Preliminary Design Review of Serialising ASIC

(ATLAS Level-1 Calorimeter Trigger)

Introduction

The PDR for the Serialising ASIC (SerASIC) took place between March and June 1999. Its purpose was to assess the functionality, features and feasibility of the design. The review was conducted largely by e-mail, with a final series of meetings with the designers during which the collated list of comments was discussed.

It was intended that the review panel should assess the system specifications from several different aspects, so the reviewers were chosen for having expertise in areas of electronics and system engineering, data acquisition and software. The panel consisted of:

Tony Gillman (RAL)

Paul Bright-Thomas (Birmingham)

Paul Hanke (Heidelberg)

Ullrich Pfeiffer (Heidelberg)

Richard Staley (Birmingham)

Comments from several others in the level-1 calorimeter trigger community were also received, and were channelled via these people.

The designers provided the review panel with a single specifications document. However, as a major function of the SerASIC is concerned with interfacing to the DSS and prototype ROD modules it was found helpful also to read their specifications documents. Three of the panel members were also involved in the DSS and ROD PDRs, so the panel as a whole could assess the SerASIC in its proper context.

All documentation was available on the Web, at:

http://hepnts1.rl.ac.uk/atlas-l1/Modules/Modules.html

General comments

The SerASICs provide many of the core functions of the level-1 calorimeter trigger processor. Located in the e/γ Cluster Processor modules (CPMs), their primary purpose is to time-multiplex and fan out the data arriving from the PreProcessor system into 160 Mbit/s serial bitstreams for transmission to the Cluster Processing ASICs (CPASICs). A secondary role is to capture and assemble raw data samples for transmission to the ROD modules following level-1 Accept signals.

Numerous issues emerged during the review process, many of which concerned fundamental performance limitations. The more serious of these will be summarised in the following section. In general, the panel identified several areas in the specifications document where the text and figures would be improved by some careful revision, and it was recommended that this be done when the other major changes were made. It was requested that unusual acronyms (e.g. BIST – Built-In Self Test) should always be defined.

During the period of the review, continuing simulation studies indicated that the smallest of the Xilinx Virtex FPGA family would provide the equivalent functionality and performance of an ASIC, but at a significantly reduced cost and with the benefits of flexibility.

Furthermore, the next larger device in the family would allow the channel density to be doubled, with only a modest increase in cost. The possibility of operating the serialising system at 200 Mbit/s would be precluded by the use of these particular FPGAs; nevertheless the review panel strongly recommended that the device be implemented in the most appropriate FPGA technology and packaging, and renamed Serialiser. The designers are reminded that minimal latency remains an important goal in the design process.

Specific comments

The issues raised by the reviewers fell into three main groups.

a) Real-time data path

It was noted that the specifications of the incoming data from the PreProcessor system allow for operation with two different link technologies, G-link and LVDS. The reviewers felt that this added unnecessary complexity to the specifications and figures, as these two technologies differ in their data path widths and in the number of their error flags. They therefore recommended that the specifications should be revised to assume that the links from the PreProcessor system use the preferred LVDS technology (which in fact is the most likely scenario). If G-link technology were to be used here the re-programmable nature of the FPGA implementation would allow a modified front-end design. Note that G-links will remain the as technology of choice for the relatively small number of links transporting DAQ data to the RODs, as discussed in **b)** below.

The panel expressed concern that the proposed clocking scheme, using only the recovered strobes from the incoming serial links from the PreProcessor system, would result in the CPASICs receiving their 160 Mbit/s bitstreams with a relatively large phase spread. This would produce maximal crosstalk on the backplane. It was therefore recommended that a two-stage clocking scheme be adopted, whereby the incoming data are first captured by their recovered strobes but are then re-synchronised and serialised using the system-wide 40 MHz clock. This will ensure that all 160 Mbit/s bitstreams are synchronous. The associated potential race conditions in the DAQ data path are also avoided by this change. The panel would like to see a full description of this new timing structure appear in the revised specifications document, with simulation estimates of the worst-case timing margins for data capture.

It was considered important that link errors, indicated by a drop of the Link-Ready flag, should automatically zero the data associated with that particular time-slice. It was therefore recommended that this functionality be specified and provided. In addition, the reviewers suggested that the addition of counters could allow link failure rates to be monitored for diagnostic and statistical purposes.

The panel remarked that no clear specifications were given about the data format on the outgoing 160 Mbit/s bitstreams. They recommended that care should be taken to specify exactly how the 20-bit fields are grouped into four-bit nibbles, taking into consideration the fanout requirements. Close interaction with the CPM designer(s) will be essential.

b) DAQ data path

The panel addressed the issue of the minimum performance level that the Serialiser must achieve. The worst-case scenario would require transferring data from five successive time-slices at the maximum level-1 Accept rate of 75 kHz. To determine the feasibility of this data

rate it was necessary to understand the format of the data being assembled, and the maximum bandwidth of the proposed G-links transporting the data out to the RODs (as distinct from the LVDS links mentioned above, which transport data in from the PreProcessor system).

Configured as a pair of cascaded Serialiser chips as described, it would take 3.2 μS to transfer one time-slice of data from eight trigger towers to a ROD, allowing a maximum of four time-slices at a sustained 75 kHz rate (or three time-slices at the extended 100 kHz rate). In order to transfer five time-slices (to match the capability of the PreProcessor system) it was thought it would be necessary to zero-suppress the data in the Serialiser, at the expense of further complexity. The panel therefore explored other options, and another solution was finally recommended, the details of which are as follows.

Although reading five time-slices might on occasion be useful as a diagnostic tool, for routine data-taking a smaller number (probably only one) would be used. One suggestion was therefore to design for a programmable number of time-slices, which could be modified before each run if wished.

Another important parameter to be considered is the depth of the de-randomising buffers in relation to the maximum bandwidth permitted on the outgoing data links. The designers were asked to justify the specified 32-deep FIFOs.

The specifications document assumes that the data are assembled from the real-time path without any processing, i.e. still in their bunch-crossing multiplexed form. The panel considered very carefully the implications of this proposal, and concluded that it required considerable complexity in order to achieve fail-safe synchronisation. The recommendation therefore was to perform the de-multiplexing function before assembling the data for transport to the RODs. This identical function will in any case be needed in the CPASIC.

The panel noted that a significant reduction in the width of the assembled data word could easily be achieved. Firstly, removal of the bunch-crossing multiplexing flags, tidying the parity fields and stripping out spare bits can remove 16 bits. Secondly, the six-bit chip ID can be omitted (the “hard-wiring” to the G-link data-field already uniquely identifies the chips). Finally, doubling the channel density means that the bunch-crossing number is common, hence effectively saving two bits. The assembled data word will then be only 80 bits wide, compared to the 128 bits originally proposed for two cascaded chips. However, it was noted that the three-bit Hit multiplicity formed from the CPASICs must be included in this word (it had been omitted from the specifications document), which with one global parity bit takes the total word width to 84 bits.

It will therefore take a total of 2.1 μS to transfer each time-slice to the ROD, thereby allowing up to six time-slices (for a total of eight trigger-towers) to be transferred at the maximum level-1 Accept rate of 75 kHz (or five time-slices at almost the extended 100 kHz rate). The panel also noted that the G-links will be operating asynchronously, and so are not limited to frame rates of 40 MHz. Increasing this rate to 45 MHz (900 Mbit/s serial bitstream rate) would comfortably allow five time-slices to be transferred at a sustained 100 kHz rate if wished.

c) Miscellaneous issues

The panel considered that recent advances had considerably enhanced the performance of the Serialiser. Doubling the channel density per package, together with the choice of low-power

LVDS link technology (if confirmed), had almost certainly removed the need for expensive MCM technology on the CPMs.

In the area of system compatibility, the panel recommended that there should be close interaction with the designers of modules both upstream and downstream of the CPM (e.g. the PreProcessor and the ROD modules), in order to ensure that data formatting and interfacing specifications are always compatible.

It was also suggested that it would be useful to define limits for some of the key system performance parameters (e.g. Bit-Error rates, time jitter on incoming data, latency, etc.) where possible.

Finally, the panel would like to see a little more detail of the proposed test procedures for evaluating the performance of the Serialiser as a component before loading on to the CPMs, e.g. what the minimum upper limit of clocking frequency is, and how it will be determined.

A.R.Gillman 30 June 1999