Excerpt from the full proposal available at:

https://drive.google.com/file/d/0BxWfo2ViJ6r5MlpkbUpjbEIybUk/view?usp=sharing

See excerpts from letters of endorsement and appreciation of Crosetto's inventions and copies of a few full letters.

1. Final official report by the FERMILab review panel dated January 31, 1994 recognizing Crosetto's invention valuable and assigning him $150,000 to complete the current development and leave the project in a state where it could be continued

Page 1 - Final official report by the review committee of Crosetto's invention

January 31, 1994

Fermilab

Fermi National Accelerator Laboratory
P.O.Box 500 • Batavia, Illinois • 60510

Project Review Report on:

Digital Programmable Level-1 Trigger with 3D-Flow Assembly

Held: 14 December 1993

Fermi National Accelerator Laboratory

This is the report of the Committee charged to review the "Digital Programmable Level-1 Trigger with 3D-Flow Assembly" and to make recommendations concerning the project to the SSC management. The report consists of a list of the committee members, the summary and recommendations, and a collection of comments and concerns by the committee members that may be of use in guiding future work. The committee charge is appended to the report.

Committee:

Roger Bartley, ASIC Designs Inc.
Robert Downing, University of Illinois (Chair)
Dan Edmunds, Michigan State University
Don Machen, Scientific Systems International
Livio Mapelli, Lawrence Berkeley Laboratory
Vivian O'Dell, Fermi National Accelerator Laboratory
Greg Sullivan, University of Chicago
Silvio Turrini, Digital Equipment Corporation

1

Page 2 - Final official report by the review committee of Crosetto's invention - January 31, 1994

Summary

The committee finds this project an interesting and a unique concept for constructing programmable level-1 trigger systems. We believe the concept will work for calorimetry and may work for level 1 tracking. We see no technical reason why the proposed ASIC processor could not be built in approximately one year. We do not believe that there are any major flaws in the proposed system, however, we believe that more work need to be done in several areas. The committee is, however, concerned with funding a hardware development unless it can be used in an experiment. Since there is currently no candidate "customer", we believe that the current development should be brought to a conclusion and carefully documented for potential future use. Further development should be contingent upon a "customer" from either High Energy Physics or another discipline.

To complete the current development and leave the project in a state where it could be continued, the committee recommends that the 3-D Trigger receive funding of $150,000 for at least six months. This money should fund Mr. Crosetto, a consultant expert in the design of ASICs and cover costs of consulting with foundry engineers. The deliverables from this funding should be:

1. A detailed executable model in VHDL for the proposed ASIC. This model should be available to anyone desiring to test it as a component their system. The model should reflect the specification in item 3 as closely as possible. This VHDL code should be written for easy portability among various VHDL simulators. Documentation should be sufficient for use by others familiar with the art.
2. A review by Mr. Crosetto, based on this committees comments, of the chip specifications should be undertaken to consider such items as register file size, program memory size, applicability for use in pipelined configurations, I/O suitability for passing data to other trigger elements such as global decision hardware, etc. This review should be an iterative process with the consultant since tradeoffs will need to be made between the hardware and the chips capability.
3. Mr. Crosetto, along with the consultant, should develop a specification for the ASIC based on a 0.5 p. - 0.6 p. standard cell device with 100 MHz operation at 3.3 volts. This specification along with the VHDL description should be in enough detail that the ASIC could be taken to silicon if an experiment desired to build the device. It is important that this ASIC specification try and anticipate system issues as much as possible. Diagnostics (JTAG), observability, programmability, DAQ interfacing, etc. should be carefully considered and hooks for them specified in the design.
4. The overall system concepts should be well documented at the current state of development. This will serve as a conceptual design if/when the 3-D trigger is built. If hardware is constructed at a later time, clocking, cooling, etc. can be defined in detail. We see little risk in this approach to the processor chip design itself. The documentation should be relatively easy since most of what is required currently exists.

Page 3 - Final official report by the review committee of Crosetto's invention - January 31, 1994

5. Paper(s) detailing the conceptual system design and the processor design should be submitted to publications such as NIM or appropriate IEEE journals. Funding should be provided for oral presentations of accepted papers at two professional society sponsored conferences.

As stated above, we see no technical reason why the proposed ASIC will be a problem to develop. The design is straight forward and is in the mainstream of today's technology. If the chip were developed the cost should be about $300,000. This cost includes a foundry NRE of about $150,000 and another $150,000 for engineering work by an outside consultant and/or the foundry. The committee feels strongly that the extra engineering help is necessary and prudent to insure the success of the chip. Too many pitfalls abound in the design of any like device even though ASIC vendors advertise to the contrary.

The committee was impressed with the work already completed by an essentially one person operation. It is clear that if or when this project is committed to hardware a larger group needs to be assembled with expertise in project management, mechanics and cooling, electronics, software, etc. More work needs to be done on algorithms for other triggering such as tracking. The current design does not address the hardware necessary for global triggering and data readout issues. We see nothing fundamentally wrong, just not detailed in these areas.

Page 4 - Final official report by the review committee of Crosetto's invention - January 31, 1994

Committee Comments

The following are a compilation of detailed committee comments and concerns during meeting. These comments can be used as background material and basis for guiding future work. All comments are not necessarily self-consistent as some differing views were expressed.

ASIC Specification

Some of this section's comments assume issues of executing various trigger algorithms are answered satisfactorily.

The committee reviewed the proposed sea-of-gates ASIC and was concerned that the technology proposed would not perform at the desired speed. There was also some concern on register file size and program memory size. It was felt that a standard cell approach would result in a better design at less risk in achieving the desired function. The standard cell approach would also result in more useful equivalent gates for use as larger register stacks and program memory.

The VHDL Model for specifying the chip is an excellent approach. This technique, however, has difficulty in describing device timing variations. Another document will be needed for a complete specification that can be used to design real silicon. The VHDL model can be used by others simulating systems using the 3-D processor.

The committee discussed other chip issues. Some concern was expressed on the register size but the standard cell approach will allow easy expansion here. One concern was the saving of previous results for later calculations. The chip architecture for passing of data to non-adjacent processors in level 1 tracking needs investigating.

The design need to be carefully looked at for observability. This design technique helps in diagnostics by making all registers available for interrogation. Additional diagnostic features such as IEEE 1149 (JTAG) are included in the current chip and should be carried over into the standard cell design.

Since a large number of outputs can switch at the same time, ground bounce issues need investigation and the design adjusted to minimize the effect.

The committee felt that a change to standard cell could allow the design to be specified at 100 MHz or slightly higher. The proposed CMOS Gate Array was estimated by the committee to operate at 43 MHz not the 60 MHz advertised by the chip vendor. Also, it was felt that the design should be done for 3.3 volt logic. A major shift is taking place in the industry and designing for future use with 5 volt logic is no longer wise. Some power saving will result from this approach.

Although considerable work has been completed on the ASIC design we feel that it is only at the 20% completion point. About 60% of the effort will go into test vector generation with the remaining 20% for detailed schematic and layout.

Since this chip will be a low volume device by industry standards any contract should specify that the foundry will make it at that low volume.

Page 5 - Final official report by the review committee of Crosetto's invention - January 31, 1994

Consultant

The committee felt strongly that an outside consultant familiar with the chip process proposed should be hired for the project. Many traps are lurking for the designer which are undocumented by the vendors. After three or more designs the need for a consultant is less, however, when changing foundry processes a consultant should be hired again.

Trigger Simulator

A simulator for the trigger processor will be necessary if the device is to become a generally used architecture. This simulator will allow experimenters to test triggering algorithms before committing the hardware. This simulator model is in addition to the VHDL model of the processor chip.

No Current Experiment

The major concern of the committee with committing large amounts of money to this project is the lack of a customer. Both CDF and DO are targeting 1996 as the beginning of the next run The time scale seems to be too short for either experiment to commit to this development.

If an experiment could be found to use this prototype, the funding levels proposed herein should be seriously reconsidered.

It is possible a fixed target experiment could use a simple version of the system. A longer range project would be a prototype for LHC. Currently some LHC R&D projects exist and this would be another approach to the triggering.

The committee also felt that without a system to go into the fabrication of a few chips would yield little more information on the 3-D Processor than that which can be gained from simulation of the chip.

Work for general purpose

To make this device a truly general purpose device software tools need to be developed for the experimenter. Without these tools a similar situation would result as that with the first processor farm at FNAL. Software has made the later farms very successful. Some of this development could be contracted to software houses which specialize in this type of code development.

A system simulator needs to be developed for experimenters to test trigger algorithms. This simulator should take Monte Carlo generated event data as input.

Page 6 - Final official report by the review committee of Crosetto's invention - January 31, 1994

Architecture

The architecture appears to work for calorimeters, assuming there are no unforeseen problems with the global event decision hardware. Further work needs to be done on tracking algorithms as they relate to first level triggers. The best test of the design would be to model the logic of an existing experiment trigger in this architecture. Using real data, one could determine trigger efficiencies and delay times.

One unknown with this architecture has to do with the added flexibility provided by the programmability of this system. It is hard to crystal ball gaze, however, there was some feeling that given this feature experimenters would probably think of clever uses not now possible. Better level one triggering will reduce the data rate into level two. If a large enough reduction could be achieved, level two triggers could be replaced by a processor farm. It is possible the farm is the same as used for final event processing before storage. The rapid advances in computing power appear to make this approach feasible, if not now, at least within a few years.

The committee felt that in some cases a pipeline architecture might prove useful. The current proposal assumes all processing in one plane. Additional processing planes add more computing when necessary. Pipelining needs to be architected carefully so additional bandwidth is not needed between processors. An intelligent compiler should be able to do this easily by insuring that no extra intermediate values have to be passed. The interplane bandwidth needs to be no more than the data input rates.

System Design

The committee believes there are no major flaws in the conceptual design. More detailed work is needed in several areas:

Cooling: Many air path 'leaks' occur between cables. These become low resistance paths for air flow causing it to shunt around processors which are not adjacent to the fans. The problem seems solvable with metal work and foam rubber air dams at appropriate places. The cooling design should be checked for higher power chips at higher speeds.

Clock tolerance and distribution: Some ideas are presented but a detailed study of setup/hold times, skew, temperature effects has not been done. Nothing fatal in the general concept was seen, however, as the clock rates increase this issue will become more critical.

The mechanics are of necessity not conventional. Some special tools need development to ease assembly and disassembly. These are not complex, just not done. A larger mockup with adjacent units on all sides will answer these questions.

Page 7 - Final official report by the review committee of Crosetto's invention - January 31, 1994

The hardware diagnostic system is based on IEEE 1149. This is a common industry standard. The main issue here is if the "correct" points are being monitored. The ASIC consultant will be able to advise on how best to do this. Other diagnostic tools are needed to monitor and maintain the system. These tools will require a significant software effort if the system is to be used and maintained by others (not the designers). The serial system is also available for this use in addition to controlling the processors.

There was a question if RS232 is the appropriate serial technology to control many 3-D processors. A conventional microprocessor in each card cage might improve communication by using IEEE802. This micro could take on additional tasks of control and monitoring.

A useful project would be the design on an interface to actual detector outputs. A place is provided in the hardware for such boards, however, one would be more comfortable if a straw design was done on several existing front ends.

Some more work is required for the output data and DAQ connections from the level 1 trigger. Triggers requiring a global view need to gather information form many places in the system. We don't see any serious problems, just not done yet. Data paths to the DAQ need investigating. There should be enough output ports and bandwidth - needs checking to confirm.

Software

Potentially the most work will be done in writing support code. The prototype system can be brought up by the design groups software people. Their intimate knowledge will let them test and debug the system without expert tools. These tools will not provide the software environment necessary for wide spread acceptance.

As stated earlier a simulator seems necessary for trigger code testing. Diagnostic software will

be needed.

Although the instruction set is simple, a compiler which understands concurrency will be useful to the non-hardware user.

A runtime monitor is needed. The complexity of any trigger system demands such a monitor.

Run time tests patterns needed. Could be pre loaded data in special registers that can be triggered for periodic checks on the system integrity.

Page 8 - Final official report by the review committee of Crosetto's invention - January 31, 1994

Cost

The committee couldn't agree on the overall system costs relative to current triggers in design, e.g. CDF upgrade. Some felt that this system could cost a factor of two more, some the same. We did not have enough time or data to resolve this issue. LHC comparisons presented were hard to evaluate. The general feeling seemed to be that it could be cost competitive.

Management

The committee felt that a larger group will be necessary to make the system viable and credible to many potential users. There was some feeling expressed that a larger group will bring up issues of management. Organization will be important and require the project leader to have these skills for a timely and successful conclusion. These issues will need to be solved up front to avoid potential problems downstream.

3. The PI, Dario Crosetto met or exceeded all requirements made by the FERMILab review panel going beyond the imagination of future science

Crosetto's innovative idea was recognized valuable and endorsed by Andrew Lankford who wrote a letter of endorsement and defined the charges to the reviewers (see below) of his 3D-Flow invention.

Department of Physics
University of California
Irvine, CA 92717

November 5, 1993

Dr. Antony Montgomery
Director of the Office of Research and Technology Applications
Superconducting Super Collider Laboratory
2550 Beckleymeade Avenue
Dallas, TX 75237

Dear Dr. Montgomery:

I strongly recommend that work of Dario Crosetto on the "3D-Flow" processor project be supported as part of the closeout of the SSC Project.

Crosetto's R&D project has aroused considerable interest in the high energy physics community as a technique to perfonn fast, programmable triggers. He has done a tremendous job over the last couple years to develop this concept to the point of convincing the community of its feasibility. In particular in recent months, Crosetto has established the technical feasibility of his processor chip, through discussions with expert commercial processor IC designers, and has established the technical feasibility of his system architecture, by assembling a 3D-Flow Mini-Rack with a number of prototype cards. The "3D-Flow" concept has potential for application in nearly any future high energy physics experiment. It is now important that this project be supported at least to the stage of a prototype of the processor chip, and preferably to a prototype of a small system in order to demonstrate the architecture as well as the processor. Without support from the SSC closeout, this interesting work risks being lost to the HEP community.

Thank you for your consideration of support for this interesting and very important project.

Sincerely yours.

Andrew J. Lankford
Professor and
SDC Electronics Subsystem Manager

xc: Fred Gilman
James L. Siegrist

One year after the review the Director of the Computing Division at FERMILAB, Joel Butler also confirmed the positive outcome of the review and the value of Crosetto's invention stating: "***the 3D-Flow project is the only detailed study demonstrating the feasibility of executing several level-1 trigger algorithms of different experiments***."

Crosetto's 3D-Flow invention is still today the only solution demonstrating feasibility of executing several level-1 trigger algorithms of different experiments because none of the level-1 triggers developed during the past 23 years have the capability to execute L-1 Trigger algorithms as the 3D-Flow OPRA does.

FEB-27-'95 MON 12:24 ID:FNAL COMPUTING WH TEL NO: 708-840-2783 #110 P02

Fermilab

Fermi National Accelerator Laboratory
P.O. Box 500 • Batavia, Illinois • 60510
708-840-3148 Butler@FNAL.GOV

Computing Division
Joel Butler

February 27, 1995

To whom it may concern,

About one year ago, Fermilab conducted a review of the proposal by Dario Crosetto for the development of a 3D-Flow Architecture for triggering in high energy physics experiments. The review committee consisted of five electronics experts and three physicists, all acknowledged to be experts and actively involved in trigger issues. This review committee found the design to be promising for its potential application in HEP triggering and possibly elsewhere, and to be a technically sound and feasible approach.

The 3D-Flow Architecture offers the possibility of performing decision-making, image-processing, and pattern recognition in a flexible manner - because of its inherent programmability - and at higher speeds than conventional approaches - due to parallelism. The high degree of connectivity between processors is an especially significant advantage of the proposed system. These elements make it a promising approach to solve many problems, ranging from high speed triggering applications in High Energy Physics to image processing applications of significance in the commercial sectors.

The model developed so far, which is a fully documented design and a simulation, can be 'plugged in' to a larger simulation framework to test the applicability of this approach. The model is, therefore, an important step in itself which allows designers to make informed choices. Using this model, Dario has demonstrated the versatility of the approach. Work to advance the realization of the model in actual hardware would also be of great interest and benefit not only to the High Energy Physics community but to several other application areas. Prototype development would result in the solution of several system and integration problems and would reassure potential users that the technology was viable. Similarly, the development of a support environment is an important step in producing a usable product. As far as applications to High Energy Physics, at present the 3D-Flow project is the only detailed study demonstrating the feasibility of executing several level-1 trigger algorithms of different experiments. I would like to strongly endorse funding to build this chip for the benefits that the HEP community will stand to gain. If this chip was to be available 'off-the-shelf', it would undoubtably be of interest to the scientific community and will save the cost of designing several separate ASICs.

Joel Butler

Head, Computing Division

cc: T. Nash
J. Venard

More information about the major public scientific review of Crosetto's invention at FERMILAB

4. SSC Director's request for an international summit and evaluation of Crosetto's invention is reported in the email dated November 19, 1993, specifying that " *A physicist from the European HEP* [High Energy Physics] *community should be considered*."

From: To: CC: Subj:
FNAL::IN%"lankford@lankford.ps.uci.edu" 19-NOV-1993 17:32:39.54 IN%"BUTLER@fnalv"
IN%"KIRK@sscvxl.ssc.gov"

-

3d flow review

Return-path: <lankford@lankford.ps.uci.edu>
Received: from FNAL.FNAL.GOV by FNAL.FNAL.GOV (PMDF V4.2-12 #3998) id <01H5IAMCP2XCOOJB45@FNAL.FNAL.GOV>: Fri, 19 Nov 1993 17:32:28 CDT Date: Fri, 19 Nov 1993 15:11:59 -0800

From: lankford@lankford.ps.uci.edu (Andrew *J.* Lankford)
Subject: 3d flow review To: BUTLER@fnalv
Cc: KIRK@sscvxl.ssc.gov
Resent-message-id: <01H5IAMCRHQA00JB45@FNAL.FNAL.GOV>
Message-id: <9311192311.AA04381@ lankford.ps.uci.edu.ps.uci.edu > X-VMS-To: IN%"BUTLER@FNAL.FNAL.GOV"
X-VMS-Cc: IN%"KIRK@sscvxl.ssc.gov"
Content-type: TEXT/PLAIN; CHARSET=US-ASCII
Content-transfer-encoding: 7BIT

Joel,

Please review the following charge and let me know if it is acceptable to you. Thanks,

Andy

November 19, 1993

Dr. Joel Butler Fermilab, M.S. 120 P.O. Box 500
Batavia, IL 60510

Dear Joel,

Tom Kirk, acting in his capacity as Deputy Director of SSCL, has asked me to help arrange a technical review of Dario Crosetto's "3D-Flow" processor project. This project has aroused considerable interest among high energy experiments; however, because of the forward-looking nature of the project, high energy physicists are often uncertain of its detailed technical feasibility, particularly of the processing elements. Via this letter, I am charging you with organizing the needed technical review. The results of the review will be used by SSCL in determining the level of support which this project should be given as part of the SSC closeout.

You may formulate the detailed charge to the review committee. I recommend the following charge:

Perform a detailed technical review of the "3D-Flow" processor project being conducted by Dario Crosetto of SSCL. The review should address the following issues:

a) The technical feasibility of an integrated circuit implementation of the high-speed processing elements, including an estimate of the required development time and development costs,

b) The technical feasibility of the proposed system implementation, including an estimate of the required development time and development costs,

c) The general suitability of this approach for implementationas a trigger processor for future high energy physics experiments, including practical issues such as flexibility in application to specific experiments, maintainability,

and projected cost.

d) Progress to date in developing this approach.Provide to the Deputy Director of SSCL a-written summary of the review addressing the above issues. If appropriate, the summary may also recommend appropriate goals for continued work on the project as part of the SSC closeout.

The review panel should consist of at minimum the following:

- a technical expert on digital IC design, preferably microprocessor design,
- an additional digital electronics engineer,
- two physicists expert on triggers.

It may consist of other experts, for instance a technical expert who is also knowledgeable about commercialization of the processor for the high energy physics market. A physicist from the European HEP community should be considered. The committee may be expanded as convenient; however, I suggest that the review committee be kept sufficiently compact so as to facilitate a review date in the near future.

The review should be held as soon as possible. It must be held before the new year in order to be useful for SSC closeout. It may be held at a location convenient to you and the reviewers in order to facilitate its timeliness.

I suggest a one-day review, including time to draft the written summary. I do not believe that the review can be performed in sufficient depth in less time. In order to facilitate the in-depth nature of the review, detailed technical material on the project should be circulated to the reviewers well in advance of the review date. I believe that Crosetto already has prepared suitable materials. Sufficient time should be allotted for detailed technical questions from the review committee. In fact, I suggest only a very short formal presentation by Crosetto at the review, in order to allow time to adequately address the committee's questions.

Since this review may be of interest to a number of members of our community, you may choose to publicize it to at least the leaders of trigger groups of current and future experiments. Such interested parties could submit their comments to you, as head of the review, and to Dr. Antony Montgomery, Director of the Office of Research and Technology Applications, SSCL.

Thank you for conducting this important review. I regret that I will be unable to participate.

Sincerely,

Andy

Andrew J. Lankford

SDC Electronics Subsystem Manager

Dear Bob,

Below is an excerpt of the letter from Andy Lankford, head of the SDC electronics group, to me asking me to set up the review.

It contains the basic elements of the charge. After the charge, I include several concerns that have been raised, and which are in the nature of personal communication and not part of the charge,

Best wishes

Joel

The Charge:

(From Andy Lankford, via Joel Butler)

Tom Kirk, acting in his capacity as Deputy Director of SSCL, has asked us to help arrange a technical review of Dario Crosetto's "3D-Flow" processor project. This project has aroused considerable interest among high energy experiments; however, because of the forward-looking nature of the project, high energy physicists are often uncertain of its detailed technical feasibility, particularly of the processing elements. Via this letter, I am charging you with organizing the needed technical review. The results of the review will be used by SSCL in determining the level of support which this project should be given as part of the SSC closeout.

The detailed charge to the review committee is as follows:

Perform a detailed technical review of the "3D-Flow" processor project being conducted by Dario Crosetto of SSCL. The review should address the following issues:

a) The technical feasibility of an integrated circuit implementation of the high-speed processing elements, including an estimate of the required development time and development costs,
b) The technical feasibility of the proposed system implementation, including an estimate of the required development time and development costs,
c) The general suitability of this approach for implementation as a trigger processor for future high energy physics experiments, including practical issues such as flexibility in application to specific experiments, maintainability, and projected cost.
d) Progress to date in developing this approach.

Provide to the Deputy Director of SSCL a written summary of the review addressing the above issues. If appropriate, the summary may also recommend appropriate goals for continued work on the project as part of the SSC closeout.

5. Agenda of the formal, International, PUBLIC, scientific review of Crosetto's invention held at FERMILab on December 14, 1993 and outline of Crosetto's presentations to the scientists at the FERMILab auditorium, the public Q&A with scientists and with the review panel and follow up

On December 14, 1993, eight reviewers and Dario Crosetto gathered at FERMIlab from different parts of the world to review Crosetto's invention.

Crosetto presented his invention in FERMIlab's auditorium before hundreds of scientists followed by a Q & A session. All questions posed to Crosetto were answered satisfactorily. The in-depth review continued for the entire day in several sessions with the review panel and a smaller group of FERMIlab scientists.

Agenda of the Scientific Review on Crosetto's:

Digital Programmable Level-1 Trigger with 3D-Flow Assembly

Held: 14 December, 1993 at Fermi National Accelerator Laboratory

		Time	Place	
1)	Committee	9:30 – 10:00	7th floor NE	(closed session)
2)	Crosetto's Presentation	10:00 – 11:30	1 West	(open session)
3)	Discussion	11:30 - 12:30	1 West	(open session)
4)	Lunch	12:30 – 01:15	Small Dining Room	
5)	Crosetto Q/A	01:30 – 02:30	7th floor NE	(open session)
6)	Committee	02:30 – 05:00	7th floor	(closed session).

Description of the Presentation to the Review Committee

Since the design is complex and could not be fully covered in one day, an overview was presented using the simultaneous display of large pictures, described below. Additional details were supplied to reviewers upon their request. Concrete backup material including documentation from VLSI chip manufacturers, quotes from industries, and reports on other HEP trigger systems were included as attachments.

To accomplished this, the following were displayed simultaneously:

1) Twelve picture-posters on the right side of the room describing the architecture applied to physics (calorimeter and tracking Level-1 trigger algorithms versus the 3D-Flow architecture).

2) Nine picture-posters on the left side of the room describing the 3D-Flow assembly of the system, including details of selected portions.

3) Six picture-posters at the front of the room describing:

   a) A calorimeter Level-1 trigger algorithm subdivided in steps executable by the 3D-Flow (3D-Flow mnemonic instruction versus calorimeter element operation);

   b) Routing of information and results within the 3D-Flow processor array to an exit point;

   c) A block schematic of the internal architecture of the 3D-Flow processor;

   d) A detailed schematic of the 3D-Flow processor internal architecture. The logic symbols (register, Multiplexer, etc.) had a one-to-one correspondence to the VHDL code displayed on a monitor from a personal computer;

   e) The microcode of all instructions of the 3D-Flow processor;

   f) Commented microcode of each instruction of the algorithm described in poster 3a).

With the above pictures displayed, the user could follow the flow of data in the different paths of the 3D-Flow processor (from the input FIFO, to the internal buses, internal units, and output registers) on the i) detailed logic schematic of the 3D-Flow and at the same time ii) on the timing diagram on the computer generated by the VHDL model of the 3D-Flow processor.

4) On the overhead projector there was the hierarchy of the VHDL files that described the 3D-Flow processor. The execution of these files (the complete set of them make the 3D-Flow processor simulator) was then displayed on the monitor from a personal computer.

5) On the table there was the hardware prototype of a 3D-Flow mini-rack with motherboards, daughterboards, receiverboards and flexible printed circuits interconnection

6) A stack of documents, as attachments, (quotes from industries, VLSI functions data-sheets, HEP reports on trigger systems).

With the above simultaneous display of the project aspects, one could obtain a complete picture of the project from bottom-up or from top-down design, and follow the logical path.

The conclusions of the review panel were published in a report signed by the Chairman on January 31, 1994. It states that it is a valuable and feasible invention with the following words:

1. an invention (stated using the word "*unique*," see first line of the Summary on page 2)
2. valuable and feasible (stated in several sections of the report, some of which are highlighted in yellow).

   E.g. on page 2 "*We believe the concept will work… We see no technical reasons why the proposed ASIC processor could not be built… We do not believe that there are any major flaws… We see little risk in this approach to the processor chip design itself. The documentation should be relatively easy since most of what is required currently exists.*"

   On page 3 "*As stated above, we see no technical reason why the proposed ASIC will be a problem to develop… The committee was impressed with the work already completed by an essentially one person operation* [Crosetto]."

   On page 6: "*…this architecture has to do with the added flexibility provided by the programmability of the system… given these features experimenters would probably think of clever uses not now possible… The committee believes there are no major flaws in the conceptual design*"

   On page 8: "*The general feeling seemed to be that it could be cost competitive*"

In recognition of finding Crosetto's invention valuable and feasible, the FERMIlab review panel assigned $150,000 to Crosetto (all the money available during the SSC closeout), with the expectation that another "*customer*" in HEP or other discipline would further fund its the development and extended his contract for six months "*To complete the current development and leave the project in a state where it could be continued…*" This six month period was extended to nine months, so Crosetto continued to work at the SSC until October 1994 while the majority of the 3,500 SSC employees were terminated in 1993.

No statements in the report expressed any concern, only pointing out that some procedures still needed to be developed. Crosetto developed all of them in the following years; he used much of his own money or used donations from friends to buy components to prove the feasibility and functionality of his invention in hardware.

In summary: All requirements from the Fermilab review panel have been delivered and their anticipated technological advantages and benefits to advance science confirmed.

Crosetto addressed and solved satisfactorily all items pointed out by the review panel as needing further work and surpassed those requirements in developing new innovative concepts that further facilitate experimenters in discovering new particles, reducing the cost of the experiments and by developing new inventions that enable the implementation of an effective early cancer detection using a safe, low-cost screening procedure with the 3D-CBS technology.

He defined the specifications of the 3D-Flow processor, created a detailed executable model in VHDL, anticipated system issues that permitted a working hardware version to be built, implemented the diagnostic (JTAG), observability, programmability, DAQ interface.

He designed the data path to the DAQ verifying that the number of ports and bandwidth can satisfy the expected data rate based on the occupancy of the detector.

He defined in detail clocking, setup/hold times, skew, temperature, cooling, etc. that permitted functional hardware to be built later.

He designed the necessary nonconventional mechanics and tools that permit easy assembly and disassembly of the system.

He addressed the hardware necessary for the global triggering and data readout.

He solved the problem of a large number of possible ground bounce issues due to the large number of outputs switching at the same time with an effective ground distribution, several on-board voltage regulators and by using LVDS current drivers that avoid voltage swings.

He completed in the following year (in part at his own expense) 80% of the effort in generating test vectors and creating the detailed schematics.

He created the assembler to allow users to program trigger algorithms in mnemonic instructions.

He implemented the simulator in C++ that is replicating many times the VHDL model of the 3D-Flow processor for hundreds of thousands of processors working in parallel in a trigger system simulating all the chain from the data arriving in parallel from thousands of detector channels to the filtered output data in a single line to the global trigger processor.

He implemented the fault tolerant software that could detect and provide the ID of the faulty component or connection in the system for easy maintenance and repair. He implemented run-time software that constantly monitors and maintains the system, detecting anomalies, running test patterns that assure performance and checks for system integrity.

He implemented tools to create a system environment of any size, for any type of detector, simulating the trigger algorithms of different experiments (CMS, ATLAS, Alice and LHCb, and any other level-1 and part of level-2 trigger algorithm) that experimenters could test before construction of the hardware. Crosetto published as requested by the review committee results of his study in NIM peer-reviewed scientific journal (NIM references) and presented them to several IEEE Nuclear Science Symposium and Medical Imaging conferences (provide references).

Furthermore, he provided the results of his study to CERN scientists in communications via email and seminars that ECP group leader organized.

Going beyond what was requested by the FERMIlab review panel, Crosetto, in part at his own expense, compiled four units of the 3D-Flow processors in a single component from three different FPGA (Field Programmable Gate Array) technologies (ORCA, from Lucent, Xilinx and Altera) and to a standard cell CMOS technology using Synopsys tools and also as a consultant (they are recognized having among the best software tools and being among the best ASIC Design companies).

He was able to do this because his innovative 3D-Flow system architecture is technology-independent so it can migrate to any future technologies giving it the advantage of lowering the price and increasing the performance as technology improves.

All these studies permitted the implementation of the hardware without a problem. First, Crosetto built a demonstrator of the proof of concept using two Altera prototypes boards and showed its functionality on the oscilloscope at the IEEE-NSS-MIC conference in 2001 in San Diego (CA). He then built an industrialized, modular board to build a system that can satisfy the requirements of detectors of any dimension and event rate of any radiation dose (or luminosity in HEP). The modular board that he designed and built has several testability features like 64 LED to monitor signals and 128 test points for monitoring fast signals on the oscilloscope. This forethought in advance testability of the board during the design phase made it easier to test and have the board and the system working at once.

The statement in the letter of the Head of the Computing Division of FERMIlab reported on page 9: "…*promising approach to solve many problem, ranging from high speed triggering applications in High Energy Physics to image processing applications of significance in the commercial sectors*" after Crosetto built his system in hardware, demonstrating that the trigger systems of different experiments costing hundreds of millions of dollars to develop and implement, can ALL be replaced by a box with several modular boards of his programmable 3D-Flow parallel processing system. Not only can each experiment implement its programmable algorithm, but they can change it after some data taking in order to better filter the background noise and be more selective in finding the Higgs Boson and other subatomic particles.

However, the greatest benefit of Crosetto's invention is that the same "box" with those modular boards can extract information from the radiation administered during a screening process on cancer patients, achieving major benefits in a) lowering the radiation dose to the patient, b) detecting cancer at an early stage and c) lowering the examination cost.

Starting from the year 2000, Crosetto developed new applications in the field of Medical Imaging providing a 3D-CBS technology (3-D Complete Body Screening) that makes use of the 3D-Flow system. The 3D-CBS technology hundreds of times more efficient than current PET (Positron Emission Tomography) devices (currently over 6000) for the first time makes possible an effective early cancer detection at an affordable examination cost.