STAR DAQ Requirements

Revised, February 19, 1997

M.J. LeVine, A. Ljubicic, M. Schulz
Brookhaven National Laboratory

E. Schäfer
MPI, Munich

1. Introduction

The DAQ system of the STAR detector is responsible for accumulating data, on an event-by-event basis from each of the detector sub-systems and assembling the resultant data into large data structures that constitute a complete view of the detector response to a particular experimental interaction. The resulting data sets (i.e. events) are then moved to archival storage as well as to on-line computing resources for varying degrees of event reconstruction in near real-time. Although various DAQ system components perform a diverse set of tasks, the principal DAQ system functions (and thus, its specific requirements) lie in the following areas:

- Creation of the electrical and logical interface with various detector system electronics that allows detector data to be manipulated in a standard computing environment.

- Generation and application of the Level 3 software trigger capable of reducing the input data stream by a factor of 100.

- Reduction in size of accepted event data stream, through various compression techniques, by a factor of up to 10.

- Recording of resulting data stream on an archival-quality medium for subsequent off-line analysis.

- Creation and distribution of sufficient event data for on-line monitoring of both DAQ and detector system performance.

- Integration of all these functions with other major components of the STAR experiment (e.g. controls, trigger system, etc.).

The DAQ system functional requirements that follow address each of these areas.

2. Executive Overview

DAQ system functional requirements have been broken down into the following categories:

Operational modes
Trigger
TPC
SVT
Other Detectors
Level 3 Trigger

Calibration
System Control
On-line Monitoring
Data Archiving
Reliability
Maintenance and Mechanical
Documentation

3. Requirements Overview

3.1 Operational modes

3.1.1 Archiving stream

Requirement:

An archiving stream shall be provided which ensures the quality of archived data, i.e., that the data for all detectors are acquired under a set of identical conditions.

Justification:

The principal task of DAQ is to deliver data from one or more coupled detectors to a high-speed logging stream. These data must be acquired under strictly controlled circumstances; i.e., the conditions under which they are acquired must not be allowed to change during a run.

Before a main archiving run is started, a test of all concerned detector system must be carried out in order to determine that they are capable of proper operation.

No detector which is participating in such a run may change its configuration or any of its operational parameters, for example, high voltages. Experiment control must guarantee that this is the case.

The coupled stream must guarantee that each trigger accepted will end up on the archiving tape. Stated otherwise, flow control must back up from the main archiving tape to the detector-specific BUSY to the trigger system.

Current Status:

3.1.2 Test streams

Requirement:

There will be ancillary partitions of DAQ which will allow acquisition from detectors under less stringent conditions than those described under **3.1.1**.

Justification:

Other requirements dictate that less formal data streams should be available for commissioning, testing, and calibration. These *test streams* should be capable of starting and stopping without affecting the archiving stream's operation. In contrast to the archiving stream, the test streams impose no constraints on the participating detectors regarding change of parameters.

Test stream data may be delivered to the online system via the network or to low performance tape systems. In either case, events will be delivered for these streams on a best effort basis: non-availability of the output device (tape or network) should not result in additional detector dead time.

Independent operation for partitions smaller than an entire detector need not be provided.

Current Status:

3.2 Trigger

3.2.1 DAQ<=>Trigger interaction

Requirement:

The Data Acquisition is at all times in one of two major system states: **halted** (stopped, paused) or **running**. Whenever DAQ is in state *running*, the trigger system can issue a new signal, called readout request (Remark: in this document the term trigger is used in two contexts: in place of readout request, or simply to refer to the trigger subsystem). The readout request is accompanied by the issuance of a unique identifier, or token, whose numerical value token shall be identical to a token distributed by trigger via the detector front end electronic systems to the DAQ receiver cards; the DAQ system will use this identifier to verify the correct assembly of fragments into an event. In addition, this token corresponds to a data buffer shared by trigger and DAQ, which contains the trigger detector data for the event in question. From the time of issue of a readout request until the token is returned by DAQ, the contents of this buffer may not be modified by trigger.

Upon receipt of a trigger input, the DAQ system will begin processing the event and indicate so to the trigger system. Further triggers will be accepted by the DAQ system up to the limit of its internal resources. When finished processing the current trigger, the DAQ system will indicate so to the trigger system by returning the relevant token.

The trigger system will not present any triggers to the DAQ system that have not passed the level 2 acceptance criteria of the trigger system.

In addition to the trigger-DAQ handshake discussed above, there will be a detector-specific busy signal for each detector subsystem; this reflects the fact that these detectors require readout times which are dissimilar and may be running in an uncoupled mode (see **3.1.2**). DAQ does not issue a global busy signal.

Justification:

Flow control allows the unavailability of any resource (e.g., tape drives, buffers) to be signaled to all upstream data sources, reaching as far upstream as necessary to prevent any loss of data.

Under most circumstances, DAQ system performance will be governed by data archiving speeds. Given this fact, a flow-controlled system design presents a consistent view of detector dead time as the *or* of trigger, relevant detector, and DAQ dead time contributions.

Current Status:

3.2.2 Multiple Trigger Inputs

Requirement:

The DAQ system will respond to multiple trigger types within a single run configuration.

Justification:

The DAQ system must be able to respond to multiple experimental triggers within the same data-taking run. Detector data from individual triggers must remain correlated for later analysis. Limiting trigger types to a modest number results from the requirement that cost and complexity must be limited.

For any given trigger, hardware signals from the trigger subsystem to various detectors must agree with the list of detectors to be read out by DAQ.

Current Status:

Each trigger type will represent a trigger sub-system request for readout of a specific set of detectors using a specific set of readout algorithms. Detector data read out in response to a particular trigger will be collected together and will be tagged as a single event. The DAQ system will accommodate at most 16 trigger types. The acceptance by DAQ of a trigger will result in at most one event being built. Since the trigger system will provide a trigger word of 16 bits, the trigger system will map this word onto one of 16 possible trigger types and will distribute this trigger type to all of the detector front ends and to DAQ. The trigger system 16-bit word for each event will be available to DAQ level 3 algorithms as well as for further event processing and/or rejection. This word will include an indication of the detectors participating in the readout. The current DAQ-trigger interface specification is in agreement with these requirements.

The DAQ system and trigger system will maintain a common description and/or database of the mapping of 4-bit DAQ command word onto DAQ readout sequences. This information will be easily accessible by the DAQ system and the DAQ system will be notified of any alterations. The information contained within this database will be sufficient to insure complete coordination between trigger system readout requests and DAQ system responses.

Mapping of these DAQ command words onto DAQ readout sequences, once agreed upon by DAQ and Trigger, will not be changed while DAQ is in state running.

3.2.3 Multiple Trigger Streams

Requirement:

DAQ system response to multiple triggers will not be solely governed by the slowest detector readout time. A slow detector (e.g., TPC) which is currently being read out shall not prevent a trigger being issued to another detector. The DAQ system will provide individual detector busy signals that reflect the state of each detector's DAQ readout sequence.

DAQ has no requirement to read event fragments in the temporal order in which their triggers occurred. Event fragments can overtake each other, i.e., an older event in the DAQ pipeline can leave the event builder later than a younger event. This is true whether the event fragments have the same or different trigger type.

Justification:

Trigger readout sequences that involve lengthy readout sequences of detectors like the TPC should not completely inhibit readout of smaller, independent triggers involving other detectors (e.g., EMC) (STAR physics requirement).

Current Status:

The current trigger-DAQ interface design accommodates the de-coupling of independent trigger streams. This is accomplished by requiring that *individual detector sub-systems include sufficient event buffering to meet their own data rate requirements*. Data from fast detectors corresponding to an event which involves a slow detector are buffered for later readout.

It must be pointed out that DAQ performance is determined by a number of factors. Although the current design attempts to minimize the effects of DAQ dead time due to the readout of large events (e.g. TPC events), additional factors, such as event building overhead, will strongly affect data rates from all trigger types.

3.2.4 Effects of High Trigger Rates

Requirement:

In *no* case will any trigger system behavior swamp the DAQ system and make it unresponsive to command/control inputs.

Justification:

Required for proper control of DAQ under adverse circumstances.

Current Status:

3.2.5 Event numbering

Requirement:

DAQ shall label each event with a unique number distinguishing that trigger from any other trigger in the same run.

Justification:

Labeling all fragments which will ultimately form an event allows detection of software and/or hardware errors which might result in the incorrect assembly of an event, in particular assembling fragments from distinct triggers into the same event. In order to accomplish this, DAQ needs a trigger sequence number which contains enough bits to distinguish it from all other trigger sequence numbers which might have occurred during the period in which an event is in the DAQ pipeline up to the event building process.

Current Status:

This information is contained in the trigger token which is supplied by the trigger system and which must arrive at the receiver boards following the same physical path as the detector data. The 12-bit token issued by trigger is adequate to meet this need.

In addition, an incrementing number will be assigned to each event built by DAQ. This number will be unique for each run.

3.2.6 Trigger Rates

Requirement:

When presented individually, the DAQ system will accept triggers requesting individual detector readout at the following rates:

- TPC: 100 events/sec.
- EMC: 1000 events/sec.
- SVT: 100 events/sec.
- FTPC: 100 events/sec.

When triggers require the readout of multiple detectors for each trigger, the DAQ will perform at the following rates determined by the slowest detector in the configuration.

Justification:

STAR requirements.

Current Status:

3.3 TPC

3.3.1 TPC Readout Rate

Requirement:

The DAQ system is required to read TPC data at rates up to 100 events/sec. At these rates, the portion of the DAQ system responsible for the TPC readout will be 100% busy. These dead times include partial (or complete) event movement from sector readout boards to DAQ hardware, level 3 trigger processing for up to 100 events/sec. and archiving at least one compressed TPC event per second.

Justification:

Stated STAR requirement.

Current Status:

3.3.2 DAQ<=>TPC Cable Plant

Requirement:

The physical connection between the DAQ system and the TPC detector shall not allow any ground current return paths (i.e. ground loops) between the DAQ and TPC electronics.

Justification:

Electrical noise generated by such ground currents is unacceptable.

Current Status:

The use of fiber optic cables for DAQ<=>TPC interconnections satisfies this requirement.

3.3.3 DAQ System Data Processing

Requirement:

The DAQ system will be capable of performing the following operations on the 10-bit TPC data stream:

- Subtraction of a programmable, per channel, per time bucket offset of up to 8 bits in magnitude
- Per channel, programmable compression of each time bucket 10 bit ADC value into an 8 bit quantity.
- Programmable zero suppression of time bucket data that fall outside of programmable threshold criteria.

These corrections will all be performed in real time with no adverse affect on the data rates referred to in the triggering requirements section.

Justification:

These operations are required for reducing the magnitude of the event data early in the acquisition cycle. These reductions are necessary to meet data archiving performance requirements later in the system.

Current Status:

3.3.4 TPC Raw Data Forms

Requirement:

The TPC readout boards produce 10-bit ADC values; early in the DAQ readout sequence these 10-bit values are converted into an 8-bit compressed value. This conversion is called 8-bit compression and is performed on all TPC data.

The DAQ system will be able to produce TPC raw data events that incorporate a variety of levels of compression. That is, the DAQ system will be able to produce events that have been:

- 8-bit compressed data, without pedestal subtraction or zero-suppression.
- 8-bit compressed data with pedestal subtraction, no zero suppression.
- 8-bit compressed data with pedestal subtraction and zero-suppression.
- Any of the above formats with overall +/- 5% gain correction performed by scaling of 8-bit compression conversion algorithm.

Justification:

Hardware and system debugging require access to data at varying levels of processing. Although data will ultimately be acquired in its most efficient, compressed form, experiment staging and debugging will require some or all of the above formats.

Current Status:

Current plans call for 8-bit compression algorithm to be performed via downloadable lookup table.

3.3.5 Raw Data Buffering

Requirement:

DAQ will provide buffering for multiple TPC events. All events within this buffer will be available for level 3 trigger processing.

Justification:

STAR requirement.

Current Status:

TPC receiver cards provide for 12 event buffers.

3.3.6 TPC Level 3 Track Reconstruction Performance

Requirement:

Level 3 trigger decisions are based on real-time processing of the data produced by each detector, separately. For the TPC, this corresponds to track reconstruction.

Justification:

Overall performance of this stage of the event selection process is critical in achieving the 100 to 1 rejection required to achieve trigger performance requirements.

Current Status:

There are, as yet, no detailed TPC level 3 track reconstruction algorithm performance requirements. There is, however, a known time budget for this process. The TPC receiver boards will provide sufficient buffering for up to twelve events. Therefore, given the maximum TPC event rate of 100 events per second (see **3.2.6**), any *average* event processing that exceeds 120 msec. will result in increased DAQ system dead time. Although this does not yield a requirement for the processing power needed for level 3 algorithm processing, it does provide a time budget within which system performance will not be degraded. Stated differently, independent of the number of level 3 processors operating in parallel on separate events, the total elapsed time for dealing with a single event in all stages of the level 3 processing cannot exceed 120 msec.

Under current plans, the DAQ system will provide a minimum of 40000 Spec92Int of processing power distributed within the TPC sector processing level. Level 3 processing performance will be proportional to the processing power available and decreasing according to the complexity of the level 3 algorithm.

3.4 SVT

3.4.1 SVT Readout Rate

Requirement:

DAQ will accept triggers requesting either SVT individual readout or in combination with other STAR detectors at a rate of 100 events per sec. DAQ will archive compressed SVT events at a rate of at least one per sec.

Justification:

Stated STAR requirement. This is the same rate required by the TPC. Since SVT and TPC FEE RDO can produce up to 100 events per sec. due to SCA limitations, similar performance requirements fit both subsystems.

Current Status:

3.4.2 DAQ<=>SVT Cable Plant

Requirement:

The physical connection between the DAQ system and the SVT detector shall not allow any ground current return paths (i.e. ground loops) between the DAQ and SVT electronics.

Justification:

Electrical noise generated by such ground currents is unacceptable.

Current Status:

The use of fiber optic cables for DAQ<=>SVT interconnections satisfies this requirement.

3.4.3 DAQ System Data Processing

Requirement:

The DAQ system will be capable of performing the following operations on the 10-bit SVT data stream:

_ Subtraction of a programmable, per channel, per time bucket offset of up to 8 bits in magnitude

_ Per channel, programmable compression of each time bucket 10 bit ADC value into an 8 bit quantity.

_ Programmable zero suppression of time bucket data that fall outside of programmable threshold criteria.

These corrections will all be performed in real time with no adverse affect on the data rates referred to in the triggering requirements section.

Justification:

These operations are required for reducing the magnitude of the event data early in the acquisition cycle. These reductions are necessary to meet data archiving performance requirements later in the system.

Current Status:

3.4.4 SVT Raw Data Forms

Requirement:

The SVT readout system produces 10-bit ADC values. Early in the DAQ readout sequence, these 10-bit values are converted into 8-bit compressed values. The DAQ system will be able to produce SVT raw data events that incorporate several degrees of compression. In particular, the DAQ system will be able to produce events with:

_ 8-bit compressed data, without pedestal subtraction or zero-suppression.

_ 8-bit compressed data with pedestal subtraction, no zero suppression.

_ 8-bit compressed data with pedestal subtraction and zero-suppression.

_ Any of the above formats with overall +/- 5% gain correction performed by scaling of 8-bit compression conversion algorithm.

Justification:

Hardware and system debugging require access to data at varying levels of processing. Although data will ultimately be acquired in its most efficient, compressed form, experiment staging and debugging will require some or all of the above formats.

Current Status:

Current plans call for 8-bit compression algorithm to be performed via downloadable lookup table.

3.4.5 Raw Data Buffering

Requirement:

DAQ will provide buffering for multiple SVT events. All events within this buffer will be available for level 3 trigger processing.

Justification:

STAR requirement.

Current Status:

SVT receiver cards provide 12 event buffers.

3.4.6 SVT Level 3 Track Reconstruction Performance

Requirement:

Level 3 trigger decisions are based on real-time processing of the data produced by each detector, separately. For the SVT, this corresponds to track reconstruction.

Justification:

Overall performance of this stage of the event selection process is critical in achieving the 100 to 1 rejection required to achieve trigger performance requirements.

Current Status:

There are, as yet, no detailed SVT level 3 track reconstruction algorithm performance requirements. There is, however, a known time budget for this process. The SVT receiver boards will provide sufficient buffering for up to twelve events. Therefore, given the maximum SVT event rate of 100 events per second (see **3.2.6**), any *average* event processing that exceeds 120 msec. will result in increased DAQ system dead time. Although this does not yield a requirement for the processing power needed for level 3 algorithm processing, it does provide a time budget within which system performance will not be degraded. . Stated differently, independent of the number of level 3 processors operating in parallel on separate events, the total elapsed time for dealing with a single event in all stages of the level 3 processing cannot exceed 120 msec.

Under current plans, the DAQ system will provide a minimum of 8000 Spec92Int of processing power distributed within the SVT sector processing level. Level 3 processing performance will be proportional to the processing power available and decreasing according to the complexity of the level 3 algorithm.

3.5 Other Detectors

3.5.1 Readout Rate

Requirement:

The DAQ system is required to read detectors at a rate equal to their digitizing rate, but not to exceed 1000 events/sec. This rate is not identical to the event building rate, which will be considerably lower.

Justification:

Readout of partial detector information should be possible at the maximum digitizing rate of the detector. Summary information based on these data will be used to select events to be passed to the event builder, at a lower rate.

The event building rate will be subject to several limits, including archiving rate and a maximum event handling rate.

Current Status:

The limit of 1000 Hz is the current best estimate of readout rate, based on little knowledge of the detector readout mechanisms.

3.5.2 Physical Interface

Requirement:

The physical interface for other detectors will be similar to that of the TPC and SVT, utilizing optical fiber for serial transmission.

Justification:

See **3.3.2**

Current Status:

3.5.3 DAQ System Data Processing

Requirement:

Currently only specified for TPC (see section **3.3.3**) and SVT (section **3.3.3**).

Justification:

Current Status:

Input from other detectors is required.

3.5.4 Raw Data Formats

Requirement:

Except for TPC and SVT, no information is available. No current requirements.

Justification:

Current Status:

3.5.5 Level 3 Performance

Requirement:

Except for TPC and SVT, no information is available. No current requirements.

Justification:

Current Status:

3.6 Level 3 Trigger

3.6.1 Event Selection Mechanism

Requirement:

Level 3 will allow event selection before event building, based on summary information extracted from each detector separately. Summary information consists of, for example, tracks found in the TPC. In addition to the information derived from the detector itself, each level 3 processor will have available summary information derived from the trigger detectors which may be useful in interpreting its own results.

DAQ will make available, on a per-event basis, histograms of distributions of any quantities on which cuts are made in Level 3; this information will be included in the assembled event.

Justification:

Level 3 is an opportunity to select events before incurring the enormous overhead of moving large amounts of data. Well-designed algorithms to summarize the data from each detector will allow characterization of that detector. However, cross-correlation among different detectors must be done keeping in mind the strict time budget imposed on Level 3 processing.

Current Status:

In order to allow the maximum possible cross-correlation between detector subsystems, other detectors (TOF, EMC) will send their data to the TPC sector crates, where the geometry of the detector justifies this. This is done with the intention that the cross-correlation can be done in parallel in the sector crates.

Other detectors whose organization is not along TPC sector boundaries (e.g., SVT) will send their data to the global Level 3 crate, where available time decrees that a much less ambitious cross-correlation is possible.

3.6.2 Level 3 Processing Requirements

Requirement:

DAQ will provide required pre-processing of raw detector data to generate the summary information to be used in the level 3 CPUs, as well as the level 3 CPUs themselves.

Justification:

The raw data are found in buffers on DAQ receiver boards, or equivalent. In some cases (TPC, SVT), massive amounts of processing are required to convert pad (anode) responses to space points to be passed to tracking code in the level 3 CPUs; in these cases, the pre-processing can occur in real time only with help from custom-built hardware.

Current Status:

For the case of the TPC and the SVT this requirement represents the basis for the DAQ design.

3.6.3 Level 3 Software Development

Requirement:

DAQ will provide a mechanism to port and test software which has been developed in a simulation environment to the embedded environment of the level 3 CPU. High level language will be supported.

Justification:

Conditions in the simulation environment will be quite different from that encountered in the level 3 CPU. First, it is unclear that STAF tables can be supported in level 3. Second, it should not be necessary that the level 3 CPU have an attached keyboard and screen for debugging. Finally, support for virtual memory larger than the physical memory cannot be provided in a real-time environment.

Current Status:

It is expected that DAQ will provide two different mechanisms to aid in the migration of code between the two environments. The first will be a level 3 emulator, running in a standard UNIX environment, which would transform between STAF tables and the embedded level 3 format. The second will be virtual terminal facilities which will logically attach a workstation screen to standard I/O facilities provided by the language supported in the level 3 environment.

3.6.4 Detector Data Delivery

Requirement:

Detector subsystems will deliver data to DAQ receiver boards in a manner which facilitates preprocessing needs specific to the detector. For the TPC, SVT, and FTPC, this means that, to the largest extent possible, neighboring pads (anodes) will be delivered to the same receiver board.

Justification:

For the TPC, SVT, and FTPC conversion of ADC data to space points requires that responses from all adjoining pads (anodes) be available to the same CPU. This is essential for recognition of boundaries of clusters formed by a single particle trajectory passing through a pad row (TPC) or layer (SVT).

While limited information can be passed between receiver boards, a requirement for large-scale sharing of information will result in overheads which will adversely impact system performance.

For other detectors, similar constraints may arise.

Current Status:

This requirement is currently satisfied for the TPC and SVT.

3.6.5 Level 3 Performance

Requirement:

Level 3 will be able to handle the preprocessing of data and transport of summary information to level 3 CPUs at a rate of 100 events/sec.

Justification:

STAR requirement

Current Status:

DAQ will provide transport consistent with this requirement. Present budget allocations, however, do not provide for full level 3 implementation.

3.6.6 Buffering of raw data

Requirement:

Each detector will provide, or be provided with, buffering for 12 events which are in the level 3 pipeline. This is in addition to buffering required for independent multiple trigger streams (see Sec. **3.2.2**).

Justification:

It is necessary that all detectors have buffering which results in a consistent pipeline length.

Current Status:

Requirement is satisfied for TPC and SVT. For other detectors, no information is available.

3.7 Calibration

3.7.1 TPC, SVT, and FTPC Channel Pedestal Value Calibration

Requirement:

The DAQ system will perform channel pedestal value calibrations at the receiver board level.

Justification:

Pedestal calibration events for the TPC and SVT are very large. During normal operation, there is little need to inspect or save data from individual pedestal calibration events. They are primarily used as inputs to a simple algorithm that produces individual channel offset averages that can be loaded into local pedestal memory. By placing the pedestal calibration algorithm as far upstream as possible, we lessen the need to move these large black events through the system. Thus, pedestal calibration is faster, easier and, hopefully, performed more frequently.

Current Status:

DAQ will be capable of accumulating and calculating ADC pedestal offset values for each channel's time bucket and loading internal pedestal memory. The resulting pedestal values, if requested, will be sent through the normal data stream for archiving or inspection. Pedestal value calibration should take no more than approximately 2 minutes for 64 samples.

During debugging or when individual channel sigma's appear unacceptably large, the results of the pedestal calibration algorithm can be record or delivered to a network client (see **3.9.2**).

Calibration requires DAQ system operation in a dedicated mode. This implies that normal data taking will have to stopped for the duration of the calibration procedure.

3.7.2 Geographical Pad Address Verification

Requirement:

The DAQ system will, as part of the TPC, SVT, and FTPC calibration process, be able to probe and verify the presence and geographical pad location of all FEE boards.

Justification:

Correct DAQ system operation requires an exact understanding of individual pad data ordering within the data stream emitted by each readout board. An automatic probe and verification of this data stream is required to provide this information in a concise and reliable manner.

This capability is essential to identify errors resulting from incorrect connection of fibers carrying the detector raw data stream as well as errors in arrival sequence resulting from incorrectly configured multiplexer-demultiplexer combinations.

Current Status:

3.7.3 Geographical Anode Cabling Verification

Requirement:

The DAQ system will, as part of the SVT calibration process, be able to probe and verify the connectivity between components of the SVT and DAQ readout board channels.

Justification:

Correct DAQ system operation requires an exact understanding of individual SVT anode data ordering within the data stream. An automatic probe and verification of this data stream is required to provide this information in a concise and reliable manner.

This capability is essential to identify errors resulting from incorrect connection of fibers carrying the detector raw data stream as well as errors in arrival sequence resulting from incorrectly configured multiplexer-demultiplexer combinations.

Current Status:

3.8 System Control

3.8.1 DAQ Control Interface

Requirement:

At any given time, the DAQ system will be controllable through a single interface. This interface will be selectable, under supervisory control. There will be a mechanism for determining which interface is capable of controlling the DAQ system at a given time. Both graphical user interface and support for running command files will be provided.

Justification:

Flexibility in debugging, combined with reliability.

Current Status:

3.8.2 DAQ Data Bases

Requirement:

DAQ will maintain appropriate data bases for DAQ system parameters.

Justification:

Successful analysis of recorded data will depend on accurate information of the experimental run circumstances. This includes details of how the DAQ system has been configured and what operations it is performing on the recorded data stream. Therefore, any parameters that can impact later analysis must be archived.

DAQ must be able to run, for maintenance and software development, independent of both the RHIC computing facility and Online.

These logging operations are not time critical as they should all take place prior to actually beginning a data run.

Current Status:

All DAQ system parameters that affect either the quality or quantity of data in the data stream will be archived into a local data base by DAQ system begin-of-run time. These parameters will be available to other sub-systems and on-line analysis programs at any time. These parameters will describe DAQ system behavior in the following areas:

- Enabled/disabled triggers.
- Enabled/disabled detectors.
- Enabled/disabled detector channels or features.
- Level 3 trigger control parameters.
- Data archiving mode.
- Archive data set name and description.
- Data formatting parameters (compression, pedestal subtraction, etc.).
- DAQ system operating mode (see below).

Parameter changes will be logged at the beginning of every data-taking run.

3.8.3 DAQ Database Archive

Requirement:

All critical DAQ parameters (see above) will be logged to the archive at the beginning of every data taking run.

Justification:

The DAQ system must be able to include critical parameters with the archived data set. Local DAQ parameter storage, however, is only intended to facilitate DAQ configuration and operating mode changes. It is not intended to act as long term archival storage for later analysis purposes.

Current Status:

Local DAQ parameter storage will contain only those parameter values that are current. Additional parameter storage will be provided on a named parameter set basis (see **3.8.4** below).

3.8.4 Save/Restore Functions

Requirement:

It will be possible to save and restore the DAQ system operating modes and configurations.

Justification:

The number of user-alterable parameters that fully describe the operations of the DAQ system is expected to be substantial. There should be a simple mechanism for saving the state of these parameters, altering several for testing purposes, and then restoring the originals. Furthermore, the saved parameter sets should be cataloged and recallable by a simple naming mechanism. Care will be taken to not allow this facility to subvert the normal DAQ system checks that prohibit begin run state changes into illegal configurations (e.g. requiring SVT data when SVT is missing or disabled).

Current Status:

These configurations will be stored by name with an accompanying descriptive text. The operation of saving a DAQ operating mode and configuration and replacing it with another should take no more than 30 seconds.

3.8.5 DAQ<=>Sub-system Interlocking

Requirement:

There will be a reliable mechanism that inhibits sub-system state changes while the DAQ system is transitioning between the ***idle*** and ***data-taking*** states (i.e. beginning a run). Furthermore, there will be a reliable mechanism available to the DAQ system that will, when appropriate (see **3.1.1**), inhibit sub-system state and/or parameter changes that will affect the quality of experimental data.

Justification:

The transition of the DAQ system into the data-taking state will involve a number of configuration and parameter checks. As these will take some time to perform (several seconds), there must be a mechanism of insuring that other sub-systems on which the DAQ system depends, do not alter their state in an unsynchronized way. There must be a reliable mechanism that, during several time-critical state changes, insures that all sub-systems are synchronized.

In addition, if the DAQ system in operating in a very restrictive mode (e.g. a serious data archiving run), certain sub-system operations must be inhibited (e.g. TPC HV set point changes). This also requires a reliable inter sub-system synchronization mechanism.

Current Status:

These assurances are the responsibility of the STAR online system.

3.8.6 DAQ Hardware Configuration

Requirement:

The DAQ system will be able to probe all attached FEE's and hardware data paths when explicitly requested.

Justification:

The DAQ system must not operate in a mode where unrecoverable errors are likely. If, for example, run parameters require that all TPC sector processing components are powered and ready, these components should be checked by DAQ before actually allowing a run to begin. Failure to do so might allow portions of the readout code to attempt to access parts of the DAQ system that are not present and/or ready with unpredictable results. To the extent that the DAQ system can probe its interface environment, this technique can add a great deal of robustness in a fairly straight-forward way.

Current Status:

Based on operating mode and parameter information, the DAQ system will not successfully transition from the ***idle*** to the ***data-taking*** state if any required sub-system is not present or, in some detectable way, inoperative (e.g. requiring tape archiving with no tape present). As this procedure may be time-consuming, a ***quick begin run*** option will be available.

3.8.7 Error Logging

Requirement:

Exceptional conditions that occur during DAQ system operation will be logged to an on-line error logging facility. Detection and logging of exceptional conditions will be the explicit responsibility of those writing DAQ system code. The error logging facility will provide a mechanism for displaying errors in a timely fashion as they are detected and logging them to a disk file. All entries will be time-stamped by the originating operating system. There will be provision for indication of the severity of the error encountered as well as its source. System clocks (as opposed to experimental time stamps) will be synchronized to within about 1 second. The logging of errors may adversely affect DAQ system performance.

Justification:

Many portions of the DAQ system will contain embedded processors with no direct connection to an interactive screen. Moreover, the sheer number of such embedded processors makes the attachment of each (or even several) of these processors to terminals (or workstation windows) completely unworkable. There is a real need for a centralized, general purpose mechanism that allows operator notification when error conditions arise in remote, difficult to reach DAQ components.

Since exceptional condition are just that (i.e., unusual behavior, exceptions to normal operation), degradation of DAQ system performance is acceptable.

Current Status:

3.9 On Line Monitoring

3.9.1 On Line Data Interconnect

Requirement:

The DAQ system will be connected to the on-line analysis system through a commercially-supported, high-speed data communications channel. Although the technology in this area is rapidly evolving, this communications channel will support, at a minimum, the equivalent of a 100 Mbits/sec. transfer rate.

Justification:

Most data analysis will be performed in commercially-built workstations and processor farms. By using a widely-accepted high-speed communications standard consistent with our data rate requirements, we will save in purchase and software development costs.

Current Status:

3.9.2 On Line Data Channel

Requirement:

The DAQ system will distribute data to consumers via a single independent data distribution channel. The number of independent data channels will be kept to one in order to minimize the potential for DAQ system latency and to minimize the complexity of this portion of the DAQ system. Distribution of data to online consumers will, at most, reduce DAQ taping performance by 5%.

The channel will interface to a single online process. Further distribution of these events to multiple consumers is not a DAQ requirement.

Justification:

On line data will be required by a large and, as yet, undetermined number of monitoring programs. Requiring the DAQ system to support and manage all of the connections to these monitoring programs places a large and unnecessary burden on DAQ system resources. The DAQ system will support a single on line data channel to event consumer programs. This consumer may, in turn, re-distribute events to various monitoring processes. These consumer programs are not considered to be part of the DAQ system and its responsibilities.

Current Status:

3.9.3 On Line Data Formatting

Requirement:

The DAQ system will provide on line data in a format that allows recipient processors to efficiently convert event data into its correct native representation.

Justification:

Over the course of the STAR experiment, on line data will be sent to a wide variety of computing platforms-some of which will differ in their internal data representations (*e.g.* big-endian *vs.* little-endian). The format of the data stream emanating from the DAQ system must allow the recipient system to efficiently convert (if necessary) the data into the representation that is most natural and therefore, most efficient, for its architecture.

Current Status:

3.9.4 On Line Monitoring Functions

Requirement:

The DAQ system will provide status and statistical information on system performance in several areas. These include:

- Trigger rates for all enabled trigger types.
- Summary of aggregate event size before and after data compression.
- Summary of detector readout board time-outs or errors.
- Summary of Level 3 hit-finding algorithm performance-including statistics on number of hits per event by both readout board and sector.
- Summary of Level 3 track-finding performance and statistics (categorized by rejection criteria) for Level 3 event rejection.
- Summary of event builder activity and number and size of events passed on to the data distribution channel (see above).
- Summary of trigger information provided by trigger sub-system.

Justification:

In order to monitor both DAQ system performance and detector performance, we need access to the above quantities. The above is only an initial list. It is expected that additional quantities will be added as necessary.

Current Status:

This information will be available to DAQ and, on demand, to Online.

3.9.5 On Line Data Filtering

Requirement:

The DAQ system data distribution channel will consist of a mechanism to accept requests for events from the online system. Each request may specify a trigger type, as well as a scope (see **3.9.6**). The events delivered will be those that arrive at the event builder based on the DAQ action assigned to their trigger types.

Justification:

The DAQ system data distribution channel will provide a particular class of events. These classes can include statistical events describing DAQ system performance, collections of raw, unprocessed readout board data, partial or complete events containing multiple detector sub-systems.

In no case will events be built more than once for any given trigger type. For example, a given trigger which results in a zero-suppressed TPC readout cannot also give a non-suppressed version to the online system.

Current Status:

Requirements for the content of each of these data distribution channels have been requested from other STAR groups. These requirements will drive the exact method used to build or generate the events to be provided through these channels.

3.9.6 Data Transfer Characteristics

Requirement:

The DAQ system will transfer complete events-as defined by the event builder-over these channels, or partial events as defined by a scope specification. The scope will determine the size of the subevent that will be delivered. Scopes down to the TPC sector level or equivalent will be honored.

Justification:

The restrictions on event granularity are meant to keep the design of the event builder relatively straightforward. On the other hand, it is recognized that there are valid reasons for requesting a series of events comprised of a partial detector.

Current Status:

3.10 Data Archiving

3.10.1 Start of Run Information

Requirement:

Information necessary to define the conditions of the run will be archived to the logging tape.

Justification:

Appropriate data must be acquired, from Controls, to allow data or pointers to data to be included on the archiving tape at the beginning of each run.

It is expected that voluminous data (e.g., pedestals or EMC HV values) may be stored in a database and only a reference to this dataset will be stored on the tape itself.

Current Status:

3.10.2 Archiving Data Rates

Requirement:

The DAQ system will be capable of archiving event data to a recording channel or channels at sustained rates of up to 16 Mbytes/sec. Although the DAQ system will support multiple archiving streams (see below), there is no requirement for aggregate logging rates in excess of this number. The presence of additional logging channels will not adversely affect the performance of this single channel.

Justification:

The required rate of 16 Mbytes/sec. corresponds to the expected size of compressed TPC events logged at a rate of one per second. The one event per second data logging rate is a STAR requirement..

Current Status:

This requirement may be substantially altered by the introduction of non-lossy compression to the TPC and SVT data. Preliminary studies indicate that these techniques can save up to 40% of the TPC event size. No estimates are available on the impact on CPU resources. Based on these estimates, the 16 MB/sec could be reduced to 12 MB/sec.

3.10.3 Archiving Controls Data

Requirement:

Controls data concerning detector conditions will be archived to tape.

Justification:

It may be necessary to have information concerning detector operating conditions interspersed in a timely fashion during the course of the run, e.g., it may be necessary to know the drift velocity in the TPC at intervals shorter than the duration of a run, in order to analyze the data.

Current Status:

DAQ will respond to special triggers generated by Controls by acquiring data from Controls, which will be formatted as a special event, to be included in the archiving stream.

3.10.4 Multiple Archiving Streams

Requirement:

The DAQ system will support multiple archiving output streams.

Justification:

A given trigger configuration may yield multiple event streams that are completely uncorrelated. During initial detector testing it may be desirable to direct these streams to different logging devices.

Current Status:

In addition to the main archiving stream, one or more low-performance taping streams will be supported for test or calibration purposes.

3.10.5 Archiving Media

Requirement:

The DAQ system will archive data onto industry standard, archive quality media.

Justification:

Due to the large volume of data recorded, transcription of data from non-standard media will become expensive and prohibitively time consuming.

Current Status:

Industry standards in the high-capacity storage market are still evolving. We are watching several developing standards and will make a choice at a suitable time.

3.10.6 Archiving Modes

Requirement:

The DAQ system will allow data archiving to be performed in one of several modes. These will include, but not be limited to, record all data, sample data for recording with a programmable selection rate, record events identified as originating from a given trigger type, and record events that pass a selectable programmable filter function. In all cases, complete events-as defined by the event builder- will be recorded.

Justification:

Since data acquisition performance will, in most cases, be determined by the performance of the data archiving sub-system, a great deal of operational flexibility is required.

Current Status:

Current options listed above are preliminary. Subsequent requirements must be determined by various detector groups.

3.10.7 Data Archiving Format

Requirement:

Data archiving output will be formatted in such a manner as to facilitate its translation into appropriate internal representations when read into computer systems of differing architectures. This translation will not require access to external event description databases or files. The selected output format will be reasonably compact and will not increase data structure sizes by more than 10%.

Justification:

Data archiving will have data representation problems similar to those found in the DAQ<=>on line analysis area. Data written to an archiving medium will, in some cases, be read by computer systems of different architectures-with differing internal data representations. To facilitate the distribution of data for analysis, the DAQ system will archive experimental data in a format which facilitates the translation, if necessary, of elements of the data stream into an appropriate internal representation. This implies that, in regards to primitive data types and representations appearing in the data stream, the DAQ archiving data stream will be self-describing.

It must be recalled, however, that DAQ will be able to archive data at a fixed rate (16 Mbyte/s). Increases in event size due to formatting will result in fewer events on tape. Attempts to structure data in an arbitrary fashion may result in burdens imposed by self-described data which cannot be tolerated.

Current Status:

3.11 Reliability

3.11.1 System Reliability

Requirement:

Once fully operational, the DAQ system shall be unavailable due to failures and/or repairs for not more than 5% of the time when beam is available.

This requirement translates into a combined constraint on frequency of repairs and time-to-repair a malfunctioning (hardware or software) component.

Justification:

DAQ is a common element in all measurement activities. In general, DAQ being unavailable results in inability to utilize any of the detector systems.

The requirement can be met by a combination of efforts: good diagnostic software facilitates quick localization of faulty components, hardware design which allows for easy module replacement speeds repairs, and careful design of both hardware and software will minimize frequency of repair.

Current Status:

3.11.2 Archiving Error Rate

Requirement:

Tape recording shall not encounter a bit error rate exceeding 10^{-9} uncorrectable write errors.

Tape recording shall not encounter a bit error rate exceeding 10^{-12} undetected write errors.

Justification:

Tape archiving systems use error-detection-correction schemes which correct the vast majority of errors in real-time, without having to notify the client.

A detected, uncorrectable tape write error is not in itself disastrous; the client can rewrite the event. At some level, however, it begins to impact system performance. The limit given here represents a rewrite once per every 10 events.

The assumption is that only a single record has to be rewritten, so that the overhead is negligible.

Undetected write errors are much more serious. If such an error affects a single ADC value, its effect may be unnoticeable. On the other hand, if the error occurs in a pointer to another structure, the entire event may never be readable! The limit given corresponds to one such error in 10,000 events.

Current Status:

Depending on the taping hardware chosen, rewriting a record with a detected error may not be possible. In that case, the numbers cited above are not directly applicable.

Simple strategies such as writing pointers 3 times in combination with checksums on every event will be used; it is expected that these methods will allow salvage of all but the most pernicious errors. Care has to be taken at every stage of formatting an event in order to limit the amount of data within an event which is compromised by a single bit error.

3.11.3 System boot time

Requirement:

DAQ will guarantee that the system can be brought to full operational status from a hardware reset in 5 minutes.

Justification:

In the event of disastrous failure resulting in the need to perform a hardware reset, or in the event of a power dip, DAQ should not be the limiting factor in returning to full operational status.

Current Status:

Hard disks may be located in each VME crate to provide boot capability consistent with this requirement; use of 100 Mbit/s Ethernet connections to the DAQ fileserver are also under study.

3.12 Maintenance and Mechanical

3.12.1 System Functional Test Suite

Requirement:

It is required that the DAQ system be able to test and exercise all processors and data paths within the DAQ system in the absence of or without control of any detectors or sub-systems. These tests will verify overall system functionality and will allow isolation of gross hardware failures. These tests will not take the place of detailed, board-level diagnostics and are not intended to produce data rate loadings equal to those encountered during an experiment. These tests will not require any specific detectors or slow- control subsystems.

Justification:

The DAQ system will be a large and complex collection of processors and interconnect hardware. In systems of this kind, repairs are most often performed by substitution-boards, cable, whole sub-systems, etc. However, in systems of this size, it is not always possible to trivially identify the failed component. Without some mechanism of identifying a failed component, blind replacement repair schemes will be lengthy and potentially costly in lost beam time. There must be a suite of diagnostic programs that will be first used in staging the DAQ system and subsequently assist DAQ personnel in locating and repairing failed hardware (*and software*) DAQ components.

Current Status:

3.12.2 Quick System Functional Test

Requirement:

There will be a single test program that will indicate, to a very high level of certainty, that the DAQ system hardware and software is functional. Unlike the above test suite, this test is not intended to help identify failing components. It will result in a simple go/nogo response that will indicate the system's readiness to take data. This test will not require any specific detectors or the Controls subsystem.

Justification:

Ultimately, STAR detector operation will take on the characteristics of a facility. That is to say, it will be staffed and used by a wide variety of users who may have little detailed knowledge of system designs. It will be important to allow relatively DAQ-unskilled personnel to verify the DAQ system early in the run preparation cycle, probably well before all detectors are powered up. If the DAQ system fails this functionality test, members of the DAQ group can then use more detailed tests to affect repairs.

Current Status:

3.12.3 Component Diagnostics

Requirement:

Diagnostic programs and tools will be available to debug and repair any non-commercial hardware DAQ system components. These programs and tools should not require the entire DAQ system to function. They will come into use after failed hardware components have been identified through the use of a system function test (*see above*) and verified by hardware component swaps.

Justification:

Some hardware components of the DAQ system will be custom designs. During use, some of these components will fail. When this occurs, there must be a set of dedicated hardware and software tools available for locating failed components and re-testing these assemblies prior to being re-installed in the DAQ system. The full DAQ system is not an appropriate test and repair station for PC boards.

Current Status:

After manufacture, these components need to undergo some form of acceptance testing. Thus the tools need to exist for manufacturing purposes. The requirement translates into the need to reserve VME crates and CPUs to carry out repairs without disturbing the experiment.

3.12.4 Environmental and Power Monitoring

Requirement:

Wherever possible, internal hardware ventilation temperatures will be monitored. Wherever possible, power supply voltages and currents for non-commercial DAQ system hardware will be monitored. It will be the Control system's responsibility to monitor temperatures as well as power supply voltages and currents.

Justification:

As noted earlier, portions of the DAQ infrastructure will, on occasion, fail. While it is very desirable to be able to quickly replace failed fans and power supplies, it is of even greater importance to know that they have, in fact, failed.

Current Status:

3.13 Documentation

3.13.1 DAQ System Context Diagram

Requirement:

There will be a context diagram that shows the relationship between the DAQ system and all the sub-systems and program interfaces with which it interacts.

Justification:

Good design practice and consistency with other STAR sub-system documentation.

Current Status:

3.13.2 DAQ System User's Guide

Requirement:

There will be a DAQ system user's guide available in both printed and WWW readable form. Applicable portions of this guide will be available for use during early detector tests involving the DAQ system.

Justification:

Good design practice.

Current Status:

3.13.3 DAQ Functional Description

Requirement:

There will be a DAQ functional description document that will describe, in some detail, the design and implementation specifics of the DAQ system.

Justification:

Although a requirements and user's document go a long way towards describing a system's capabilities, they do not give any real feel for how a system is internally organized. During the course of the STAR experiment, personnel will certainly change. Maintaining a clear understanding of the system's internals is critical for future system enhancements and long term maintenance.

Current Status:

3.13.4 DAQ Code Development Environment

Requirement:

DAQ system code will be developed and supported within a suitable source, distribution and version control system.

Justification:

Good software engineering practice.

Current Status:

DAQ code development presents a slightly different set of development environment issues and concerns than off-line or monitoring program development. In addition to the obvious requirements of tracking operating system, compiler and utility library versions, the DAQ system must track several other components that may currently have individual development control systems in place. These may include the EPICS system, i960 development tools, vxWorks (or other RTOS), etc. Wherever possible, available and supported development environments will be utilized.