The Matching Unit *Version 2* for the HADES Trigger System

Revision History

Revision 0.1	20.01.2003
Revision 0.2	17.02.2003
Revision 0.3	03.03.2003
	Suggestions of Jörg Lehnert and Jan Hoffmann have been integrated
Revision 0.4	24.03.2003
	Suggestions of Ingo Fröhlich have been integrated

Author: M. Traxler

Purpose: Description of the Matching Unit Version 2 and concentrator boards for the HADES LVL2 Trigger

Date: 2003-03-24 12:25:27

1 Concept

The overall concept and the physics background of the Matching Unit (*MU*) in the framework of the Second Level Trigger of the HADES Trigger System can be found by following the link:

http://www.physik.uni-giessen.de/hades/groups/matching/Matching_unit.concept1.pdf

2 Experience with version 1

The version 1 of the Matching Unit is a prototype and was only build to prove that the concept is feasible. It has not the computing power and data paths to fulfill the requirements of the HADES-proposal (100 kHz LVL1 trigger rate).

2.1 Results achieved with the MU V1

The MU V1 proved that the concept of a central LVL2 decision taking device to be feasible. The module is currently (January 2003) being used in the HADES experiment. It is running absolutely stable (no DAQ-stops due to the MU are known). The proposed algorithm for the LVL2 trigger decision is implemented and also checked by an off-line analysis (except invariant mass determination, which is implemented but was not checked until now).

2.2 Problems found in the prototype

The only severe problem of the MU V1 is the performance. The maximum rate achieved with this

module is 30 kHz LVL2 trigger rate, which is expected to be reduced to 20 kHz under beam conditions (could not be tested up to now, due to the limitations of the other detector systems).

The bottlenecks of the MU V1 were determined by actual timing measurements of the whole system. The following issues have been determined:

a) The time from the addressing of an IPU to arriving data from the IPU to the MU is approx. 300 ns due to pure traveling time of the signal along the long cables. This adds to 1.2 µs total for 6 IPUs in the sequential addressing mode.

b) As the MU is programmed in *'C'* to achieve greater maintainability, an interrupt for pending data must at first save the local variables. This takes with the current 40 MHz DSP approx. 2 µs, which then also adds up to at least 20 µs (10 IPUs in total, each addressing results in an interrupt, can be optimized by polling) dead time.

c) The idea of the MU V1 with the central scheduler, which receives all the data and then distributes the data to one of the "working-DSPs" has not been successful, due to the fact, that the receiving and checking (trigger tag, length, consistency) of the data is on average the most time consuming part of the work the MU is doing. This is only true if the multiplicity in the detector is low, but for light systems (carbon) this is on average the case. Therefore, it would then take even more time to transport this data to an other DSP, which then on average has "nothing" to do.

3 Concept of MU V2 to accept higher LVL1 rates

Several ideas have come up while using the MU V1 to improve performance and architecture.

3.1 Concentrators

As DSPs are designed for Digital Signal Processing algorithms, they are capable to process larger sets of data very efficiently. Therefore, one should concentrate the small data pieces of the different IPUs and send them in one block to the DSPs, which then doesn't have to react on to many interrupts, which slow down the system (saving of the register set, which gets bigger and bigger for newer DSPs).

Concentrator boards do the addressing and receiving of data from the IPUs. This is done in parallel and not sequential anymore to reduce the propagation delays (refer to 2.2). These are controlled by CPLDs to achieve fastest response time. The concentrator stores as many events from one IPU system (or just a subset) as there are resources available. The bulk-data is then sent to the MU.

For higher immunity to EM distortions and Common Mode Rejection in a hash environment it is preferable to use optical links. These optical links will be used in the new concentrator boards, but as an backup-solution, the usual twisted pair connector will still be available, but as an "Mini Ribbon" connector, which is more robust and half the size than a standard IDC-connector.

3.2 Faster MU V2

A new MU with a fully parallel architecture and a new generation of DSPs (Tiger SHARC) will allow to run with the required 100 kHz LVL1 rate. This can be estimated, due to the following changes in speed and the architecture:

a) Tiger SHARC DSPs TS101 [Ana02a] run with an internal frequency of 250 MHz, so without using the new 3 fold parallel buses with 128 bit width and the two parallel working ALUs, this will speed up the algorithm by a factor of 6 compared to the SHARC21060 DSP at 40 MHz clock.

b) The parallel concentrator concept will reduce the data transport time by a factor of 3, to approx.

3 µs. This will be “bought” by an additional latency of 5 µs, which is of no concern, due to the LVL1 pipes of the trigger system.

c) The parallel distribution of one whole event to one DSP will remove the bottleneck of the central DSP which receives every event and checks it for consistency.

4 Implementation Details of the MU V2

4.1 Block diagram of the MU V2 concept

In Figure 1 the block diagram of the MU V2 concept is shown. The concept is configurable and scalable, which is important, as there are planned extensions of the TOF system as well as a new detector subsystem, the RPC (Resistive **P**late **C**hamber), which will also be used in the LVL2 trigger algorithm.

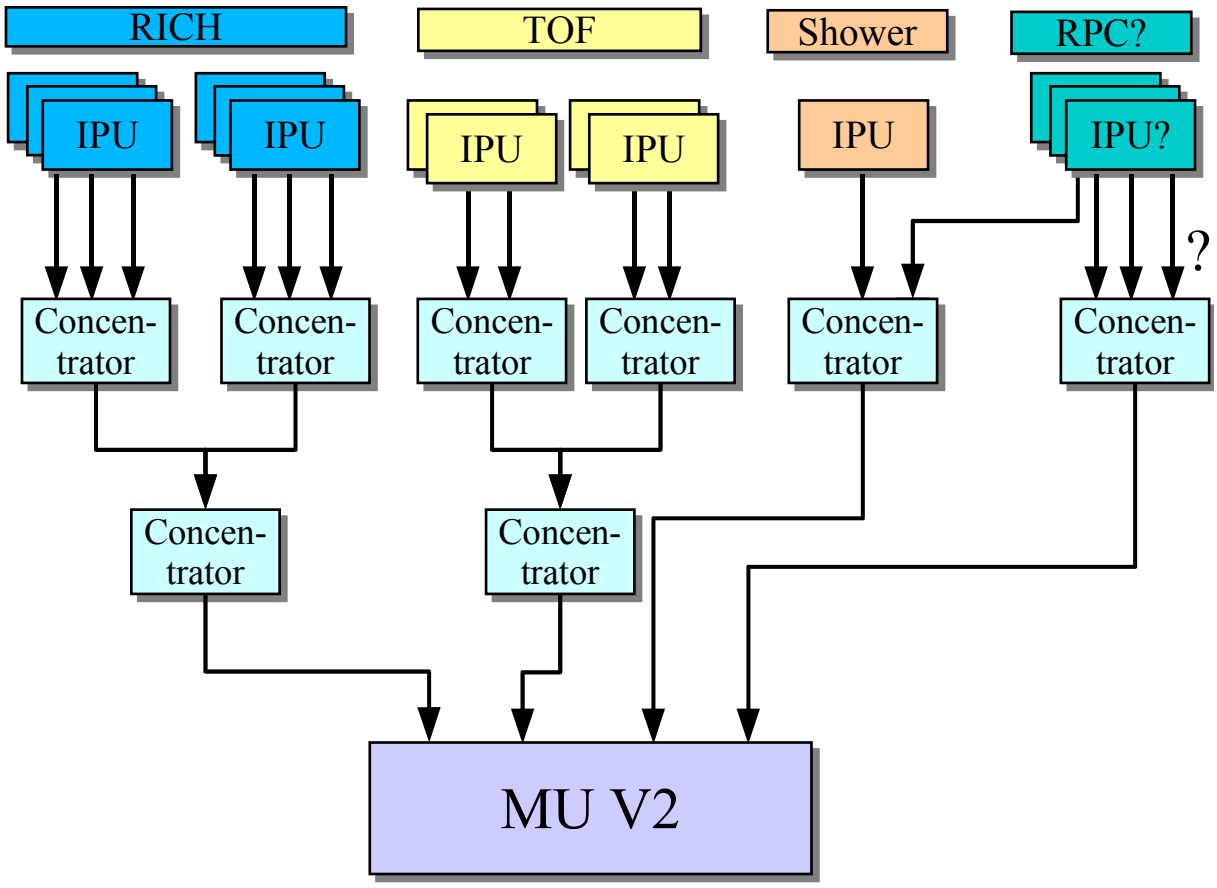

Figure 1: Block Diagram of Trigger System with MU V2

4.2 The concentrator

The concentrator block diagram is shown in Figure 2. The concentrator uses 3 (or 4) IPU-Buses to address and transfer the data form the IPUs in parallel. The IPU-Bus is the existing IPU-Bus with differential transmission of 8-bit data with a 20 MHz clock and up to 7 individual addresses. The input Fifo must be able to store one complete event of one IPU, because the data transmission is event driven and without stop-signal during transmission of the data of one event.

The collected information is then stored in an output Fifo, which must be able to store many events.

The MU then addresses the concentrator board over the MU-link and receives the stored data of one complete IPU-subsystem in one packet. This improves the speed compared to the serial addressing and transmission scheme currently used in the MU V1 scheme.

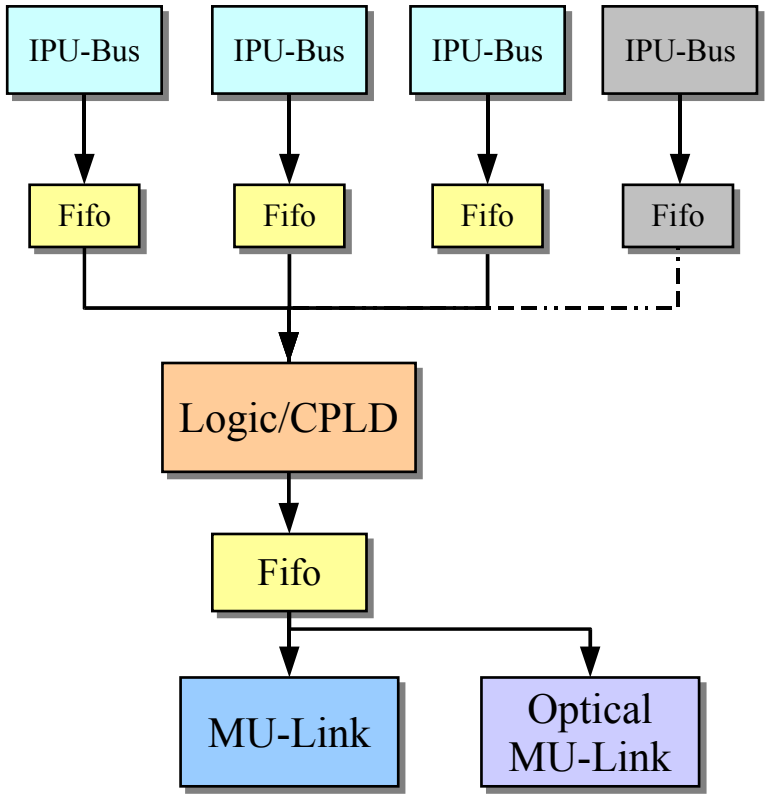

Figure 2 : Block Diagram of the MU Concentrator Module

The inputs may also be from different IPUs, which would have the following advantages:

a) less modules needed

b) improved load balancing

c) no strong restriction in the number N of final MU input ports to N different IPU types

Additionally, this concentrator has a logic core (large CPLD, Cypress Delta39K, [Cyp02]), which allows to make some simple trigger decisions. As a first step in this direction, it will count the number of words received by the RICH-IPU and is able just with this information to determine if there is at least a single hit found by the IPU. In experiments with small systems (e.g. C+C) the number of rings found in the RICH IPU is on average very low (1:10) and thus this information can be used to suppress the transmission of the IPU data to the MU. In the case of an empty RICH-IPU event, the concentrator will send the MU a special command, which signals the MU not to read out the other IPUs, but just discard their data, which is then done by a special address sent to the concentrator.

4.3 The MU Version 2

The Matching Unit Version 2 block diagram is shown in Figure 3. It has 4 ”conventional” IPU-link inputs and an optical input, capable of 1GBit/s transfer rate (upgradeable to 2Gbit/s just by

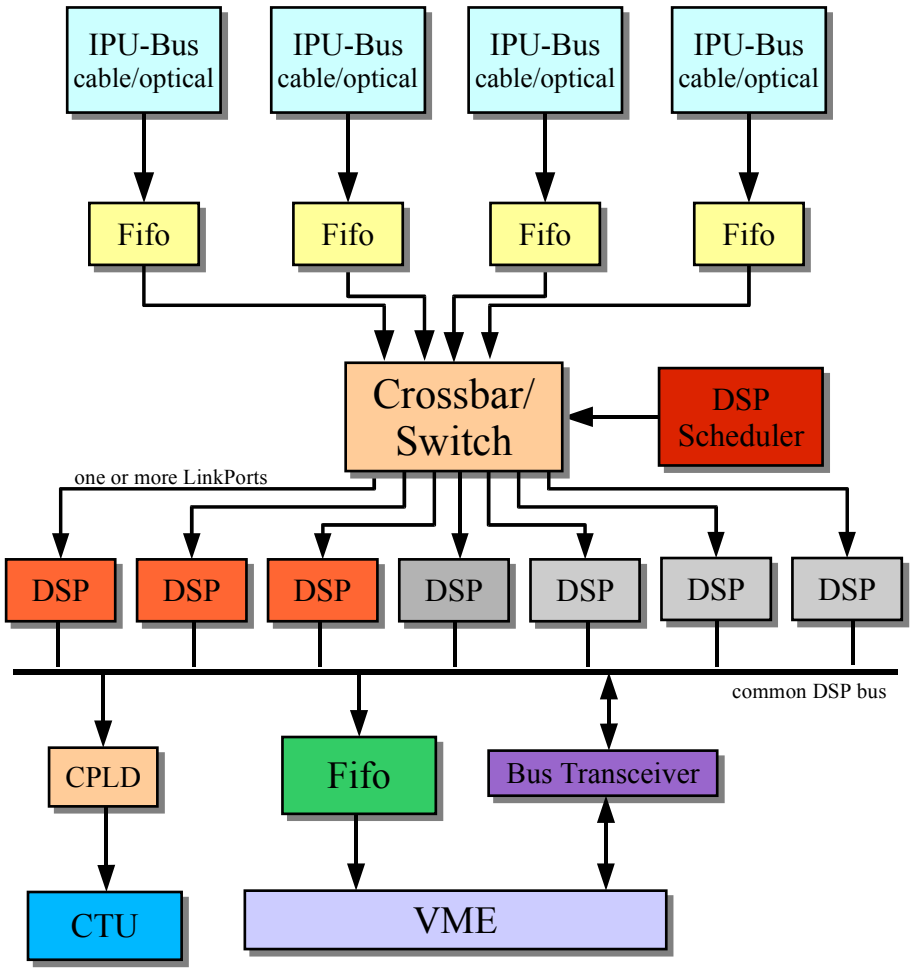

Figure 3 : Block Diagram of the MU Version 2

exchanging the optical SFP transceiver). The optical link can be used, when connected to an concentrator board. The data of the IPUs and/or the concentrator-boards are buffered in input-Fifo's. Then, the data is multiplexed and sent by the crossbar to one of the 7 computing DSPs. The crossbar is controlled by an eighth “scheduler” DSP, which keeps track of the free and occupied resources of the MU computing DSPs and determines where the next event it sent to. Alternatively, one can implement a token algorithm, which simplifies the sorting of the events. One token for the trigger decision, which has to be sent to the CTU as fast as possible, and one token for the transfer of the MU algorithm data to the readout-fifo [Fro03]. Each of the DSPs is performing the full MU trigger algorithm, with all the necessary data consistency checks. After the trigger decision of one DSP is made, it can send its internal data (found leptons/di-leptons, invariant mass, IPU-raw data for off line efficiency determination, etc.) to an large Fifo, which is connected to the common DSP bus. This Fifo can then be read out by VME, without acquiring the DSP bus, which is very important, as VME cycles are very long (approx. 1 µs) compared the 10 ns bus-cycle of the DSPs. As the events in this fifo have to be in sequence (to avoid resorting on the eventbuilder), the events are first stored in the local memory (6 Mbit on TS101) until the DSP which runs the matching-algorithm for the event which follows in sequence is finished. This will allow under “normal” (no events with unreasonable high multiplicities are accepted) conditions, that all DSPs on the board process an event in parallel and no idle time is introduced.

Due to the high demand of signal integrity for the common DSP bus for a bus-cycle of 10 ns, it is planned to have the possibility to split the multiprocessing DSP-bus into two segments, of only 4 DSPs each. Simulations have shown [Plex02], that it is very problematic to keep the signal timing in the specifications when building a 8 processor system at 100 MHz bus frequency.

The arbitration of the computing DSPs receiving data from the crossbar and sending the data to the fifo is managed by the scheduler-DSP, ensuring that the event sequence in the readout fifo is adhered. The fifo has to be very deep, to allow the readout to be very late, which is needed in the current scheme, as the TOF-concentrator as well as the MU are read out in the same crate by one readout process.

References

[Ana02a]: *TigerSHARC DSP Microcomputer*, Rev. .0, 10/02, Analog Devices, Inc., 2002

[Cyp02]: *Delta39K ISR*, Rev. F, Cypress Semiconductor Corporation, Nov. 2002

[Fro03]: Proposed by Ingo Fröhlich, 24.03.2003

[Plex02]: *ADSP-TS101S MP System Simulation and Analysis*, Revision 1.2, Plexus Corp. March 12, 2002