

Contents

8	Rates, Trigger and Data Acquisition	3
8.1	Expected rates	3
8.1.1	Overview	3
8.1.2	Trigger elements	4
8.1.3	Accidental rates	5
8.1.4	Rates in tracking chambers	7
8.2	Trigger	9
8.2.1	Overview	9
8.2.2	Level 1 trigger	12
8.2.3	Trigger simulation	14
8.3	Data acquisition	15
8.3.1	Overview	15
8.3.2	Data flow and rates	16
8.3.3	Level 3 trigger	17
8.3.4	Monitoring and Control	18

Chapter 8

Rates, Trigger and Data Acquisition

8.1 Expected rates

8.1.1 Overview

We estimate trigger and background rates in GLUEX using measurements of the hadronic cross section combined with the CLAS experience. The hadronic rate between any two photon energies E_1 and E_2 can be written as

$$R = \int_{E_1}^{E_2} n\sigma(E) \frac{dN}{dE} dE$$

where n is the number of target protons per unit area, $\sigma(E)$ is the hadronic cross section as a function of energy, and dN/dE is the photon energy spectrum. The photon flux is composed of a coherent and incoherent sum as detailed in Chapter 4. Background rates are dominated by the broad-band incoherent flux. The signal rates result from the photon flux in the coherent peak, which will depend on the radiator crystal structure and its orientation. The coherent peak will be optimized to the specific physics program. For our rate estimates, we use the typical case for the flux computed on a diamond radiator with the coherent peak at $E_\gamma = 9 \text{ GeV}$.

Both coherent and incoherent fluxes are proportional to the electron beam current and radiator thickness. Multiplying the number of electrons per second by the radiator thickness in radiation lengths gives the product N_0 which we will use in the following calculations. For conditions which we will refer to as “low intensity” (300 nA beam on a 10^{-4} radiator), $N_0 = 1.9 \times 10^8/\text{s}$. For the coherent peak at 9 GeV the tagged photon flux between 8.4 and 9.0 GeV is

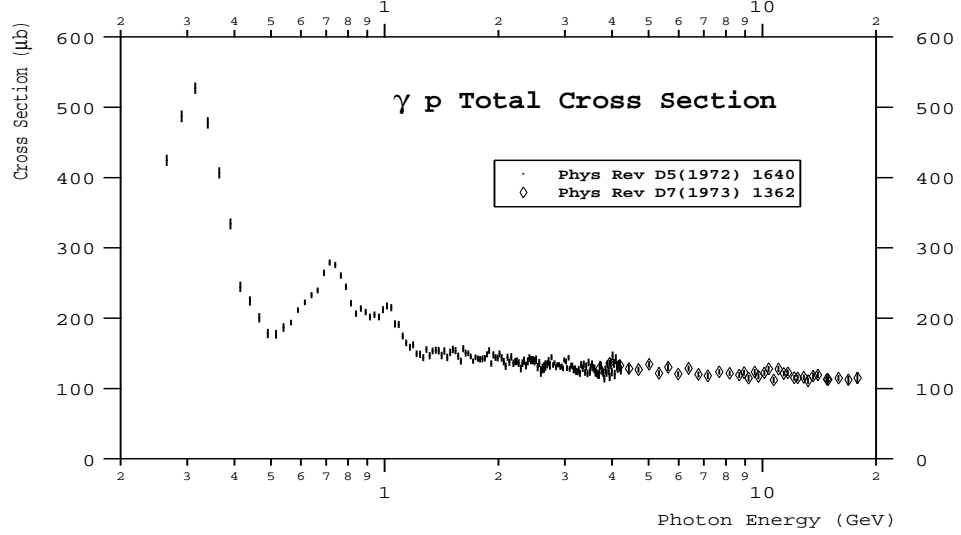


Figure 8.1: Total cross section for $\gamma p \rightarrow \text{hadrons}$ as a function of photon energy.

$R_{tag} = 0.14 N_0$. The average tagging efficiency over this interval is 0.375, so the tagged photon flux on target is $1.0 \times 10^7/\text{s}$. “High intensity” running, where the tagger becomes ineffective as part of the level 1 trigger, nominally corresponds to $N_0 = 1.9 \times 10^9/\text{s}$ and yields $10^8/\text{sec}$ tagged photons on target.

The total hadronic γp cross section¹ is plotted in Fig. 8.1. For the experimental conditions defined above and a 30 cm liquid hydrogen target, the total hadronic rate in the detector is

$$R_0 = 2 \times 10^{-4} N_0 \quad (8.1)$$

and a tagged hadronic rate

$$R_T = 7.4 \times 10^{-6} N_0 \quad (8.2)$$

For low intensity, the expected total hadronic rate is 37 kHz and the tagged hadronic rate is 1.4 kHz.

8.1.2 Trigger elements

We make some rudimentary assumptions about the trigger elements in order to estimate various rates. These assumptions are discussed further in the

¹We use measured cross sections [1, 2] with actual data obtained from the Durham Data Base [3]

Trigger section of this document. Initial commissioning of the detector at low rates will use a level 1 trigger to select events of interest. At higher rates a sophisticated level 3 software trigger² is required. We concentrate here on discussion of rates at lower photon beam fluxes.

The trigger consists of coincidences between several counter elements of the detector. It must select the tagged hadronic rate in the presence of accidentally coincident backgrounds. The first trigger element is the photon tagger, essentially a segmented scintillation counter. The rate in this counter is determined by N_0 , which is controllable (within limits) by the experiment.

The second trigger element is the start counter/vertex chamber. This detector package will provide position and timing information with sufficient resolution for track reconstruction. In comparing the demands of the GLUEX start counter to the CLAS experience, it is useful to note that the GLUEX target is inside a solenoidal magnetic field which will protect the start counter from the flux of low-energy Compton scattered electrons emerging from the target. The CLAS start counter does not enjoy this protection.

The tagger and start counters are small, and are therefore the best candidates for determining the precise event timing. For this discussion, we will assume that coincidences between them can be identified within a time window $\Delta T_1 = 15 \text{ ns}$.

Interesting events will have particles in the final state other than the one that satisfied the start counter requirement. These particles may be energetic, forward-going charged particles, forward or large angle photons, and/or charged particles with sufficient transverse momentum to reach the bore of the solenoid. Any particles of this type will be registered in other elements of the detector and these signals can be used as further requirements in the trigger. This refines the loose interaction definition given above. We refer to this collection of signals as the global level 1 trigger. As its elements are counters of extended size, we take a coincidence time window $\Delta T_2 = 100 \text{ ns}$ when the global level 1 trigger is required.

8.1.3 Accidental rates

The rate of interesting events given by Eq. 8.2 is 1.4 kHz ($N_0 = 1.9 \times 10^8/\text{s}$) and 14 kHz ($N_0 = 1.9 \times 10^9/\text{s}$) for low and high intensity beams respectively. However, various other processes will form accidental coincidences at the different trigger stages, and we need to recognize these. It is most important that these do not form the bottleneck for the data acquisition system, regardless of

²We are reserving level 2 for a possible intermediate level hardware trigger

our ability to reject them offline.

We consider two sources of accidental background. They are not entirely mutually exclusive, but we consider them separately for ease of explanation. The first (A_1) of these comes from purely random time coincidences between the trigger elements, in which case we compute the time overlap based on the various counter singles rates. The second (A_2) is more “physical”, considering hadronic photoproduction that is outside the tagging range, but in accidental coincidence with the tagging system.

First consider purely random coincidence events. A coincidence between the tagger and start counter loosely defines an interaction in the target. The rate A_0 of this coincidence is given by

$$A_0 = SR_{tag}\Delta T_1 \quad (8.3)$$

where S is the total rate in the start counter. Based on the experience in CLAS, we take $S = 0.03N_0$, scaled using appropriate factors for collimation and beam intensity. This is most certainly an upper limit because of the solenoidal shielding effect. For $R_{tag} = 2.7 \times 10^7/s$ we find $A_0 = 2.3 \times 10^6/s$, considerably larger than the tagged hadronic rate $R_T = 1400/s$. Further refinements are achieved by the global level 1 trigger.

The rate of the global level 1 trigger, $f_{L1} \times R_0$, is taken to be the total hadronic rate³ reduced by the rate for single pion production for $E_\gamma \leq 0.5$ GeV ($f_{L1} = 0.5$). A loose trigger which uses a charged particle track count in the start counter and requires neutral energy in the barrel and/or forward calorimeter (see Section 8.2.2 below) should easily be able to eliminate these low energy events. The accidental rate using both the interaction and global level 1 triggers is

$$A_1 = A_0 f_{L1} R_0 \Delta T_2 = S f_{L1} (0.28 \times 10^{-4}) N_0^2 \Delta T_1 \Delta T_2 \quad (8.4)$$

where we have substituted from Eqs. 8.1 and 8.3 and used $R_{tag} = 0.14 \times N_0$. The second accidental background comes from true hadronic events, and therefore would pass the global level 1 trigger. They are out of time with the precise RF signal, but that is much smaller than the online resolving time ΔT_1 of the interaction coincidence. Ignoring the “true” events that are part of this rate, one calculates

$$A_2 = f_{L1} R_0 R_{tag} \Delta T_1 = f_{L1} (0.28 \times 10^{-4}) N_0^2 \Delta T_1 \quad (8.5)$$

In order to evaluate the total accidental contribution numerically, correlations must be taken into account. This reduces the sum of the above estimates.

³The cosmic-ray rate is small and has been neglected.

For $N_0 = 1.9 \times 10^8/\text{s}$, the accidental contribution to the trigger is 7.3 kHz , and the tagged hadronic rate is $R_T = 1.4 \text{ kHz}$. We note that as the photon flux increases, the start counter and tagger lose their effectiveness in reducing trigger rates, so the trigger rate asymptotically becomes proportional to the hadronic rate. At higher currents, a DAQ system with a software level 3 trigger is required. A summary of the rates is shown in Fig. 8.2 as a function of electron beam current.

8.1.4 Rates in tracking chambers

At the high photon flux anticipated for GLUEX, one concern is that the occupancy rates in the drift chambers may be too high to allow reconstruction. In order to estimate these occupancies, a test of high intensity running with photons was performed in the CLAS detector.⁴ Measurements were taken at 10, 80, 250 and 320 nA with a 10^{-4} radiator, and rates were measured in the forward TOF scintillators (7.5-12.5 deg), the electromagnetic calorimeter (8-45 deg), and the drift chambers. The drift chamber occupancies at the highest current (320 nA) are given in Table 8.1.

Beam	Region 1 S1	Region 1 S2	Region 2 S3	Region 2 S4	Region 3 S5	Region 3 S6
Photon	2.3%	2.3%	0.3%	0.4%	0.7%	0.9%

Table 8.1: Drift chamber occupancies for each superlayer (in percent) for run 21998 at the maximum beam current of 320 nA (logbook entry #7031).

Although the conditions of the test did not duplicate precisely the conditions expected in GLUEX, reasonable estimates can be made by appropriate scaling. In Table 8.2 we compare the differences in target, collimation and beam energy. As the majority of background results from lower energy photons, we assume the energy dependence of the measured rates is small. The rates are scaled by a factor of 1.7 (ratio of target lengths) and the beam current is scaled up by a factor of 5.33, which is the expected collimation ratio. The drift chambers in region 1 are completely unshielded by any magnetic field in CLAS, whereas the drift chambers in region 2 are shielded by the field of the CLAS torus. The 2.2 T solenoidal field for GLUEX is expected to be at least as effective as a shield as the CLAS torus. Therefore, we expect the

⁴This information is taken from CLAS-NOTE-2000-004 High-Rate-Test by Elton Smith.

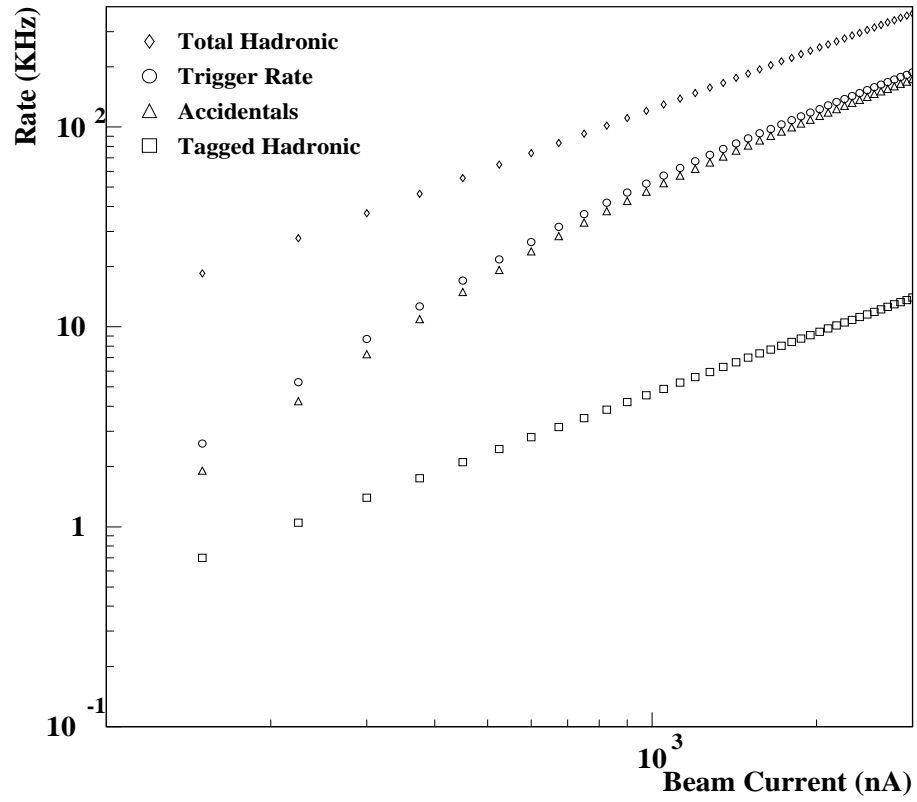


Figure 8.2: Estimated rates as a function of electron beam current. Plotted is the total hadronic rate and the estimated trigger rate, which is the sum of accidental coincidences and the tagged hadronic signal.

occupancies in the GLUEX drift chambers to be as low or lower than those in CLAS for comparable granularity.

	Hall B Test	Hall D
Beam Current	80 nA \rightarrow 320 nA	300 nA \rightarrow 3 μ A
Radiator	10^{-4}	10^{-4}
Collimation keeps	80%	15%
Target Length	18 cm	30 cm
Beam Energy	2.4 GeV	12 GeV
Trigger	Restricted	Open

Table 8.2: Comparison between conditions in Hall B during high rate test and anticipated running parameters for GLUEX. A current of 3 μ A in GLUEX corresponds to 10^8 photons/s in the coherent peak.

Extrapolating measured occupancies in region 2 to a current of 3 μ A (GLUEX with 10^8 photons/s in the coherent peak), we expect an occupancy of 0.6%. The rates are plotted versus electron current scaled to GLUEX in Figure 8.1.4. This is well below the typical operational limits of 2.3% imposed for the region 1 drift chambers in CLAS during electron beam running, a rate at which tracks can still be reconstructed with reasonable efficiency. We note that the extrapolated rates in region 1 for a beam current of 3 μ A is approximately 5%, exceeding usual operational limits by a factor of 2, but this figure is for a configuration which is completely unshielded by any magnetic field whatsoever and thereby represents an absolute maximum to the expected rates. We note that the operation of a polarized target in Hall B (which replaces the mini-toroid with a solenoidal field) allows running at twice the normal luminosity. Thus we expect that for comparable segmentation, raw rates in the GLUEX detector at the maximum design current will be similar to the current experience with CLAS. The conclusion is that the GLUEX detector should be able to handle rates up to 10^8 photons/s.

8.2 Trigger

8.2.1 Overview

In order to achieve the roughly 20-1 reduction in event rate, GLUEX will use a two-stage trigger, combining a hardware-based level 1 trigger with a software

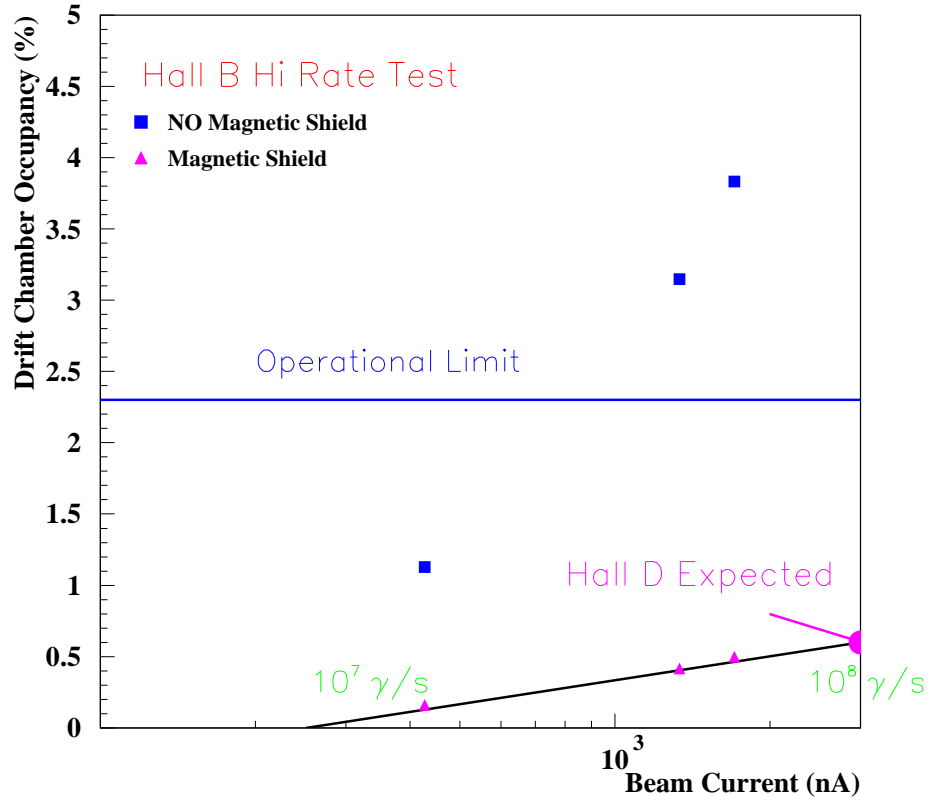


Figure 8.3: Drift chamber occupancies (scaled by target thickness = 1.7) plotted versus beam current (scaled by collimation factors = 5.33) expected for HALL D operation. The drift chambers in region 1 (squares) are completely unshielded by any magnetic field in CLAS, whereas the drift chambers in region 2 (triangles) are shielded from backgrounds by the main torus field. The nominal low current operation in GLUEX (10^7 photons/s in the coherent peak) corresponds to 300 nA. The 2.2 T solenoidal field for GLUEX is expected to be at least as effective as a shield as the CLAS torus.

(reconstruction) based level 3 trigger. An essential feature of the GLUEX design is to build pipelining into the entire trigger, digitizer, and data acquisition systems at the outset. This has the twin virtues of allowing adequate time for the level 1 trigger to do its job, while eliminating signal degradation involved in delaying the signals while the trigger operates. Pipelining in this way also allows us to upgrade from initial photon fluxes of 10^7 photons/sec to eventual fluxes of 10^8 photons/sec without any significant changes to the trigger/DAQ architecture. Eliminating conversion deadtimes will allow us to acquire events which occur very close together in time.

Figure 8.4 shows a schematic of the implementation of the GLUEX level 1 trigger. The level 1 trigger makes a decision based on detector elements which

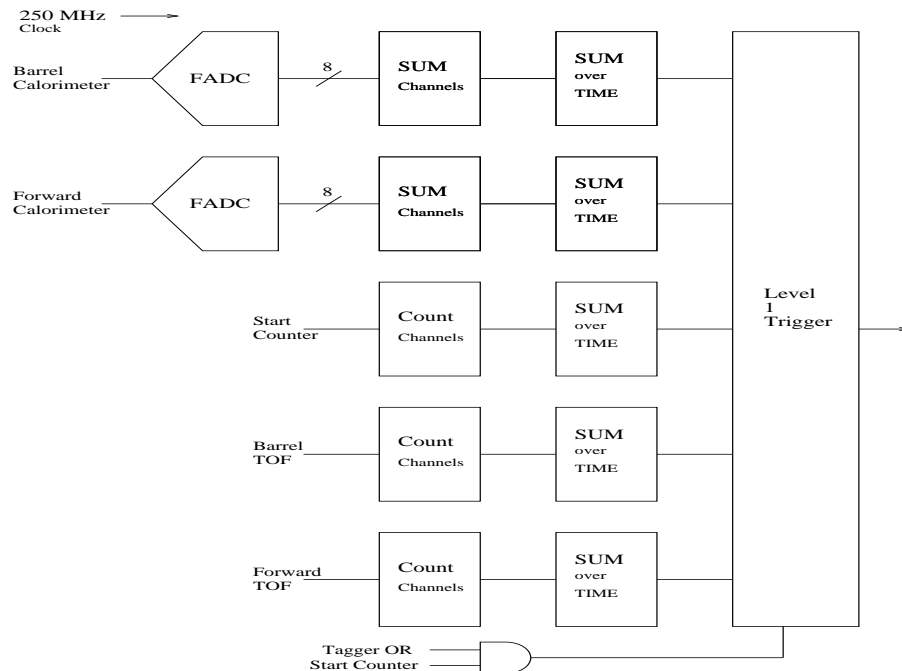


Figure 8.4: A schematic diagram of the GLUEX trigger.

measure hadronic multiplicities (track counts) and energies. In the schematic shown, the start counter and barrel calorimeter and forward TOF detectors provide the track count while the barrel and forward calorimeters determine the energy. A tight tagger OR/start counter coincidence is used as input to the level 1 trigger for low photon fluxes of $\approx 10^7$ photons/sec.

For high photon fluxes ($\approx 10^8$ photons/sec), the tagger OR/start counter coincidence is not a useful requirement, and the level 1 trigger will probably

only be able to cut the rate down from 385 KHz to around 180 KHz. Most of this background comes from multi-pion events caused by untagged (low energy) photons. In order to reduce this rate by a factor of 10, a very accurate reconstruction of the photon energy is required. Because of the complexities involved in accurately determining track momenta and then linking information from the different detectors, we believe the best approach is to use a software level 3 trigger embedded in the DAQ architecture, rather than to build a series of specialized level 3 trigger processors. This level 3 trigger will do a simplified full reconstruction of the event, using all of the data, in order to throw out events from low energy photons.

8.2.2 Level 1 trigger

The level 1 trigger consists of five subsystems, and a global trigger processor (GTP) which combines these five outputs into the global level 1 trigger. Each of the subsystems continuously (via a digital pipeline) computes a parameter, then compares it against a number of programmed value/function pairs. The trigger pipeline would sample input data at the rate of the FADC clock (250 *MHz*) or possibly at half that rate (125 *MHz*). A value function pair might be an energy value and a $<$, $=$, or $>$ function. When any of the value/function requirements is satisfied, the subsystem sends a timestamped subsystem event report (SER) to the GTP. The GTP is programmed with a number of different level 1 trigger configurations, each combining different value/function pairs from the subsystems, along with a trigger coincidence window (TCW) specifying the maximum time window for coincidence of the different trigger requirements.

The five level 1 trigger subsystems are:

1. **A track count** - obtained from the start counter. The start counter discriminator signals are used to create the prompt OR for coincidence with the tagger, but are also sent into a *track count pipeline* to determine the number of tracks. Two different track counts may be programmed, each with a $<$, $=$, or $>$ criterion attached.
2. **A track count** - obtained from the barrel calorimeter. The discriminator signals from the central calorimeter are sent into another *track count pipeline* which determines the number of tracks. This pipeline runs synchronously with the start counter track count pipeline. Two different track counts may be programmed, each with a $<$, $=$, or $>$ criterion attached.

3. **An energy sum** - obtained from the barrel calorimeter. The barrel calorimeter will be digitized by 8 bit, 250 MHz flash ADCs (FADC). All channels are then digitally added together (in a pipeline tree) to form the barrel calorimeter energy sum. The energy sum then passes through a shift register thus making available a time window. Successive samples within this time window are added together. This is analogous to the gate width in a conventional charge sensitive ADC. Two different energy values may be programmed, each with a $<$, $=$, or $>$ criterion attached.
4. **A track count** - obtained from the forward TOF. Discriminator signals from the forward TOF are sent into a *track count pipeline* which determines the number of tracks. This pipeline runs synchronously with all the other level 1 pipelines. Two different track counts may be programmed, each with a $<$, $=$, or $>$ criterion attached.
5. **An energy sum** - obtained from the forward calorimeter. This sum is constructed in the same manner as for the central calorimeter, except that the selection of which digitized analog sums are added together to form the forward energy sum, is programmable. Two different energy values may be programmed, each with a $<$, $=$, or $>$ criterion attached.

As mentioned above, the GTP may be programmed with several different triggers. Programming a single trigger means selection of

1. Either a minimum, maximum, or exact number of tracks in the start counter.
2. A minimum, maximum, or exact number of tracks in the barrel calorimeter.
3. A minimum, maximum, or exact number of tracks in the forward TOF.
4. A minimum or maximum for the global energy in the barrel calorimeter.
5. A minimum or maximum for the global energy in the forward calorimeter. Certain areas might be programmed out of this sum.
6. The appropriate boolean combination of elements 1-5.

The trigger will have the capability to have at least eight simultaneously defined triggers. This trigger is very flexible and can be programmed to be very loose (say one track in the start counter) or very tight and complex (specific track counts and energy thresholds in each detector). Examples of triggers which can be programmed in this model include:

1. At least two tracks in the start counter AND at least one track in the downstream TOF.
2. At least one track in the start counter AND a minimum energy in the downstream calorimeter.
3. At least two tracks in the start counter AND at least one track in the barrel calorimeter AND a minimum requirement of energy in the barrel calorimeter AND a minimum requirement in the forward calorimeter.

All subsystems will run synchronously and will be timed so that the time stamps from average momentum tracks ($\sim E_{\text{beam}}/3$) will match at the GTP. Higher and lower momentum tracks will be slightly out of time, but this effect should be less than 20 ns , and this is compensated for by programming the TCW value. The synchronous output of the level 1 trigger will then be ANDed with the coincidence of the tagger OR and the start counter OR. This allows the timing to be determined by the tagger and start counter, and removes the synchronous nature of the trigger.

The rate of the global level 1 trigger, $f_{L1} \times R_0$, is taken to be the total hadronic rate⁵ reduced by the rate for single pion production for $E_\gamma \leq 0.5 \text{ GeV}$ ($f_{L1} = 0.5$). A loose trigger which uses a charged particle track count in the start counter and requires neutral energy in the barrel and/or forward calorimeter should easily be able to eliminate these low energy events. The resultant level 1 trigger rate is about 180 kHz . We note, however, that 80% of the hadronic rate comes from photons with energies below 2 GeV . This energy cut, which would require a more sophisticated trigger, would reduce the level 1 rate to 70 kHz .

8.2.3 Trigger simulation

As mentioned above, background events are typically due to low energy photons, resulting in low energy events. Not only are these background events lower in energy, but they are also less forward, due to reduced Lorentz boost. Thus, good events typically deposit a larger fraction of their energy in the forward calorimeter, and have more tracks and hits in the forward time-of-flight. The goal of the Level 1 trigger is to use these differences to cut as large a fraction as possible of the background events, while minimizing the number of good events lost. The goal for the data reduction in the level 1 trigger is to

⁵The cosmic-ray rate is small and has been neglected.

remove at least 50% of the background events, without losing more than 0.5% of the good events.

In order to test the the trigger, the six reactions listed in Table 8.3 were simulated and studied. The simulated events include two low energy delta production channels, and four interesting physics channels at low (background) and high energies. Reaction events were generated using *Genr8* [4]. After generation the events were then run through *HDGeant* [5] for simulation. This provided the necessary data needed. For each reaction 10,000 events were generated giving 120,000 events.

A function of the form given in Eq. 8.6 was used as the basis for deciding cuts. When the calculation is less than Z the event is cut. A genetic algorithm was used to optimize the coefficients and Z. The fitness function was weighted such that keeping good events was given a higher score than cutting background events. If good events were cut then it would be penalized and if it cut too many then the score received was zero. As shown in Table 8.3 the best set of coefficients cut nearly all of the delta's and most of the low energy background events. On average 72% of the background events are cut, while no single good event channel lost more than 0.5%.

$$\begin{aligned}
 Z \geq & A * [NumberTracksForwardTOF] & (8.6) \\
 & + B * [EnergyForwardCal] \\
 & + C * \frac{[EnergyForwardCal + 1]}{[EnergyBarrelCal + 1]}
 \end{aligned}$$

8.3 Data acquisition

8.3.1 Overview

The GLUEX data acquisition system is being designed to accept a 200 KHz Level 1 input rate, and will be pipelined so as to incur no deadtime. Front-end boards will continually digitize and store several microseconds worth of data to allow time for the Level 1 trigger decision. When a Level 1 accept arrives the boards will extract the appropriate time slice of data from the digitizing memory and move it into a large secondary memory store. Readout controllers will collect data from many boards over a backplane, then transmit the data to event building processors over a network. Note that the readout controllers likely will not need to run a hard real-time operating system (e.g.

Reaction	Energy(GeV)	Percent Cut
$\gamma p \rightarrow \rho^0 \pi^+ n \rightarrow n \pi^+ \pi^- \pi^+$	1	67.99%
	2	41.68%
	9	0.05%
$\gamma p \rightarrow \rho p \rightarrow p \pi^+ \pi^-$	1	70.48%
	2	54.82%
	9	0.50%
$\gamma p \rightarrow X^*(1600) n \rightarrow (\eta^0 \pi^+) n \rightarrow n \pi^+ \gamma \gamma$	1	90.10%
	2	56.24%
	9	0.11%
$\gamma p \rightarrow X^+(1600) \Delta^0 \rightarrow (\pi^+ \pi^+ \pi^-) (n \pi^0) \rightarrow \pi^+ \pi^+ \pi^- n \gamma \gamma$	9	0.23%
$\gamma p \rightarrow \Delta \rightarrow n \pi^+$	0.337	99.99%
$\gamma p \rightarrow \Delta \rightarrow p \pi^0$	0.337	98.75%

Table 8.3: Trigger cut rates for reactions and their energies.

VxWorks) due to the large memories on the digitizing boards, an important simplification.

Complete events will be shipped from the event builders via a network to a large farm of Level 3 processors. The Level 3 farm will reduce the event rate by approximately a factor of 10 before shipping the remaining events to event recording processors, which will then write the events to a staging disk in preparation for transfer to tape. We are designing the system to handle a recording rate of 100 Mb/s. During initial running at low luminosity (10^7) this system will be able to record all events to disk, and no Level 3 rejection will be needed.

Most of the hardware components needed to build the DAQ system described above are available now or will be available soon, so there should be no problem finding hardware a few years from now. The main challenge will be to develop the DAQ, online, monitoring, and controls software.

8.3.2 Data flow and rates

GLUEX will have approximately 12500 FADC channels. Assuming a typical occupancy of 2%, a 250 MHz, 8-bit FADC, a time window of 100 nanoseconds, and readout of the full time window, the total amount of FADC data would potentially be: $12500 \text{ channels} * 0.02$ (occupancy) $* 25$ bytes/channel = 6.25 Kbytes per event.

The 25 bytes/FADC channel will be used to extract an energy and a time signal. Previous work [6, 7] indicates that a time resolution better than the FADC sampling interval can be achieved by fitting the FADC waveform (see also Chapter 7). Thus we plan to reduce the FADC data to an energy, time, and channel identifier in real-time using special on-board hardware. The amount of data per hit will drop from 25 bytes to 10 bytes per FADC channel, thereby lowering the total FADC data to a more manageable 2.5 Kbytes per event.

In GLUEX there will be approximately 8000 TDC channels so the data volume for the TDCs will be: 8000 channels * 0.02 (occupancy) * 4 bytes/channel = 640 bytes per event.

There will be little data from devices other than TDCs and ADCs (scalars, latches, etc.) so the total event size will be about 4 Kbytes per event. Taking 5 Kbytes per event as the design goal gives 5 Kbyte/event * 200 KHz = 1 Gbyte/sec off the detector. Assuming 100 front-end VME crates (cPCI will need more) gives a backplane rate of 10 Mbytes per second, easily handled by current technologies.

Event building will be done in parallel on 8-16 event building processors. Event analysis will be performed in parallel on 50-200 Level 3 farm processors (see below). Event recording will be done in parallel on 2 to 8 event recording processors. In all cases existing network switches can easily route the volume of data between stages. Note that we are investigating use of advanced (e.g. layer 7 routing) network switches to further simplify transfer of data between stages.

8.3.3 Level 3 trigger

If the Level 1 trigger rate for low intensity running (10^7 tagged photons/s) is less than 20 KHz, or 100 Mbytes/sec, the Level 3 trigger farm will not have to cut any events since the DAQ system is being designed to handle this rate to disk. In high intensity mode, where the Level 1 rate may be as high as 200 KHz, the Level 3 trigger must be able to reduce the event rate by a factor of ten.

Most of the unwanted events result from an untagged, mostly lower energy photon interacting in coincidence with a tagged photon. To reject these events Level 3 must be able to estimate the energy of the photon which produced the event. This involves reconstructing tracks, matching them with the calorimeters, and adding additional energy deposited by neutral particles in the calorimeters. This is most simply and easily done in a commodity processor Level 3 farm, rather than in specialized hardware.

We estimate the required processing power required as follows. The Hall B

	Low Rate	High Rate
Event Size	5 KB	5 KB
Event Rate to Farm	20 KHz	200 KHz
Data Rate to Farm	100 Mbytes/s	1000 Mbytes/s
Num Links to Farm	1	10
Data Rate per Link	100 Mbytes/s	100 Mbytes/s
Link Technology	Gigabit Ethernet	Gigabit Ethernet
Events/s per Link	20000	20000
SPECints/ev for L3	0.1	0.1
Num SPECints/link	2000 SPECints	2000 SPECints
Num SPECints/link x 2	4000 SPECints	4000 SPECints
Num 200 SPECint processors/link	20	20
Total Num 200 SPECint processors	20	200

Table 8.4: Rates, sizes, and processing requirements for the Level 3 trigger.

online hit-based event reconstruction system obtains 3% momentum resolution using about 5 milliseconds of cpu time on a 20 SPECint processor, or about 0.1 SPECint per event (full reconstruction with better than 1% resolution takes about 45 milliseconds). Assuming the same for *GLUEX* gives 20000 SPECints total for the full Level 3 farm at 200 KHz event rate. Assuming 50% processor utilization (due to I/O overhead, etc.), approximately 40000 SPECints or 200 processor boxes at 200 SPECint each are needed (150 SPECint boxes are currently running in the JLab farm system). Depending on the improvement in cpu performance over the next few years, far fewer boxes will likely be required, perhaps 1/4 as many.

Table 8.4 shows the rates, sizes, and processing requirements for the Level 3 trigger.

8.3.4 Monitoring and Control

Monitoring and control tasks include hardware configuration and control (“slow controls”), bookkeeping, online event monitoring, alarm systems, and messaging systems. These are less demanding tasks than data acquisition in *GLUEX*, and should not present unusual challenges. We plan to follow some examples from Hall B, but to also make use of lessons learned there. In particular, we

plan to integrate offline data analysis tools with the online software at the outset to reduce the total cost of software development.

The framework for slow controls will be uniform for all subsystems in GLUEX, but the framework choice is not obvious. VME-based EPICS works in Hall B, but does not mesh well with the online requirements and has proven to be manpower intensive. In fact, a number of Hall B systems do not use EPICS or VME, but instead resort to CAMAC or other options. We believe that an open, message-based system that takes advantage of commodity hardware and software, and that implements a uniform user interface to diverse underlying hardware is best. The JLab Data Acquisition group is currently developing an agent-based system meeting these requirements.

Bookkeeping tasks include all recordable activities of the experiment other than raw and calibration data. We expect this will be done using object/relational databases. Current commercial and public domain database technology should be adequate.

The alarm and messaging framework allows sub-systems to communicate their state to monitoring programs and operators. This system needs to be integrated across the entire online, DAQ, and database systems in a simple, uniform manner. The scale and performance requirements of this system are modest, and similar to other systems running or in development at Jefferson lab.

List of Figures

8.1	Total γp cross section	4
8.2	Estimated rates as a function of electron beam current.	8
8.3	Drift Chamber Occupancies	10
8.4	A Schematic of the GLUEX Trigger	11

List of Tables

8.1	Drift chamber occupancies	7
8.2	Comparison to Hall B	9
8.3	Trigger cut rates	16
8.4	Rates and Processing Requirements	18

Bibliography

- [1] Armstrong *et al.* *Phys. Rev.*, D**5**:1640, 1972.
- [2] Caldwell *et al.* *Phys. Rev.*, D**7**:1362, 1973.
- [3] The Durham online database is located on the WWW at <http://durpdg.dur.ac.uk/HEPDATA>.
- [4] P. Eugenio. *Genr8*: A general monte carlo event generator. Technical report, Carnegie Mellon University, 1998.
- [5] R. Jones, 2001. The HDGeant Monte Carlo Program.
- [6] J. Va'Vra. *Nuclear Inst. and Meth.*, A**244**:391–415, 1986.
- [7] D.G.Cussans. *Nuclear Inst. and Meth.*, A**244**:277, 1995.