CERN/DRDC/90-56
DRDC/P12
30 October 1990

R&D Proposal

Embedded Architectures for Second-level Triggering in LHC Experiments (EAST)

M.v.d.Guchte, L.O.Herzberger, A.v.Inge
Computer Systems Group, University of Amsterdam

K.Bos, J.Vermeulen, L.Wiggers
NIKHEF-H Amsterdam

V.Buzuloiu, M.Nelu, R.Ruscu
Institute of Atomic Physics, Bucharest

J.Koch, T.Maroti, S.Molnar, J.Toth
Central Research Institute for Physics, Budapest

R.K.Bock (*), A.Gheorghe, W.Krischer, R.McLaren, J.Pfennig
CERN, Geneva

Z.Hajduk, M.Kajetanowicz, P.Malecki, Z.Natkamiec, A.Sobala
Institute of Nuclear Physics, Cracow

B.J.Green, T.Medcalf, J.Strong, A.J.Wildish
Royal Holloway and Bedford New College, London

J.Badier, C.Charlot
Ecole Polytechnique, Palaiseau

F.Marzano, L.Zanello
Dipartimento di Fisica, Universita di Roma "La Sapienza", Roma

W.Lourens, A.Taal
University of Utrecht

H.Leich
Institute of High Energy Physics, Zeuthen

Note: Proposal is made subject to acceptance by national funding agencies

(*) Spokesperson

Abstract

The next generation of hadronic colliders will confront our community with extraordinary problems of building highly *time-sensitive* and *selective* detectors. Both problems are not solvable simply by extrapolations from the past. Some solutions will have to be found in strongly detector-dependent areas, e.g. signal processing for separating optimally individual bunch crossings in the detector and its associated front-end electronics, or synchronization of events whose transit time through detector and electronics is several bunch crossings. Other problems will have solutions that are less detector-specific; they simply result from the high data flow.

In this proposal, we present a coherent plan for exploring one of the critical aspects of the implementation of detector selectivity, that of embedding 'intelligent' devices for triggering. We break down the interesting components of this problem area into the categories: triggering architectures, data collection and switching components, and system integration and modelling. Building on existing experience (like the studies done as part of the LAA project), and drawing as much as possible on concepts and implementations existing outside High-energy Physics, we propose to evaluate possible solutions and critical components in the laboratory, and to demonstrate their functioning in realistic test devices and in real detector prototypes, to the extent that the developments of detectors (in particular front-end electronics) allow us to do so. Our proposal covers a time span of no more than three years, in order to demonstrate the viability of the more promising technologies, in time for planning future full-scale LHC detectors.

1) Introduction

Future high-luminosity colliders will provide our community with unique opportunities for physics discovery by experiments. In order to take advantage of these possibilities, new standards for building particle detectors and associated data collection electronics will have to be set. Indeed, not only will the collision energy be boosted by more than an order of magnitude compared to existing collider machines, but the cross sections at which interesting phenomena are predicted to be found will be the smallest ever. The machine parameters of the LHC are optimized for pushing the discovery limits of physics phenomena towards the lowest possible cross sections.

At TeV energies, however, total cross sections of proton-proton collisions are high. Comparing the *interesting cross sections* with the *total rate* of physics events in which they are embedded, presents us with an environment that is unprecedented in its technical difficulty. The combination of high initial collision rates with the scarceness of interesting phenomena translates into new challenges to detector builders. Detectors will have to be sensitive to extremely short time scales, viz. at the level of individual bunch crossings. They will also be required to reduce, in real time, an extraordinary mass of data. At least six orders of magnitude must be gained, using the 'interesting' signatures against the very high cross section 'background' of hadronic physics. Sophisticated local data processing will be required, using the best information available, to achieve such factors. We advocate in the following an approach to this problem area, which is based as much as possible on partial solutions imported from quite different applications, and which relies, for quick results, on an inter-laboratory collaboration with a loose connection to detector developments.

LHC opens a new era for detector building

Today's hadron colliders, at CERN and FNAL, operate at bunch crossing rates of several microseconds, and with maximum rates of tens to one hundred thousand per second. This allows practically dead-time free operation, and comparatively simple decisions which reduce the rate for high-level general purpose computers to take over (real time as farms, or after mass storage in central mainframes). The operating conditions for e-p machines are, mostly due to lower cross sections, much more relaxed, despite the 96 nsec bunch crossing period of HERA. Electron positron colliders like LEP are in a different category altogether, due to the very much smaller cross sections encountered there.

The hadron-hadron future will look quite different. The design parameters for the LHC [1], even ignoring the highest-luminosity option, show that bunches will cross at 15 nsec intervals, and that multiple collisions will regularly occur in every bunch crossing: fluctuations will oblige physicists to cope routinely with tens of collisions in each 'event'. It is no exaggeration to call the data collection environment of the LHC a hostile one. Fast detector response, corresponding front-end electronics, an internal synchronisation scheme and signal processing to pick the correct 'time slice', and extraordinary selectivity thereafter, will thus be vital detector qualities. Failure to give them the necessary attention and allocate sufficient resources will *impair the physics output* of LHC most seriously.

Interesting signatures require real-time intelligence

At the core of high selectivity is a problem of *physics signature.* There exists a number of 'gold-plated' signatures in the predicted physics, like in the key search for the particle of the electroweak symmetry breaking, viz. the Higgs particle, which may manifest itself as two Z and eventually as 4 charged leptons. These signatures are low-cross section 'paths of luck', expected to be available when LHC has reached full luminosity. The majority of forecasted physics, and the much higher cross section signatures, are found in partially leptonic channels, with compound signatures including QCD jets. They will require carefully tuned cuts in pt and/or lepton isolation, and association between several of these phenomena [2], [3], [4]. While it is likely that some experiments will be built using detectors tuned for one or the other of the exceptionally clean signatures, which may require only comparatively simple triggers, it is our belief that many studies at the LHC will rely on detectors with associated 'intelligence': intelligence extracts the relevant physics information from the flood of data, and provides flexibility which is needed to find clean signatures under varying accelerator and physics conditions. In particular, decisions will have to be based on correlations between phenomena in different parts of the detector(s), and not simply rely on isolated inclusive signatures. Some of the necessary cuts, to be efficient, are likely to require careful and permanent tuning and optimal information including detector calibration constants; a maximum of access to intermediate results also has to be part of a triggering system. Using the term *intelligence*, in this context, does *not* imply *full programmability* in high-level language; rather, we refer to devices providing the *required flexibility,* be it by programming at any level or by easy reconfiguring.

The ultimate limits in what can be achieved in real time decision making should be set by our understanding of the detector components and their calibration at time of decision, but not by the difficulties of access to technologies.

Coupling with detector developments

We maintain that many detector subsystems now in discussion or proposal state, share common data collection and triggering characteristics vital for extracting physics, and far from trivial in their technology. They will eventually require a major fraction of a detector's budget, and hence need resources for studying and pilot projects in parallel with detector developments, and at a comparable level. This is true for any of the non-trivial physics signatures, and unrelated to whether a 'general-purpose' or a 'special-purpose' detector is intended. The R&D activities proposed below are complementary to the various detector developments, and will eventually be crucial stepping stones for most of the collaborations that will form around LHC. They will be a vital preparation of the decisions future experiments will have to take a few years hence, and should be seen, due to the long lead time, to be on the critical path for collaborations.

If the discovery potential of future detectors is believed to be directly coupled to the capability of detectors to extract high-level information from fine-grain information, then a sizeable fraction of future detector budgets will have to be spent on this electronics/architecture aspect, and a development plan as proposed here is an obvious necessity.

We elaborate below how we intend to couple our work to several of the proposed detector subsystems. Some detector prototypes or hardware emulators will be needed to provide an environment for the desired proof of principle, i.e. for demonstrating system parts in real time. In longer term, detector association should become a more and more stringent constraint, until all or most detector-independent work will be absorbed by future collaborations in, say, three years from now.

We would also like to stress that, beyond the necessary association with detector R&D projects, we will seek to stay in close contact with other generic work in data acquisition and/or triggering that may be proposed and undertaken in the framework of R&D for LHC detectors (e.g. the development of frontend electronics). Common development tools are an obvious area where harmonization is essential, but we also foresee that common interest in certain industrial components will lead to local collaborations on some subtopics.

Inter-laboratory approach based on developments in industry

The developments proposed concern 'embedded systems', multiple components of some programmability, coupled to detector electronics. Industry contributions are likely to play an essential role in implementing practical solutions. The systems that will eventually be part of experiments, will have a useful life of a decade or more, and scalability will, therefore, be a major concern as much as the inevitable evolution of technology.

The detector-oriented developments proposed in view of LHC today can not reasonably be expected to separately support developments of a purely technical nature. On the other hand, any one of our HEP laboratories is hard put, by itself, to provide the resources for a comprehensive R&D program: the existing knowhow in our laboratories combined, is barely sufficient for a serious development program in this domain. While the intended developments will need serious adaptation to individual detector designs, there is enough overlap in the problems to ensure that a coordinated program for complementary developments in several laboratories starting now, will at least stand a chance of introducing the more important long-lived concepts in several future experiments, and avoid duplication.

It is, therefore, our suggestion to support, in the short term, the more critical LHC-related projects in electronics and real-time computing as inter-laboratory collaborations, and independently of their firm and definitive association with specific detector developments.

2) Overview of Data Collection and Triggering

We show in fig.1 a schematic view of data collection in a high-rate experiment, useful mostly to identify our terminology. Front-end electronics, after shaping and preamplification, holds information in analog form over the time it takes to form a first-level trigger (this 'latency' is assumed to be of the order of a microsecond). Some channel signals are split and participate in forming the first level trigger. When available, this trigger, assumed to have a 1/1000 selectivity, broadcasts its decision to all local pipelines. For retained events, signals are now digitized and regrouped. Again, they have to be held in a pipeline for the latency of the second-level trigger, and a fraction of the signals is used in forming the second-level trigger algorithms. Eventually, trigger results selectively let events continue to the final phase, full digitisation and collection of all data pertaining to the same event ('event building'), whence they can enter full (again selective) analysis in a real-time farm ('third-level trigger').

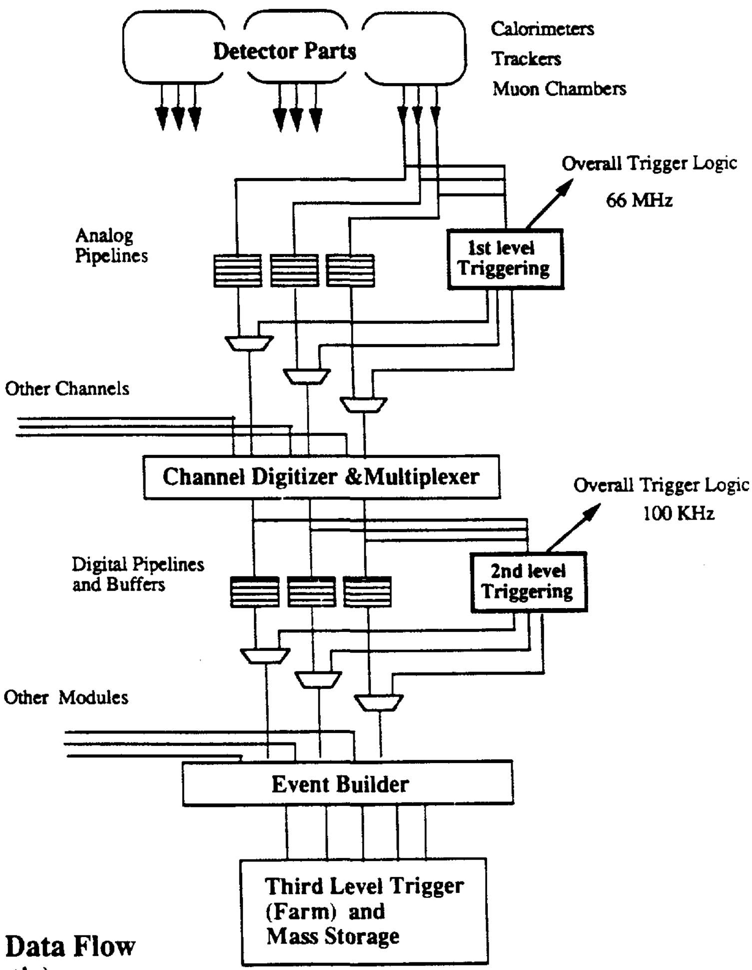

Fig.1:
Overall Data Flow (Schematic)

Note that several critical choices are not apparent in the diagram, but could have a major influence on the second-level trigger architectures and their implementation. As an example, the availability of on-detector per-channel digitizing at bunch crossing speed, as is likely to be proposed for calorimetry, or the success of analog signal transport, for individual channels, over comparatively long distances by optical fibres, as researched in the context of TRD detectors, can not fail to influence deeply the subsequent data flow.

In this proposal, we concentrate on a 'generic' second-level trigger stage, i.e. on data flow components to be used by detector components contributing to decision making using digital processing of some complexity. Two relevant boundary conditions have to be observed for this layer: the data rate expected to flow into the subsystem from the detector ('input'), and the rate which can be unloaded onto the event-building stage ('ouput'). In between these two limits, the specific contribution which a detector component can be expected to make to the selection process can be expressed in terms of a data flow diagram and a triggering algorithm. We have thus reduced the scope of our proposal by two relevant assumptions, as expressed in reference to fig.1.

a) The data at the initial bunch crossing rate will have been synchronized, pre-processed for a bunch crossing in analog or digital form, and locally held, by (largely detector-specific) fast electronics. The arising problems and the technologies involved in this front-end electronics part have also many overlapping aspects like VLSI development, radiation hardness, power consumption, cable or fibre plant design etc. We make the assumption here that other general proposals and detector-specific developments will cover this area. In agreement with discussions over the last years and backed up by Monte Carlo studies [5], we assume 'second-level' subsystems to be fed, after a first-level trigger, with digital data at an average rate of one event every 10 μsec.

b) We take for granted a final link in every experiment's data flow, a massive battery of general-purpose computing elements. We expect it to be backed up by extensive data bases, and to support the full range of high-level tools available. In such 'farms', the final high-precision cuts on full event data will be implemented as high-level language algorithms, and data flow monitoring and calibrations will be executed. We assume that this area is being investigated by other groups, and in particular that its progress is largely dictated by the evolution of commercial offerings. We make the assumption here, that full digitizing (even of passive detector parts) and the 'event building function' of the system, will be able to sustain an influx of at most 1000 events per second, expected to be treated in a 'third-level trigger' farm of microprocessors. Note that the third level trigger is not necessarily implemented as a physically localized unit; it may well be a network of processors, with access to data memory units spread throughout the upstream electronics.

With these assumptions, we have situated our proposal domain of 'second-level triggering' between fast electronics for front-end and first-level triggering on the one, and overall processing in powerful computer farms on the other hand, without firm assumptions about siting, i.e. about the physical location of the components. We have further identified, between the input and output rates defined above, a critical factor of at least two orders of magnitude, to be gained by what we call a 'second-level trigger'. In the same simplified language, the 'first-level trigger' is assumed to have reduced the data rate from bunch crossing rate by a factor of typically 1000.

Quite generally, second-level ('intelligence') triggering is distinguished from first-level ('brute force', not synonymous with simple) triggering by the decisiion frequency and the amount of time available for decision taking (latency). We see several further characteristics of second-level triggers:

The use of *fine-grain information* i.e. of locally fully digitized data is likely to be necessary. Obviously, better information will allow finer cuts, and algorithms close to optimal. Hence, where distribution functions are steep (as is the case for most pt distribution functions), an important gain in signal/background ratio can be expected, and complicated studies of trigger efficiencies can be avoided. Important examples are high-precision tracking for muons, electron identification by transition radiation or preshower counters, or detailed cluster analysis in the lateral and longitudinal direction of calorimetric showers;

The accurate association between *different detector parts*; may be achieved. This will be a necessary ingredient in event selection, even though our assumption is that we perform such association locally and leave global event building (i.e. preparing data such that any association algorithm becomes possible) to a later stage. Examples are the association between tracking or TRD (or preshower counters) and calorimeter, tracker and vertex chamber, and muon tracker and calorimeter.

The *association of first-level trigger information* for working on partial data in the second level, is another characteristic. First-level 'pointers' may be obtained from the same or a different detector; examples are energy accumulations in calorimetry for finer cluster definition or for electron identification as in the previous paragraph;

We believe that *commercially available computer-like architectures* can be applied to implement the necessary second-level algorithms. In view of this, we have to assess carefully and in detector dependence which of the architectures existing in applications outside high-energy physics, special or general-purpose, can be adapted to satisfy economically the requirements of our future detectors, and how they can be interfaced to the rest of the system.

It is the last argument, commercial availability, which has prompted most strongly the present proposal: the availability, outside high-energy physics of hardware and software which we can put to our use. Hence there is a need to familiarize ourselves with promising components and to choose the ones most suitable for our purpose. Special-purpose processors, for instance, which have been developed for applications in image processing in space, medicine, even in the wide consumer market (High-Definition TV, [6]), can deliver extraordinary computing power equivalents if used for a suitable task. A hard demonstration that such devices can indeed be used in the context of particle detectors, and the assessment of the competition they are undergoing from general-purpose devices, require sizeable resources. The constant and, in many fields, turbulent evolution of technology requires, furthermore, good contacts between our community and industry (or industrial technology research), which will contribute to establishing the expertise required for making decisions of considerable importance in the future.

3) Projects Areas in More Detail

3.1 Trigger Processor Architectures

It is our intention to continue preliminary work done in the context of the LAA project at CERN [7, 8], which has put into prominence two commercially available architectures that seem able to execute non-trivial algorithms logically close to the detector and at the required speed: Pipelined Image Processing modules and 'SIMD'-type Massively Parallel Processors.

a) Pipelined Image Processing systems are special-purpose 'computers' made up from a variety of specific architectural modules, developed for quite different applications, like in television, aerial surveillance, medical scanning. These devices can be purchased off-the-shelf from industry; their performance is evolving over time with a rate similar to that of microprocessors, and their internal and external interfacing possibilities are under rapid evolution, too. They need detector-specific interfaces, but no internal changes or additions, in principle. Pipelined image processors execute e.g. neighborhood operations like convolutions or morphological operations, useful for feature extraction. They develop for these tasks an enormous equivalent of computing power. Being two-dimensional devices, their most promising application is in second-level triggering on fine-grain information from calorimeters, i.e. e/π separation and pileup rejection with variable isolation criteria, also hadronic jet definition, e.g. for quark tagging. There exists a model of this architecture at CERN, the MaxVideo system [8], running, amongst others, algorithm tests on recently taken SPACAL data. Such systems will have to be interfaced to suitable detector prototypes as they are being built, using commercial data connections and flexible transmission controllers as much as possible. Further algorithm research on these modules will have to be done, and new types of higher performance modules, with faster internal bus, and easy interfaces should continuously be understood, i.e. evaluated and purchased (e.g. the next generation of MaxVideo, or the future Data Wave Processor [9]).

b) 'SIMD'-type Massively Parallel Processors: Here we deal with an open architecture of many thousands of mesh- or string-connected processors, working in lockstep ('SIMD system', for 'single-instruction-multiple-data'). Individual processors are far from general-purpose, typically, they are 1-bit or 4-bit processors, with few basic instructions implemented, but with the architecture being fully open for programmed acces, and hence of considerable sophistication. For suitably parallelized applications, unprecedented computing power can be extracted. The recently launched collaboration (MPPC) between Aspex, CERN, Orsay, and Saclay [10], to build several 8Kprocessor boards and program them for tracking (TRD and muons) and calorimeter applications, targets one possible implementation of SIMD parallelism based on an industrial design, the ASP (Associative String of Processors). Like pipelined image processors, such systems need to be complemented by interfacing electronics to various detector prototypes, in order to demonstrate the feasibility of the approach. In the case of the ASP this is done in collaboration with the industrial partner (Aspex Ltd.). In case of a positive conclusion of the ASP feasibility project, it can be expected that a detector-specific processor and communication design, based on the same basic concept of many thousands of very simple processing elements, will be proposed in the context of some future detectors.

It is very likely that these specialized pipelined and parallel architectures will have to be complemented by processing capabilities of a more conventional type, i.e. programmable general processors. We intend, therefore, to complement the study of specialized architectures with a study of 'MIMD'-type architectures, based on networks of transputers and/or relatively cheap digital signal (or other) processors, augmented by facilities for inter-processor communication. Given future processor performance and the possibility to reduce data rates locally by partitioning, networks of commercially available processors may have a major role to play. We expect to make use of the existing experience in our collaboration with DSP-s and MIMD (loosely coupled processors) architectures, e.g. with transputers [11], or take advantage of the work on FDPP [18].

The exploration of the capabilities of neural networks ([12], [13]) will also be part of the collaborations's work . This includes the exploration of the neural net architectures most suitable for the implementation of real-time algorithms for generic peakfinding, calorimeter cluster amnalysis, and track finding. The study will deal with the simulation of different architectures [14], like feed-forward neural nets, and with their possible hardware implementation [15]. We will also study the possibility of realizing second-level trigger architectures entirely as custom-made systems, like a contiguity processor [16], or devices based on general products like content-addressable memory [17]. We would also keep in mind that the complications due to special architectures and their programming are to be checked permanently against the yardstick of what general-purpose processors, so much more familiar to handle, could achieve in their place. Conventional processor architectures may be sufficient, given future performance of processors, if the locality of decision making can be put to use in reducing the data rate (locally).

3.2 Connection and Switching Technology:

In high-rate data acquisition and in any second-level trigger system, the fast buffered data flow from frontend electronics to the event collection and/or triggering devices requires multiple rearrangement and possibly time multiplexing of data, schematically shown in fig.2.

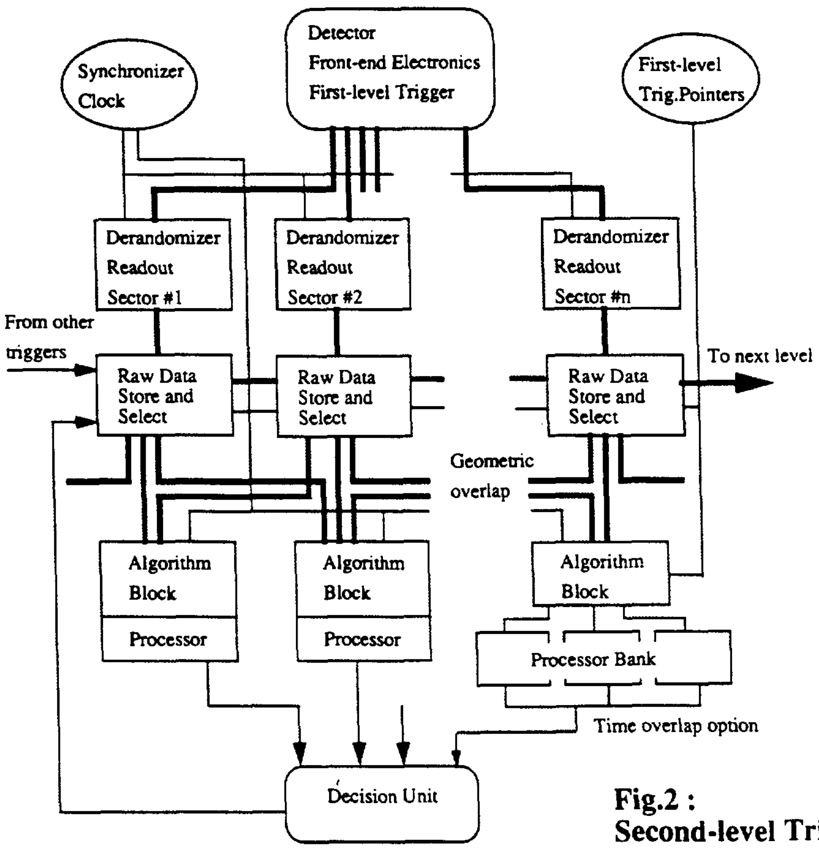

Fig.2 :
Second-level Trigger
Typical Data Flow

General buses as widely used today will not have the required bandwidth, and appear expensive because largely overspecified: locally or globally, high-bandwidth point-to-point links with minimal protocol and (largely synchronous) flexible switching devices will instead have to be studied and physically evaluated. Switching may, of course, be done in special hardware, but could also lead to programmable or LUT-driven devices issuing addresses for block DMA transfers, and thus allow more flexibility. Programmable data transformers (e.g. DSP-s for data compaction) must be seen as part of such studies, a possible starting point might be the high-bandwidth intelligent crate controller card developed at UCL [19]. Flexible switching will be needed because data has to be selected and rearranged *spatially,* in order to cope with locally defined algorithms, and distributed *temporally*, because operations may be pipelined over several events, to overcome performance limits (at the expense of latency, of course). The advantages and limits of locally synchronous systems should further be explored seriously: synchronous operation is a solution simplifying considerably control and data transfer protocols. A sufficient condition for synchronous operation is the data independence of all operations. The use of components developed for cache memories in hierarchical computer memories is also to be looked into.

3.3 System Integration and Modelling

With increasing complexity of the data flow system on and around a detector, the sound engineering of fully predictable solutions takes an increased importance. The modelling, at behavioral level, of mixed hardware/software systems and of system parts, satisfying the system's requirements, is a major aspect of such an engineered solution. Modelling proceeds by stepwise refinement, adapting the model permanently to improved (because prototyped) understanding of system parts, and 'synthesizing' down as automatically as possible to hardware and software design. We intend to use commercial software tools ('platforms'), in common with other hardware and software activities. Local interfaces between modelling tools, proprietary simulation packages (e.g. for specific trigger processor architectures) and physics and detector simulators may well become part of our activity.

4) Initial Workplan

The collaboration is driven simultaneously by the intention to evaluate various *commercially available components*, and by the desire to judge them for *specific applications*. We have therefore made an initial choice of triggering architectures and connection components for study and implementation, and of target detector projects with whom we want to collaborate most closely. Our planning concerns primarily the subprojects of the first year. It should be understood that planning beyond these initial activities is part of the proposal; other subprojects will be added as suitable. We assume that we report to the DRDC after the first year of activity, and update our workplan at that time.

4.1 Collaboration with detector developments

We have agreed to use the following detector developments as realistic test environements for architectural and algorithmic developments. They cover a wide range in the relevant aspects like occupancy, locality of algorithms, and, probably, siting of frontend electronics.

a) the SPACAL calorimeter [20], whose second-level trigger will consist of using local digitized data of finest granularity, to obtain e/hadron discrimination, jet definition, and pileup rejection;

b) the TRD detector [21], whose second-level trigger is invoked by a first-level electron candidate from a (electromagnetic) calorimeter, and whose role is local track finding with statistical pulse height analysis, for e/hadron discrimination, possibly in presence of jet or other close-by background tracks; we are also considering using a model for more general tracking, like the microstrip gas detector [22];

c) the muon detector [23], where second-level triggering constitutes a task of high-precision tracking in the projection defining the transverse momentum.

4.2 Trigger algorithms and architectures

We want to concentrate on studying

a) the pipelined image processing devices, presently characterized by the MaxVideo system. New components for this device, with faster internal and external line speed, and improved (commercial) software will be required. Some more recent chips in this market, with yet another major step in link speed, should also be evaluated. The application of these devices could be primarily in calorimetry, hence tests will be done in collaboration with SPACAL.

b) the interfacing and embedding of ASP boards; track triggering using the SIMD architecture ASP is part of the feasibility study in the existing MPPC collaboration. To the extent that interfacing to the detector and additional architectural components may be required, more resources may be needed than foreseen in MPPC, for application-dependent parts. Our target detectors for the ASP are the tracking devices, i.e. the TRD and the muon detector.

c) model solutions based on DSP-s for SPACAL and for the TRD. Hardware implementations and benchmarks with cheap digital signal processing devices will involve Motorola 56000s and the FDPP-introduced AT&T DSP32C.

d) solutions based on networks of general-purpose processors, in particular of the next generation of transputers [24]. Each transputer may eventually be coupled to one or more special-purpose IC-s, to speed up critical parts of the algorithm.

e) evaluation of the use of CAMs and of he suitability of a simple contiguity processor analog (both custom-built processors), for the fine tracking needed as part of the muon detector.

f) the possibilities of using neural network techniques as a realistic second-level trigger, although hardware implementations are not initially planned.

4.3 Connection and switching technology

In connection technology, we want to build connecting hardware using the low-level hardware features offered on the market for HIPPI [25], starting with a detector-independent testing interface and a workstation as monitoring device. Specifications for the test setup and HIPPI as a first transmission test are being drawn up. Suitable alternatives to HIPPI like SCI and fast data switches will be added as our problem understanding evolves.

4.4 System integration and modelling

The modelling activity will concentrate first on obtaining more experience with existing tools (SIMSCRIPT, Verilog, object-oriented languages like Modsim, SMALLTALK or C++), using our own emerging developments as models. As detector and system parameters evolve, our models will expand and include more complete systems, closer to the future reality. The evaluation of any further possible tools for modelling may become part of our activity, if not supported under other programs in our laboratories.

4.5 Funding and requests from CERN

Although it is premature to present a detailed budget plan involving all our institutes, we can estimate the following cost for the first year, which includes a fraction of local buildup of necessary infrastructure.

	(in KSF)
Test bench for 2nd level, prototype	100
Test and work platforms (workstations, test benches, general software) appr.70 KSF per major institute, some covered by existing infrastructure	280
Detector-specific interfacing electronics	100
New architectural components (image processors, DSP-s, switches, fast connections, CAM-s, etc)	200
Software tools, mostly for modelling	250
Collaboration metabolism (common training, travel etc.)	200
Total	1 130

We estimate that the CERN share of this total is about 25 to 30%, i.e. 300 KSF. Estimates for subsequent years can be expected to be similar or lower. This proposal does not include requests for test beam or computing time, as these are covered by the detector development teams with whom we collaborate.

5) Summary

Goals

The aim of this proposal is to demonstrate the most suitable alternatives for architectures running flexible second-level trigger algorithms for calorimetry, tracking, and track identification devices. This includes economic ways to channel massive amounts of data from the detector to these architectures, and useful toolsets for integrating these communication and processing elements into a coherent overall system design.

Strategy

The collaboration proposes to identify a coherent set of components needed in implementing second-level trigger systems in LHC detector subsystems, choosing a small number of target detectors. These components are subsequently to be implemented in hardware and software, using agreed communication protocols and sets of tools for design, integration and simulation. Implemented steps will be tested in the laboratory and, together with the target detectors, in suitable test beam environments.

6) References

[1] G.Brianti, The Large Hadron Collider in the LEP Tunnel, 2nd European Particle Accelerator Conference, Nice, June 1990, CERN/AC/FA/90-01.

[2] Workshop Proceedings: Large Hadron Collider in the LEP Tunnel, ECFA-CERN Workshop, Lausanne/Geneve, M.Jacob ed.,CERN 84-10.

[3] Workshop on Physics at Future Accelerators, Proceedings, La Thuile/Geneve 1987, J.H.Mulvey ed.,CERN 87-07.

[4] Summer Study on High-Energy Physics in the 1990-s, Proceedings, Snowmass 1988, Sh.Jensen ed., World Scientific 1989.

[5] A.Lankford et al.,Nucl.Inst.Meth.Phys.Res.A289 (1990),597.

[6] ECFA Working Group for Triggering and Data Acquisition, Aachen Workshop October 1990, to be published.

[7] A.Zichichi,The LAA Project, CERN-LAA/90-01

[8] S.Lone et al., Fine-grain Parallel Computer Architectures in Future Triggers, Nucl.Inst.Meth.Phys.Res. A288 (1990) 507.

[9] U.Schmidt, Data-Driven Array Processor for Video Signal Processing, IEEE Transactions on Consumer Electronics, 8/90

[10] M.Lea, ASP: A Cost-effective Parallel Microcomputer, IEEE Micro Oct.1988, and MPPC, Progress Report, to be submitted to the DRDC.

[11] L.W.Wiggers and J.C.Vermeulen, The Use of Transputers in the ZEUS online System, Comp.Phys.Comm. 57 (1989) 316.

[12] D.Cutts et al., Applications of Neural Networks in High Energy Physics, Santa Fe Conference 'Computing in High-Energy Physics', April 1990, Proceedings to be published.

[13] B.Denby, Neural Networks and Cellular Automata in Experimental High-energy Physics, Comp.Phys.Comm.49 (1988) 429; and B.Denby and S.Linn, Status of HEP Neural Net Research in the USA, Comp.Phys.Comm. 57 (1989) 297.

[14] B.Humpert, A Comparative Study of Neural Network Architectures, Comp.Phys.Comm. 58 (1990) 223.

[15] Silicon Neural Networks, in IEEE MICRO, 9,6 (Dec.1989)

[16] G.Darbo and B.Heck, The Contiguity Trigger for the DELPHI TPC, Nucl.Inst.Meth.A257 (1987) 567.

[17] L.Ristori, Nucl.Inst.Meth. A278 (1989) 436.

[18] D.Crosetto, Fast Digital Parallel Processing Module FDPP, CERN-DD/89-33 (AC) and S.Buona, D.Crosetto, Test Results of Real-Time Algorithms executed on FDPP with SPACAL data, CERN/ECP 90-06.

[19] R.Belusevic, G.Nixon, and D.Shaw, An 80 Mbyte/s Data Transfer and Processing System, RAL Report 90-028 (May 1990).

[20] Scintillating fibre calorimetry at the LHC, Proposal, R.Wigmans, spokesman, CERN - DRDC/90-23

[21] Integrated High-Rate Transition Radiation Detector and Tracking Chamber for the LHC, Proposal, B.Dolgoshein, spokesman, CERN - DRDC/90-38

[22] F.Hartjes et al., Recent tests with a gaseous microstrip chamber, ECFA Study Week Barcelona, Sept. 1989, CERN 89-10.

[23] Study of muon triggers and momentum reconstruction in a strong magnetic field for a muon detector at LHC,
Proposal, M.Della Negra and K.Eggert, spokesmen, CERN - DRDC / 90-36

[24] D.Pountain, Virtual Channels: The Next Generation of Transputers, Byte, April 1990, 'Europe and World' supplement, p.3.

[25] HPPI, High-Performance Parallel Interface, Specifications, ANSI X3T9.3/88-023 (February 1990)

Imprimé au CERN. Prix : 0.90 Fr.