Notes on Pilot Project modelling meeting

NIKHEF, January 27 + 28, 1999, Amsterdam

Notes by R. Cranfield, F. Wickens, J. Vermeulen

Present

First day :

A. Amadon, R. Blair, R. Bock (part-time), P. Clarke, R. Cranfield, N. Ellis, R. Hughes-Jones, K. Korcyl, L. Levinson, I. Mandjavidze, F. Wickens, R. Scholte, J. Vermeulen

Second day :

A. Amadon (part-time), R. Blair, R. Bock (part-time), P. Clarke, R. Cranfield, R. Hughes-Jones, K. Korcyl, F. Wickens, R. Scholte, J. Vermeulen

Agenda

<u>First day, morning, chair F. Wickens</u>

1. Parameters and models :
    - 1.1 ATLAS specific parameters (e.g. number of ROBIns, ROB mapping, etc.)
    - 1.2 High-level full system models (building blocks : ROBComplexes, switches, procesor farms, supervisor, etc. and relevant parameters (e.g. network bandwidth))
    - 1.3 High-level models of test set-ups and relevant parameters
    - 1.4 Component models for the full system (e.g. model of the ROBComplex) and relevant parameters
    - 1.5 Component models of test set-ups and relevant parameters

<u>First day, afternoon, chair N. Ellis</u>

2. Implementations of models discussed in 1.1, 1.2, 1.3 and 1.4 and results

3. Workplan, including discussion of writing a Master Working Document (MWD)

<u>Second day</u>

4. Further discussions on parameters and models

1. Parameters & Models

1.1 ATLAS specific parameters

Update to DAQ Note 62

P. Clarke is working on an update of DAQ Note 62 concerning the organization of the read-out of the experiment and of the amount of and structure of the data generated by the various detectors. For the pixels + SCT + TRT the information is complete, for the calorimeter the information is essentially available – P. le Dû has some updates and has promised P. Clarke to provide these before he is away. A. Amadon is also involved.

P. Clarke prepared a table for the MDTs, S. Falciano has said that changes are needed, but that the necessary information can be provided quickly. However, more recently she has said that this may take a while. P. Clarke will urge her to provide the information required soon.

It was agreed that a snapshot is urgently needed, to be available by the February ATLAS week! It is less important that it is completely up-todate..

Update to DAQ Note 70

Information on a number of items is necessary for an update of DAQ note 70.

<u>RoI positions</u>

The positions of the RoIs provided by the LVL1 trigger are clear and lie on a 0.1 eta & phi grid for calorimeter and muon barrel, for the muon endcaps L. Levinson has provided the information required. This information is approximate but probably sufficiently accurate, it is not clear how to improve the accuracy.

<u>Trigger menus</u>

The trigger menus are problematic - it was suggested that S. George & A. Amadon work together to produce something as soon as possible (see also 3.1). P. Clarke queried why exclusive lists are needed, i.e. why the items in the menus in the TPR etc. are not adequate. N. Ellis explained that these menus are inclusive, but that an exclusive list is needed to specify the rate for each possible combination of RoI's. R. Böck pointed out that the trigger is dominated by background so the exclusive list for modelling can be quite short.

A. Amadon showed a single (large) page menu (mainly inclusive rates, but these were used to estimate exclusive rates). He also showed a similar menu for high luminosity (the spreadsheet with these menus is available from the modelling page). Some details need further work (e.g. assumptions about which items are independent/exclusive)

N. Ellis pointed out that the list contains jets with 15 GeV/c (off-line) threshold. J. Baines + D. Hubbard have found that a rule of thumb is that each extra jet reduces the rate by a factor root 2. This would change when going to higher momentum jets but no one has these numbers yet. It was asked whether J. Bystricki can help with this. Several people agreed that it would be better to use the jet threshold as given in the TPSR.

J. Vermeulen remarked that there is also a need to agree on reduction factors for sequential processing. A. Amadon has some of these (for the calorimeter from Saul Gonzalez).

N. Ellis suggested that Monte-Carlo figures could be used where available and that gaps could be filled in by (less reliable) interpolation and guess.

It was urged that the information should be available by the February ATLAS week.

<u>ROB mappings</u>

J. Vermeulen showed an eta-phi mapping for the SCT coming from P. Clarke (J. Vermeulen has a ROB's per RoI program). There are two extremes – mapping by layer and mapping by towers. Sarah Wheeler (UCL) is studying the more complex intermediate scenario's. There was much discussion about what is the important thing to optimize – it was agreed that this is probably the ROB RoI request hit rate. It was also agreed to use a simplified mapping for modelling for now (in particular for tower read-out of the SCT and pixel detectors), to let Sarah Wheeler present her results during the ATLAS week and next to iterate with the detector and trigger groups while keeping the mapping used for modelling studies fixed.

J. Vermeulen worried that the same problems we had with the calorimeter could be hit if a simple layer model is used, i.e. that we could mislead ourselves by taking a simple average for the RoI request rate in paper models.

A. Amadon showed numbers for numbers of ROB's and RSI's per RoI for both tower and layer mapping for the calorimeter (see his spreadsheet available from the modelling page). For the calorimeter there is also the question whether not using all of the calorimeter layers in the trigger brings advantages.

J. Vermeulen was concerned about the small numbers of the ROBs for the SCT and the pixel detectors (46 for pixels), independent of whether a tower or layer based read-out is used.

It was concluded that it is necessary to prepare the relevant numbers for both layer and tower options and to provide the results to the Detector Interface Group.

<u>Level-2 RoI sizes</u>

What assumptions should be used for the size of the RoI's ? Are the LVL1 RoI sizes used, or are they refined as L2 processing continues ? P. Clarke reported that T. Hansen had promised to provide some better numbers based on simulations. N. Ellis remarked that for muons it is necessary to allow for curvature of the tracks. Once the momentum is known a more restricted RoI can be used. However, its dimensions are no longer constant in eta and phi, but change with R!

What is the RoI data size ? N. Ellis : maybe this information is known in the detector groups. L. Levinson : what about RoI data selection from ROB? R. Böck : for the SCT whole wafers will be selected.

1.2 - 1.5 Component and system models and parameters

J. Vermeulen showed as an example parameters for the ROBin : 8 time parameters (can be set to 0 for simpler model). Ideally probability density functions should be used for each parameter (supported by simdaq). For the ROBout 6 time parameters are used.

A discussion on the desirability and usefulness of detailed models followed. F. Wickens remarked that we are not choosing between technologies and J. Vermeulen said that complex models can be transformed into simpler ones by setting time parameters to 0. This discussion gave rise to a discussion of the aims of the work on modelling. Persons present were asked to give their views.

- J. Vermeulen : the aim is to make a full generic model of the Pilot Project architecture and see if it works well. The paper model gives answers but it needs to be shown that these are supported by dynamic model. P. Clarke then asked to clarify the term Pilot Project architecture, answer : as shown, basically model C. F. Wickens changed this in terms rather of all processors connected to all ROBs.

- R. Blair : by the time of the TP it is necessary to show that we understand what we are proposing, so it is necessary to show that the model also reproduces the behavior of a small scale system in the laboratory. Also show that averages agree with paper model results.

- N. Ellis agreed with R. Blair. An existence proof of a working system is needed, complemented by results on the scaling behavior. He also sees modelling as a tool to make choices between sub-options. E.g. number of sequential steps, mappings etc – can all be done with generic models. Modelling provides also a tool that can help to understand the behavior of test-bed systems.

- P. Clarke : the aim is simulation of the full system, but not with more detail than is needed to show that ATLAS can handle the traffic patterns. Modelling also is a tool for making choices with respect to optimisation of the system. P. Clarke disagreed with the assumption that everything built has to be modelled : e.g. modelling ROBin's is a waste of time as this is not believed to form a bottle neck. Modelling of networks is however important. There is a problem with the testbeds – what needs to be modelled should be as close as possible to what finally will be done. However, for understanding the behavior of the test-beds e.g. details of Windows NT and/or LINUX have to be modelled, will this eventually be needed ?

- R. Hughes-Jones is interested in modelling of the details of components and mentioned that the DAQ should also be added to the full model.

- I. Mandjavidze : paper models should work at say 60% to allow for queues, then it is necessary to ascertain that the dataflow can pass through the system - look at individual components, especially the switch (to show that there is no data loss, contention, blocking, etc) - this cannot be technology independent. However, approaches used for a certain technology, for example use of congestion avoidance of ATM, could be used in a generic model. One cannot hope to get close to proper modelling of the testbeds in view of the complexity of the behavior of the operating systems (if a model is found to work it will only show what we could calculate anyway). The need to determine the overall latency very precisely is not seen since this will not be a critical issue. DAQ should be inside the model! R. Böck: the model should be kept as simple as possible : many parameters can be set to 0. J. Vermeulen : the DS link based test setup provided an example of how an initial set of parameters (for a model in which handshaking between data sources and destinations was ignored) were not good enough - it is better to have a fuller set of parameters with variable subsets set to 0. I. Mandjavidze : any important queues must be part of the simulation. Several people agreed. There was some discussion as to whether this includes queues in processors – at least whether they are important.

- R. Böck : sees a separate activity emerging – have a number of people interested in smaller items which need to be done in more detail. E.g. no one has mentioned the need to model processor farms. There are constraints on the architecture for which a full blown study is not required, e.g. a study of the ROB mappings for the SCT and the calorimeter. J. Vermeulen : this can be done with SIMDAQ, R. Böck : for me this is a problem (Note by JV : see also discussion on Ptolemy).

- K. Korcyl : what cannot be constructed for demonstration needs to be modelled. More emphasis needs to be put on reliability of modelling and scalability properties of models. Can the application test-bed be modelled with the paper model to check scalability ? The switch should be run at a load considerably below 100%. I. Mandjavidze : at 60% load a switch can block depending on the nature of the traffic flowing through it.

- F. Wickens found that the level of divergence of opinion was less than feared. He agrees that the major aim is to have a generic model of the full system – including important technology specific items that address scaling such as congestion avoidance. Also tools are needed to study individual items (mappings and detailed components). R. Cranfield noticed a divergence of opinions on how important it is to model the test-beds. F. Wickens suggested to draw up a list of tasks with priorities assigned to them and associated with these tasks, where the people involved should decide which tool should be used. N. Ellis remarked that more effort into paper modelling to study mapping options is required. He also noticed that several people have suggested that the model should include DAQ and that it is not agreed that DAQ traffic is included on the same switch used for the LVL2 traffic. F. Wickens said that L. Mapelli has invited modelling of the DAQ system by the LVL2 group – so we can proceed, ideally to model both the DAQ and LVL2 views. Nevertheless one has to be sensitive for the DAQ view in models run and in how results are presented. Modelling is not a high priority item for the DAQ group (R. Böck : L. Mapelli places less weight on generic models), but is also not considered to be a waste of time.

A discussion followed on the prospects of arriving at a common view with respect to the design and implementation of the DAQ and higher-level trigger systems and on possible delays in the completion of the TP. P. Clarke remarked that in his view the TP should contain a common view. N. Ellis commented that an important issue for the co-ordinators and the TDSG is to arrive at a common view. There are movements already as illustrated by the minimal coupling view and testbed (with two switches) proposed by L. Mapelli.

J. Vermeulen said that SIMDAQ has hooks already for modelling of DAQ aspects and that he is interested to model event building, both with a combined and separate switch.

2. Implementations of models

<u>A. Amadon : Paper modelling</u>

A. Amadon presented an update on the paper model for the pilot project architecture. He showed results on RoI rates for each subdetector, data request rates, data volumes rates and occupancies of ROBs, RSIs and switch (the sheets presented as well as the Excel spreadsheet are available via the modelling page). For low luminosity triggering the TRT full-scan presents a problem : if the TRT scan data goes across the network the RSI's of the TRT become very heavily loaded.

High luminosity triggering is not critical. Low luminosity triggering with B-physics with TRT full scan in a special-purpose processor (FPGA based) places stringent requirements on the SCT and pixel detector ROBs, but looks to be feasible if the number of tracks found by the scan is smaller than 20. If the TRT scan is done in the processors 15 times more processors are needed than for high luminosity !

It was discussed how firm the 20 tracks per scan is. Current studies indicate that the number is indeed high and hard to cut down. Although the B-physics trigger implies a large increase in number of processors, the cost is still within the current budget models, assuming the availability of cheap 2000 MIPS processor. P. Clarke remarked that the possibility of making the TRT full-scan within the ROB complex should also be investigated. Partial histograms could be filled within each ROB complex. This only works if the data is not spread across too many complexes (since histograms bins above threshold have to be transmitted between complexes).

L. Levinson commented that FPGA processors are also becoming more powerful and also that one can say that processors are becoming I/O starved. F. Wickens : these overheads will reduce for higher processor speeds- even if not scaling simply. R. Blair expected the difference between 200 and 400 MHz Pentium's to be known. J. Vermeulen noticed that the change in bus speed (from 66 MHz to 100 MHz) also has to be taken into account

K. Korcyl : Ethernet Switch Model

K. Korcyl described work on modelling (in OPTNET) of an Ethernet store and forward switch. Ports are attached to 256 MByte/s busses (vertical lines at the left and right in the figure below). These connect to intermediate 40 MByte/s busses. There are 8 32 Mbit/s links connecting two 40 MByte/s busses via the switch matrix to each other.

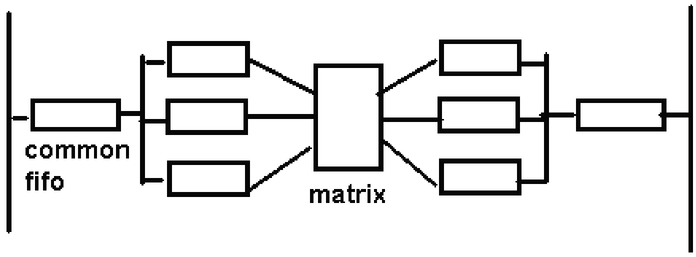

The optimised Linux software driver has an overhead of 28 µs. The latency on a packet through the switch is 4.24 µs. Two latency distributions are needed : one for the latency of messages staying within a switch board and one for those that have to pass through the matrix. The switch internals had to be deduced and were confirmed afterward to be correct by the manufaturer.

I. Mandjavidze : the results shown are for random traffic. What happens if a buffer in the switch starts to fill up. How does the model handle this ? Modelling should also be done with ATLAS like traffic.

P. Clarke : proposal for Ptolemy modelling

P. Clarke discussed the specification of a modelling framework in Ptolemy. The boundaries of nodes and switches need to be specified at a high-level. The message passing mechanism should refer only to simple messages. Each type of node should be specified by a document which defines the external boundaries (including defintion of the information content of messages + responsibilities) – note that this is tool independent. For the switches the initial idea was that they should have a common simple interface at the boundary. A detailed draft document is being worked on by K. Korcyl and R. Hughes-Jones)

All currently available documents can be found at :

http://www.hep.ucl.ac.uk/atlas/simulation

Ptolemy code is written for some of the models specified in these documents. B. Rensch was invited but declined to participate in this work (the meeting concerning the use of Ptolemy for modelling planned for the beginning of January in Copenhagen did not take place).

P. Clarke said that existing work should be looked at before embarking on the most difficult parts – e.g. look at the configuration etc. .

R. Blair asked whether this will lead to a wider effort. Answer : to be seen.

I. Mandjavidze asked about the performance of Ptolemy. J. Vermeulen estimates from B. Rensch's report a factor 20 more CPU performance required and a factor 7-10 more memory needed than for SIMDAQ (n.b. for sufficient accuracy for models with a high load of its components it typically is necessary to run for a day with SIMDAQ, B. Rensch needs about 128 MByte of RAM for the full model).

K. Korcyl has no intention of producing a full ATLAS model in Optnet- the current model contains 8 nodes and takes 2 hours CPU time. R. Böck said that he is interested in a simple model for an HPCN based LVL2 system. J. Vermeulen offered to try to incorporate into SIMDAQ (Note by JV : in an off-line discussion the next day it was agreed that R. Böck will specify properties of the HPCN system he is looking at and that R. Scholte and J. Vermeulen will look at how this can be turned into a configuration file for SIMDAQ. It is expected that, at most, only small changes in the code are necessary).

A. Amadon asked whether architecture C can now be modelled with SIMDAQ. Answer by J. Vermeulen : yes, the model consists of single switch talking to all processors, sequential processing is included, all processing for a single event is done in a single processor. Checking against paper model results (parameters of DAQ note 70) is being done : currently some discrepancies have been found but these are being resolved. The switch at present is a single switch, but SIMDAQ contains code for a switch fabric built from up to 3 layers of switching elements

A discussion followed on the scope of the work and the utility of using Ptolemy. P. Clarke : Ptolemy gives a comprehensive set of tools and documentation, but it does not have the waiting concepts of MODSIM.

R. Hughes-Jones got interested in Ptolemy to study switches and does not have a strong feeling on a wider use. He remarked that it is a lot of work to cross-check different tools and that this has to be done carefully. The question now is do we continue with the wider work and how different are the framework tools.

J. Vermeulen : If Ptolemy models for certain technologies provide state transition models coding these models for SIMDAQ is not so hard.

K. Korcyl is not planning to model the whole ATLAS in OptNet. Once the probability density functions are obtained these could be used in some other model.

3. Workplan, Master Working Document

3.1 Paper Model Update

The October 1998 milestone (parameter specification) is still only partially met, once all parameters are known updating the spreadsheet (milestone December 1998) is cumbersome due to symbolic names not being used (it was assumed that this is not possible) for referring to cells. L. Levinson commented that it is possible to use symbolic names.

It was decided to base modelling in the coming months on :

- the trigger menus found in A. Amadon's spreadsheet (to be updated later by A. Amadon together with S. George),
- RoI sizes to be provided by A. Amadon and T. Hansen,
- for ROB mapping a "regular" (simplified) tower mapping or layer mapping will be used for the SCT and pixel detectors, to be updated later; for the calorimeter the mappings used by A. Amadon are to be used,
- data volumes : fixed numbers are to be used first, then distributions: for the SCT and pixels these are to be provided by P. Clarke and F. Wickens; for the TRT by D.Froidevaux(?) ; for the calorimeter the numbers are fixed (note added by JV : but will be different for different ROBs) ; for the muon detectors S. Falciano has to be asked (but a Poisson distribution can be used), in view of work on the NIMROD J. Vermeulen may be able to find the required information at NIKHEF.

3.2 Full generic models

J. Vermeulen reported that in SIMDAQ all components of the full generic system have been implemented and tested, as specified as milestone for December 1998. Also the documentation has been improved (as agreed on in Chamonix to be done before February 1 1999). An original milestone was October 1998 for sequential processing .: 1/1999 : the RoI rates agree with the paper model. It is now possible to work with the program for a variety of system architectures, work on comparisons with the paper model is in progress. Version 4 of the program and associated documentation are available from the modelling page.

N. Ellis asked how to include the DAQ in this model and how far to go. F. Wickens commented that it is necessary to distinguish between the LVL2-SFI and EF-SFI. J. Vermeulen : they have different names in the program.

3.3 Generic test system models

The original milestone was October 1998 for having a model and results available for the DS-link based model B demonstrator system : this has been done but has not yet been documented (this was put aside to get SIMDAQ documented).

It was discussed whether this (modelling the model B and model C demonstrator systems) should still be done. P. Clarke pointed out that some people who were now saying that this should be dropped (which he agrees with) argued strongly against him a year ago and he would warn against such wasted effort in future.

J. Vermeulen suggested the use of Ptolemy to model the test-beds, but N. Ellis was concerned about using different tools for the full scaling and checking against the test-beds. P. Clarke questioned the value of modelling the test-beds since they are so far from optimised hardware plus software. In his absence the meeting was unanimous in agreeing that they should be modelled. I. Mandjavidze had earlier also questioned whether the test-beds could be modelled accurately enough. R. Böck was concerned that if these could not be modelled how one can be confident about modelling the full system.

P. Clarke undertook to evaluate the application test-bed set-ups with respect to what aspects can sensibly be modelled. K. Korcyl expressed an interest in modelling the application testbeds.

R. Blair is still interested in modelling the ATM based model C demonstrator (Note added by JV : after the meeting this was discussed by RB and JV and it was found that everything needed - due to the work on the demo B demonstrator setup - is available in SIMDAQ. A configuration file specifying the system to be modelled was produced and debugged).

3.4 Technology models

K. Korcyl and R. Hughes-Jones are at present the main players for Ethernet switching, I. Mandjavidze for ATM. J. Vermeulen likes to have information to be included in the Master Working Document to be produced (see next item) by the end of March 1999.

3.5 Full models

To be planned, it is still too early for this point

3.6 Master Working Document

J. Vermeulen suggested to produce and maintain a Master Working Document (1st iteration end March 1999, 2nd iteration 2 weeks before next modelling meeting (combined with ROBComplex meeting ?).

4. Further discussions on parameters and models

During the second day of the meeting the components and associated parameters to be used for a generic modelling of the Pilot Project system were discussed. It has been attempted to define a baseline for all relevant items, so that initial modelling can be done. It is accepted that the model is only an approximate one and that refinements will have to be included later. However, it is hoped that critical components / parameters can be found. The choices made are conservative ones, so that the model results may be too pessimistic and in reality a better performance may be found. If the performance with the model as defined is satisfactory it therefore should be likely that in reality the performance should also be as required.

4.1 Switch models

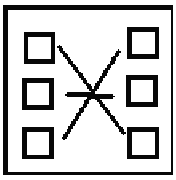

Currently SIMDAQ has the above concept for a switch. The crossbar switch has input and output buffers for each port and arbitration for access to the output buffers. Switches can be put together into a fabric of at maximum 3 layers.

It was decided that a single large store and forward switch as currently available in SIMDAQ will be used initially, in case of problems a fabric built from 64 port switches can be used. The ports are bi-directional: this is emulated in SIMDAQ by using pairs of switches where the ports on one of the switches are paired with ports on the other switch in input-output pairs. The underlying assumption is that the interaction between data streams flowing in opposite directions is negligible.

For the baseline link speed 15 MByte/s will be used, as is the case in A. Amadon's paper model. The speed for an internal connection is also 15 MByte/s.

4.2 Supervisor

It was decided to model the supervisor with 5 independent RoI processors. There is no "steward" handling decisions and processor assignment as in previous designs. Assignment of an event to one of the processors in the LVL2 farm (SFIs are not used, see below, but also if they would be used processors rather than SFIs would be assigned) is done on a round-robin basis by each RoI processor independently. An LVL2 decision is sent from the LVL2 processor to the RoI processor which assigned the LVL2 processor. Decision blocks are sent by the RoI processor to the ROBOuts as a multicast message (i.e. the switch broadcasts the decision block to all ROBOuts connected to it). The size of the decision blocks is determined by requiring that there is on average one accept message per block. This leads to smaller decision blocks than used until now, but reduces the effect on the average time needed for a decision to reach the ROBs caused by having 5 processors instead of a single processor collecting decisions.

4.3 SFI

In the paper model of A. Amadon there is one SFI per processor (to handle I/O and merging of data into single RoI fragments ready for handling by the FEX algorithm). It was decided to assign the SFI function to the processor taking care of the LVL2 algorithms and to model without SFI’s.

4.4 LVL2 farm processors

The LVL2 processors take care of preparing input for feature extraction (originally part of the SFI function, see above), FEX + global processing and of RoI request formulation (or "steering" in the sense of the reference software). In SIMDAQ the "ModelCProcessor" object has the required functionality. The various processing tasks are modelled in SIMDAQ as separate processes connected via queues. Each process, once handling an event, can only be interrupted by a higher priority process becoming executable (i.e. there is no round-robin scheduling). The alternative model where each event is assigned to its own thread makes no difference if queues take no time, as the required amount of context switching is the same for both models when all processing tasks in the first model are given the same priority.

It was decided to use the processing times of DAQ note 70 for the different processing steps, but scaled to 1000 MIPS, i.e. reduced with a factor of 2. A. Amadon scaled to 2000 MIPs in his spreadsheet, but although this may prove to be realistic it was preferred to be somewhat more conservative. In DAQ note 70 the total processing time was considered to include an overhead of 10 %, therefore the nominal processing times were increased by 11.1111 %. It was decided to use the same convention. The processor time needed for fragment merging is determined by a merging speed of 80 MByte/s. For the context switching time 10 microseconds is to be used. For SIMDAQ two numbers have to be specified (e.g. 5+5) to allow for cases when a process finishes, but nothing else replaces it or when a process starts, but another process does not need to be descheduled : one number specifies the time needed for scheduling a process, the other number the time needed for descheduling.

4.5 ROBComplex

The ROBComplex consists of a number of ROBIns connected to a single ROBOut. It has been decided that in the model each ROBIn has a direct connection to the ROBOut (i.e. the ROBIns and ROBOut do not connect to a shared bus). These connections transfer data with a speed of 80 MByte/s. In the ROBOut data from different ROBIns but corresponding to the same RoI are merged into a RoI fragment with a merging speed of 80 MByte/s. For each input and output for theROBOut and for the ROBIn there is an overhead of 10 µs processing time (except for event data arriving in the ROBIn), this can be treated as context switching time. For fanning out RoI requests to ROBIns 10 µs processing time (independent of the number of ROBIns)+ 10 µs I/O overhead per RoI request is needed. The model should allow to specify a processing time per destination, which has been decided to be set to 0 µs initially. For distribution of decision blocks the same numbers (10 µs processing time (independent of the number of ROBIns) + 10 µs I/O overhead per decision block) are also to be used.

The numbers of ROBIns per ROBOut are the same as used by A. Amadon in his spreadsheet. Indexing an event fragment received via the ROL (Read Out Link) in the ROBIn takes 5 µs per fragment. For decision handling the ROBIn needs 1 µs per decision + 10 µs I/O overhead for receiving the decision block. The event data will be transferred under DMA control from the buffer memory to the output link with a transfer speed of 80 MByte/s. Support for this type of DMA transfer needs to be included in SIMDAQ. For each RoI request received by the ROBIn 10 µs processing time + 10 µs I/O overhead is used. It is assumed that there is no pre-processing and no data selection in the ROBIns. Also for jet RoIs the full data is transmitted (this is different from what has been specified in DAQ note 70).

Data is transmitted to the Event Builder by the ROBOut after receiving a LVL2 accept decision. This has to be adapted later to handling requests from the Event Filter

As is the case for the processors 10% has to be added to the nominal processing times because of overhead (for comparison with paper model - unfortunately added as 10% of total!).