

REIHE INFORMATIK

25/95

**Enable++: A Second Generation FPGA Processor**

H. Högl, A. Kugel, J. Ludvig, R. Männer, K.-H. Noffz, R. Zoz

Universität Mannheim

Seminargebäude A5

D-68131 Mannheim

# Enable++: A second generation FPGA processor\*

H. Högl, A. Kugel, J. Ludvig, R. Männer, K.H. Noffz, R. Zoz  
Lehrstuhl für Informatik V  
Universität Mannheim

## Abstract

*In the computing community field programmable processors are going to fill the niche for special purpose computing devices. A typical example is ultra-fast pattern recognition in experimental particle physics – a task for which we constructed two years ago Enable-1, an FPGA processor rather specialized for pattern recognition algorithms in  $\mu\text{s}$  domain, but also provided with modest features for coping with more general applications.*

*This paper presents the follow-up model Enable++, a 2nd generation FPGA processor that offers several substantial enhancements over the previous system for a wider range of applications:*

*Enable++ is structured into three different state-of-the-art modules for providing computing power, flexible and high-speed I/O communication and powerful intermodule communication with a raw bandwidth of 3.2 GByte/s by an active backplane. The technical realization of all three modules is guided by the maximum usage of field programmable logic. The actual demand of computing- and I/O-power can be satisfied by the number of modules plugged into the crate.*

*Enhanced features of Enable++ comprise the configurable processor topology provided by programmable crossbar switches. In combination with the  $4 \times 4$  FPGA array and 12 MByte distributed RAM the Enable++ computing core offers a strongly increased and scalable computing power. For building new applications the system offers a comfortable programming and debugging environment consisting of a compiler for the C-like hardware description language spC, a simulator and a source level debugger for hardware design. The goal in planning the hardware design environment for Enable++ from scratch is to transfer established methodologies in software design to the design of digital logic.*

\*This work has been performed within the RD-11 and ATLAS collaborations and has been supported by the Bundesministerium für Forschung und Technologie (BMFT) under grant PH/-RS6-93/09 and the Gesellschaft für Schwerionenforschung (GSI) in Darmstadt, Germany, under grant MAMÄK.

*Concerning pattern recognition tasks, we estimate that Enable++ surpasses modern RISC processors by a factor of 100 to 1000.*

## 1 Introduction

Reconfigurable hardware built with field programmable components is an ideal concept for constructing special-purpose processors combining both the speed of a hardware and the flexibility of a software solution. This type of processor consists of a computing core using a set of Field Programmable Gate Arrays (FPGAs) which are configurable by software to form a hardware implementation of an algorithm. Two exemplary implementations of this concept are Enable-1<sup>1</sup> and DecPeRLe-1<sup>2</sup>.

Both processors were designed with different goals in mind. Enable-1 was tailored for a specific pattern recognition task (called triggering in high energy physics) at the ATLAS/LHC experiment and realized an architecture optimized for a general class of systolic algorithms with similarities in dataflow. The high degree of parallelism and the execution of the algorithms in hardware, both provided by two FPGA matrices consisting of 36 FPGAs, were the major advantages that made the Enable-1 machine superior to others at a benchmark arranged from the EAST RD/11<sup>3</sup> collaboration [BBB<sup>+</sup>93].

DecPeRLe-1 was introduced later and is based on a more general architecture and thus has a broader range of applicability than Enable-1. It is built around a computational core of  $4 \times 4$  XC 3090 FPGAs. Surrounding the matrix there is  $4 \times 1$  MByte of static RAM. For interfacing to the external world DecPeRLe-1 uses three general purpose real-time links (e.g.

<sup>1</sup>Built at Lehrstuhl für Informatik V, Universität Mannheim [NZK<sup>+</sup>93].

<sup>2</sup>DecPeRLe-1 has been developed at the Paris Research Laboratory (PRL), a research laboratory of Digital Equipment (DEC) [BRV93].

<sup>3</sup>The RD11/EAST collaboration within CERN is responsible for the ATLAS realtime pattern recognition (trigger) system.

HIPPI) and one TurboChannel adaptor.

This paper is the result of two years of in-depth experience with FPGA processors. It presents Enable++ as a genuine 2nd generation of FPGA processors which is intended for multiple-purpose applications. The new processor combines advantages from Enable-1 and DecPeRLe-1 and introduces several new features. Section 2 gives a brief overview of the hardware architecture of Enable++. The most important improvement is that Enable++ is programmable in a high-level language similar to C, that makes application development more convenient. This programming environment is discussed in section 3. Section 4 deals with the expected performance of Enable++ as a real time pattern recognition system. The last chapter closes with an overview over the current status of the project.

## 2 The Enable++-System

The general view of the whole system will be a standard workstation which is connected to the Enable++ hardware. All hardware parts are contained in one (or up to three) 9HU sized VME crates. The workstation will remain in any respect an ordinary 'working environment', enhanced with additional capabilities by the fact that it is connected to a powerful FPGA processor. For this purpose one can use either a standalone model or a VMEbus based plug-in board. The physical connection depends on the type of frontend: If a VMEbus type workstation is used, one has the advantage of connecting both by a high bandwidth system bus (e.g. SBus) and/or by a serial bus which is part of the systems *service network* and is realized by a transputer link. Using a standalone workstation one is restricted to use only the serial bus because of the distance between both systems. Both buses offer the same functionality, which means that both have the ability to access all system resources.

After connecting the frontend workstation to Enable++, the user will concentrate on writing applications. For this purpose he will deal with the Enable++ Development Environment (EDE) which offers all the functionality for breaking down a high-level specification of a new application to an Enable++ configuration. The user can choose among different functions as there are *compiling* of a system description, hardware *source level debugging* and *simulation*. EDE offers with spC its own brand of hardware description language. It is derived from both C and Hardware-C and flavoured with additional constructs for describing systolic data processing. spC is the attempt to realize a

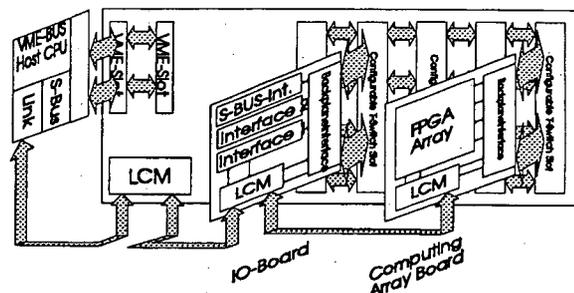


Figure 1: Enable++ system overview

hardware description language which is intermediary with respect to abstraction level and hardware control. Because EDE is planned to fit tightly in a UNIX environment, the user will have in addition dozens of useful and well established tools for free.

After outlining the main 'look-and-feel' of the whole Enable++ system, this section will proceed with a closer look at the hardware components of Enable++. As depicted in figure 1 there are three different types of components:

- Computing Array Boards (CA)
- I/O Boards (I/O)
- an active Backplane (AB)

Computing power is mainly located on the Computing Array Boards. On each board a matrix of FPGAs is connected with high-speed RAM and a network of I-Cube<sup>4</sup> crossbar-switches (FPIDs) building an extremely high-powered computing core which can be adapted to a broad range of different topologies.

The I/O Boards deal with incoming and outgoing data from different sources and sinks (SBus<sup>5</sup>, FibreChannel<sup>6</sup>, HIPPI<sup>7</sup>, SCI<sup>8</sup>, etc.) and may use different types of receivers and transmitters. A single board uses four different ports for data transmitters or receivers. By the use of FPIDs another convenient usage of this board is that of a router. Source and destination of all four ports can be other ports or the active backplane.

The concept of an active backplane was realized after investigating generalized backplane concepts. The backplane now constitutes a field programmable module in the same way as the other two components. It

<sup>4</sup>Field Programmable Interconnect Devices (FPIDs). We use the IQ240, which offers 240 connectable I/O pins.

<sup>5</sup>SUN System I/O Expansion Bus

<sup>6</sup>FibreChannel is a high-speed optical connection, 1062 MBd

<sup>7</sup>High Performance Parallel Interface with 100 MByte/s data rate.

<sup>8</sup>Scalabel Coherent Interface

connects I/O boards and CA boards and is mainly responsible for high-speed data transfer between them. Its open interconnection scheme has the ability to implement a variety of communication models in multiprocessing applications. A possible protocol could be token-ring based with a raw bandwidth of 3.2 GByte/s.

Furthermore both CA and I/O boards are supplied with a novel bus interface architecture based on FPGAs and local accessible dual-ported RAM. This interface is capable of doing high-speed pre- and post-processing of data flowing to and from the board (e.g. formatting) and of synchronizing the data flow between the on-board logic and the backplane.

For the purpose of system maintenance we introduced the concept of Local Control Modules (LCMs) which are connected via an autonomous service network. All of the modules introduced above contain a LCM module which serves for standalone board level test, board initialization, board monitoring and debugging support.

## 2.1 Computing Array Boards

The main functional parts of the Computing Array Board (CA) are the FPGA array and the bus interface. In terms of absolute FPGA resources each CAB has a complexity of approx. 288,000 gates (using the Xilinx figures for sixteen XC4013- and eight XC4010-type FPGAs). The FPGA array consists of 4x4 XC4013-type FPGAs which are connected as shown in the left part of figure 2 (for the sake of clarity we omitted the connections to the lower half of the bus interface).

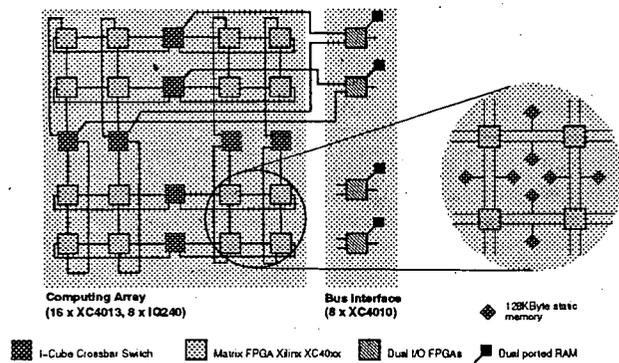


Figure 2: Major building blocks of a Computing Array Board.

Besides nearest neighbour connections, FPGAs are connected through I-Cube crossbars so that different global topologies can be achieved via configuration. Possible topologies are for example linear arrays, rings,

cubes and hypercubes (see figure 3). This highly regular connection scheme allows for very fast and broad pipelined flows of data as needed with systolic algorithms. We believe that with this architecture high utilization is achievable for many regular systolic designs and that even random logic can be mapped on this topology with acceptable losses. In comparison with Enable-1 we have nearly the same count of configurable logic blocks (CLBs) due to less but higher integrated FPGAs. By the new design we end up with improved routing possibilities, higher I/O bandwidth and extended interconnect possibilities.

In addition there will be 96 high-speed 128 kByte synchronous SRAM chips on board. This is in total 12 MByte of extremely fast local RAM. These RAMs are located at the direct interconnections between neighbouring FPGAs as shown in the right part of figure 2. A scheme like that is suitable for pipelined lookup tables (FPGA1-RAM-FPGA2) with up to 17 input bits and 16 output bits and general read/write RAM. In both cases these RAMs have no wait state and one clock cycle latency at 50 MHz system clock. The total throughput is thus 4.8 GBytes/s.

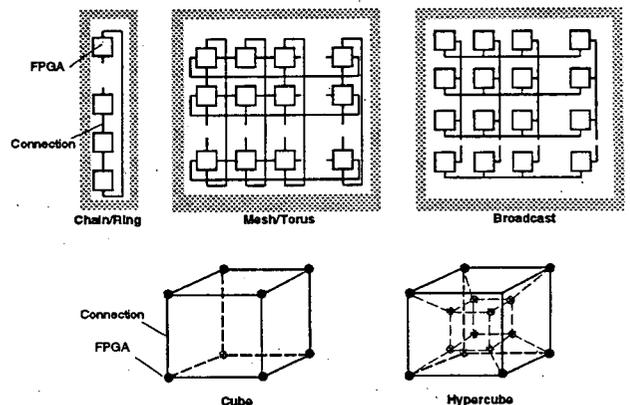


Figure 3: Examples of possible configurations.

The computing array is connected via the crossbars to a bus interface (shown in the middle part of figure 2), comprising eight XC4010-type FPGAs and fast dual-ported RAMs. This circuit decouples the synchronously operating array from asynchronous bus protocols on the backplane and buffers incoming or outgoing data. Connection to the backplane is done by eight 40-bit buses (32 bit plus 8 bit control) which can be operated synchronously up to the full clock speed of 50 MHz. This equals a data rate of 1.6GByte/s.

## 2.2 Enable++ I/O subsystems

The Enable++ I/O subsystems connects the computing core of Enable++ to its real-world environment via high-speed datapaths. It is scalable to achieve an overall data throughput capability equivalent to the installed computing power. To provide a maximum flexibility in connecting to external data sources or destinations a modular design of the I/O boards was selected. Every I/O motherboard can carry up to four I/O daughterboards, each one implementing a specific transmission protocol and media access. The LCM module described in section 2.4 can be used as a fifth virtual I/O port in addition to its primary functions. Together with the four buffered ports connecting to the backplane bus, there is a total of 9 data channels on one I/O motherboard. Every channel can be arbitrarily configured as input or output. Programmable logic (FPGAs and FPGAs) serves for routing and pre-processing of the data. All 9 data channels can deliberately be connected to any one or more of the other channels.

With these features the I/O system is well suited to perform complex routing operations which go far beyond a trivial point-to-point connection. Possible operations are interleaving of two external channels to one of the backplane's subbuses, replication of incoming or outgoing data, even format or protocol conversion can be done. The on-board bandwidth for every one of the 9 channels is 50 MHz by 32 bits (200 MByte/s). Although the 4 backplane ports can keep up with this speed, this value is obviously limited in the case of the outbound channels by the properties of the specific I/O daughterboard. I/O daughterboards to be manufactured in the nearest future will connect to the HIPPI channel and to the SBus of SUN-type workstations. There are also plans for ATM<sup>9</sup> and FibreChannel interfaces. For special purposes there will also be a module providing several Megabytes (approx. 32 ... 128) of buffer memory with a bandwidth of 200 Mbyte/s in block access mode.

## 2.3 Backplane

Handling as high data rates as done by Enable++ is viable only with a special kind of bus system. Since Enable++ is intended to be a multiprocessing system care must be taken to avoid any kind of bottleneck. For this reason a configurable active backplane bus is integral part of the system. It will have a multiple-T

<sup>9</sup> ATM is a network protocol used in high-speed serial connections (155 MBd, 622 MBd available)

topology where at every slot the same kind of router circuit will buffer, split, bridge and route signals coming from neighbouring slots. Thus static and dynamic reconfiguration of the system is possible. It furthermore enhances overall data rates since an active routing scheme makes it feasible to use wider buses on the backplane than on the Computing Array Boards and have other than nearest-neighbour connections. Every slot is splitted horizontally into eight 40 bit wide subbuses, each having 32 bit of data and 8 bits for control. With this architecture it is straightforward to implement systolic datapaths with the full system speed of  $8 \times 200$  MByte/s data-rate in total.

## 2.4 Enable++ Local Control

One important feature of the Enable++ system is the implementation of a uniform local control module (LCM) on each of its module boards. The main elements of this LCM are a T425 transputer with 4 to 32 MByte dynamic RAM and 4 transputer links plus an XC4005-type FPGA. Mechanically the LCM is a piggy-back daughterboard to be plugged onto the specific Enable++ motherboard: Computing Array, Backplane or I/O board respectively. An additional LCM may be used in conjunction with a special interface board in the host computer, if needed. Each LCM can be connected via its transputer links to up to 4 other LCM modules or other link interfaces. Typically all LCM modules in the system form a daisy-chain by connecting via two links each to its left and right neighbour. At least one link connects from the daisy-chain to the host computer. Through the FPGA the transputer gains access to all hardware resources of the motherboard it is controlling. Additionally the FPGA has the capability to interrupt the transputer on an internal event or on request from the motherboard, and to perform direct-memory-access to the dynamic RAM.

The tasks of the LCM are

- standalone board level test
- board initialization
- board monitoring
- debugging support

### 2.4.1 Standalone board level test

In this basic mode the LCM is connected to a host computer via one of its links. The transputer performs JTAG/boundary-scan compliant board level testing plus special test algorithms, if necessary. Our goal is

to achieve 100% test of all connections and functions of the motherboard.

#### 2.4.2 System initialization

The host computer transfers all configuration data to the LCM module which first stores it in its internal memory. For a full-scale Enable++ CA board this will take about 1.5 Mbyte of memory. Depending on the memory resources on the LCM, alternate configuration data sets may be transmitted subsequently to provide means for dynamic configuration changes. The LCM transputer performs configuration of the motherboard by selecting any group of up to 32 configurable logic chips and transferring the data either under program control or via direct-memory-access. Typical configuration or reconfiguration times will be less than 50 ms for a full-scale Enable++ CA board. With a large memory to store different configuration data sets the LCM may perform complete or partial dynamic reconfigurations of the motherboard, either on demand by the host computer or after an interrupt generated by the motherboard. Algorithms may take advantage of this self-modifying-hardware capability.

#### 2.4.3 Board monitoring

The FPGA chip on the LCM can be connected to some user defined status lines on the motherboard. Application specific designs can be developed for the LCM to trigger on certain conditions on these status lines. Various timers, counters or logical equation checks may be implemented. As the transputer is not fast enough to keep up with the speed of the motherboard, the FPGA brings the real-time capability to the LCM module.

Any kind of error handling and logging procedures can be run on the transputer. Samples of data can be taken for online monitoring, either small data sets at full speed or arbitrary data sets with a reduced speed. Thus data can be transferred to the host computer via the transputer link without interfering with the main datapath on the backplane bus.

#### 2.4.4 Debugging support

Support for debugging relies on several properties of the components used on all boards or modules and on special features of the LCM:

- The clock system of each motherboard can be controlled by the LCM
- All FPGA and FPID components support full JTAG boundary scan capability

- All FPGA components support a non-destructive readback capability of all internal states and memory locations/registers
- Each motherboard provides a set of programmable status lines
- Each LCM runs a special debug task communicating with the Enable++ system debugger
- High-speed trigger conditions on the motherboard can be handled by the LCM
- All LCM modules form a link-based network independent from the main datapath.

The LCM performs the mapping of source level node names to local hardware resources with the help of the debugging database (see chapter 3). The average access time from the LCM to any given resource on its motherboard varies with the location of that resource (memory, I/O buffer, FPGA internal register) but will not exceed a few milliseconds - fast enough for interactive debugging. As the data flow between the host computer and the LCM modules consists mainly in control information the speed of 20 MBd on the link network will be fairly sufficient.

Debugging of the Enable++ system can be performed not only board-by-board but on system level and with the comfort of a high-level debugging tool. The LCM assists the main debugging task running on the host computer by virtualizing the hardware on the Enable++ board components.

### 3 The Design Environment

One major criterion in the design of the Enable++ machine is the underlying general hardware architecture which allows a broad range of applications in different fields such as pattern recognition, number crunching, fast networking and logic simulation. Because the system's configuration space is enormous, it is most important to supply the Enable++ hardware environment with powerful tools for designing and debugging new applications. Our goal is to establish a toolset which shifts the design process for hardware as far as possible to a standard which has been well established for years in software design.

The EDE (Enable++ Development Environment) we are currently building is designed with the following principles in mind:

- The procedure of building an application for Enable++ should be very similar to building an ordinary software application.

- The operation of Enable++ is tightly bound into a program running on the host workstation. Both host and Enable++ functionality is specified in one common set of source files.
- A source level debugger allows comfortable testing and debugging of Enable++ applications *while* the hardware is actually operating.
- A simulation tool provides the possibility to test Enable++ algorithms completely prior to any hardware implementation.

### 3.1 Overall structure of the design environment

The usefulness of a large FPGA processor like Enable++ depends heavily on how easily the system can be used and adapted to new applications. In an extreme scenario a human operator would solely write ordinary software without being aware of using application specific reconfigurable logic. An intelligent design system would decompose the software specification and automatically map parts of the software to hardware. With Enable++ we will provide a design environment which offers an intermediate solution to the design problem which is much more useful in application design than the above 'radical' solution: Both hardware and software is explicitly specified in one common set of source files. In case of the software part on the host computer the language is C, in case of the hardware part it is spC (systolic parallel C, a derivative of both C and Hardware-C), which is designed for describing systolic and parallel logic systems.

As shown in figure 4, the EDE consists mainly of four different tools. The spC compiler `spcc`, the host C compiler `gcc`, the hardware debugger `spcdebug` and the hardware simulator `spcsim`. The principal processing of the two paths (host software/Enable++ hardware) of an application program is done by the compilers `gcc` and `spcc`.

The paradigm shift of specifying hardware by software makes well established rules in traditional software design applicable in hardware design. One of the most important rules states that a programming system is only as useful as its debugging facilities. We adopted this rule and plan the source level debugger `spcdebug` to have complete control of the hardware design in the debugging phase similar to the debugging capabilities of C source level debuggers (`gdb`, `ups`) in traditional software design. For this purpose `spcdebug` makes use of the LCM service network to access Enable++ hardware.

With the simulator `spcsim` we want to enhance the development environment with a tool to verify the

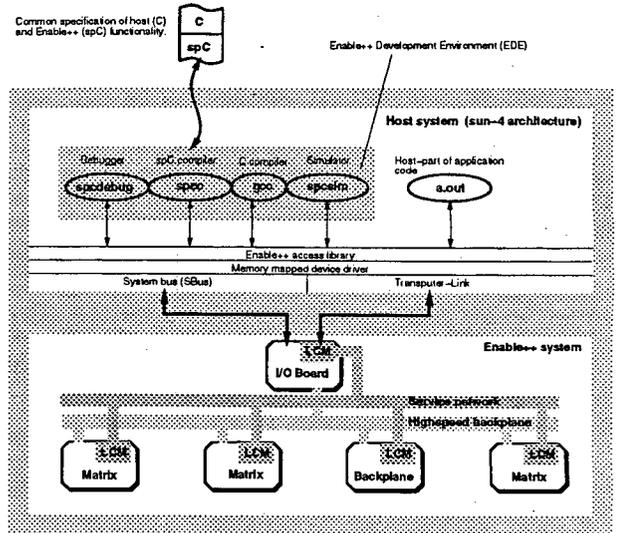


Figure 4: Overview of Enable++ Development Environment EDE.

behavior of algorithms prior to any mapping in Enable++ hardware.

Figure 4 also shows the approach we chose in communicating between host and FPGA processor. There are in principle two possibilities for connecting a host workstation to the hardware. First, if the host computer is a standalone workstation, then the connection is established via a Transputer Link to the Service Network. Second, a VMEbus based host machine can be plugged into the Enable++ crate with either Link or SBUS connection or both to an I/O board. A layered communication software provides an abstract and structured way of accessing and exchanging data between both entities. The access point for services like EDE and application programs is the Enable++ access library. The low-level device driver and physical I/O are in this way hidden from the application programmer.

### 3.2 The development cycle

As shown in figures 4 and 5, a mixed description consisting of C and spC code is the starting point for a new design. A preprocessor handles the separation in C and spC code.

C path: The C code output by the preprocessor constitutes the host computer part of the application program. For accessing the FPGA processor the code has furthermode to be linked with the Enable++ access library.

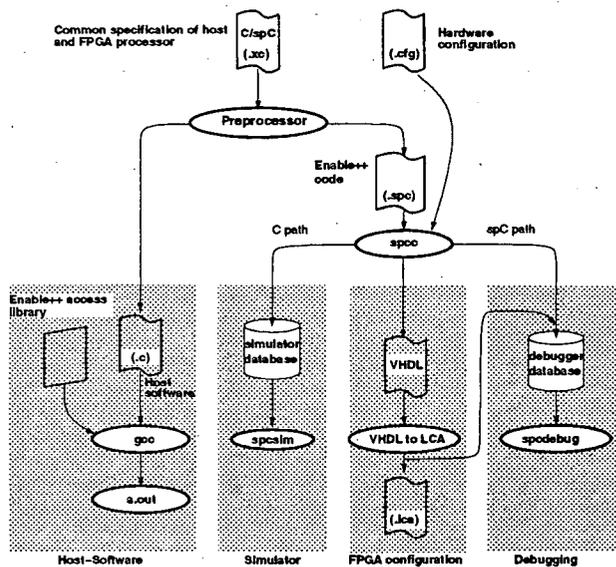


Figure 5: Design paths for building Enable++ applications.

**spC path:** The spC code is processed with the spcc compiler which uses a configuration file for controlling various task specific hardware parameters (e.g. configuration of the FPIDs). spcc outputs code for three different purposes:

- VHDL code for specifying the functionality of the FPGAs.

spcc converts the spC description to VHDL code. We chose this format because of the recent availability of commercial CAD tool which map VHDL code to different FPGA technologies, e.g., Synopsys and Asyl+. With one of these tools the VHDL code is further processed to Xilinx configurations (LCA files).

- A simulator database which can be "executed" with the simulator kernel spccsim. The simulator database contains data structures which are built during compilation with spcc.
- A debugger database which allows source level debugging of the Enable++ hardware with spccdebug (see also section 2.4.4). For source level debugging it is essential, that the elements of the synthesized logic are mapped to corresponding code fragments of the initial spC description.

Relevant parts of the debug database are distributed across the LCM network. The debugger task on each LCM module transparently performs the accesses to the hardware resources, setting and checking

of breakpoints and inspecting/modifying of variable registers.

## 4 Applications

Although Enable++ is designed as a general purpose processor the first applications will be real-time pattern recognition in High Energy Physics. The high data rate of today's particle physics detectors requires a powerful online pattern recognition system for data reduction called the trigger system. Typical trigger tasks for the subdetectors of the ATLAS experiment, one of the two experiment at the world's largest particle accelerator LHC/CERN (planned for 2004) has to be handled within  $10\mu s$ . The data rates are in the order of some 100 Mbyte/s. While in the past particle physicists were forced to build dedicated electronics for each subdetector trigger, Enable++ allows to use one system for all subdetectors just by reprogramming.

Three of the ATLAS trigger tasks are specified as C programs and implemented on modern RISC processors (PowerPC, Alpha, Sparc10) [BHL94]. This allows a performance comparison between an FPGA processor and a standard CPU solution. As long as the Enable++ hard- and software is not completed and direct measurements therefore impossible the benchmark algorithms are formulated in spC, and speed as well as resources are manually estimated. This method is simplified by the pipelined structure and the lack of data dependencies in the algorithms. The accuracy of the performance estimation is acceptable (between 10% and 20%) as experiences with a similar method for Enable-1 have shown.

The following section gives a short overview about the three pattern recognition algorithms and summarizes the performance of an implementation on standard workstations and Enable++. Detailed descriptions of the algorithms, which are named after the corresponding subdetectors (Transition Radiation Tracker TRT, Silicon Tracker SCT and, Calorimeter) can be found in [Leg94] and [KM94].

The trigger task of the TRT is to determine the most significant electron track in the  $\phi/z$  projection of the detector. Due to the detector geometry particle tracks appear as zick zack lines of mean slope  $d\phi/dz$  in the 96 times 16 2-bit image. In order to cope with the track geometry the TRT algorithm consists in a Hough transform for track recognition and a weighted maximum finding for the electron identification.

The Hough transform is implemented as a lookup table ( $64k \times 256$  channels) distributed over the FPGA

matrix RAM. By asserting the pixel coordinates onto the RAM in a systolic way a very effective implementation is reached. The FPGAs histogram the lookup table outputs and perform a weighted maximum search of the histogram channels.

Similar to the TRT algorithm a Hough transform is used to find tracks in the  $r$ - $\phi$ -projection of the SCT (image size is 4 layers of 1000 pixels each). Because both TRT and SCT use a zero suppressed pixel coordinate list as input, the principle implementation of the SCT is similar to that of the TRT. The major differences to the TRT are the 8 times wider (16 times less deep) Hough space ( $4k \times 2048$  channels) and a different histogram evaluation algorithm.

In contrast to the previous algorithms the Calorimeter requires a huge amount of arithmetic circuitry. The center of gravity has to be evaluated within two 16-bit  $20 \times 20$  data windows. Depending on the found center of gravity the 2nd moment of one of the data windows is calculated.

Because of that dependency the algorithm has to be divided into two steps. In a first step the input data are accepted by Enable++ and stored in a dual-port memory. During this step the center of gravity is calculated using bit-serial (but data parallel) arithmetic circuitry. The found center of gravity serves as base pointer for the 2nd moment coefficient tables. In a second step the input data are read from the dual-port memory and the 2nd moment convolution is performed. At the same time the 2nd moments are calculated, the center of gravity for the next event can be evaluated reducing the required frequency by a factor of two.

Each algorithm implemented on Enable++ requires less than  $10\mu s$ , in some cases more than one image can be executed concurrently. The execution of the trigger tasks onto Enable++ provides a speedup of 100 to 1000 compared with traditional processors. Table 1 summarizes the benchmark results for Enable++ and compares the results with the performance on three modern RISC processors (DEC-Alpha, Sparc10, PowerPC). Note that for the workstation implementation no I/O is considered (in contrast to the implementation on Enable++). Because none of the CPUs shows significant performance advantages over the others only a mean speedup (above the mean CPU performance) is listed.

## 5 Conclusions and Status

The Enable-1 machine and DecPeRLe-1 have proved that FPGA processors are well suited for high-speed

algorithm	$t_e$ [ $\mu s$ ] <sup>a</sup>	$n_{IM}$ <sup>b</sup>	$T_e$ <sup>c</sup> [ $\mu s$ ]	$\langle S \rangle$ <sup>d</sup>
Calorimeter	10	4	2.5	430
TRT	9	2	4.5	180
SCT (2.5%) <sup>e</sup>	8.5	1	8.5	580
SCT (1.0%)	4	1	4	600

Table 1: Benchmark results for Enable++

<sup>a</sup>Execution time

<sup>b</sup>Number of images processed in parallel

<sup>c</sup> $t_e/n_{IM}$

<sup>d</sup>Speedup

<sup>e</sup>Value in percent is the ratio of set to all pixels

processing requirements, e.g. in feature extraction algorithms. Enable++ is the result of rethinking both architectures for a more general applicability offering the following enhancements:

- A new feature of Enable++ is the flexible processor topology provided by a computing array configurable by field programmable crossbar switches. In combination with the state-of-the-art FPGA resources and 12 MByte distributed RAM, Enable++ offers strongly enhanced computing power.
- A major part of Enable++ focuses on a flexible and extremely powerful I/O system. The division of the system in Computing Array Boards and I/O Boards connected by a programmable FPGA bus provides a high degree of scalability.
- For building new applications the Enable++ Development Environment EDE has been planned and up to now partly realized which consists in the final version of the following tools: A compiler for the high-level hardware description language spC, a simulator and a source level debugger for hardware designs.

Once realized this concept will be an extreme improvement in our FPGA-design capabilities, since it breaks up the design process in small steps which are well defined and independent. This gives programmers freedom in developing complex applications because they can work most of the time on an abstract level without thinking about all pitfalls of hardware design. Debugging of logical errors and first tests can be done with spC and VHDL code while optimisation and hardware debugging are done afterwards. Very fast designs can be implemented using the low-level tools if appropriate. Hardware debugging and system monitoring allows verification of designs for real-time data and real system environments. Together with its inherent scalability Enable++ will be a very powerful, universal and user-friendly FPGA processor system.

All type of boards are under construction and will be operating in a-prototype version until summer/fall

1995. An  $\alpha$ -release of the spC compiler is already available now and for 9/1995 we scheduled the release of an  $\alpha$ -version of the EDE. A full featured stable version of the EDE with complete debugger and simulator support is expected to be available at 6/1996.

## References

- [BBB<sup>+</sup>93] J. Badier, R.K. Bock, Ph. Busson, S. Centro, C. Charlot, E.W. Davis, E. Denes, A. Ghorghé, F. Klefenz, W. Krischer, I. Legrand, W. Lourens, P. Malecki, R. Männer, Z. Natkaniec, P. Ni, K.-H. Noffz, G. Odor, D. Pascoli, R. Zoz, A. Sobala, A. Taal, N. Tchamov, A. Thielmann, J. Vermeulen, and G. Vesztergombi. Evaluating Parallel Architectures for two Real-Time Application with 100KHz Repetition Rate. *IEEE Tr. Nucl. Sci.*, 40(1):45-55, 1993.
- [BHL94] R.K. Bock, R. Hauser, and I. Legrand. Algorithms in econd-level triggers for ATLAS and benchmark results. ATLAS DAQ note 94-27, EAST note 94-37, 1994.
- [BRV93] P. Bertin, D. Rincin, and J. Vuillemin. Programmable Active Memories: A Performance Assessment. Technical report, Digital Equipment Corporation. Paris Research Laboratory, 1993.
- [KM94] W. Krischer and L. Moll. Implementation of a pattern recognition algorithm for the Si tracker on DecPerLe-1. EAST note 94-21, 1994.
- [Leg94] I. Legrand. Data collection and preprocessing for the ATLAS second-level trigger. EAST note 94-30, 1994.
- [NZK<sup>+</sup>93] K.-H. Noffz, R. Zoz, A. Kugel, F. Klefenz, and R. Männer. Results of On-Line Tests of the ENABLE Prototype, a 2nd Level Trigger Processor for the TRD of ATLAS/LHC. EAST note 94-03, 1993.