in a PC possible
- Summary
  - up to 256 MB/s feasible
  - different applications possible
  - remote control and configuration via PMC
  - local Linux on MCU

    - LEP crate environment
    - system cost 0.5 MCHF for 100 RUs
    - 48 RU/rack
    - IO rate beyond 1 MHz for events from 50 B to 32 KB
    - 100 RU potential: 1600 links, > 5 GB/s
    - DAQ Gbit technology: max. output 40-60 MB/s, 1.5 KB packets
    - VELO
        - requires 64 bit @ 66 MHz PCI
        - 25 RU, 100 input links 5 GB/s to SCI
    - 9U version ready and fully tested
    - optimization in progress

- J. Christiansen: Impression of full test was not given.
    - H. Muller: You cannot test everything. Only one problem found so far.
- J. Christiansen: Has CMS a parallel development?
    - H. Muller: PCI board would be a common development with CMS. Unique in LHCb: write-only model, with only throttle to control flow.

V. Network Processor-based Readout-Unit

A. Introduction to Network Processors (B. Jost)

- today 20 companies offer or announced to develop NPs
- small companies dominate market, but big companies catch up
- large number of RISC processors
- hardware assistants for e.g. traffic shaping
- on-chip buffer space
- integrated MACs
- integrated general purpose processor on-chip insuring communication e.g. to controls
- fully programmable
    - flexibility
- Summary
    - useable for data merging
    - next generation will deal with 10 Gbit/s and 40 Gbit/s link speeds
        - lower price for current devices
    - risks
        - economical failure of the technology
        - short lifetime for special product
        - cost

Q&A:

- J. Christiansen: How does the market look like?
    - B. Jost: Estimated revenues ~ 100 M$ to rise to ~ 3 B$ within next year.

B. IBM Network Processor, Software Development environment and LHCb Software (N. Neufeld)

- IBM NP4GS3
    - 4 full duplex Gbit Ethernet MACs
    - 16 CPUs with 2 hardware threads each
        - 2128 MIPS
        - up to 4.1 Mpackets/s
    - 128 KB on-chip input buffer

    - up to 128 MB DDR RAM on output buffer
    - 2 switch interfaces DASL 4 Gbit/s
    - embedded PPC405
    - in production since beginning of 2001
- RU design on 9U board
    - 1 or 2 mezzanine cards with NP and all memory
        - connections: PCI, DASL, physical network, JTAG, power, clock
    - physical layer connectors
    - L0 throttle output
    - power and clock generation
    - LHCb standard ECD interface (CCPC)
- chip dataflow
    - physical layer -> ingress event building -> DASL -> egress event building -> physical layer
- sub-event merging
    - main task
    - two locations possible which leads to two different algorithms
        - ingress: 1.5 MHz for small frames
            - on-chip memory (no wait)
        - egress: large frames (9000 B) at lower rate
            - large buffer space
            - external memory (wait states)
    - SW development environment
        - from IBM, very rich
        - documentation available with examples
        - comfortable user interface
        - simulator available
- performance
    - 4:1 event building
        - limited by output link
    - 3:1 event building
        - NP performance for all fragment sizes
- test
    - setup
        - reference kit HARDWARE
        - 4 Tigon2 NIC 1000SX
            - limited to 620kHz
        - 1 µs timing resolution
    - procedure
    - results:
        - single thread and single source
            - 4.9 µs/fragment
        - single thread and 4 sources
            - 2.8 µs/fragment
        - single source 16 treads
            - 0.0 µs/fragment (not measurable)
- requirements met

Q&A:

- J. Christiansen: Throttling needed?
    - N. Neufeld/B. Jost: Threshold can be set in software. Track is kept how much memory is used independently from fragment size.
- M. Schulz: Fragment size may not be larger than 30 bytes for Level 1 application?

    - N. Neufeld: May be optimized for header overhead. Used header not needed for this application, just as convenience for test the same header was used as for DAQ application.
    - B. Jost: Ethernet frame compliance not needed. So all Ethernet address space can be used for data.
- J. Christiansen: What is the enforced framing to be used?
    - N. Neufeld: Only few words needed for frame identification.
    - B. Jost: Only 4 bytes CRC at frame end is enforced. Rest may be redefined.
    - N. Neufeld: Concept of independence from physical layer even more enforced in successor chip.
- F. Formenti: 16 CPUs needed for fix purposes?
    - N. Neufeld: One thread for booting, one thread for PPC communication fix. Rest free.
    - F. Formenti: Task of management unit to steer threads?
    - N. Neufeld: Configurable.
- J. Christiansen: Shared program code possible?
    - N. Neufeld: Implementation as 4 different program memory blocks to facilitate access during execution.

C. IMPLEMENTATION, APPLICABILITY AND SUMMARY (B. JOST)

- Architecture
    - see above
- Mezzanine boards
    - 14 layer board
    - most complex parts are confined
    - fewer pins than using NP on main-board
    - modularity
- Features
    - generic, no bias towards an application
    - 8 fully connected Gbit Ethernet ports
        - bias towards Gbit Ethernet for readout network
        - needed Gbit Ethernet – SLink interface for Level 1 electronics
            - worked on in ATLAS
        - no need for NICs in RU
        - possibility to plug standard PCs for diagnostics everywhere
- Applications in LHCb
    - DAQ
        - FEM/RU
            - FEM and RU are equivalent
            - allows for multiplexing N:M with $N+M \leq 8$
            - performance good enough for any trigger rate up to 100 kHz
        - building block for readout network
            - intrinsically an 8 port switch
            - any-size network possible
            - benefits
                - full control and knowledge over switching process
                - full control over flow control
                - jumbo frames
                - full monitoring capabilities
        - final event building
            - baseline solution smart NICs in subfarm controllers
                - off-load of SFCs
            - NP module
                - used as 4:4 MUX/data merger

        - software port-independent
            - small scale lab setup
    - Level 1 (Proposal)
        - same basic architecture as for DAQ
        - one unit acts as 2x 3:1 MUX
        - one unit as final event building unit
        - performance sufficient (see above)
- Design and Production
    - mezzanine on basis of IBM kit
    - main board would be a separate effort
        - CERN: cheap
        - outside: less cheap
    - before production of prototype review with IBM preferable
    - production clearly external
    - acceptance test in LHCb
- Cost
    - mezzanine:
        - 3 k$ each for ~100 cards
        - complete RU NP
            - 6 k$
            - dominated by NP
    - motherboard
        - CCPS 150 $
        - power/clock ???
        - physical layer 8x 90 $
        - total 2 k$
    - complete unit
        - < 8 k$ each for ~100 modules
        - ATLAS has shown interest
- numbers used
    - DAQ
        - all applications 259 RUs: 2.07 M$
        - only FEM/RU 140 RUs: 1.12 M$
    - Level 1
        - all applications 80 RUs: 0.64 M$
        - only FEM/RU 32 RUs: 0.256 M$
- Summary
    - software driven: easily configurable
    - own strategies to be implemented easily
    - elaborate debugging and diagnostic capabilities
    - connection via PPC towards ECS
    - debugging of code in-situ
    - NP RU fulfils requirements in speed and functionality
    - functional equivalent reference kit available
    - simulation and measurements in agreement and show sufficient results
    - no bias towards specific network
    - error detection and reporting
        - CRC32 hardware supported
        - PPC/CCPC allow efficient monitoring
        - debugging capabilities
    - plugging of standard PC possible in any stage
- Planning
    - hardware design
        - offer: 300 k$ for design and production tools

        - 1 man year for infrastructure development (online team)
- Environment
    - ~ 60 W
- Conclusion
    - promising technology
    - performance sufficient
    - cost high but not prohibitive, will drop with new generations of NPs
    - strong points: flexibility, support, single module for all applications

Q&A:

- Ph. Charpentier: Is the mux presented in the first talk fix or subject to change? What is the view of the subdetectors? Is partitioning taken into account? Granularity of partitioning is an important input.
    - J. Christiansen: Only assessment by electronics coordinator. Only fixed value is the number of links.
    - B. Jost: High mux meant for flat architecture. Now things have evolved in subdetectors. Granularity is mainly defined by the TTC system. No data-specific handling between FE and trigger CPUs.
- H. Muller: What about Crates?
    - B. Jost: Aim at single width VME-module. ~15 crates
- J. Christiansen: Many subdetectors start building specific FEMs. This may lead to high load.
- F. Formenti: Does programming in assembler limit more fancy features?
    - N. Neufeld: IBM is working on a C compiler. A GNU compiler port seems to be possible in a short time.
- M. Schulz: Are CPUs pipelined?
    - N. Neufeld: Yes, one should avoid branches. Program memory sufficiently big to put code without subroutine codes (16k 32bit).
- M. Schulz: How should the Level 1 scenario work? Switch will be needed.
    - B. Jost: May be done with 30 more NP RUs as 1:7 demux.
- M. Schulz: $/CHF ratio over time has to be taken into account.
- F. Formenti: Is this a single-vendor solution? Are there future developments already foreseen?
    - B. Jost: Around 10 different NPs available on the market. No manpower to look at all available solutions. ATLAS had a look at the INTEL implementation. IBM solution uses MAC in hardware, so the processing is far simplified. Cost for NPs at the same level for all manufacturers.
    - B. Jost: IBM will produce version 2. ATM support will be in. Next generation should be available in I/2002 and support 10 Gbit Ethernet.

VI. QUESTIONS AND ANSWERS SESSION

- J. Christiansen: A new version will be built on a PCI card after finalizing the existing one?
  - H. Muller: After October the new version will be attacked if requested.
  - F. Formenti: That should be discussed in advance.
  - H. Muller: The NIC could then be put on PCI, not as mezzanine.
  - B. Jost: This would limit the setup to one RU per PC.
  - M. Schulz: New PCs have three independent 64bit/66MHz PCI busses with two connectors each. This would mean 3 RU applications inside one "server"-PC.
- Ph. Gavillet: How can the module be used as a simple test system e.g. in the lab?
  - B. Jost: Only a tree has to be implemented with a sizeable quantity of NP RUs and a PC with Gbit interface. The module can easily be used for any test setup within LHCb.
  - H. Muller: Exerciser was shown. This would do the job to test the RU system. For only four channels, a RU is the full system including event building.
  - B. Jost: This was done for the performance tests, except for the final transfer of the built events to a destination like a PC. It should not make a difference if one would connect FE hardware instead of test generators.
  - J. Christiansen: Should the RUs be used in real beam tests, too?
    - B. Jost: This implies that there exists real FE electronics.
    - St. Wotton: One does not learn much by plugging real Level 1 electronics instead of a generator.
    - J. Christiansen: There will not be final FE electronics before 1-1.5 year's time.
    - B. Jost: The bulk production of electronics will take place in 2003.
    - F. Formenti: For a production in 2003 one should exercise the whole system in small setups at the latest in 2002.
    - J. Christiansen: The subdetectors should be asked when they want to have a RU for their testing.
    - B. Jost: To build a readout system with RU for the testbeam one would need a complete DAQ system with farm. That does not seem to be feasible.
    - Ph. Charpentier: At a certain level of detector evolution, the subdetectors will have devices that can produce a high data rate. Do we want to build a readout network especially for this?
    - J. Christiansen: The quality of communication between FE and RU has to be tested early enough to allow for correction of serious problems.
    - Ph. Gavillet: Do we want to produce a "vertical slice" of LHCb, i.e. to test the interoperability of the full system?
    - B. Jost: One or two RUs are sufficient for these tests.
    - M. Schulz: There will be an integration test Level 1 electronics / FPGA RU late 2001. Then a complete slice will be available.
    - H. Muller: A complete documentation/support set has to be done for the exerciser. A special test will show if the correction possibilities of STF are sufficient or have to be improved.
  - J. Christiansen: When can we have a full slice available for testing?
    - B. Jost: End of 2002.
    - B. Jost: The LHCb DAQ will be different from LEP systems, as there will be no cross talk. Every part of the system will be a complete system, scaling plays a role in configuration, but not in operation.
    - H. Muller: This is not true, as in the end there will be a main event builder.
    - B. Jost: This test cannot be done, as you will not be able to overload a switch with only a limited number of sources.
- Ph. Gavillet: Any remarks on the cost?

- B. Jost: NP RU has a high cost. Given numbers are rather pessimistic than optimistic.
- J. Harvey: ~ 3MCHF for all RU applications including event building, 2 MCHF for RU including links and infrastructure, 125 kCHF TFC, 1.5 MCHF farm, 1.4 MCHF event builder.
- J. Christiansen: subdetectors should pay for FEMs.
- C. Gaspar: What would a normal switch for the event building cost?
    - B. Jost: The full speed Gbit Ethernet switch 4x4 cost is not known. The test setup Foundry switch, which is functionally equivalent, was ~18 k$, so we assume ~1000$/port for the NP RU.
- J. Christiansen: What are the prices for IBM NP based switches?
    - B. Jost: The available product looks like a WAN router.
- F. Formenti: Why can the evaluation kit not be used as RU itself? What is planned in addition?
    - B. Jost: The form factor is not feasible. In addition it has more features than needed. The cost is ~25 k$.
    - H. Muller: Why don't you ask IBM to build the whole system?
        - B. Jost: That was done. The answer was yes, 2.5 M$ for the design. Another company was asked, the answer was ~300 k$ for the mezzanine.
        - B. Jost: To get a customized router, from e.g. Alcatel, would steel the possibility of loading own code.
    - H. Muller: LHCb cannot be the only customer for something like this.
        - B. Jost: CISCO bought 9 evaluation kits and will build switches with the NP.
    - F. Formenti: Unfortunately the NP RU is in competition with industry which will definitely come up with something similar.
    - J. Christiansen: F. Formenti's group should also look at the cost and possibilities to build the module at CERN.
        - F. Formenti: To build such a design at CERN will take long.
        - H. Muller: To take schematics and produce a layout will not be a generic job for CERN, this should be done outside.
        - F. Formenti: To build a first prototype will take at least 6 months. There are many aspects to be evaluated.
        - J. Christiansen: Only PCB routing is not the main issue.
        - H. Muller: For PCB one would go to a company, like Mont ford, and get a completely tested device.
        - J. Christiansen: One needs a designer taking care of the full process, design, testing, signal integrity, …
        - B. Jost: If the project should start, one would not hesitate to buy knowledge from IBM.