Hyper-K DAQ: Supernova burst monitoring Technical Note

B. Richards, H. M. O'Keeffe, T. Dealtry, T. Doyle

April 2019

1 Introduction

The data acquisition (DAQ) systems for Hyper Kamiokande (Hyper-K) undertake a broad number of tasks including data transfer, buffering, persistification, monitoring, slow control, cluster control, electronics and networking. At present, there are 3.35 FTE working on the Supernova DAQ and burst identification task. These are supported by additional FTEs from undergraduate students.

This note will describe the software, networking and hardware systems that are responsible for the flow of data from electronics readout to persistification on disk. To facilitate this description 3 scenarios of different data types will be tracked through the system as they express the most likely and edge cases. These data types will be normal data taking, a near supernova and a far supernova. As detailed in the Hyper-K design report [1], a near supernova is defined as being at 0.2 kpc and far supernovae are those up to distances of 4 MPc.

The DAQ system will be built using ToolDAQ [2] a modular, highly scalable and fault tolerant DAQ framework, which will be installed on each of the commercial computing server nodes that form the DAQ system. These servers will be connected to commercial network switches for transferring data between them. Layers of network, networking and messaging protocols as well as redundant connections and buffers are used within ToolDAQ along with dynamic service discovery to increase fault resilience, tolerance and to dynamically re-route and correct faults as they appear. A simplified model of the DAQ system can be seen in Fig.1, which shows the front end electronic (FEE) boards and four of the major nodes that exist in the DAQ. These are the readout buffer units (RBUs), the trigger processor units (TPUs), the event builder units (EBUs) and the Brokers.

The basic data flow of the system starts by buffering the data as rapidly as possible on the RBUs from the FEEs, where it is than catalogued. From here the Brokers manage the data flow but do not act as intermediaries in the transfer. They assign jobs to the TPU and EBU worker node to analyse or save data respectively, maintaining records of the results and job assignments. The TPU and EBU workers then request the data for their assigned jobs independently from the RBUs.

The networking layers between the each of the nodes and the FEEs (shown in Fig.1) form two physically separate networks. One network will be referred to as the front end network and it connects only the FEEs to the RBUs via network switches. The other network, connects the RBUs, TPUs, EBUs, Brokers and other archiving, monitoring and control systems. The main role of the front end network and computing resources of the RBUs is to readout and store the data from the FEEs, buffering it on disk and in memory for as long as possible (~10 hours for normal data, longer if required, supernova data more permanent). The back end network is responsible for processing the hit data stored in the RBUs and saving it to disk via the TPUs, EBUs and Brokers.

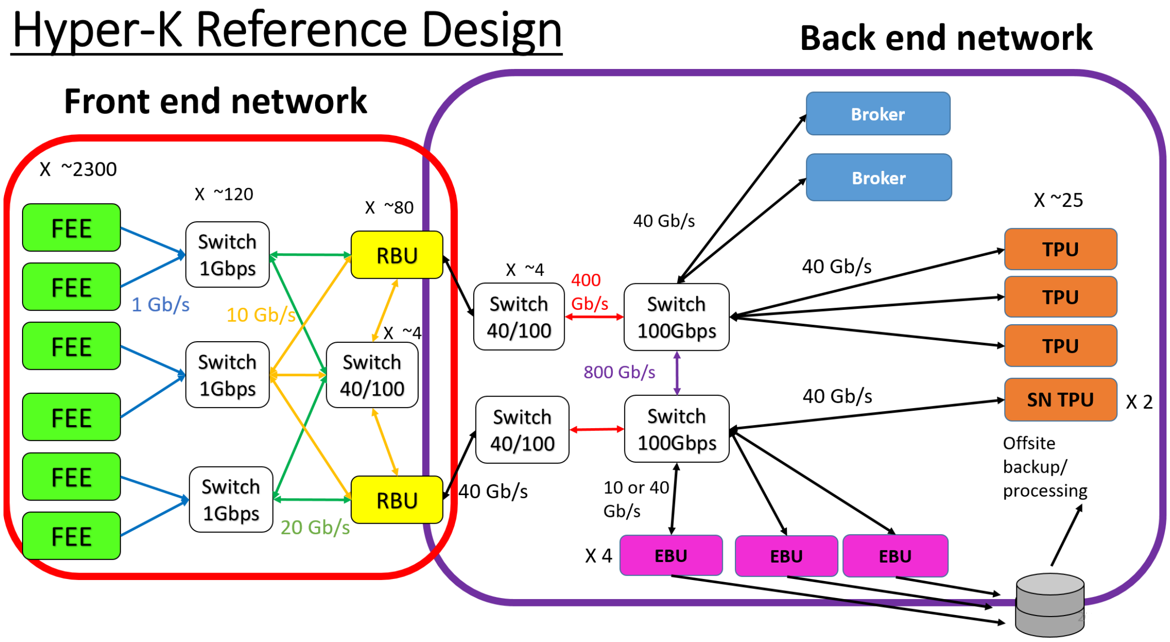

Figure 1: Network reference diagram of the DAQ system.

2 Data Readout and the Front End Network

As mentioned previously, all of the challenges to readout the raw data from the electronics to disk are encompassed by the front end network. This has to have a dead-timeless readout of the FEEs and store the data on disk as quickly as possible where it can stay for a longer period of time to allow processing in the back end network. At present, multiple configurations of detector PMT geometry currently exist, which require anything from 2000-2500 FEEs depending on the numbers of PMTs. However for this note only two configurations that represent the maximal hardware and data challenge will be shown and are known as the 'normal' variant with 40,000 20" PMTs for the inner detector (ID) and the 'investigated alternative' variant or 'hybrid configuration' with 30,000 20" PMTs and 7,700 MPMTs for the ID (whose values will be shown in square brackets). Both variants also include 15,000 PMTs for the outer detector (OD) and the ToolDAQ system is built from modular and scalable blocks so the configuration can be expanded or contracted for any chosen configuration.

The FEEs are designed to provide readout for 24 channels each for the ID and OD and 456 channels for the MPMTs (separated into 24 MPMT units), this gives ~2300 FEE boards in total for the configurations. Each of these boards is fitted with six 1 Gb/s optical fibre connections, of which 4 are used for data transfer and slow control combined, and 2 for clock and synchronisation (alternate designs using seperate data and slow control connections have also been considered). The front end DAQ network only connects to the slow control and the data connections, with the clock and sync provided by another, independent, system. Of these four connections to the front end network, only two are connected with one active and the other as backup. The other two are disconnected and would only to be used in case of connection failure for redundancy. This gives ~4600 1 Gb/s optical fibre (SFP) connections. The amount of data per FEE is estimated based on a single hit requiring 12 bytes (Tag/Flags[16 bits], Channel[16 bits], Time[31 bits], Charge[16bits], Fit Quality[16bits]), where the fit quality is a number defining the accuracy of the electronics on the FEE to sit the wave form with a standard distribution to determine the time and charge correctly. If this fails, the full wave form will be sent from the FEE to the RBU which will be 24 Bytes (Tag/Flags[16 bits], Channel[16bits], Time[31bits], Waveform[12 bits x 10 samples]), we would expect this to occur for no greater than 5% of hits. Scope for this hit data to be reduced (both waveform and T&C) exists with data compression occurring on the FEEs. However studies are underway to determine how successfully this can be done in a lossless way and what are the processing overheads to decompress it.

During each of the three data scenarios (normal data taking, near supernova and far supernova) the data rates from each FEE are expected to be ~8 MB/s for both normal data taking and for a far supernova and for a near supernova ~117 MB/s in the first second and ~189 MB/s for the next 10 s. This gives a total data output per FEE of ~306 MB for a near supernova and a total detector output of ~560 GB [530 GB]. All data is sent via TCP/IP protocol eliminating the loss of data packets in the network and switch hardware. During normal and far supernova data taking periodic calibration data will be sent that will not effect the data flow, however during a near supernova event when data rates increase, this calibration data will be neglected or demoted to prioritise the supernova.

Near supernova therefore present the greatest data rate challenge and as such the system must been designed to cope with this. The 1 Gb/s data link from each FEE will support a maximum of 125 MB/s of data throughput, which is u98umore than the 117 MB/s peak rate expected from a near supernova (this peak rate increases to 235 MB/s depending on PMT location, so buffering on the FEEs exist to handle this momentary saturation in the first second of a near supernova). However the rest of the front end system, the networking layer and RBUs also have to be able to handle the peak connection rate and continue to write data to disk in a dead-timeless way. In order to achieve this, roughly 80 RBU servers are required, with each RBU responsible for the readout of up to 30 FEE boards. The RBUs will therefore be expected to deal with input data rates of ~240 MB/s data during normal and far supernova operation and a peak rate of ~3.5 GB/s during a near supernova. To achieve this each RBU will have four 10 Gb/s optical SFP connections connected to the front end network. This gives a total connection speed of 40 Gb/s and therefore total data rate of 5 GB/s (roughly 1.5 times the peak supernova rate), in fact the maximum output of the 30 FEEs to saturate their 1 Gb/s links is 30 Gb/s or 3.75 GB/s. This means that the RBU links have been designed such that as fast as the FEEs can physically supply data, that data can reach the RBUs without bottle neck. However the lack of bottlenecks is not the only factor in the design of the network layers and links. The front end network also has to have redundancy and be as resistant as possible to hardware failure. Standard networking wiring tree schemes, where links are aggregated through a pyramid of ever increasing performance switches would therefore not be suitable. Their single point of failure could causes massive and catastrophic data losses, as much as 50% on a single switch failure. Therefore a scheme has been designed which eliminates single switch failure (or minimises it to a maximum of 0.8% if separate slow control and data connections), has redundancy for RBU link failure and can support the data rates required to reroute data to alternative RBUs if an RBU server fails. This rerouting of data under server failure is a software feature of ToolDAQ made possible by the dynamic service discovery and fault tolerant communication layers. This will allow the system to discover a failure in the hardware automatically and send instructions to rewire and reroute the data to other RBUs to cope with the hardware loss. The network wiring scheme can be seen in Fig.2.

To achieve the redundancy and single switch failure prevention, modular blocks of 2 RBUs and three commercial network switches of 48 x 1 Gb/s SFP ports with 4 x 10 Gb/s SFP+ ports are used. Each of the three switches have SFP connections of 20 FEE boards and redundant connections from the next set of three switches (40 connections) connected to them giving 60 actives FEEs across three switches. As each RBU reads 30 FEEs, each RBU has sole use of one of the switches and share half of the 3rd switch. They connect to these switches using the 10 Gb/s SFP+ ports with two links to the dedicates switch and one to the shared switch allowing full speed readout of each of the FEEs

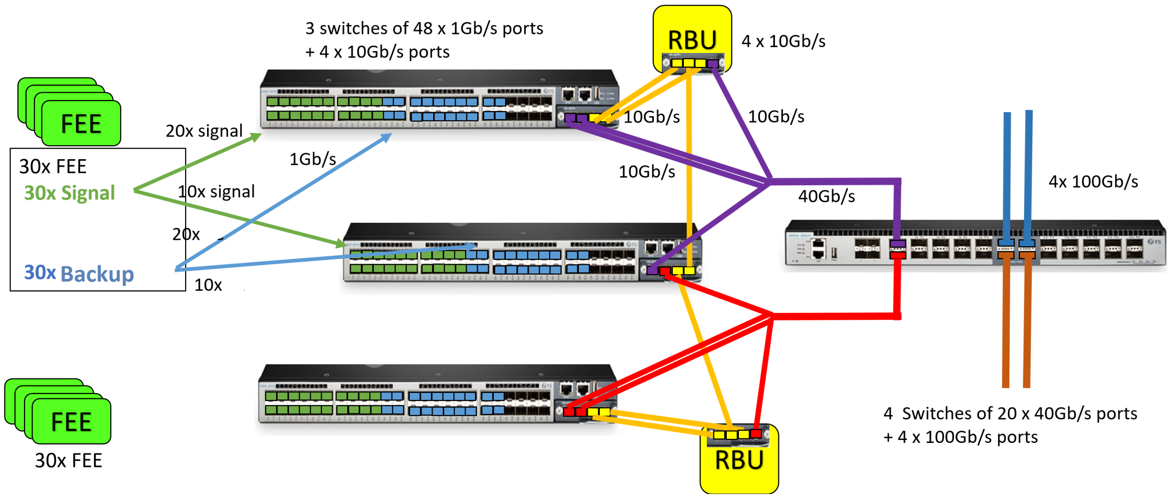

Figure 2: Front end wiring scheme showing readout of 60 FEEs via three switches[centre] ($\frac{1}{40}th$ of full detector) and one of 4 switches [right] which provide the redundant network for the front end data rerouting

connected. In total there would be 40 such blocks of RBUs and switches to read out the full detector. The remaining two SFP+ ports on each of the three switches and the 1 SFP+ port each of the RBUs are then connected via a 40 Gb/s (QSFP+) → 4 x 10 Gb/s (SFP+) breakout cable which is attached to another network switch which has 20 x 40 Gb/s (QSFP+) ports and 4 x 100 Gb/s (QSFP28) ports. There would be four of these higher spec switches which forms the redundant backbone of the front end readout network. They allow for the 30 FEEs connected to an RBU to be shared between 20 other RBU on failure of an RBU. If link failure to an RBU occurs then it can also act as an alternate route for data from the FEEs as well. These extra switches allows for 20 RBUs to be interconnected for redundancy as it stands, which is $\frac{1}{4}$ of the detector readout. This should be sufficient as it is unlikely that more than a $\frac{1}{4}$ of the detectors systems would fail, however further redundancy could be achieved by either inter connecting the 4 x 100 Gb/s (QSFP28) ports in a ring between the 4 switches with 2 in either direction aggregated, or by the addition of another 32 port 100 Gb/s switch. This would allow any RBU to read out any FEE in the entire detector.

Once the data reaches the RBUs it then needs to be buffered, organised and catalogued for processing and long term persistification and archiving. The RBUs achieve this via threaded dead-timeless readout of the FEEs sockets to the front end networks. Dedicated threads are used to simply pull the data from the buffered network sockets straight into RAM with very low latency. At this point the path of normal data and some supernova data will bifurcate.

During normal data taking the incoming data will contain both time and charge information as well as possibly occasional waveforms (this is dependent on the choice of electronics). This data will be organised and buffered in RAM for ~100 seconds (at 240 MB/s) which requires >24 GB of RAM. The data then will be shunted to magnetic hard drive storage for longer buffering of multiple hours (e.g. 10 hours = 8.7 TB). Magnetic drives can typically store 250–300 MB/s, so a single drive is sufficient for this. However with a second drive we have added redundancy and/or greater headroom for the write speed.

Most processing and event readout will occur within this 100 s window and so the RAM storage will be sufficient, however for the ability to look back at data outside of this range which may occur with a delayed external supernova trigger or if extended periods before/after a supernova are later required the magnetic drives are used. Additionally, if trigger reprocessing or processing delays occur further downstream in the data path, this storage can act as an extra buffer. It also serves as a way to recover data for a period where an RBU might fail. In the event of an RBU failure, the data for that period would not be lost; it would be stored on the drives and could be recovered at reboot or rerunning of the process.

During a supernova, data will take a different path. For a distant supernova detection of the supernova will not occur until the TPUs have analysed the data and the amount of data is an insignificant increase over that of normal data taking. Therefore the data will follow the same path as described for normal data, however when the trigger processing systems determine a supernova has occurred a message will be sent to the RBUs to record a copy of that period around the far supernova in a separate more long term storage on an SSD.

For a near supernova, the data rates will increase massively as previously described in [1]. This should be apparent without further trigger processing at the RBU level. FEE boards will stop forwarding waveform data if loads get too high and the RBU will send a message out to the Brokers alerting them to the situation of high rate on their channels. This could be cross checked between RBUs by the Broker and a validation sent back to the RBU that a supernova is occurring as well as with the SNTPUs. Whether it receives validation or not, the RBU will be able to cope with and respond to the increase in data rate. At peak rate for a near supernova the RBUs will receive 3.5 GB/s from the 30 connected FEEs during a near supernova, with a total supernova data size of 9.2 GB per RBU for the full ten seconds. Therefore the minimum ~22 GB of 100 s buffer in RAM is more than sufficient to hold all of this data which would take priority. The RBUs will have a spec with a minimum of 48 GB of RAM, which actually is sufficient to cover both the 100 s buffer for the near supernova data and other processes of the server. However despite this RAM buffer the supernova data will also be copied to disk as soon as it arrives for a persistent storage. The peak rate of 3.5 GB/s is beyond the performance of magnetic drives storage but NVME M.2 SSDs can write at 2-2.5+ GB/s and with each RBU only required to store 9.2 GB per supernova, a 1 TB or 500 GB drive which are reasonably priced and readily available could store 54 or 108 near supernova bursts on a single drive. So one of these would therefore provide a dedicated supernova storage disk and have the speed required to record them for indefinably long periods (this would be the same SSD used for the far supernova after detection). If extra backup overhead speed is required two of these drives could be used RAIDed giving 4-5+ GB/s write performance which could not only cope with the very peak near supernova rate, but also cope with full saturation of the 30 FEE links (3.75 GB/s max) for 3-6 mins. Important to note however is that due to the link speed and fast writing abilities of the RBU and network, a near supernova will cause no dead time on data recording from the detector by the RBU, so data directly after is still being saved and buffered, even if two near by supernova where to occur one after another.

3 Data Processing, Archiving and the Back End Network

Once data is received in the RBUs it is somewhat secure for a semi-permanent multi-hour time period (with supernova data permanently stored on the SSD). Despite this, the goal of the TPUs, EBUs and brokers are to analyse and write out the useful data in real time and certainly within the 100 s RAM buffer. The design of the system however means that the RBUs operation is not dependant on anything in the back end network, so any failures or problems within it be they TPU, EBU, broker, hardware or network issues the RBUs will continue to run and buffer and store data without loss.

In order to organise and coordinate the processing and storage of data in the back end network we have decided to use a broker assisted paradigm with two server nodes known as the brokers. As with the front end network each node (TPU, EBU or broker) is independent, so failures of any node in the system will not cause cascade failures in software or processes running on other nodes. This means that on most occasions single or multiple nodes failing will just cause the ToolDAQ system to rewire itself to re-adjust to the loss seamlessly and nodes can be rebooted and will automatically join back in with the system without problem, human input, re-configuration or loss of data. This is achieved for the TPUs and EBUs by having them operate as worker nodes in a pool managed by the brokers that are assigned time windows for trigger processing and event building and writing to disk respectively. The built in dynamic service discovery allow for each node to know when others on the system go down and broadcast their presence when they restart to be re-added to the pools of workers by the brokers. The Brokers therefore, despite not doing any of the processing work, are the organisers and managers of the system keeping track of which jobs have been assigned to which nodes and then collating results of those jobs and re-issuing them if need be to other nodes if failure occurs. This does however mean that loss of a broker in the system is a larger problem, which is why two identical brokers exist for redundancy with a master slave relationship. The slave broker keeps track of all decisions of the master and is ready to jump in and switch roles to the master if the current master falls over. The time widows for which the TPUs are assigned are overlapping to allow for edge cases and are decided on in advance of the data being produced and sent out to the TPUs so they are ready to process as soon as the data is captured. The broker handling the job distribution for the TPUs and EBUs mean that the system is load balancing with new jobs being sent out to workers that finish their current jobs, this automatically allows for the fact that some trigger decisions are a lot more complicated and require more processing time than others.

3.1 TPUs and Trigger Processing

Upon receiving a job in the form of a trigger window of ~1 ms to process from the broker, the TPU then will send a request to all of the RBUs to send data for that relevant period to the TPU. Not all data will be sent to the TPU only the charge and time information not the waveforms as this data will require further processing offline and is not useful to trigger decisions. This data will be ~5 MB [4.2 MB] per 1 ms time slice and with some ~25 TPUs that works out to 197 MB/s [169 MB/s] for real time processing. 197 MB/s requires a link speed of 1.6 Gb/s, but to allow redundancy for TPU failure and to increase the transfer speed so more time is available for processing 40 Gb/s links to TPUs have been incorporated meaning that 1 ms of data can be transferred in <1 ms.

Once the data has arrived at the TPUs it is then passed through a sequential list of trigger algorithms and reconstruction of increasing complexity; when a positive trigger condition from one of these algorithms is met, the sub time window within the 1 ms responsible for it is flagged and later algorithms may not run on that section. These algorithms make use of GPUs for accelerated processing and are organised using the ToolDAQ framework which they are written as modular blocks within.

The baseline for identifying supernova bursts in Hyper-Kamiokande will be the same as that used in Super-Kamiokande. This will be outlined below and is based upon that detailed in [3]. Extensive testing of the Super-K strategy is currently taking place and results should be known in advance of, and therefore included in, the next draft of the technical report. Some minor changes to the Super-K strategy are anticipated for Hyper-K and these will be discussed below.

3.1.1 Super-K real-time burst monitor

The Super-K real-time neutrino burst monitor was developed using three supernova models

- Wilson Model [4]

- Nakazato Model (NK1) with $M = 20\,M_{\odot}$, $t_{\text{revive}} = 200\,\text{msec}$ and $Z = 0.02$ [5]
- Nakazato Model (NK2) with $M = 13\,M_{\odot}$, $t_{\text{revive}} = 100\,\text{msec}$ and $Z = 0.004$ [5]

where $M_{\odot}$ is the progenitor mass in units of the solar mass, t_{revive} is the shock revival time and Z is the metallicity. Neutrino oscillations were taken into account using models from [6] and values from [7]. Three possible neutrino oscillation cases were simulated for each of the three supernova generators, leading to a total of 9 Monte Carlo samples. The oscillation cases considered were:

- No oscillation
- Normal hierarchy
- Inverted hierarchy.

These samples were used to test the performance of the real-time supernova burst monitoring system.

Data recorded by the data acquisition systems is processed by the event builder, packed and the data stored in a sub-run file. Each sub-run file corresponds to approximately one minute of detector activity. Copies of each sub-run file are sent offline and to the supernova burst monitoring system. The offline files are used for standard analyses and are retained indefinitely. The supernova (SN) monitoring system runs on a single computer. Each sub-run file received by the SN monitoring machine is first converted to the format used by the offline analysis and this reformatted file is used as the input to the event reconstruction process, which outputs the vertex position, vertex time, direction and energy for each event. The reformat and reconstruction takes approximately 2 minutes per sub-run.

Using the reconstructed information from the sub-run, events with $E_{\text{total}} > 7\,\text{MeV}$ and with a vertex position inside the 22.5 kton fiducial volume are selected. Cosmic ray and Michel electron events are removed. For each event that remains, a 20 second time window extending backwards from the event's time stamp is opened. The number of events within this 20 second window, N_{cluster}, are counted, taking into account any sub-run boundaries.

For each 20 second window, a second variable, D is computed. The D variable characterises the dimensionality of the distribution of vertices in the time window and assigns an integer value of $D = 0, 1, 2$ or 3, depending on whether the distribution is point-, line-, plane- or volume-like respectively. The D parameter is computed as follows

1. Collect positions of event vertices in the window
2. Calculate the average position of the events
3. A covariance matrix can be computed using the position of each vertex and the average vertex position
4. The eigenvalues and Trace of the covariance matrix are calculated.
5. The χ^2 for each of the line, plane and volume hypotheses can be calculated using the eigenvalues and the trace [8]
6. From these three values, the minimum value of χ^2 is chosen and the value of D is assigned accordingly.

Once the D value and the number of selected events in the 20 s window (N_{cluster}) have been calculated, they are compared to three levels of supernova warning.

- Golden: $N_{\text{cluster}} \geq 60$ and $D = 3$
- Normal: $60 > N_{\text{cluster}} \geq 25$ and $D = 3$
- Silent: More than 13 events in 10 s

In the event of a golden warning, experts are called by phone and meet in order to make a world-wide announcement within 1 hour. For a normal warning, a notification is sent to the supernova monitoring service (SNEWS) and experts via email. If a silent warning is issued, notification is sent to a few experts only and no notifications are issues outside of the collaboration.

The golden warning is set such that there is 100% supernova detection efficiency at the Large Magellanic Cloud and the normal warning is such that there is 100% detection efficiency at the Small Magellanic Cloud. The efficiencies were derived from the Monte Carlo samples described at the beginning of this subsection. Results are detailed in [3].

In the Super-Kamiokande burst monitor, the supernova direction is determined using a likelihood fit. As the direction of the outgoing electron is strongly correlated with the direction of the incoming neutrino, elastic scattering events are used for this purpose. A set of probability density functions (pdfs) are computed using Super-K Monte Carlo simulations. These pdfs contain information about the reaction channel type, energy of the neutrino, direction of the event and direction of the supernova. There are a total of 14 parameters used in the fit. The initial value of the supernova direction used in the likelihood fit is obtained using a coarse grid search. At each step of the search, the number of events with $\cos\theta_{SN} > 0.8$ are computed, where $\cos\theta_{SN}$ is obtained from the dot product of the event and test supernova direction. Details of the performance of this algorithm are outlined in [3].

3.1.2 Towards a burst monitor for Hyper-K

As stated in section 3.1, the Hyper-K supernova burst monitor will use the Super-K monitor as a baseline. Studies are in progress to ascertain whether any major changes to the Super-K method are required. As the Wilson model is outdated, the studies for Hyper-K will focus on the use of the Nakazato model. For the triggering studies, we use the model parameters that give the lowest event rate in order to be conservative, namely:

- Nakazato Model with $M = 13\,M_{\odot}$, $t_{\mathrm{revive}} = 100\,\mathrm{msec}$ and $Z = 0.004$ [5] and no oscillations. This will be referred to as the NK2NO model.

It should be noted that for calculating data rates throughout this note, the model parameters were chosen such that the largest event rate were chosen:

- Nakazato Model with $M = 30\,M_{\odot}$, $t_{\mathrm{revive}} = 300\,\mathrm{msec}$ and $Z = 0.02$ [5], with normal hierarchy oscillations for the highest rate in a 1 ms window;
- Nakazato Model with $M = 30\,M_{\odot}$, $t_{\mathrm{revive}} = 300\,\mathrm{msec}$ and $Z = 0.02$ [5], with no oscillations for the highest rate in a 1 s window;
- Wilson Model [4] with inverted hierarchy oscillations for the highest rate in a 10 s window.

Once positive sub-windows have been determined by the TPU these will be communicated to the Brokers along with a completion statement which will trigger the next time window to be sent for processing. As well as communicating with the Brokers, the TPUs will also communicate with a pair of dedicated supernova TPUs (SNTPUs). The normal trigger algorithms that run on the generic TPUs will look for evidence of a supernova on the 1 ms time slices processed, by examining the rate of triggers in the incoming data. If the rate of triggers observed in a nominal 1 ms slice of data exceeds the usual rate in a statistically significant manner, a warning will be issued and the RBUs will begin to write all data to their SSDs. The threshold for making such a decision will be determined using Monte Carlo studies. To permit monitoring of supernova (both distant and near) whose data rate is below the RBU threshold, the TPUs will pass the 1 ms time slice to a reconstruction programme which runs as part of the normal TPU trigger algorithm. Whilst it is possible to have longer time slices, there will be

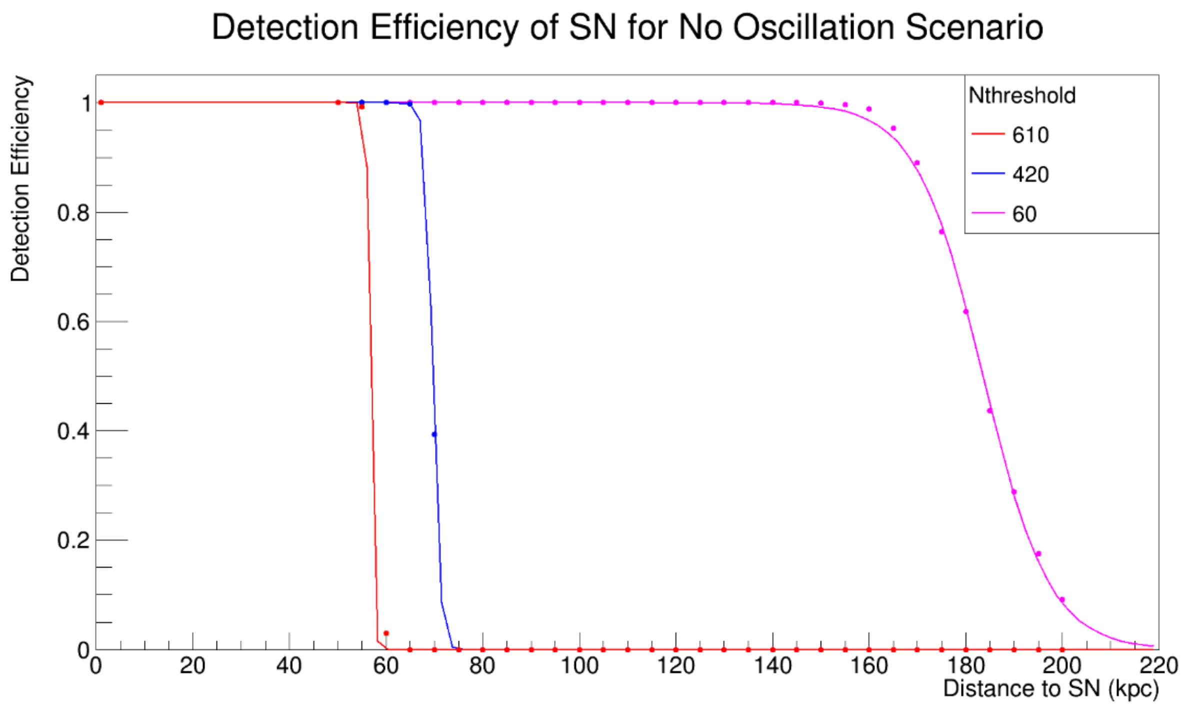

Figure 3: Detection efficiency for a population of 3000 simulated SN at each distance. The NK2NO model, corresponding to a conservative estimate of the event rates. Fiducial volume and $E_{\mathrm{true}} > 7\,\mathrm{MeV}$ cuts are used. The red, blue, magenta lines are shown for N_{cluster} (Nthreshold) values corresponding to 100% efficiency at the LMC, SMC, and LMC in SK, respectively.

a point at which the graphics card memory will not be sufficient to store the time slice data. This optimisation will be performed a full simulation of the time and performance of complete collection of trigger algorithms is available.

The direction, vertex position, energy, and time of the event will be returned and passed from the TPU to the SNTPUs. The SNTPUs, will collate multiple time slices of data to form a 20 s monitoring window. Two dedicated supernova nodes exist for redundancy and both will be active so identical cross checks can be made to ensure the data is not missed. Upon a positive supernova determination messages will be sent to all nodes including the RBUs to transfer the data to their long term SSD storage.

On the dedicated supernova TPUs, the total number of events with energy greater than 7 MeV in the fiducial volume and in the monitoring time window (nominally 20 s in Super-K) will be calculated.

3000 SN were simulated using the NK2NO model every 5 kpc from 50 to 200 kpc with sntools [9]. Using this information, values of the N_{cluster} cut required to provide a 100% detection efficiency at 50 kpc (LMC) and 60 kpc (SMC), corresponding to how SK set thresholds for their golden and normal warnings respectively. The efficiency distribution is shown in Figure 3. The SK golden warning threshold ($N_{\mathrm{cluster}} = 60$) is shown for reference.

This figure uses a cut on $E_{\mathrm{true}} > 7\,\mathrm{MeV}$. In order to assess the impact of energy reconstruction, 1.2 million SN events (corresponding to 3000 SN @ 50 kpc) were ran through WCSim [10], with dark noise turned off. A straight line fit of E_{true} vs N_{hits} is used to determine the N_{hits} cut corresponding to $E_{\mathrm{true}} > 7\,\mathrm{MeV}$ for each geometry configuration. Qualitatively, there is not a large effect. A quantitative measure of this will provided in the next version of this note.

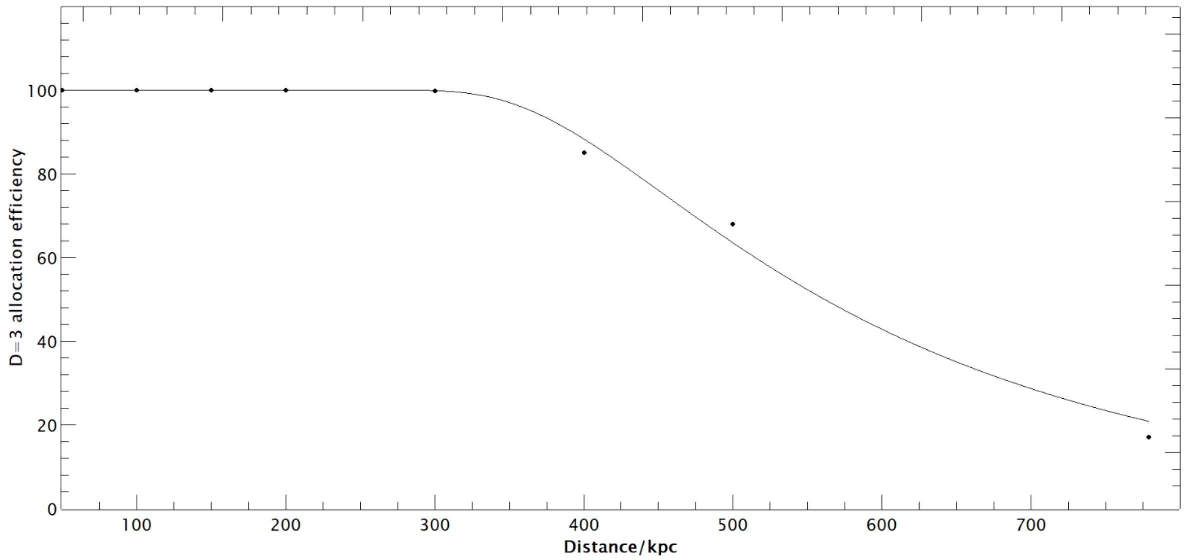

Figure 4: $D = 3$ allocation efficiency for a population of 3000 simulated SN at each distance. The NK2NO model is used, corresponding to a conservative estimate of the event rates. True vertex positions are used.

As with Super-K, the D value will be calculated for all events within the monitoring time window. 3000 SN were simulated using the NK2NO model across distances from 50 kpc (LMC) to 779 kpc (M31) with sntools [9]. True vertex positions were passed to dimfit, the code used by SK to allocate D values. The efficiency to allocate $D = 3$ is shown in Figure 4. The efficiency is 100% out to 300 kpc. Results using a reconstructed vertices will be produced for the next version of this note.

The determination of D from the χ^2 values depends on some geometry-dependent thresholds. These thresholds will need to be modified for HK, due to its different size and height/diameter relative to SK. These new thresholds will be set in the next version of this note.

These results show that the SK real-time burst monitor method will work for HK for SN out to at least the LMC (60 kpc). The studies mentioned above will be performed in order to set precise trigger thresholds.

In the event of a near supernova (0.1 kpc), 54,000 to 90,000 events are expected to be detected [1]. In the event of such a supernova, a large event rate will be experienced by an RBU. Each RBU will immediately start writing all data to their SSD. It is anticipated that for near supernova, the detection and alert system will be fully automated. For supernova further away, a similar procedure to that used in Super-K will be adopted. Experts will be contacted and a meeting convened while the TPUs continue to process 1 ms time slices that will be used to decide whether the increased event rate is due to a supernova,and also to determine the direction of the supernova. It is expected that a decision on the supernova and its direction would be known as quickly as possible, though tests are ongoing to calculate this time more precisely.

The inclusion of information from the outer detector (OD) in the supernova burst monitor is being considered. If it is possible to incorporate the OD into the supernova burst monitor, it may allow Hyper-K to retain some supernova detection capability during periods of detector calibration. This will be pursued further once tests of the inner detector burst monitor are complete.

3.2 Event Builder Units

Once a positive sub-window from the TPUs is received by the Brokers it will then assign it to one of the 4 EBUs communicating the time window of interest and the trigger type/ other useful info. The EBU receiving the job will then contact all the RBUs to request the relevant data (this time including waveforms) and it will be saved to both the local magnetic disk and then later the off-site copies. The EBUs will then send a confirmation to the Brokers of the job being complete which will trigger the next job being sent to them. Assuming a ~50 kHz expected event rate and a time slice of 500ns this will mean 80 MB/s [60 MB/s] data saved data rates, 7 TB [66 TB] per EBU per day and 28 TB [24 TB] total per day. This quantity of storage per EBU could be achieved with one magnetic drive providing 2 days of local storage. However for redundancy you would want a minimum of two drives. The current system design has 4 drives per EBU allowing 4 days of local storage. Added to this it is probably a good idea to have a second onsite backup of this data running a disk serve that can offload all the EBUs and also take care of the transfer to offsite. This data could then be transferred off site for archiving with a 10 Gb/s connection.

In the event of a near supernova, the Brokers will broadcast a node wide message and assign each EBU a collection of RBUs. The EBUs will then contact the RBUs and request all their supernova data. The supernova data which will total ~140 GB [130 GB] per EBU (~7 GB per RBU) will be sent to directly and they will begin to save it to local magnetic drive storage. This will take ~5 mins for each EBU to record the data to disk based on two drives RAIDed per EBU. However If 4 drives are used this will take 2.5 mins. Once completed, normal processing will resume and the added capacity will catch back up to real time quickly.

In the event of a near supernova, the TPUs and SNTPUs may not need to process all of the incoming data. Positive identification of the supernova will be achieved by the significant rate increase observed over a relatively short time period and decrease to nominal rates shortly thereafter. Studies will be performed to determine the absolute number of neutrino events required to get a direction resolution close to 1–1.3° @ 0.1 kpc [1] in a short (< 2 minute) time period.

3.3 Back End Network

In order to transfer all the data between RBUs, TPUs, EBUs and brokers, the back end network needs to be quite fast. As mentioned in the previous section the TPUs are designed with 40 Gb/s QSFP+ optical connections and the RBUs need similar 40 Gb/s QSFP+ to cope with the data rates. The EBUs are somewhat rate limited by the speed of the magnetic drives within them, so could make do with 10 Gb/s SFP+ ports but could also also use of 40 Gb/s QSFP+ ports as switch capacity is available. The Brokers despite their low data through put demands require fast connections so will use 40 Gb/s QSFP+ ports. To route all these connections 4 switches of 20 x 40 Gb/s (QSFP+) and 4 x 100 Gb/s (QSFP28) ports will be used to connect all of 80 RBUS with the 100 Gb/s connections aggregated to 2 switches of 32 x 100 Gb/s (QSFP28) ports (Fig.5). Eight of the ports on each 32 x 100 Gb/s will go to the two sets of 20 x 40 Gb/s (QSFP+) switches with a further 8 ports used as aggregate interconnects to the other 32 x 100 Gb/s switch. This allows 16 free ports per switch for TPUS, EBUs and brokers so any node can readout any RBU at the same speed.

4 Conclusions and future work

Despite a small number of open issues regarding the specific thresholds used in each of the warning levels (section 3.1.2), a baseline supernova identification and readout strategy has been presented in

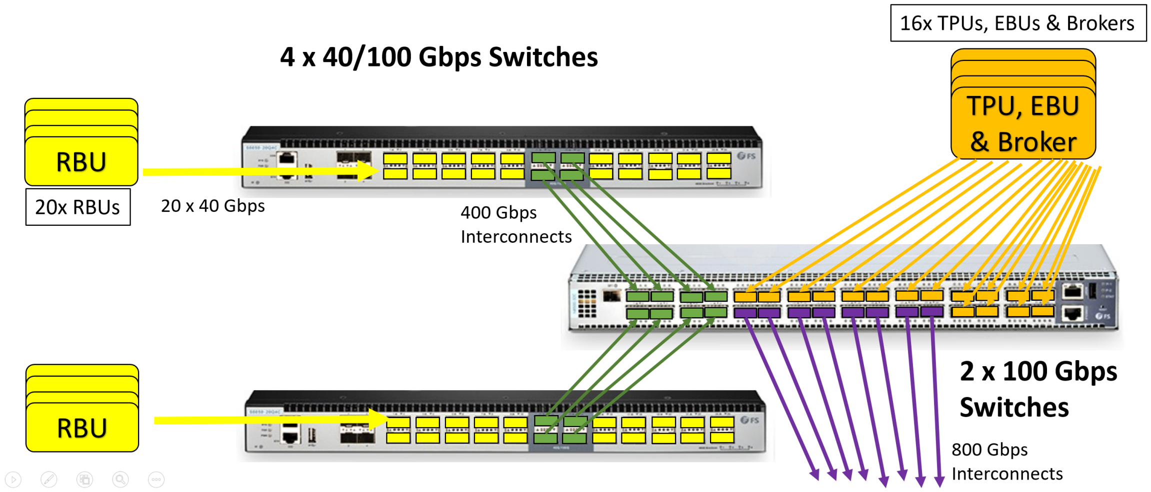

Figure 5: Back end wiring scheme

this report. In the coming months, these will be extensively tested and finalised for inclusion in the next draft of the technical report.

The main areas for us still to address are:

- Confirmation of the warning level thresholds for gold, silver, etc.;
- DIMFIT code to confirm $D = 3$ for supernova using reconstructed vertex positions;
- DIMFIT code to confirm $D \neq 3$ for non-supernova using reconstructed vertex positions;
- Simulation of backgrounds (most notably spallation) to assess the false-positive rate;
- Reconstruction of supernova direction.

References

[1] K. Abe et al. Hyper-Kamiokande design report. 2018.

[2] Benjamin Richards. ToolDAQ framework v2.1.1. https://doi.org/10.5281/zenodo.1482767, November 2018.

[3] K. Abe et al. Real-time supernova neutrino burst monitoring at Super-Kamiokande. 2016.

[4] T. Totani et al. *Astrophys. J.*, 216(496), 1998.

[5] K. Nakazato et al. *J. Supp.*, 3(205), 2013.

[6] A.S. Dighe and A.Y. Smirnov. *Phys. Rev. D*, 0330007(62), 2000.

[7] Particle Data Group. 2014.

[8] M.B. Smy. Fitting a point, line or plane to a collection of three-dimensional points. 2004.

[9] sntools. `https://github.com/JostMigenda/sntools/`.

[10] Wcsim. `https://github.com/WCSim/WCSim/`.