

CERN LIBRARIES, GENEVA



SC00000253

CERN/DRDC 93-12  
RD-11 Status Report  
1 May 1993  
(EAST note 93-08)

SCP  
CERN-DRDC  
93-12

## Status Report: Embedded Architectures for Second-level Triggering (EAST)

J.Vermeulen (*NIKHEF-H Amsterdam*)

F.Constantin, A.Gheorghe

(*Institute of Atomic Physics and Polytechnic Institute, Bucharest*)

E.Denes, G.Odor (*Central Research Institute for Physics, Budapest*)

R.K.Bock (\*), J.Carter, W.Krischer, R.McLaren, I.Legrand (*CERN, Geneva*)

J.Renner-Hansen (*NBI Copenhagen*)

St.Fisher, R.Middleton, F.Wickens (*Rutherford Appleton Laboratory, Didcot*)

D.Belosludtsev, N.V.Gorbunov, V.Karjavin, S.V.Khabarov (*LHSE, JINR, Dubna*)

V.Dörsing, A.Reinsch, H.U.Zühlke

(*Institut für Angewandte und Technische Informatik, Universität Jena*)

K.Cetnar, Z.Hajduk, W.Iwanski, K.Korcyl, P.Malecki, A.Sobala

(*Institute of Nuclear Physics, Krakow*)

R.Nobrega, J.Varela (*LIP, Lisbon*)

B.J.Green, T.Medcalf, J.Strong, A.J.Wildish

(*Royal Holloway and Bedford New College, London*)

R.Hughes-Jones (*Manchester University*)

F.Klevenz, A.Kugel, R.Männer, K.H.Noffz, R.Zoz (*Institute of  
Computer Science V, University of Mannheim, Mannheim*)

J.Badier, Ph.Busson (*Ecole Polytechnique, Palaiseau*)

P.Bitzan, M.Novak (*Inst. of Computer and Information Science, Prague*)

L.Levinson, M.Sidi (*Weizmann Inst. of Science, Rehovot*)

Ch.Balke, J.Haveman, W.Lourens, A.Taal (*Utrecht University*)

U.Gensch, H.Leich, U.Schwendicke, P.Wegner

(*Institute of High Energy Physics, Zeuthen*)

(\* ) Spokesperson

### 1. Summary

The EAST Collaboration (RD-11) has been exploring one of the aspects of the critical problem of event selectivity in LHC detectors, second-level 'intelligent' triggering. The context, assumed initially, and now borne out by quite detailed physics simulations for the LHC proposals, foresees a level-1 event passing rate of the order of ~100 KHz, requiring a further reduction in level 2. Reduction 'algorithms' of some complexity will be required, using physics features extracted from fine-grain local data in multiple detector windows, combining them into global decisions by correlating data from different subdetectors. Implementation using commercially available components as much as possible is a primary goal.

The collaboration has used as methods both simulation studies and prototyping of hardware implementations, concentrating first on the feature extraction task, i.e. the local conversion of raw detector data into estimators meaningful for physics.

In reference to the milestones for **1992/93** as defined at the time of the last Status Report (DRDC 92-11), the collaboration has attained all major goals. The following results were obtained over the period of reporting (details in chapter 3):

- The *feature extraction benchmark* exercise was concluded by publishing results.
- New *improved benchmark data*, most importantly QCD background jets filtered by a first-level electron trigger, were made available for calorimeter and TRT algorithm studies by mid-1992, and gave increased confidence in our assumptions. The implementation of an algorithm implementing a track trigger based on Si detectors was also studied, in collaboration with ATLAS physicists.
  - A system integrating the SLATE *emulator and MaxVideo* via the 'HiMax' interface board, has been *demonstrated* in 1992, using a TRT trigger algorithm. Tests with the TRT prototype and RD-13's data acquisition system are now planned on the timescale of RD-6 (1993/4). A Router unit, arranging raw TRT data in suitable iconic form, has been designed and is *under construction*.
  - A TRT trigger based on a prototype of the *Enable machine*, is also *under construction*. Enable is a feature extraction alternative to the MaxVideo system, based on field-programmable gate arrays.
  - The collaboration between EAST and FERMI (RD-16) has resulted in detailed *simulations* of the data collection/feature extraction in a FERMI board, based on a DSP (TMS320C40). Prototype boards for the implementation are under design. A prototype board for *optical high-bandwidth links* from FERMI to FEAST has been built and shown to run at 1 GHz.
  - A *Decperle system* has been installed and is *under evaluation*. Porting the benchmark feature extraction algorithms to this system has already been done in part; performance and ease of use will be compared against other systolic solutions.
  - *Model algorithms* for the collection of feature data from multiple subdetectors and regions of interest (global trigger) have been defined, and corresponding *data sets* have been generated.
  - Various alternatives for implementing the global trigger have been selected, and *system simulations started*. Several components will be implemented in pilot projects, to measure the critical parameters of algorithm processing and communication times.
  - System specifications in a potentially commercial *object-oriented specification language* (VDM++) have started; a first report has been made to the AFRODITE project (Esprit) of the pipelined feature extraction problem as encountered in the TRT.

In **1993/94**, the third project year, the collaboration intends to (details in chapter 4):

- *demonstrate* two pipelined architectural candidates for *feature extraction* in a beam test of a *detector prototype* (RD-6/ATLAS);
- *demonstrate* a working *FEAST* prototype board and measure its *characteristics*;
- *evaluate* the suitability of the *Decperle* system for the implementation of feature extraction algorithms, most importantly under the criteria of programmability and cost;
- *remain in contact* with new and relevant developments in industry;
- *evaluate* competitively several overall (*global*) *second-level trigger architectures* by simulation, using standard test data and available tools; these architectures are characterised by general-purpose *computing nodes* and *switching networks*; both can be implemented in quite different ways;
  - *implement* some essential parts of global architectures in *demonstration hardware*, to provide better input for simulation and improve the evaluation procedure;
  - refine system *specifications* using VDM++
  - in general, move activities *closer to future LHC collaborations*, where embedding in an overall readout and data acquisition system, and hence system aspects like interface standardisation, error handling, maintainability, evolution, will be increasingly relevant.

## 2. Overall principles in the second-level trigger

The past two years of activity in EAST have permitted to decompose the problem of second-level triggering in some detail, under the constraints of detector design as it is being pursued in the LHC pre-proposal work. This structuring of the problem is vital to introduce the choices of transmission and processing technology, and of parallelism, based on which a final system will eventually be costed and chosen. Being part of the achievements of work in EAST, we find it useful to explain shortly this decomposition and the major implementation options.

### *Phase 1: Routing and feature extraction, front-end buffering*

Raw detector data are naturally collected close to the detector, in local non-overlapping modules, typically chips, boards, possibly crates. Nearly all detector electronics under design today intends to implement the storage over the latency period of a level-1 (L1) trigger, usually assumed to be of the order of 2  $\mu$ sec, on these modules. For detectors that contribute to the L1 trigger, transmission into the system, and possibly local aggregation for the trigger, is also done by the same modules.

It is one of our assumptions that the basis for any second-level (L2) trigger has to be laid by an algorithm that converts local data from a subdetector into variables containing the relevant physics message ('features'). This conversion ('feature extraction') can best be achieved by restricting the algorithm to work on local subsets of data only ('regions of interest'), as indicated by the results of L1. Thereby the bandwidth, critical for most detectors, can be kept low, the algorithms remain simple, and the natural parallelism of the detector can be put to use.

Locality of data requires that a device realized outside the L2 implementation indicates the whereabouts of 'regions of interest' (RoIs). It is another of our basic assumptions that all RoIs submitted to scrutiny in L2 are pointed out by what is essentially a L1-guided unit (a 'RoI-builder'), which 'drives' all L2 operations. This is true even for RoIs that do not themselves actively participate in the L1 decision, like lower-threshold calorimeter clusters.


The data pertaining to regions of interest (RoIs), of course, have then to be selected by some mechanism, which we term 'routing'. The Router function is strongly dependent on the readout design of a detector. Differences appear particularly when comparing the more advanced designs for storing data for those events that are accepted by L1. Substantial variety exists in the modularity of collecting data in buffers, where they are kept over the period of L2. In the extreme case of a clear 'push' architecture, data are collected into one, or very few central buffers. From there, they can be (selectively) transmitted to a feature extraction device, without substantial loss due to non-overlapping access to data stored in different buffers.

In the opposite extreme, data remain stored locally, in multiple buffers, and can be 'read' from there by the device implementing the feature extraction algorithm. In this case of a 'pull' architecture, the overlap problem (RoIs are non-congruent with readout modularity) necessitates a data collection from (parts of) several readout modules. Further, the large number of local buffers requires many potential connections to the feature extraction devices or, if these are also implemented locally, to the global L2 processors. A switching device of some generality, custom-made or commercial, becomes then a necessary element. If raw data are switched to feature extraction devices, bandwidth requirements for this device will be an important constraint.

### *Phase 2: RoI processing and event processing*

Once raw data have locally been converted to physics features, these have to be collected from all subdetectors and from all RoIs, for forming a global decision on the entire event. It is quite apparent, if one decomposes further the part of the algorithm

dealing with multiple sets of features, that the natural and efficient order of processing is to combine first all subdetectors that have 'seen' the same physics 'object', into a decision which is again local (same RoI), and follow this by combining all RoIs into an event decision. While this seems the most inviting algorithm strategy, it is by no means necessary (although possible and possibly cost-effective) to express the strategy also by a corresponding implementation in separate and parallel processors. If this is the adopted architecture, one may want to take advantage of another natural parallelism, that of testing multiple physics hypotheses against the same set of data. That data set, i.e. all features from all RoIs, amended possibly by quantities derived from them, must then be shared by all physics processes, and the entire decomposition can be presented graphically in the following figure.



### *Pipelined or fully programmable processor farms*

We have, so far, avoided discussing the hardware on which these algorithms or algorithm parts are implemented. In fact, basic choices exist that lead to quite different architectures. A properly decomposed algorithm making use of the detector- and physics-given parallelisms can be executed on multiple hardware units in a pipelined way (viz. in functional decomposition), and maintain the imposed overall decision frequencies. This approach is particularly indicated if algorithm execution times depend little or not at all on the data content. In this case, even fully synchronous operation is conceivable, or at least a data-driven approach can be used. If, on the other hand, data determine to an important extent the algorithm execution, then processors that keep up with overall frequencies of decision have to take 'worst-case' scenarios into account. Re-synchronizing and accounting for fluctuations without redundancy will then become costly. In this situation it is likely that a more efficient approach consists in spreading the execution of algorithms onto a farm of processors, which can be scheduled according to availability. This introduces the well-known event parallelism, and totally asynchronous transmission, opening the path to (or rather necessitating) general switching networks that perform best under these circumstances.

While work in EAST has shown that by proper choice of pipelined processors a decision frequency of 100 KHz can be maintained without introducing event parallelism in farms, this should not be interpreted as an early bias towards such solutions. Pipelined systems may arguably be a simple and cost-effective solution, because using best the existing parallelisms and minimising the number of processors. However, arguments like standardizing components, homogeneity, scalability, technological evolution, ease of programming, reconfiguration in case of malfunction of components, all are cost factors in their own right, and will have to be considered. The characteristics of farm-type solutions are, therefore, also under exploration in EAST.

## *Occupancy, thresholding (zero-suppression), iconic processing*

In the above, we have indiscriminately talked about 'subdetectors', where in reality we should have taken into account the very different characteristics of devices. Among the detectors called upon to contribute towards forming L2 trigger decisions, will be low-occupancy detectors like muon chambers or Si tracking chambers, and detectors operating, at high luminosity, at the limit of acceptable occupancy, like a transition radiation tracker.

A high-occupancy detector has little or nothing to gain from data compression, and may even blow up the volume of information by thresholding (zero suppression). Such a detector can transmit information as fixed-length records, where all cells are present and the position in the transmitted record indicates the cell position. Such records are then called 'images' made up of 'pixels' (individual cells), in the language of image processing. The processing can then take advantage of existing 'iconic' processing devices, which put to use the specific parallelism inherent in the order of transmission.

In the low-occupancy case, transmission of data must take advantage of the high compression factor gained by thresholding, i.e. only those chambers that have been hit by charged particles will transmit their address. Necessarily, transmitted records will be of variable size, and synchronous transmission as well as iconic processing of such data is not possible. Instead, 'symbolic' processing (from lists) will be indicated, although the use of image processors has been shown to be still possible.

We should note that calorimeters as foreseen in the LoIs seem to give occupancies that do not fall clearly into the 'low' or 'high' category. A decision whether or not to use data compaction for calorimeters has not been taken. Consequently, in transmission models we consider both possibilities.

## *Communication and switching*

It should be clear from what has been said above, that the implementation of L2 triggers is not dominated by the question of installing large computing capacity of a general type. Although that is also needed, we are more conditioned by

- the organisation of the detector readout
- buffering
- optimal use of parallelism
- point-to-point communication lines
- commercial switching devices.

The choice of restricting L2 algorithms to an operation in regions of interest indicated from outside, alleviates the bandwidth problem, and allows structuring the problem into various phases, as discussed above. Different characteristics of detectors, however, and different choices in technology, still leave open a variety of solutions. The most characteristic property of the overall L2 system is, probably, that a fully homogeneous and a minimal-cost approach to its implementation are in contradiction, and hence neither possible nor desired. A pipelined image processing approach is near-impossible for very-low-occupancy detectors, because of imposing constraints like synchronicity, at least locally. General-purpose processor farms for detectors that lend themselves well to iconic processing are not cost-effective, because they introduce fast switched communications and schedulers where they are not needed. The cheapest (as opposed to cost-effective) system, on the other hand, will be one that does not leave open the avenues of evolution in detectors, physics demands, and technology.

As will be discussed in detail below, EAST has now started exploring the major choices in some detail, using foreseeable technological possibilities. The intended result is a catalog of remaining (along with some discarded) options, that can be costed in detail at the time when LHC experiments will have to make definitive choices.

### 3. Activities and milestones in 1992/3

During the period of reporting, work in EAST has started to polarize more clearly into two categories, which we could broadly characterize as generic R&D on the one, and experiment-related work in close collaboration with other R&D activities or proto-collaborations, on the other hand. The following report resumes point by point the milestones laid out in our previous status report, as confirmed by the DRDC (in its minutes DRDC 92-36).

#### 3.1 Completion of the benchmark exercise

The EAST benchmarks consisted of mapping defined feature extraction algorithms (tracking as in RD-6, calorimetry as in RD-1) onto various candidate architectures, and to study the possibilities and limitations of possible implementations, including cost. This work has been completed and published [1]. The number of candidate architectures for subsequent further study and for pilot implementations has been reduced substantially by the results. Lack of suitable interfacing possibilities, high cost, general unsuitability, or delayed availability, are among the causes for deciding to put some candidate architectures on hold for at least the immediate future.

Retained as possible were the pipelined architectures (**MaxVideo**, **Enable**) as potential low-cost candidates, and the fully commercial MPP elements of **iWarp** as potentially simplest homogeneous solution. Practical **implementations** of feature extraction architectures are now firmly planned on MaxVideo and Enable.

Under continuing investigation are **FEAST**, an array of DSPs with geographically fixed assignments, and **Decperle**, a board of multiple Xilinx chips with fast communication lines and buffers, allowing to implement, partly in a high-level language, multiple pipelined algorithms on identical hardware. FEAST is under development in the collaboration using as DSP the Texas Instruments TMS320C40. Decperle is a prototype series developed by Digital.

All ongoing projects are described in more detail in chapter 4 below.

#### 3.2 Improved benchmark data

The EAST benchmarks, now published, were initiated at a time when trigger-oriented Monte Carlo activities in LHC proto-collaborations had not started, and detector parameters could only be 'educated guesses'. EAST with its limited resources could not bridge this gap, and our benchmark data, while requiring relevant feature extraction algorithms, were not meant to reflect in detail the situation of a second-level trigger, which is dominated by fluctuations in physics background. This situation has meanwhile converged towards a more coherent approach. EAST has, in fact, actively participated in the Monte Carlo calculations of physics background events (QCD jets) that pass a 'reasonable' first-level trigger criterion in the calorimeter. Although algorithms were re-optimised with the new data, the results gave us increased confidence in the original algorithms used in our benchmarking exercise, and verified that corresponding triggers do indeed perform a valuable function in future experiments. The overall credibility in the work performed in EAST for future LHC collaborations thus has gained substantially. It is, however, not intended to repeat the feature extraction benchmarks with improved data and/or algorithms.

#### 3.3 Hardware implementations and demonstrations

Demonstrations of implementations of second-level trigger algorithms are possible either with SLATE (the EAST-built hardware emulator) or with detectors in beam test setups. SLATE has been presented in our previous status report, and has now been produced in a small commercial series by Transelektro Ltd., Budapest. Its HIPPI format

output stream has been interfaced to the data memories contained in the MaxVideo image processing system. The interface VME module (called HiMax) provides 50Mbytes/sec input rate. Using HiMax, algorithms already running in MaxVideo can be made to run at LHC level-2 speed, if detectors transmit in HIPPI format.

For tests with detector modules, we have found in RD-6 a prototype which will run also as a technical test for LHC-like readout electronics. We are, therefore, participating actively in preparing a **trigger test** as part of the TRT beam tests, foreseen for October 1993. The data acquisition in that test is planned to run on a system based on the pilot systems of RD-13.

### *3.4 The Enable Machine*

The Enable machine [2] provides an alternate implementation of the TRT algorithm. The device is a flexible special-purpose processor, custom-designed for the needs of the TRT feature extraction algorithm. More precisely, it is a SIMD ('single-instruction-multiple-data') machine consisting of an array of identical processing elements realized with field programmable gate arrays. In order to identify electron and/or pion tracks in binary images of limited size (regions of interest), the Enable machine is composed of two main functional building blocks: a histogram generation unit and a trigger decision unit. Like in the MaxVideo-based solution, the high and low pulse height images are sent through two separate but identical histogram generation units. The subsequent trigger decision unit analyzes the content of the histogram channels and, by comparing the bin content from the two images, achieves 'electron'/pion/'no high  $p_t$  track' classification in hardware.

Enable has been fully designed and is now under construction, equipped with a HIPPI serial-link interface to maintain the dataflow bandwidth from the detector/Router. Interfaces for Enable and Maxvideo are developed jointly and by the same institutes, to minimize the investment in specific and temporary demonstration devices.

Enable is an integral part of EAST's demonstration program around the TRT prototype planned in RD-6.

### *3.5 FEAST: geographically fixed processor/detector assignments*

In the choice of the second-level trigger processor, EAST has most thoroughly explored pipelined processor systems which work on events in monotonous sequence, keeping up with the expected decision frequency of 100 KHz. This is valid under the hypothesis that the frontend electronics 'pushes' the data into the acquisition system after first level selection, and RoIs can be selected by a special Router unit.

The FERMI project (RD-16) is developing fully digital front end electronics for calorimetric detectors. The chip now designed will include high speed analog-to-digital converters, a programmable pipeline/digital filter chain, and local buffering over the period of first-level triggering. Each FERMI board will contain 36 chips and cover a detector region of 324 channels, corresponding to a contiguous region in space.

The aims of the FERMI/EAST (FEAST) sub-project is to link the FERMI boards to feature extraction and level-2 buffering boards with the same fixed geographical assignment. This implies that the collection of data for a RoI does not need a Router; instead, data for a RoI are assembled through nearest-neighbor links. Local feature extraction processors in this architecture do not have to respond to a 100 KHz requirement, because successive events will occur at much lower rate, locally. The model also implies that many more feature extraction processors can be potentially active, and have to be connected to the global decision processors.

In the period of reporting, feasibility studies for implementations have been carried out, and first implementations of feature extraction algorithms have been made on DSPs.

The FERMI board has been simulated in VHDL, and the communication requirements for the FEAST board explored.

### 3.6 Decperle

The Enable machine, described above, is a custom implementation of the TRT trigger algorithm, in field-programmable gate arrays (FPGA-s). Although highly performant and cost-effective, a custom-made board suitable only for a given subdetector poses potential problems of long-term support and maintenance, in a future LHC experiment.

Digital had announced, in 1991, a Research Initiative around a generalized concept based on FPGA-s (Xilinx), in this context called 'Programmable Active Memories' (PAM-s). In the 'Decperle' system, multiple (presently 16) FPGA-s are surrounded by additional memories, switching elements, and four high-bandwidth I/O ports that can be interfaced to fast point-to-point links like HIPPI. One of these links connects directly, via Turbochannel, into a high-level processor (Decstation 5000). The fact that the PAM-s possess ample and fast memory coupled to active switching elements, qualifies them for different systolic or data-driven feature extraction tasks. High-level software and advanced debugging tools exist to speed up development at the gate level, and to make system changes more transparent.

A Decperle system has been installed at CERN in January 1993, as part of a joint project Digital/EAST. Several key algorithms are presently being implemented on this device, first results look very encouraging.

### 3.7 High-level decision network

The systolic or data-driven feature extraction devices (MaxVideo, Enable, Decperle) together with a Router, or the FEAST system (combined Router/feature extraction) operate on local (RoI) data from a single subdetector. As outlined in chapter 2 above, they have the function of converting raw detector data to physics variables, from which a trigger can be constructed. In doing so, they reduce the packet size and thus bandwidth requirements for further communication.

Decisions about accepting/rejecting full events ('global decisions') will have to be based on the output of multiple low-level feature extraction systems. RoI-wise, phenomena (electrons, muons, jets, hadrons...) have to be classified using information from *all subdetectors*; subsequently, three-dimensional *geometrical correlation* of physics objects (e.g. effective masses) will be needed for taking an overall decision. More flexibility and scalability to adapt to evolving physics and detectors, random access to feature data, and higher precision, possibly floating point, are needed. Processors with full high-level programmability and floating point hardware seem unavoidable, and switching devices with high bandwidth will have to channel data from the feature extraction devices to the global decision processors.

To study possible connection topologies, i.e. processors and connection networks, has, therefore, become an important goal of EAST. In order to approach systematically the understanding of the interplay and functioning of various candidates for processing and switching, we have defined, in a series of EAST meetings and notes, characteristic global decision algorithms, corresponding test data, and a variety of hardware models on which these global decisions could be implemented. These models are being studied, on a behavioral level, in detailed simulation. Our goal is to achieve a comparison in performance and cost, as input to future decisions in LHC collaborations. The extreme readout and buffering models of an event-parallel farm behind a general switching device for features, and of a systolic decision array built from various processors (but with a fixed data route) are both part of the spectrum. The initial model components considered for simulation have been defined: general-purpose processors (TMS320C40, Alpha, i860, Transputers) and switching devices (ATM, SCI, Transputer switch).



Some of these model components will be implemented in hardware, in order to obtain realistic parameters for modelling. Preparatory work for these implementations has started.

### *3.8 Fibre optics high-bandwidth connections*

High-bandwidth point-to-point connections are a critical part of the parallel data transmission close to the detector, in our terminology up to to the feature extraction device. The HIPPI standard has been introduced in our work, and in RD-6 as a useful interim standard for low-level connectivity, and seems well accepted by industry. It is the transmission standard used for the RD-6 tests later in 1993.

This standard has, however, limitations already underlined in our last DRDC status report. For a final application in LHC, hundreds of expensive HIPPI cables would be required, taking up a large volume. Cost, physical size, and cable length limitations could be a problem in future experiments.

A conversion of the 32 parallel electrical signals of HIPPI to a serial optical transmission has therefore been implemented, based on Hewlett-Packard's GLogo chip, and has been successfully demonstrated on a prototype board ('serial HIPPI'). The sender and receiver boards together make a connection appear to a user entirely like HIPPI, but avoid the large cable volume and the distance limitation: fibres transmit without repeater up to 1 km, as opposed to about 25 m for a HIPPI cable.

### *3.9 Formal system modelling using VDM++*

Being 'application partners' in the Esprit II project AFRODITE, we have invested initial work on the practical application of tools for formal specification of mixed hardware/software systems, in the context of our level-2 trigger simulations. The tool to be used in this project is an object-oriented specification language called VDM++ [3], based on the standard VDM. Members of EAST have undergone training courses in VDM and related toolsets, and have participated in discussing VDM++ specifications. The precise objective of the EAST application is the specification and modelling of a (possibly partial) acquisition system, with at least the second-level trigger steps included, requiring fault-tolerance, maintainability, and reusability of components.

We have produced, for the Esprit project, an initial report describing pipelined feature extraction in the language VDM, thus initiating a transfer of new technologies from professionals in software/hardware specifications to the high-energy physics community.

## **4. Workplan for 1993/4**

In the following, we distinguish between developments and implementations that are naturally done in the context of preparing for the LHC technical proposals of 1993/94, and developments or studies of a more generic nature. We believe such generic work is not only required to complete what we take to be our DRDC mandate of fully understanding the problems of triggers at 100 KHz: for reasons of stretching the marginal resources, some studies should be pursued jointly by future collaborations even beyond acceptance of the LHC and detector program.

## 4.1 LHC-oriented

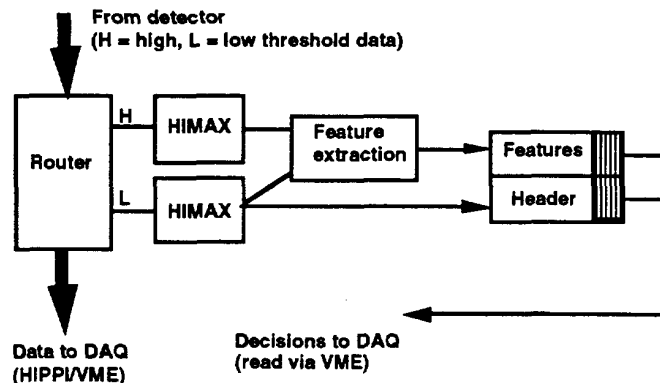
### *Beam test demonstrations of Router, MaxVideo, Enable*

Project RD-6 plans for 1993/94 tests of a TRT prototype detector equipped with LHC-like electronics, and using a pilot system under development in RD-13 for data acquisition. EAST will be participating in these tests, building up a demonstration of triggering at LHC speed with most of the LHC requirements implemented.

It is one of our key assumptions that only RoI data are sent to the trigger unit. In the 'push' architecture underlying the TRT readout model, this data selection can be done by intercepting the synchronous data stream between the detector readout electronics and the data acquisition system, in a **Router** unit. This unit implements the selection of a 'region-of-interest' (RoI) under control of a pointer. The RoI pointer is provided in the experiment by the first-level trigger, in the test setup by a register set under program control to test the corresponding Router function. The Router for the TRT prototype takes advantage of the synchronous readout, and achieves its function of data selection, with some preprocessing, using a short-term memory and a sequencer unit. It has fully standard HIPPI input and output protocols, and thus is entirely transparent to the data acquisition of the TRT, except for a constant added latency. The unit is now under construction.

For the purpose of algorithm tests with the RD-6 prototype, the already demonstrated algorithm in **MaxVideo** will be run on-line. The HIPPI interface (**HiMax**) has been prototyped and demonstrated [4], and is presently being rebuilt in a more general version for MaxVideo20. As the data acquisition will not be able to run at the speed of the trigger, all decisions will be kept and written onto tape. Matching by event number, the decisions taken for the fraction of events which has been written onto tape, can then be verified off-line.

As an alternative to MaxVideo, the **Enable** machine will also be tested on the same timescale of the RD-6 prototype beam tests.



**Milestones:** produce and test all units (Router, HiMax, MaxVideo, Enable, DAQ interface) so that they are available for beam testing in October 1993 and spring 1994, and demonstrate functioning at full speed.

### *Continued algorithm work*

Although reasonably understood now, algorithms will have to be continually adapted to the evolving definitions of the detector, adding new subdetectors as physics

requirements arise. More work will have to be invested to understand trigger algorithms and implementations for the barrel TRT of ATLAS, and for discrete trackers (inner detectors, muon chambers). In particular, the impact on physics of possible algorithm simplifications, to save on implementation cost, will have to be carefully evaluated before final decisions are taken.

In the area of algorithm development, EAST intends to become an integral part in preparing technical proposals for the LHCC, and in the optimisation of detectors, algorithms, and implementations. No specific milestones can be given.

## ***4.2 Generic work, continued***

### ***FEAST***

A complete calorimeter second-level trigger subsystem based on the FEAST idea is now approaching the stage of pilot design. The present ideas for implementation are as follows: A FERMI board contains 36 FERMI chips, each of which is connected to 9 cells. Optical links provide high-bandwidth communication from the FERMI board to a FEAST board. This board has local links to adjacent FEAST boards mediated by a DSP, and contains the local feature extraction processor. Features are then transmitted, via a switching network, to the global decision processor farm of level 2. FEAST boards also contain a data buffer, and, for second-level retained events, communicate with a different switching network implementing the event builder of the data acquisition system.

For a pilot implementation of the FEAST board we envisage to use boards from LSI (Reading, UK), containing multiple TIM modules (replaceable daughter boards). We intend to individually modify some TIM modules for our purpose. All communication (FERMI, global L2, DAQ) and feature extraction tasks are done by TMS320C40 processors from Texas Instruments. The switching technology is open, but we explore the possibility to use either the C104 transputer switch (in collaboration with the GP-MIMD project), or ATM (in collaboration with RD-31), possibly also SCI (with RD-24).

**Milestones 1993/94:** *specification of prototype based on LSI boards: July 1993; test environment prepared: December 1993; board ready for testing: February 1994.*

### ***Decperle***

A Decperle-1 system has been installed in January 1993 at CERN. The Paris Research Laboratory of Digital as producer of Decperle, Mannheim, and CERN have established a detailed workplan for this device. A port of the Enable design to Decperle and an implementation of a calorimeter trigger algorithm are under way.

Successful implementations of significant algorithms on Decperle could have a major significance: Decperle boards are presently interfaced to a standard Decstation 5000 host via a high-bandwidth Turbochannel interface. The host could be replaced by an Alpha processor. If several Decperle boards are connected to the same general-purpose node, a natural multi-detector feature extraction system for a single region of interest is obtained, which is comparatively easy to develop and adapt. If beyond that we imagine a low-latency network linking multiple nodes (see below), a homogeneous and fully scaleable overall L2 structure results, made of commercial components only. Arguably, such a system might correspond most closely to the problem structure, and hence have a chance of being among the lowest-cost commercial and homogeneous candidates.

**Milestones 1993/94:** *Calorimeter algorithm installed: May 1993; Enable design ported: September 1993; if successful, simulation of full Decperle/Alpha system: December 1993.*

### *Alpha/SCI*

In collaboration with the Digital Joint Project and with RD-24, we have started defining the implementation of a prototype platform with high computing capacity and high-bandwidth interconnection possibilities via the SCI interface. The system will be based on a Futurebus+-based board containing Digital's Alpha AXP processor, memory, a 2 MByte cache, as under development at DEC Galway. The board also contains a PCI port, to which a custom-made 'mezzanine' (daughter board) can be interfaced. The development concerns the design and construction of a mezzanine containing the SCI interface, and the development of CHORUS-based software for the board, which implies specific drivers and access to low-latency data sharing via the CHORUS microkernel. A group of Alpha/SCI nodes arranged in ringlets will constitute a prime choice for the global decision network of a level-2 trigger. If completed by a high-bandwidth turbochannel or PCI interfaces for Decperle (see previous paragraph), the nodes could be fed directly from the feature extraction device, and hence from the detector, without the need of further interfacing development.

**Milestones 1993/94:** *specification of the mezzanine: June 93; design of the mezzanine: Sept.93; test software for PCI-SCI: Sept.93; port of CHORUS to AXP: Dec 93 (external dependency); mezzanine integrated with Alpha AXP: May 94; SCI driver software: June 94; integrated field tests and demonstration: Dec 94.*

### *Global trigger, comparative evaluation by simulations*

The existing decision algorithms, and data sets generated for six different physics channels, plus QCD background, are already being used for extensive simulations. First results exist on algorithm execution times and switching latencies. The compute nodes under consideration are Texas Instrument TMS320C40 DSPs, Digital Alpha chips, Intel i860 processors, and Inmos T9000 transputers. They connect in various ways among themselves and to systolic or geographically fixed feature extraction processors, using their own fast links (C40 and T9000) or networks or switches expected to become available commercially (ATM, SCI, C104).

**Milestones 1993/94:** *simulation results to be compared in workshops in May and fall 1993, final results for publication early 1994.*

### *Other architectural possibilities*

Several critical architectural components being demonstrated, developed, or simulated, must not make us forget that technological evolution may, or even is likely to, change our conclusions with the passing of time. It is, therefore, our intention to remain in contact with new and relevant developments in industry, to the possible extent. We will continue to follow the development of the Davis chip (ITT) for image processing, which has been used as a (simulated) entry in our benchmarks. We also have close contacts with the design of the L-Neuro chip (Philips), which has properties that make it

a likely feature extraction candidate. The applicability of other commercial neural network devices like those from Adaptive Solutions, Hecht-Nielsen, or Intel, will also be looked at. We further maintain contact with the development of a second-level trigger prototype unit based on the Blitzzen SIMD architecture, undertaken by a collaboration of INFN Padova and the University of North Carolina. The CS-2 (following paragraph) is another architectural entry we want to explore.

### *CS-2 MPP machine*

As was already indicated in the discussion of EAST's feature extraction benchmarks [1], the excellent result achieved by the iWarp architecture has kept general-purpose computer systems of the massively parallel type in focus as potential candidates for implementing our challenging real-time algorithms in a fully programmable, commercial and homogeneous way. iWarp performed well because its fine-grain communication is directly accessible to the user, bypassing the system kernel and offering very low latency.

The CS-2 system will be brought to CERN as part of an Esprit project (GP-MIMD2) in the course of 1993. Compared to other MPP candidates, the characteristics (in particular the communication latency) of the CS-2 machine (made by the PCI consortium) make this system a suitable basis for attempting an implementation of such algorithms. This has been found in RD-11, by contacts with Telmat, before the GPMIMD-2 project was approved. If a fully commercial systems of this type provides the necessary performance, it is likely to be a serious candidate for implementing complete LHC-type trigger systems (second- or third-level) due to its homogeneity. Special processors, custom-designed interfaces, switches, connection mechanisms would all be obviated, although it is unlikely that the Router function would also be taken over by this system.

A thorough investigation of this possibility is, therefore, of utmost importance. It can also be justified in comparison to computer centre test cases for MPP machines: there is no other HEP-specific application that can potentially demonstrate the success and necessity of parallelism in a similar way. Present staffing at CERN, however, does not make it obvious how this can be achieved. Two manyears are estimated necessary to port and optimise algorithms for parallelized feature extraction, and for connecting multiple detectors and regions in a global decision MPP structure.

**Milestones 1993/94:** *no workplan can be made before the allocation of manpower*

### *Optical links*

More work will be done with industry to get familiar with the two types of digital optical link, serial and parallel. For the serial links we will continue to follow and implement ANSI standards. These standards will include Serial HIPPI and Fiberchannel.

Within EAST there are already many HIPPI components, and HIPPI links are also being used in RD-6 and RD-13 as well as in NA48. They may be replaced by cheap Serial HIPPI optical links in the future. Some Fiberchannel components also exist and are being used in optical links for NA48. Fiberchannel is a very relevant candidate for transmission and switching, and we must build up expertise in this standard.

For the parallel digital optical links, we will work with industry to obtain samples of arrays of transmitters, receivers, multi-fibre ribbon and connectors. We have to understand how these devices can be reliably introduced in future experiments.

**Milestones 1993/1994:** *Complete the Serial HIPPI prototype boards and test over optical fibers. Commercialise the development if appropriate. Follow Fiberchannel developments and start a small Fiberchannel test system. Test Hitachi arrays of parallel optical transmitters and receivers and measure characteristics of the MT ribbon fibre connector.*

### *Formal system modelling using VDM++*

The Esprit II project AFRODITE has as goal to develop a formal object-oriented specification language (called VDM++) and an associated toolset. The language will have constructs of concurrency (parallelism) and will allow to formulate real-time systems with hardware and software components. A semi-automated path will allow to translate specifications written in VDM++ into simulation code or software (C++), or into hardware definitions (VHDL). A major software house (Cap Gemini) is the leading partner, another software house (IFAD) and several universities share the burden of development, which starts from an existing set of VDM implementations (lacking the object-oriented quality, concurrency, and real time).

On the EAST side, Utrecht and CERN are actively involved, using level-2 trigger systems as application examples. Two more applications are partners in the project: Lloyd's Register studies a safety-critical application in ship loading, the Defense Research Agency (UK) intends to use the methodology for designing VLSI systems. The possibility to reason about the behaviour of a design, and thus to notice deficiencies early in development life cycle, should result in substantially increased quality of the designs and of the derived implementation. Complexity should be mastered better by having re-useable system modules, and contracting system parts out to industry will be made possible by using an agreed specification language. Independent of the success of the specific language and toolset, high-energy physics will be able to learn much in the area of CASE methodologies and tools, by collaborating closely with a group of professionals and experienced application partners.

**Milestones 1993/94:** *the AFRODITE workplan foresees a translation of L2 system specifications into VDM (already partially done), and into VDM++ by Sept.1993.*

## **5. Composition of the collaboration, institute responsibilities, resources**

Since the last Status Report, the collaboration has been joined by the new members Jena, Copenhagen, Lisbon, RAL, and Manchester.

The individual responsibilities of all institutes are as follows (in order of the workplan in chapter 4):

The *beam test demonstrations* are being prepared jointly by Mannheim (Enable machine), CERN and Bucharest (MaxVideo), Jena (Router), Dubna (HIPPI interfaces for Router and Enable), and Weizmann/Krakow (HIPPI interface for Maxvideo and buffer interface to the data acquisition system). Budapest is responsible for the software environment in SLATE, which is a key element in most tests.

*Algorithm work* continues at Krakow, CERN, Bucharest, and Lisbon, for optimising feature extraction algorithms of the TRT, and for defining significant global decision algorithms, including Neural Network - inspired algorithms.

The *prototype FEAST board* is being developed by RHBNC (motherboard and TIM board for local mesh connection and feature extraction), NBI Copenhagen (TIM board for data acquisition and global level-2 connections), Krakow (FEAST/FERMI board simulation and design), CERN (optical connections). Additional partners are University College London (C40 software environment) and Stockholm (overall FEAST simulation in VHDL).

Partners in the development of algorithms on the *Decperle architecture* are CERN and Mannheim, with very active help from the Digital Joint Project and the Paris Research Laboratory of Digital.

The *prototype architecture based on Alpha/SCI* is being designed jointly by RAL and Manchester University, together with industrial partners (among them the Digital Joint Project). An Esprit proposal has been submitted in April by this group. Software will be developed at RAL. CERN will provide a test environment, significant test algorithms and corresponding data.

Implementations of artificial neural networks are being studied by Amsterdam, Utrecht, Ecole Polytechnique, and Prague. Ecole Polytechnique continues to develop jointly with Philips the *L-neuro VLSI implementation* of a generalised Neural Network; this chip will then be applied for realising feature extraction algorithms, e.g. for calorimetry and/or tracking.

Simulations for the *global trigger architecture* are being done in a decentralised way at Amsterdam, Prague, CERN, Copenhagen, Lisbon, RAL, and Utrecht, with help from RHBNC, Zeuthen, the GPMIMD project at CERN, and the RD-24 and RD-31 projects.

Formal system *specifications and modelling using VDM++* is under the responsibility of CERN and Utrecht.

**Resources:** For 1993/4, despite an unchanged level of activity and envisaged hardware demonstrations, we foresee that spending will be somewhat lower than in 1992/1993. Multiple hardware investments have been made in the past year, and are not necessary at the same level in the coming year. We estimate this year's total spending at 750 KSf, of which 250 KSf for CERN. The overall EAST manpower involvement remains around 40 full-time equivalent.

We will continue to perform beam tests only in conjunction with other R&D activities (RD-6), and hence do not require a beam time allocation of our own.

We have followed, in the past, the policy of participating in the CERN computing budget of the large users (collaborations like ATLAS) for major Monte Carlo calculations. This situation will not change. For some flexibility and for general access (visitors!), we require our own budget at the level of 300 hours CERN.

## 8. Acknowledgements

We gratefully acknowledge the participation of individuals and institutes that are not formally part of the EAST collaboration. We mention the excellent collaboration from other R&D projects (RD-6, RD-13, RD-16, RD-24, RD-31) and from activities in ATLAS. The University of Oslo, the Technical University of Tampere, and the Federal University of Rio de Janeiro have participated in some aspects of our work (and continue to do so). We acknowledge excellent contacts with the development of a Blitzen-based trigger prototype in Padova. We are deeply indebted to the DEC Joint Project and the Paris Research Laboratory of Digital, for their contributions to the DecPerle investigation. G.Klyuchnikov from IHEP Protvino, supported by the DEC Joint Project, has made major contributions to physics simulation and the introduction of the Decperle system. We further acknowledge the collaboration with Cap Gemini and other partners of the Esprit project AFRODITE, and the Laboratoire d'Electronique Philips.

## 9. List of publications and internal notes generated in connection with EAST activities

### 9.1 Publications and conference proceedings

- [1] J.Badier et al., *Evaluating Parallel Architectures for Two Real-Time Applications with 100 KHz Repetition Rate*. IEEE Transactions on Nucl.Sc. 40/1 (1993) 45
  - [2] F.Klevenz, R.Zoz, K.-H.Noffz, R.Männer: *The ENABLE Machine - A Systolic Second Level Trigger Processor for Track Finding*; Proc. Comp. in High Energy Physics, Annecy; CERN Report 92-07 (1992) 799.
  - [3] E.Durr, W.Lourens, and J.van Katwijk: *A Formal Specification Language For Object-Oriented Designs*. Proceedings of the Second International Workshop on Software Engineering, Artificial Intelligence and Expert Systems in High Energy and Nuclear Physics, La Londe-les-Maures, 13-18 January (1992) 47.
  - [4] L.J.Levinson, M.Sidi, Y.Damatov: *A simple HiPPI destination interface for second level trigger prototypes*. Proc. Comp. in High Energy Physics, Annecy; CERN Report 92-07 (1992) 904.
- R.Männer, J.Gläss, F.Klevenz: *Massively Parallel Systolic Processors for High-Speed Recognition of Simple Patterns*. Mirenkov N.N.(ed.): Parallel Computing Technologies, Singapore (1991) 98.
- S.A.J.Keibek, G.T.Barkema, H.M.A.Andree, M.H.F.Savenije and A.Taal: *A Fast Partitioning Algorithm and a Comparison of Binary Feedforward Neural networks*". Europhys. Lett., 18 (1992) 555.
- R.Baur, J.Gläss, F.Klevenz, Ma Long, R.Männer: *The Development of Systolic Processor Arrays for Pattern Recognition, Image, and Signal Processing at the Universities of Heidelberg and Mannheim - A Status Report*; in: Makhaniok M., R.Männer (Eds.): High-Performance Parallel Architecture Design; Inst. of Engineering Cybernetics, Acad. Sci. of Belarus, Minsk (1992) 149.
- R.Baur, J.Gläss, F.Klevenz, R.Männer, R.Zoz: *Systolic Processors as Second Level Triggers for High Energy Physics Experiments*; accepted for publ. in Proc. Workshop on Image Proc. for Future High Energy Physics Detectors, Erice (1992)
- H.M.A.Andree, G.T.Barkema, W.Lourens, A.Taal and J.C.Vermeulen *A Comparison Study of Binary Feedforward Neural Networks and Digital Circuits*". Proceedings of the Second International Workshop on Software Engineering, Artificial Intelligence and Expert Systems in High Energy and Nuclear Physics, La Londe-les-Maures, 13-18 January (1992) 347.
- J.C.Vermeulen, W.Lourens, A.Taal and L.W.Wiggers: *Implementation of a Second-Level Neural Network Trigger*. Proceedings of the Second International Workshop on Software Engineering,



Artificial Intelligence and Expert Systems in High Energy and Nuclear Physics, La Londe-les-Maures, 13-18 January (1992) 387.

R.Männer: *Systolic Processors in High Energy Physics*; accepted for publ. in Proc. WOPLOT 92, Tutzing (1992)

S.Centro et al., *Results of Second Level Trigger Algorithms Using the Blitzen Parallel Machine*. Proc. Comp. in High Energy Physics, Annecy; CERN Report 92-07 (1992) 274

B.Green, J.Strong, T.Anguelov, R.McLaren, E.Denes: *A test system for second-level trigger and data acquisition architectures*. Proc. Comp. in High Energy Physics, Annecy; CERN Report 92-07 (1992) 701

Adrian Gheorghe (EAST collaboration): *Highly parallel signal processing architectures for second-level trigger applications*. Proc. Comp. in High Energy Physics, Annecy; CERN Report 92-07 (1992) 239

H.M.A.Andree, G.T.Barkema, A.J.Borgers, M.Kolstein, W.Lourens, A.Taal, J.C.Vermeulen and L.W.Wiggers: *Feedforward neural networks for second-level triggering on calorimeter patterns*. Proc. Comp. in High Energy Physics, Annecy; CERN Report 92-07 (1992) 654.

Pavel Bitzan and Michal Novak: *Optical Neural Networks for Second-Level Triggering*. Proc. Comp. in High Energy Physics, Proc. Comp. in High Energy Physics, Annecy; CERN Report 92-07 (1992) 901.

F.Klefenz, K.-H.Noffz, R.Zoz, R.Männer: *ENABLE - A Systolic 2nd Level Trigger Processor for Track Finding and e/p Discrimination for ATLAS/LHC*; submitted to IEEE Nucl. Sci. Symp., San Francisco, CA (1993)

H.M.A.Andree, G.T.Barkema, W.Lourens, A.Taal and J.C.Vermeulen: *A comparison study of binary feedforward neural networks and digital circuits*". Accepted for publication by Neural Networks (1993).

## 9.2 EAST notes (since March 1992)

92-04: *The MasPar Computer as second-level trigger architecture for calorimeter windows* (R.K.Bock, H.Figel, I.Legrand) 18 March 92

92-05: *Status Report to the DRDC* (CERN/DRDC 92-11) 3 March 92

92-06: *Second-level trigger: global decision structures* (R.K.Bock, J.Carter, E.Denes, I.C.Legrand, M.Novak, J.Varela) 5 November 92

92-07: *Summary of the AFRODITE Project* (R.K.Bock, E.Dürt, W.Lourens) 27 March 92

92-08: *A Systolic Cluster Identification Algorithm implemented on the iWarp System* (G.Greer, I.C.Legrand, J.Nelson) 4 May 92

92-09: *The Router unit and its implementation* (also RD-6 note #27) Rev.1 (R.K.Bock, V.Dörsing, N.Gorbunov, V.Karjavin, A.Reinsch) 1 December 92

92-10: *Benchmarking with Data from the Transition Radiation Detector - Implementation on MasPar Computer* (A.Sobala) 5 May 92

92-11: *A Parallel Algorithm for Feature Extraction from Transition Radiation Detector Data: Benchmark Results using the Blitzen Parallel Processor* (also DFPD 92/EI/28) (S.Centro, E.Davis, Ping Ni, D.Pascoli) May 92

92-12: *The ASP Benchmarks for second-level Trigger* (A.Thielmann, GEVA Ltd. Warsaw) 12 May 92

92-13: *Systolic histogramming technique for track finding on the iWarp system* (R.K.Bock, B.Greer, I.C.Legrand, J.Nelson) 8 May 92

- 92-14: *ASP algorithm for second level TRD triggering*  
(G.Vesztergombi and G.Odor) 10 May 92
- 92-15: *Benchmark Results for the Enable Machine*  
(F.Klevenz , R.Männer, K.H.Noffz, R.Zoz) 12 May 92
- 92-16: *Minutes of EAST Collaboration Meeting #6, and of Benchmark Results Workshop* (R.K.Bock) 15 May 92
- 92-17: *Second-level trigger algorithms on pipelined signal processing architectures* (A.Gheorghe, W.Krischer, Z.Natkaniec) 20 May 1992
- 92-18: *Evaluating Parallel Architectures for two Real-time Applications with 100 KHz Repetition Rate* (J.Vermeulen et al.) 1 June 1992
- 92-19: *Calorimeter Benchmark Algorithm for ASP*  
(G.Odor and G.Vesztergombi) 25 June 92
- 92-20: *SLATE program, Version 1.2, User Manual*  
(E.Denes) 25 September 1992
- 92-21: *The farm approach to second level triggering (draft)*  
(V.Buzuloiu) 25 February 1992
- 92-22: *A simple HIPPI destination interface for second-level trigger prototypes.* (L.J.Levison, M.Sidi, Y.Damatov) September 1992  
(Also in CERN 92-07, Proceedings of CHEP '92, Annecy)
- 92-23: *A second-level trigger, based on calorimetry only*  
(G.Klyuchnikov et al.) 8 October 1992
- 92-24: *EAST Collaboration meeting #7 and Workshop on Global L2 Trigger, Minutes* (R.K.Bock) 1 October 1992
- 92-25: *Level-2 triggering in ZEUS and implications for EAST*  
(J.Vermeulen) 12 October 1992
- 92-26: *Feasibility study on using the TMS320C40 DSP in implementing the LHC Level-two triggering system*  
(K.Korcyl, Z.Hajduk, J.Strong, T.Bharucha) September 1992
- 92-27: *Simulations of a calorimeter readout system using VHDL*  
(W.Iwanski, G.Applequist) November 1992. Also FERMI note #13
- 92-28: *Pattern recognition algorithms for triggering with a silicon tracker*  
(A. Gheorghe, W. Krischer) 19 November 1992
- 93-01: *Test data for the global second-level trigger*  
(R.K.Bock, J.Carter, I.C.Legrand, J.Varela) 28 January 1993
- 93-02: *EAST collaboration meeting #8, Minutes*  
(R.K.Bock) 9 February 1993
- 93-03: *Modelling of L2 global decision structures, Revision 0*  
(R.K.Bock, J.Carter, I.C.Legrand, M.Novak) 11 February 1993
- 93-04: *HIMAX (HIPPI to MaxVideo interface), User Manual,*  
(L.J.Levinson, M.Sidi, Y.Damatov, Z.Natkaniec) 1993
- 93-05: *The specification of a pipelined feature extractor in VDM*  
(J.Haveman) 4 March 1993. Also AFRODITE/UU/JH/DOC/V1