DUMAND-UWA-90-2
HDC-17-90

PROCEEDINGS

of the DUMAND 1990

TRIGGER WORKSHOP

July 12-13, 1990

Seattle, Washington

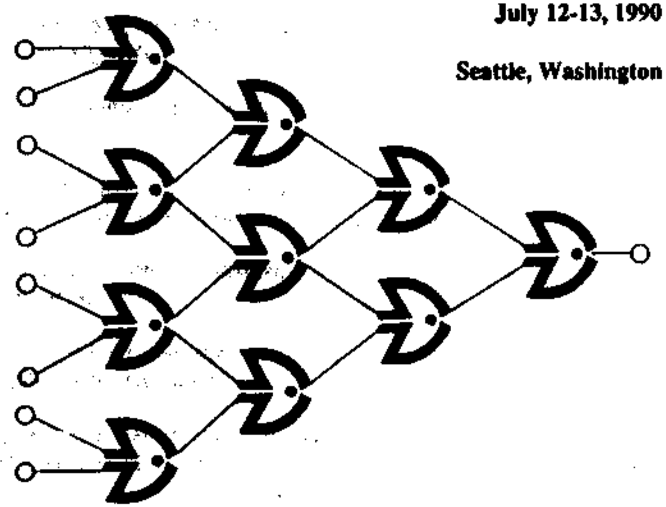

EDITORS:

KENNETH. K. YOUNG
and
R. JEFFREY WILKES

Department of Physics, FM-15,
University of Washington
Seattle, WA 98195

First published September, 1990

Department of Physics, FM-15
University of Washington
Seattle, WA 98195 USA

TABLE OF CONTENTS

FIGURES & ILLUSTRATIONS

FOREWORD

With DUMAND-II moving rapidly from proposal to reality, and deployment of the first components little more than a year away, it was important to bring together collaboration members active in planning and development of the trigger system for a brief but very productive workshop. We offered to host the meeting in Seattle so that our entire group, the newest members of the collaboration, could participate fully. The discussions brought out a number of crucial issues for clarification and decisive action, including questions as fundamental as Optical Module front-end electronics specifications and as detailed as chip selection for the String Bottom Controller interfaces. A new basic format for transmission of data to the shore station was agreed upon, and a basic window size of 2 ns for OM signal digitization was recommended. Decisions taken are summarized on pages 78-80. These actions were communicated to the collaboration by electronic mail and fax, and will result in modifications to the official DUMAND Specifications file as appropriate.

We have collected contributions from participants, which differ in some cases from initial presentations at the meeting. Questions or requests for further details should be directed to the individual authors. Viewgraphs are reproduced for presentations by participants who did not submit written contributions.

The Seattle Trigger Workshop could not have taken place without the assistance of Linda Vilett, and these Proceedings were prepared with the aid of Phoebe Bosche. We wish to give special thanks to Dr. Robert Spindel of the UW Applied Physics Laboratory for allowing us to use APL facilities.

Editors,
Kenneth K. Young,
R. Jeffrey Wilkes
Seattle, 8/15/90

LIST OF PARTICIPANTS

NAME	AFFILIATIONS
Jeff Bosel	University of Hawaii/SOEST
Paul Boynton	University of Washington
Ugo Camerini	University of Wisconsin
Dick Davisson	University of Washington
Phil Ekstrom	University of Washington
Eric Hazen	Boston University
Matt Jaworski	University of Wisconsin
John Learned	University of Hawaii/HDC
Jere Lord	University of Washington
Med Webster	Vanderbilt University
Mark Wilber	University of Washington
R. J. Wilkes	University of Washington
Akira Yamaguchi	Tohoku University
Ken Young	University of Washington

Visitors:

Bob Spindel	Applied Physics Lab, University of Washington
Steven Strauz	University of Washington
Bob Williams	University of Washington
Mike Williamson and Art Wright	Williamson Associates, Seattle

Physics Goals and Triggering DUMAND II

John G. Learned
Hawaii DUMAND Center

30 August 1990

Abstract

We discuss the options for on-line triggering for DUMAND II, beginning with the physics which is desired to be saved and estimating some efficiencies and noise rates for proposed triggers. By triggering we mean the algorithms employed for first level filtering. The suggested triggers utilize total energy deposition, and string 2 and 3 fold neighbor coincidences to efficiently capture relativistic events such as muons and cascades. This will also capture some types of exotica. Another type of trigger is needed to capture slowly moving particles, and yet another for possible supernova detection. Both the latter triggers rely upon monitoring of optical module count rates above various thresholds. It appears that all the physics goals can be met with a trigger processing scheme involving one layer of programmable processors, communicating with one data harvesting computer to achieve the necessary 3 order of magnitude front end rejection in incoming raw data rate.

DIR-12-90

1 Introduction: Where's the Physics?

The data stream coming to the shore station must be searched as it arrives for the physics data we want to save. We cannot save everything, and all not immediately recorded will be lost. Thus we must think carefully about how best to save both the expected physics (muons and neutrino interactions), try to find a few "long shots", and also to make ourselves available for fortuitous discovery. In

the following I will first list the physics we have considered as worthy and possible for exploring with DUMAND II. This is to be differentiated from the phenomena sensed by the detector: different physics may have the same requirements for triggering.

1.1 High Energy Neutrino Astronomy

The main goal is, of course, the detection of high energy muon neutrinos, made manifest as single muons. The muons produce conical light wave fronts that are detectable out to about $24m$ from the track [1]. The triggering problem addressed herein thus naturally involves use of the strong spatial and temporal clustering of the photodetector signals. High energy muons ($> 1TeV$) will give more light than lower energies, and be easier to detect, so that if we design for minimum ionizing muons we will be conservative[2]. Stopping muons are also worthwhile saving, since their rate relates to the neutrino spectral index, but produce more feeble signals. Just how efficient a trigger is at picking these out, and whether they can be fitted, needs computer simulations, now being carried out by Vic (and others I hope).

1.2 Neutrinos from Fermilab

In order to detect the $20GeV$ neutrino beam from the Main Ring Injector proposed for Fermilab, we must be sensitive to low energy muons, the lower the better. Given the muon range of roughly $5m/GeV$ one can conclude without benefit of Monte Carlo that given the $40m$ string spacing and $10m$ module spacing on the strings, that detection by a half dozen modules will require a range of at least $40m$, or an energy threshold of about $8GeV$. Note that knowing the direction and time for such events should help substantially in beating down the background.

[1] Detection probability of 1/2 for a muon at impact parameter of 23.5m, face on direction of photon arrival, see Figure 9. The probability is still 30% at 50m.

[2] See the Monte Carlo study report of A. Okada, ICRR-209-90-2, for details. He shows, in Figure 3(b) therein, that the effective area grows slowly to $1TeV$ muon energy, and then grows more nearly linearly (just as dE/dx).

1.3 Cosmic Ray Muons

Downgoing muons from interactions of cosmic rays in the atmosphere overhead are interesting for various studies, ranging from direct production to muon astronomy. (The latter only being viable if high energy muons are indeed anomalously produced by the "gamma rays" seen at around $10^{15}eV$.) Obviously, if we make a trigger that is not direction sensitive, we will collect muons of all origins with the same trigger. The downgoing atmospheric muons constitute the highest expected rate of events in the array, about 3.[illegible] in contrast to upcoming (and side going) muons from atmospheric neutrino interactions, which will occur at only about 10/*day*.

1.4 Multiple Muons

Multiple downgoing muons are useful for studies of cosmic ray composition. Since these events will produce more light per PMT and more struck PMTs with similar timing structure, one would expect also that whatever trigger is selected for single muons will work as well or better for multiple muons. Significant efforts are yet needed to understand how well we can extract the science, but triggering appears to be easy.

1.5 ν_es and Hadronic Cascades

When a neutrino interacts in or near the array, the light produced by the cascade of particles from the nuclear vertex may also be detectable. In the case of a charged current electron neutrino interaction, we may detect the electromagnetic cascade. While the rate of such events will be less than for through going muons, they are well worth studying. Not much work has yet been done with the Monte Carlo programs on this subject however, so the energy threshold is still uncertain, but as can be seen in Figure 16, the effective volume is significant even at 10*GeV*. In terms of the signal in the detector, the cascade will be a few meters long with brightness in proportion to total energy. Most of the light will travel at the Cherenkov angle to the cascade direction, but there is some in every direction. In terms of timing, the light looks somewhat like a point source, but directionally peaked in amplitude, see Figure 17.

There is, of course, also an intermediate class of events, those with showers and muons. This class would be subsumed by the above, and comes presumably

free. How well such events can be fit and distinguished from either of those classes remains to be explored in simulations.

1.6 Anomalous Interactions

We should not forget some of the exotica claimed from old underground cosmic ray experiments, because our size may reveal what has only been hinted at previously. A good example is the large electromagnetic cascades seen in the Kolar Gold Fields. Such events contained of the order of a TeV of electromagnetic energy (which makes one think of electron neutrino interactions), but with the peculiar characteristic of a large ([illegible]) opening angle between sub-cascades. Such a phenomena, if confirmed, could be the opening to preonic structure, for example, and thus is well worth making sure we can save (though fitting is not likely to be easy).

In general, any interaction involving particles moving near the speed of light in vacuum will result in something that looks like either a line source (muon) or a point source (cascade), or both. So, it seems that the triggers that catch the physics above will do for any other wonders we have not yet thought about, but which involve simultaneous relativistic particles. The main effort for triggering considerations is to push downwards on the energy threshold.

1.7 Slow Massive Particles

There remain two other classes of phenomena which we have considered, which are not caught by the above, and they are both worthy of some fair amount of effort: massive slow particles and supernovae.

By massive slow particles I mean anything slow compared to a muon, but fast compared to a fish or even the speed of sound. The likely velocities are in the range from the speed of light down to meteoritic velocities, or $1.0 \geq \beta \geq 10^{-4}$, with array crossing times between $1\mu sec$ and $10msec$. One example of such an object would be a "nuclearite" as espoused by Glashow and DeRujula to solve the dark matter problem, which would be composed of a heavy quark bag which leaves an ionized glowing trail.[3] Another example would be a magnetic monopole of the variety proposed by Rubakov, which gobbles nucleons along it's path, reradiating them as pions and such, again making much light. The latter would

[3] See Vic Stenger's note about this HDC-3-87, 1987

be less bright than the former, but would have temporally tight photon emissions. Both can be characterized as causing large light levels in PMTs distributed along their path. They may not produce tight time coicidences between modules, of the type caught by the primary muon trigger, so they form a class needing special discussion.

Such particles have been sought by existing underground detectors, and many limits published, without even a hint of a signal. However, DUMAND, being 50 times larger in area than IMB (as a benchmark), can make substantial progress with a few years of live time, moving the limits down an order magnitude or more.

1.8 Supernova Neutrinos

The final class so far identified, is due to the interactions produced by a supernova within our galaxy. On the basis of observations in Kamiokande and IMB from SN1987A we have a good idea of how many events to expect: a collapse at $5kpc$ should give roughly 1 interaction per ton of water, spread over a time of less than $10 seconds$, or about 2×10^5 events in the detector volume. Sadly they are all of energy far below (by $100\times$) our individual detection threshold, so we have no chance with the fast coincidence techniques of detecting the individual interactions. The total energy deposition is about $20TeV$, which while emminently detectable in one few nanosecond event, is not at all obviously detectable spread over 10seconds. For comparison, the total K^{40} decay energy deposited in the detector volume in the same time is about $2.6 \times 10^{16}eV$, more than $1000\times$ greater. This would seem to make the detection hopeless, except for local (at one PMT) coherence in time.

While the K^{40} events are all around $1MeV$, the mean of the supernova events is about $10MeV$, and the tip of the spectrum extends to $50MeV$ or more. One can think of each PMT as comprising a poor man's detector, which has some chance to collect a multi-PE signal from a fortuitously aimed event throughout a fairly large volume. If, for example, we have a 1 in 100 chance of getting a multi-PE hit from throughout a volume of radius $10m$, then we could have 8000 multi-PE hits in $10sec$. If the background rate above the as yet unspecified multi-PE threshold is, say, 1% of the K^{40} rate, then the increase in the multi-PE rate due to the supernova could be 20σ.

However, this handwaving argument has quite a few "ifs" in it, and should only be taken as motivation to work on the problem. For present purposes,

however, what I want to indicate is simply that (as far as I can see), the only hope we have of detecting a supernova with DUMAND II is via the rate of multi-PE events. Perhaps our best strategy will not be to trigger on superova-like events, but only to record those which are supernova like, with the intention of *a posteriori* correllation with other detectors with better low energy sensitivity.

One might object that it is not worth spending much effort on supernova detection because there are other detectors that can do far better. As we saw from the case of SN1987A, there certainly is much room for confirmatory detections. Moreover, if we are lucky and have a collapse within a kpc, then DUMAND II data could be very important in examining the time distribution, because of our high rate handling capability, where a detector such as IMB will be completely jammed by full buffers.

1.9 Types of Phenomena

Summing up, it seems that all the physics possibilities we have identified so far can be placed into 4 categories in terms of the phenomena to be sought by the triggering system:

- muon like (line source moving at c)
- point source like (localized in time and space)
- slowly moving ($\beta > 10^{-4}$) singles rate increase traversing array
- uniformly distributed multi-PE rate increase (up to $10sec$)

Perhaps we need to associate some priorities with these categories. Clearly the first two dominate our physics goals with DUMAND II. The latter are long shots, but of great physics value if found. Note that while trigger efficiency is important for the first two it is not for the latter.

2 What kind of Triggers?

The design goal for the on-line triggering equipment at the shore station is to tag, on the fly, the physics we want to save for further processing. The total data rate to shore is vastly too much to save entirely. $\sim 2 \times 10^8 bits/sec$ [4] or $\sim$ one $8mmtape/2min$! We must select the interesting physics and discard the rest; anything not saved immediately will be lost forever. (Certainly we will want to save a random sample of the unselected incoming data.) Naturally we will save as much as possible, but we cannot reasonably save more than a tiny fraction of the data stream. Restricting the data saved to one $8mm$ tape per day implies saving $\leq 10^{-3}$ of the incoming data at the shore station. Anyway, if we cannot keep up with processing on-line it is hard to imagine how we could later catch up, unless we only wanted to go back and search the archives for some event at a particular time. (In fact we probably can comfortably save all the physics data selected by the on-line analysis on disk, with tape serving only as backup, in case of disk crash, and as archive to permit the recovery of subtleties that may escape the on-line filter.)

We are thus in the uncomfortable position, typical of modern counting experiments, where the decisions made about the front end electronics will determine what phenomena we can observe. As we shall see below, that is no problem for the major goal of the project, neutrino astronomy employing high energy muons as the neutrino pointers, but we wish to be careful to attempt to allow for serendipitous discovery too.

The first level triggers must meet two criteria:

- *efficiency in extracting the physics, and*
- *suppression of background to a rate manageable by the next level*

By triggering efficiency we mean the probability of collecting events which might later be successfully fitted. The definition is thus a bit sloppy, because fittability is not easily defined, even for throughgoing muons. 1% loss is a desireable goal (99% efficiency); however as typical of such experiments, realistically we expect to live with data losses of order of 10%. If we are to do such measurements as neutrino oscillation studies with the atmospheric neutrinos, then we must work at knowing our effective area well, including the effect of triggering efficiency, which will depend upon angle.

For suppression of random background light induced triggering, the goal must be to keep the rate of triggering below that which will saturate the data harvesting computer. We do not yet know what that rate is precisely, but we shall assume

[4] 216 modules, with 60,000 counts per second, at 16 bits per count.

herein that we want to keep it to $< 10^3 triggers/sec$ for computer processing. If we take a trigger as comprising $2\mu sec$ of data (assuming that SN triggers and slow particle triggers contribute little to the total trigger rate), then $1000 triggers/sec$ would harvest 1/500 of the data stream. This is already close to the amount of data that we can afford to store, so we can contemplate saving everything passing the first level trigger[5]. Hence, once the first level processor has thrown away the obviously random data, we can make rather loose criteria on saving events thereafter, and the fitting computer should have plenty of time to run everyone's fitter, on-line.

Note that if we have a first level hardware trigger rate of $10^3/sec$, the software filtering at the second level and beyond will have to get another rejection factor of 10^6 for upcoming muon tracks in order to make the random noise background to neutrino events negligibly small ($<$ 10% of atmospheric neutrinos). (This is distinct from the fitter requirement of not reconstructing downgoing muons wrongly as upcoming, which is about a factor of 2×10^{-3}).

2.1 Optical Noise

A few words are in order about the assumptions relative to the incoherent firing of the optical modules. We learned in the SPS experiment that the background rate in the ocean would give us typically 30,000 counts per second from the Hamamatsu PMTs. We expect that with improved module efficiency, that may be nearer 60,000 counts per second in DUMAND II. The individual K^{40} decays generate so few photons (about 40) that the probability of a PMT getting 2 or more PEs from a given decay is $\sim$ 1% of the overall PE rate. The probability that 2 neighbor modules see the same K^{40} decay is totally negligible. Thus higher pulse heights than one PE will come from within the PMT, due to the sloppiness of the electron multiplication, or due to internal effects such as ion feedback.

Bioluminescence will occur sometimes, we believe. Our measurements are very limited, and not in a real equilibrium situation, but we think that random bioluminescent events may occur perhaps 1% of the time (as creature flashes). It could be that a cloud of luminescing material will drift into the array once in a while, and we assume that we will simply be off the air for that period,

[5] With simple zero suppression, 1000 events per second would be reduced to $\sim 10^6 bits/sec$ or one 8mm tape every 5 hours. Another factor of ten would then make for a comfortable recording rate.

which we expect to be infrequent. For the situation of a local flash of light (which flashes we know last for durations of typically one second) we will simply gate off the temporarily blinded PMT. In order not to overflow the buffer at the SBC digitizer, this will have to be done at the optical module. One of the jobs of the shore station circuitry will be to keep constant track of the status of every module in the array. We anticipate that the effect of such local flashes will be restricted to one module, and will not generate excess coincidences due to correllated increases in random rate, and even if they do occur, they will be almost all at the one PE level.

2.2 Simple Coincidence?

All triggering schemes make selections in four-space (and maybe more, if we include charge, for example). The simplest trigger, beloved of all experimentalists, is the time coincidence, requiring some threshold number of signals within a predetermined time window. This works very well for an experiment such as IMB, in which the major trigger is simply the observation of more than 10 photomultiplier (PMT) hits from anywhere in the detector within 100 ns. This will not work in DUMAND, as long realized, because of the substantial noise rate due to K^{40} in seawater (expected count rate of $R_m = 60,000/sec/module$*) and the large size of the array ($\tau = \sim 1\mu sec$ across). In calculating random coincidence rates (we do assume that the PMT noise rates are not correllated), the important quantity is the expected number of hits in time τ, M, which we can write as

$$M = R_m \times N_m \times \tau \approx 13,$$

where N_m is the total number of optical detectors, 216 in DUMAND II.

Thus 13 PMT coincidences will occur all the time and even 25 PMT coincidences will occur at a rate of ~ $10^1/sec$. The distribution of the number of module hits for through going muons is illustrated in Figure 1. The mean number of PMTs hit by a throughgoing muon is about 13, so such a simple coincidence is close to being useful (13 noise hits p[illegible] correllated muon induced hits). However it seems that the trigger described next is somewhat better.

I have run a Monte Carlo calculation for cascades of elementary particles, with a 1/E spectrum extending from 1GeV to 10TeV, and the result is shown

*Note that we have taken the practice of employing 100,000counts/sec as the safe design value for noise calculations.

in Figure 10 (fittable events are taken as those with more than 5 non-noise hits). One sees a rapidly falling distribution, which points to the importance of seeking efficiency for small numbers of hits. The question of whether the smaller of these events can be fitted is not addressed herein.

2.3 An Energy Trigger TE

A slightly more sophisticated trigger, which we designate TE, could employ the total number of photoelectrons (PE) in the array. Since the K^{40} noise will be almost entirely single PE uncorrelated (in time or space) noise, the summed PE rate will be close to the random PMT count rate. The mean number of PEs per through going muon event is about 23, as illustrated in Figure 2. The reason for this is, of course, that the muon is likely to pass close to one or more modules and produce large pulse heights in a few of them. The random rate for events the size of typical muon events would be, assuming Poisson statistics,

$$R_k = P(\geq (M+k) \mid M)/\tau \approx 0.33/sec.$$

for $k = 22$, $M = 13$ and $\tau = 1\mu sec$.

Since 70% of the muons producing > 6 hits in the array have more than $13PE$, a trigger threshold of $27PE$ that would include most muons would have a random triggering rate of about 452/sec. I am not very sure of this calculation, however, since I have assumed that the distribution of photoelectrons in the whole array obeys Poisson statistics.

The results of the calculation for cascades is shown in Figure 11, where once again the distribution falls rapidly with increasing numer of photoelectrons. In order to say something quantitative about the efficiency of the TE trigger for catching such events we need to use a model spectrum. An unrealistically flat spectrum was employed in the calculation (for computational ease in generating the effective volume versus energy show in Figure 16).

Implementation of such a trigger can be carried out with the specialized processors which Wisconsin plans for the trigger processor card (at least one per string). The string level trigger processors would calculate a string sum for each rollover period ($1\mu sec$), and this would be passed to the next level processor for making the sum of strings. In order to not have the rate dominated by single large hits from one module, (as normally occur in PMTs due to ions) it will probably be necessary to require a minimum number of modules as well.

While some of the specialized triggers discussed below can dip into the noise a bit better, the energy trigger has the nice property of making minimal demands upon our preconception of the event topology. It should collect any deposition of energy by relativistic particles. The energy threshold for a spherically radiating source in the poorest place for light collection, in the center of a six unit cell (wedge 10m high), would correspond to a mere 6 GeV.

We should note that in the conversion from pulse width to PE for each module, we can use a lookup table in the trigger processor. As Ralph Becker-Szendy has suggested, this need not be linear. we may find that the sum of squares, or some other power, is a better trigger. This is because the signals explored so far do tend to produce large pulse heights in a few modules. In any case, by building such conversion into the lookup table generation we can optimize the trigger experimentally.

2.4 Plane Wave, Local Coincidence $T3$

Considering possible triggers for DUMAND II, at the time of writing the proposal, Vic Stenger proposed triggers which involved coincidences on (any of) three strings. The sub-triggers within the strings would require coincidences of adjacent modules. Vic's combinations were 4-3-2, 3-3-2, 5-2-0, and 4-3-0. The individual modules were only required to have $\geq 1PE$, and no there was PE sum requirement. In June '89 it was realized that since each of the combinations above involves a triple hit on a string, we might be able to make a trigger at the string level without the extra complication of a second layer, if we could keep the random coincidence rate within bounds.

One can easily see that a plane wave will have exactly the same time difference between pairs in the 3 neighbor hits. The light wave from a muon is conical however, so subtracting the differences will not yield zero, but should be small. It is necessary to employ a Monte Carlo simulation in order to explore this question quantitatively (see description in Appendix A)

In Figure 3 we show the distribution of time differences between neighbor modules as a function of zenith angle of the muon track. The time differences are unique for vertical tracks, and most widely distributed for horizontal tracks, the overall shape being banana like. (The scattered points are due to random noise.) The time difference of neighbor coincidence pairs are plotted against each other in Figure 4, where one sees that the hits are confined below the diagonal. The difference of differences is illustrated in Figure 5, again versus

zenith angle of the track, and in the projection of this scatter plot in Figure 6, except that the latter contains only the smallest value per muon. We see that a cut of $15ns$ will keep about 85% of the muons, but there is a systematic tendancy for the large time difference events to be near horizontal. The concern about introducing a systematic bias in the trigger will be addressed below, but this example emphasizes the importance of having multiple triggers in operation simultaneously

Figures 12 through 15 present plots that are the same as Figures 3 through 6, except they are for cascades of particles. The main conclusion to draw from here is that the $T3$ works very well for cascades, as shown in Figure 15, where one sees a totally negligible number of events beyond $15ns$.

As shown in Appendix 3, the expected rate of simple triple coincidences is about a thousand per second, and scaling with the cube of the module noise rate. With the $15ns$ cut the difference of differences trigger, hereafter called $T3$, the predicted rate falls to a comfortable $54/sec$

2.5 Neighbor Pairs with High Pulse Heights $T2$

Another possibility is to use neighbor pair coincidences, but with a requirement of some minimum pulse height, which we shall designate $T2$. We realize that employing amplitude in the trigger is not generally desirable, both because it is more technically difficult, but more importantly because pulse height is poorly resolved. Nevertheless, we know that some of the muon trajectories will produce a large pulse height in some modules, for example, those trajectories that are near horizontal and pass near one module, possibly not producing a $T3$.

Some other physics phenomena may do the same, producing a large pulse height in only 2 PMTs on one string, and perhaps enough hits on other strings to make the event distinguishable from background. In a way, $T2$ can be thought of as just a local verion of the array wide TE.

We can test the efficiency of such a trigger with the Monte Carlo, but the problem is that we do not yet know the relationship between noise rate and amplitude for the real optical modules, in the real ocean, and cannot thus make reliable calculations of trigger random rates. If we assume that the Philips PMT is capable of a reduction of the noise by a factor of 50, and the Hamamatsu PMT capable of a factor of 10 similarly, then the $67k/sec$ reduces to a more tolerable $134/sec$, if we require $\geq 2PE$ in each module [7].

As we shall see below, this trigger is rather effective at extracting some events not otherwise caught, though it mostly does overlap with $T3$.

2.6 Nuclearites

Now let us discuss the case of the fast lantern, which I shall call a nuclearite generically. Such particles will raise the noise rate in the PMTs within several tens of m along the track. If the rate goes high enough it will trip the PMT into shutdown mode, as for bioluminescence. The latter takes some time to happen, being governed by the rate limiting circuit at the PMT, but probably only $10\mu sec$[8]. The fast data transmission circuit has a factor of ten or so dynamic range in terms of rate for all the PMTs, but the local digitizer will soon become saturated if we do not limit the PMT rates.

Hence we just cannot track the progress of a nuclearite through the array by watching the singles rates in the PMTs, as they arrive via the fast digitizer circuit. Since the larger end of traversal times we expect is only $10ms$, the particle will have gone by the time the fast rate limiters recover. Nor can we track the progress by the rate measurements sent ashore on the Command and Control link, because we cannot sample the rate and send it to shore many times in a second. Thus it seems that we can only observe the progress of the nuclearite by the noting times of the last pulses before the rate limiter kicks in. Moreover the relatively slow integral rate information recovered from the C^2 link can be used for total amplitude as well as fitting the trajectory spatially.

The trigger processor circuit on shore will know when a module has gone into hibernation mode by finding a *start* without a *stop* pulse. We do not know how often bioluminescence will produce such pulses, but the TTR4 data suggested

[7] Note that herein I have treated PEs as integers, while in fact the PMT output pulse charge is smeared (by fluctuations in the electron multiplier circuit) to an effectively continuous function. For this reason we need both the noise rate versus output charge, under appropriate random light illumination, and we need the distribution of in time pulses for various light levels (producing $1, 2, 3, \ldots PE$). The former is needed for the random noise calculations, the latter for efficiency assessment.

[8] The PMT rate limiter parameters have not been settled yet, but I am proposing that the optical modules will shut-off their fast data output if they receive more than 100 hits in $10\mu sec$, and that they will stay in that state until their microprocessor measures a rate back to normal, but no less than 1 second. These times may not turn out to be the final values, but they cannot be far wrong.

that it might be 1% of the time. If this were so, and the duration of the pulse were typically one second, then we could expect that there would be about 2 PMTs shutdown in the array at any given moment. It would not be too much data then to record the exact starting time of the pulses with no *stop*. Thus, one of the tasks of the data harvesting computer can be to keep a watch for unusual fluctuations in such a rate and if found to look for a progression of the hits through the array. The rate information measured at the PMT and sent to shore on the Command and Control link can then be employed to refine the fit, and measure the total light output of the track.

No one has yet studied the question of threshold and efficiency for such a particle (any volunteers?). My guess would be that we are about as efficient as for muons in catching tracks that have enough light to drive the nearest PMTs into rate limitation. The energy threshold must be something like an equivalent of 100 times the light from a muon, eg. with an equivalent dE/dx of $> 200MeV/cm$, for a $\beta \approx 10^{-2}$ particle, and scaling upwards with decreasing β. It seems that such particles might be missed by this trigger near the upper velocity range, but they probably would be caught by the TE trigger if they cause enough light to exceed the array threshold in one rollover cycle. I do not see this as a big problem if there is some loss in this region, since the expectation for heavies would be to be at least gravitationally bound to the galaxy, and travelling in the range of $\beta \approx 10^{-2 - 4}$.

Until we have more in ocean experience I have no way to calculate the rate of random triggers for such a beast. While this is a matter we will have to explore after the array is in operation, it is hard to imagine that the false event rate would be large (if it is we are in trouble for other physics with dead time from bioluminecsence!)

2.7 Monopoles

The other possibility is that the slow massive particle is visible only in bursts of light along it's path, as in the case of a monopole promoting nucleon decays along its trajectory via the Rubakov process. These remnants of the nucleon may give coincidences between at least neighbor modules. An unknown quantity in this affair is the crossection for such interactions, and the limits are usually presented with this as a parameter. The interaction length may be equivalent to a strong interaction, and thus we might see one or more decays per meter. For this case then we would have many trials for local coincidences. The light is surely more feeble than the nuclearite case, so we cannot count on picking them up by overall rate. Hence it seems that we must pick them out by $T2$ coincidence trails. Since we will trigger on $T2$ the data will go to the next level processor, which can then keep a watch for a series of $T2$'s marching through the array.

Again, we will have to employ experience to assess the false event rate. Someone who wants a nice project could, however, begin the process of calculating the range of sensitivity of DUMAND II in velocity and cross section space.

2.8 Supernova Trigger

As stated in section 1, I really am not at all convinced we have a chance for detecting supernovae, but feel it well worth some fair amount of effort. At this time the only chance I see is for us to record a statistically significant increase in the rate of individual module signals with a threshold $\sim 1PE$. As for assessing either of the important questions of efficiency and threshold of catching the signal (which would translate into how much of the galaxy we can "see"), or false trigger rate, the former needs some Monte Carlo simulations, and the latter needs the PMT noise rate versus amplitude distributions.

However, for planning the trigger we can say that it is desireable to implement the capability for each trigger card to keep a second by second tally of the number of hits per module above some individually programmable threshold (this is in addition to the total rate of triggers).

3 Rate and Efficiency of Triggers

The triggering efficiency for through going muons has been studied with the simple Monte Carlo. It would be nice to have a universal definition of fittability, but (particularly in the presence of noise) this does not seem possible. I could fit the events, but that would only test the trigger relative to my fitter, and everyone's fitter is slightly different. The following table presents some of the results for an isotropic flux of single, minimum ionizing, through going muons. Figure 7 shows the effective area for each trigger versus zenith angle, for what I define as fittable muons, $\geq 6hits$. Figure 8 shows the same plot for a definition of $\geq 12hits$ on $\geq 3strings$, demonstrating that the efficiency of the triggers does not relatively change much, though the effective area goes down a fair amount.

	$T2$	$T3$	TE	SUM
Efficiency	60.3%	63.9%	79.5%	90.8%
Excl Eff	3.4%	4.3%	14.6%	
Eff Area	$22,414m^2$	$23,750m^2$	$29,567m^2$	$33,738m^2$
Noise Rate	$134/sec$	$54/sec$	$452/sec$	$640/sec$

One might be tempted to conclude that the $T3$ trigger is fairly useless. That is not the case, though in the table it has been subsumed by the TE triggers, because it, alone among the triggers depends only upon the geometry of the coincidence and not upon pulse heights. Of the triggers, I only feel comfortable about the rate predictions for $T3$, having had to make shaky assumptions about pulse heights for calculating the others. Also, note that though the array total area, averaged over all arrival angles now totals nearly $34,000m^2$, and nearly $50,000m^2$ over the lower hemisphere, this does not mean that we can fit all of those events in practice. This gives the people writing fitting routines something for which to aim!

The effective volume of the array is illustrated in Figure 16. By effective volume I mean the fraction of generated events which produced more than 5 hits and which generated one or more triggers, times the test volume in which they were generated ($2 \times 10^7 m^3$). Note that the effective volume reached the test volume at $10TeV$, so that further calculation may reveal the effective volume continuing to rise. This is a bit amusing since the contained volume is only about $2 \times 10^6 m^3$, so almost all the events are being seen from afar! Are they fittable?

4 Conclusion

We have discussed various possible triggers for the DUMAND II array, and find that 3 easily implemented triggers ($T3$, $T2$, and TE) will work for the main goal of recording muons. The $T3$ trigger requires small second time differences between trios of modules ($15 ns$). The $T2$ trigger tags neighbor near coincidences with unusual pulse heights ($90 ns$, $\geq 4PE$). The TE trigger seeks large numbers of PEs from anywhere in the array ($1 \mu sec$, $> (13 + 14)PE$).

More simulation work is needed on cascades and low energy muons, but the proposed triggers probably work about as efficiently as we shall achieve at plucking such events out of the data stream. More work is needed in studying fitting of such low energy events. And more effort is needed in characterizing the signals and noise from the PMTs.

These triggers also seem to have a good chance to extract non-standard types of events involving relativistic particles of all types. Triggers for slow, bright particles, and for supernova detection, require special consideration, and a different type of hardware implementation. Our best approach seems to be through the implementation of means to record the times of PMT saturation, PMT noise rates each second (both single and multi-PE), and by searching for patterns in $T2$ coincidences.

While it is not clear if we have a chance to detect supernova neutrinos, the best opportunity appears to be via the monitoring the rates of pulse heights more than one PE and seeking few second increases distributed throughout the array.

Summarizing what is needed of the trigger card:

- search for single string triggers $T2$ and $T3$
- tally total PEs for each rollover and forward to next level
- tally totals of all pulses, and pulses exceeding a predetermined threshold, by module for each second
- note PMT start pulses without stops, and pass along start time, updating PMT status word to be read with scaler rates each second

Summarizing what is needed at the next level for data filtering and preliminary fitting

- collect and filter single string triggers, seek 1000× reduction
- make PE totals for TE, and test for trigger every rollover cycle
- watch for array wide increases in single PE and multi-Pe rates

- watch for ripple of PMT shutoff times

This document should be looked upon as a starting point for more detailed considerations. Of more immediate concern, for making progress on the system design, it seems that we will be able to create the triggers we need with the system as presently conceived, and that we are at least not far from optimum triggering.

Acknowledgements

The work reported herein draws upon the efforts of many people. I want particularly to acknowledge Ralph Becker-Szendy, Art Roberts and Vic Stenger for their many conversations on this matter, as well as many of the other collaborators, and particularly those at the Seattle Workshop

Appendix A Monte Carlo assumptions

The Monte Carlo program used in the studies reported herein was written not to be all encompassing, but to be reasonably fast and simple so that I could try out different trigger configurations without huge computer time. The basic muon generating routine generates tracks perpendicular to a disk centered on the array. The orientation is made uniformly random, and the tracks are taken as infinitely long. The angles and distances relative to each module are calculated, and the mean number of PE expected in each module is generated. Then the actual number of PEs are Poisson fluctuated about that number. No effects of amplitude smearing are used, and no conversion to digitized units is introduced (not needed for present purposes).

The function employed for calculating the number of PEs is

$$n_{PE}(d) = n_0(d + d_0)^{-\alpha} e^{-\beta d} F(\psi)/d,$$

where $n_0 = 277.6$, $d_o = 17.86m$, $\alpha = 0.437$, $\beta = 0.02085/m$, and the slant distance d is in m. This function is illustrated in Figure 9, where one sees that the minimum ionizing muon represented by this function will produce a mean of $1PE$ at a slant distance of $28m$. The function is for head on illumination, and the angular distribution must be inserted. I use

$$F(\psi) = A_0(0.525cos(\psi) + 0.475)$$

for the Hamamatsu PMT, and

$$F(\psi) = A_0(0.546cos(\psi) + 0.606),$$

for $cos(\psi)$ between -0.9 and 0.75, 0.0 for $Cos(\psi) < -0.9$, and A_0 otherwise, for the Philips PMT. The types of PMTs are taken to be alternating on the strings, and all PMTs taken to be face downwards. I believe that these are exactly the same functions as used by Vic Stenger, and documented in his notes about his Monte Carlo program. (The attenuation function is a result of a calculation I did years ago, updated in May 1988 for new input data. See HDC-81-10 for details.)

The cascade simulation is simplified, employing the point source approximation presented by Art Roberts in the 1978 DUMAND Workshop, 1, 103 (1978). Events are generated throughout a cylindrical volume, extending either side of the reference disk, as decribed above. For the preliminary calculations presented herein I used a disk of radius $150m$, and the cylinder was $305m$ long. It is evident (see Figure 16) that above 10 GeV we have significant contribution from events outside the array.

Noise is inserted at a rate of 60,000 hits per second in each module, in a time window of $1\mu sec$ centered on the event mean time.

Appendix B Calculation of Minimum Number of Hits

An interesting question is: what is the actual limitation on the smallest size of coincidences caused by muon-like events in the array which can be extracted from noise? Forget for the moment whether such events can be fitted. The idea is that there is some space in which these events are as tight as possible, and in which the random coincidences will be minimal. This is a sort of phase space approach.

Imagine that we continuously plot all events on a direction plot, with a number of cells equal to the best resolution we can achieve. This might be $N_D = 10^4$ for one degree resolution. Imagine also that we divide each direction into as many impact points as we have temporal resolution to support, again say $N_P = 10^4$ (about $2m^2$ grid). A through going muon would illuminate typically $l = 20$

modules, though hitting a mean of about $k = 12$ (though this is what we shall solve for soon).

For this trial track (with chosen direction and impact point) we can thus predict the relative time of arrival of photons at all modules to within the variation of geometry and the time resolution of the system. Suppose then that we ask the question every $2ns$ as to whether or not there is a coincidence of module signals amongst the 20 candidate modules along the trial track. The rate of random coincidences will then be given by the number of trials per second times the probability of a coincidence of that level.

$$R_k = (N_D N_P/\tau_m)m^k/k!,$$

where the latter term is approximately the Poisson probability of getting k coincidences when m are expected, when $m << 1$. The expected number $m = lR_m\tau_m = 0.0024$ for the stated conditions. So while the number of trials is vast, the probability drops fast with increasing k, such that for a value of $k = 8$ the rate is a negligible once in 23 years! The value of $2ns$ may be overly optimistic, but even so, if it were as bad as $8ns$ we find a rate of only once per 1.4 years for $k = 9$.

I thus conclude that we could identify physics events solely on the basis of phase space, with a negligible background from random coincidences if the coincidence threshold is as low as 8 or 9 modules. Moreover, this is conservative since the question we have asked has to do with the topology expected of muons, and we can apply pulse height criteria on top of the space time fit. Without further consideration we can also state that we should be able to employ muon fits down to the level of about 7 hits (with expected random rate of $4 \times 10^5/year$), if we accept an increase in the background due to cosmic ray muons.

It seems to me that this is telling us that we could in principle accept triggers down to the level of about 12 noise hits plus 6 real muon hits, and that an ideal fitter could extract the signals from noise. Based upon the previous discussion of the gross coincidence rate in a coarse time window of $1\mu sec$, we see that a simple trigger will not come close to this ideal goal. On the other hand, the T3 does get events down at this level, so this tells us it is being fairly efficient at getting out the maximum amount of physics.

Appendix C Random Rate of 2 and 3 Fold Neighbor Triggers

First consider the rate of neighbor coincidences. The coincidence time window, in order to accept all trajectories, must be at least the flight time of photons in water travelling from one module to the next. If the module spacing is $D_m = 10m$, the speed of light c, and the index of refraction in seawater $n_w = 1.35 \approx 4/3$, then:

$$\tau_2 = n_w \times D_m / c \approx 45ns$$

If the number of modules per string is $N_m = 24$, the number of strings $N_s = 9$, the individual module noise rate $R_m = 60{,}000/sec$ [9], then the total array rate of random pair coincidences will be

$$R_2 = (N_m - 1) \times (N_s) \times 2 \times R_m^2 \times \tau_2 \approx 67{,}000/sec$$

The random rate for three neighbor concidences is just the triple coincidence rate times the number of combinations:

$$R_3 = (N_m - 2) \times (N_s) \times 3 \times R_m^3 \times \tau_3^2,$$

where τ_3 = the coincidence window for triple neighbor coincidences. This is equal to the flight time for a photons in water to traverse the distance between the farthest modules, 20m, or about 90ns. The random rate for triples anywhere in the array is then about,

$$R_3 \approx 1040/sec,$$

which is still not tolerable, particularly since this could become $5{,}000/sec$, if the module noise rate turns out to be $100{,}000/sec$.

It was for this reason that the difference of differences trigger was developed. The question to be addressed here is what is the expected noise rate with this cut? We approach this in two steps. First consider the requirement of as close a neighbor pair coincidence as possible.

[9] We take all modules to be equal in noise rate. With some trouble one may show that for a large number of modules it works well enough to take the mean noise rate for such coincidence calculations. This has been experimentally verified in the IMB detector in 1983.

We get a factor of six right away by requiring the window to be half and the reduced combinatorics: every time there is a hit in a potential center member of a triple we ask the question of whether there was a neighbor coincidence in a time $+$ or $-$ τ_2 about that event. We can do better, however, since if the signal is early on one side, it must be late on the other, so we can gain another factor of two.

You can picture the random noise hits about the center module time as uniformly populating a square region in a scatter plot of one neighbor's time difference versus the other. This plot is shown in Figure 4 for muons, and Figure 13 for cascades. The requirement of one being early while the other is late cuts the square on the diagonal. The further requirement of the time difference being less than some value (which we determined from the Monte Carlo to be about 15 ns without losing many events), then we are further restricting the area of the plane to a parallelogram of base $\Delta = 15ns$ and height $2\tau_2$, but less a triangular region of area $\Delta^2/2$. The resulting rate is then given by:

$$R_3' = (N_m - 2) \times N_s \times R_m^3 \times \Delta \times (2\tau_2 - \Delta/2) \approx 51/sec,$$

for a net gain of a factor of 20 over the simple triple coincidence R_3.

Appendix C Random Rate of 2 and 3 Fold Neighbor Triggers

First consider the rate of neighbor coincidences. The coincidence time window, in order to accept all trajectories, must be at least the flight time of photons in water travelling from one module to the next. If the module spacing is $D_m = 10m$, the speed of light c, and the index of refraction in seawater $n_w = 1.35 \approx 4/3$, then:

$$\tau_2 = n_w \times D_m/c \approx [illegible]$$

If the number of modules per string is N_m [illegible] the number of strings $N_s = 9$, the individual module noise rate R_m [illegible] then the total array rate of random pair coincidences will be

$$R_2 \quad (N_m \quad 1) \cdot (N_s) \cdot 2 \cdot R_m^2 \cdot \tau_2 \quad [illegible]$$

The random rate for three neighbor coincidences is just the triple coincidence rate times the number of combinations.

$$R_3 \quad (N_m \quad 2) \times (N_s) \times 3 \cdot R_m^3 \times \tau_3^2,$$

where $\tau_3 =$ the coincidence window for triple neighbor coincidences. This is equal to the flight time for a photons in water to traverse the distance between the farthest modules, $20m$, or about $90ns$. The random rate for triples anywhere in the array is then about,

$$R_3 \quad [illegible]$$

which is still not tolerable particularly since this could become 5,000/sec if the module noise rate turns out to be 100,000/sec.

It was for this reason that the difference [illegible] differences trigger was developed. The question to be addressed here is what is the expected noise rate with this cut? We approach this in two steps. First consider the requirement of as close a neighbor pair coincidence as possible.

*We take all modules to be equal in noise rate. With some trouble one may show that for a large number of modules it works well enough to take the mean noise rate for such coincidence calculations. This has been experimentally verified in the IMB detector in 1983.

We get a factor of six right away by requiring the window to be half and the reduced combinatorics: every time there is a hit in a potential center member of a triple we ask the question of whether there was a neighbor coincidence in a time + or τ_2 about that event. We can do better however, since if the signal is early on one side, it must be late on the other, so we can gain another factor of two.

You can picture the rand[illegible] hits about the center module time as uniformly populating a square region in a scatter plot of one neighbor's time difference versus the other. This plot is shown in Figure 4 for muons, and Figure 13 for cascades. The requirement of one being early while the other is late cuts the square on the diagonal. The further requirement of the time difference being less than some value (which we determined from the Monte Carlo to be about 15 ns without losing many events), then we are further restricting the area of the plane to a parallelogram of base Δ [illegible] $15ns$ and height $2\tau_2$, but less a triangular region of area $\Delta^2/2$. The resulting rate is then given by

$$R_3' \quad (N_m \quad 2) \cdot N_s \cdot R_m^3 \cdot \Delta \cdot (2\tau_2 \quad \Delta/2) \approx 54/sec,$$

for a net gain of a factor of 20 over the simple triple coincidence R_3.

Figure 1: Number of muon caused hits in fittable events

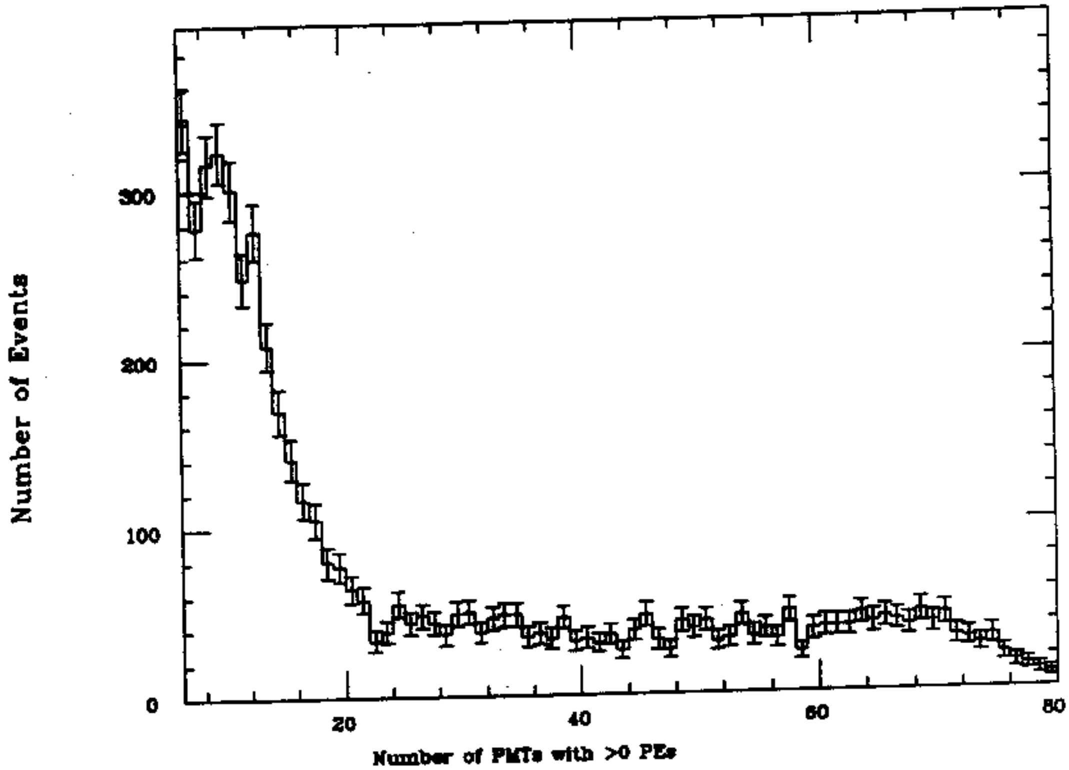

Figure 2: Muon caused PEs in fittable events.

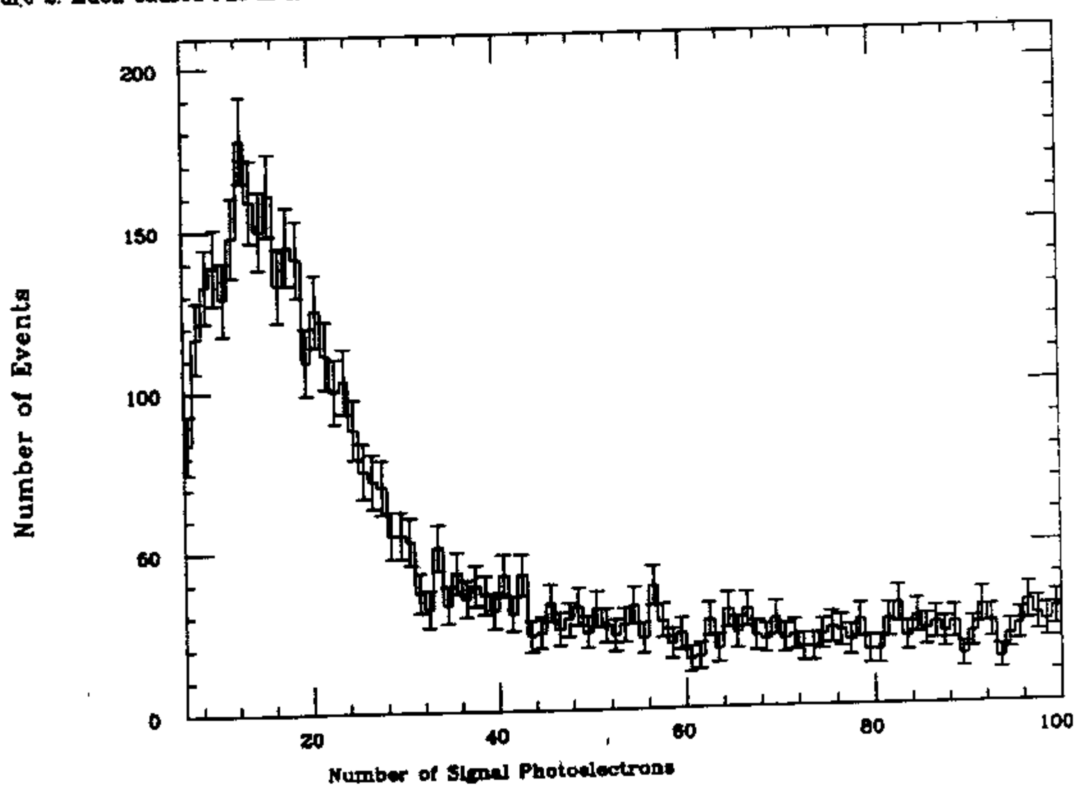

Figure 3: Pair time differences versus zenith angle.

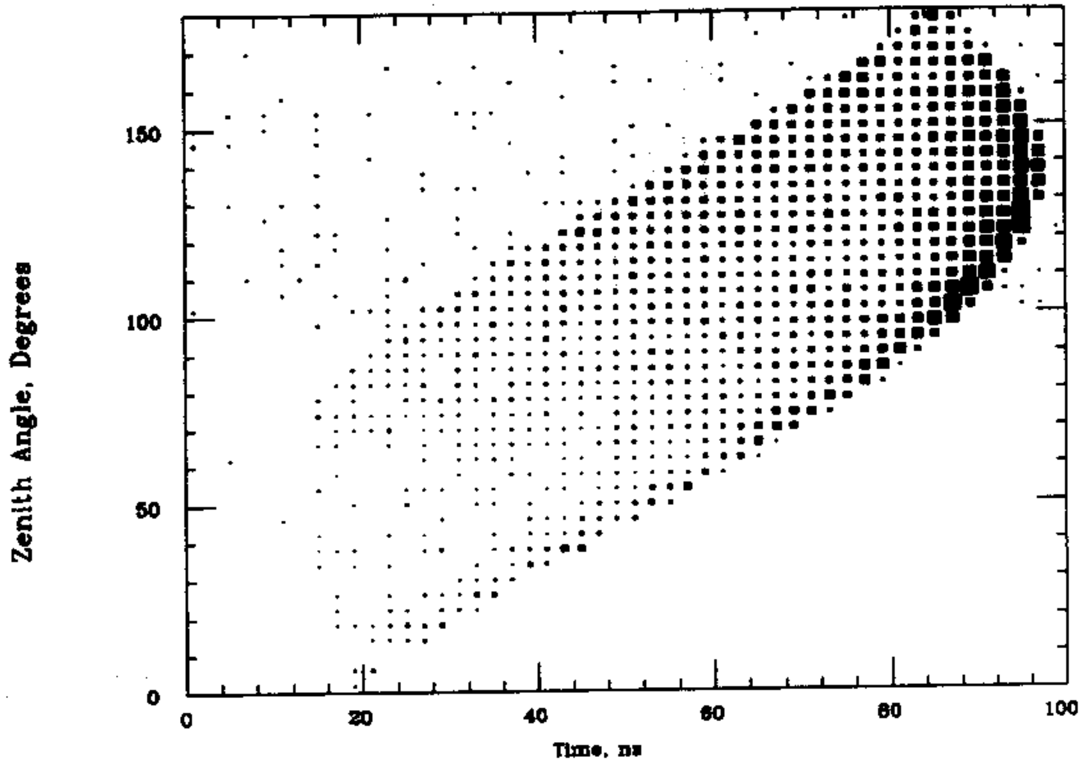

Figure 4: Neighbor pair time differences.

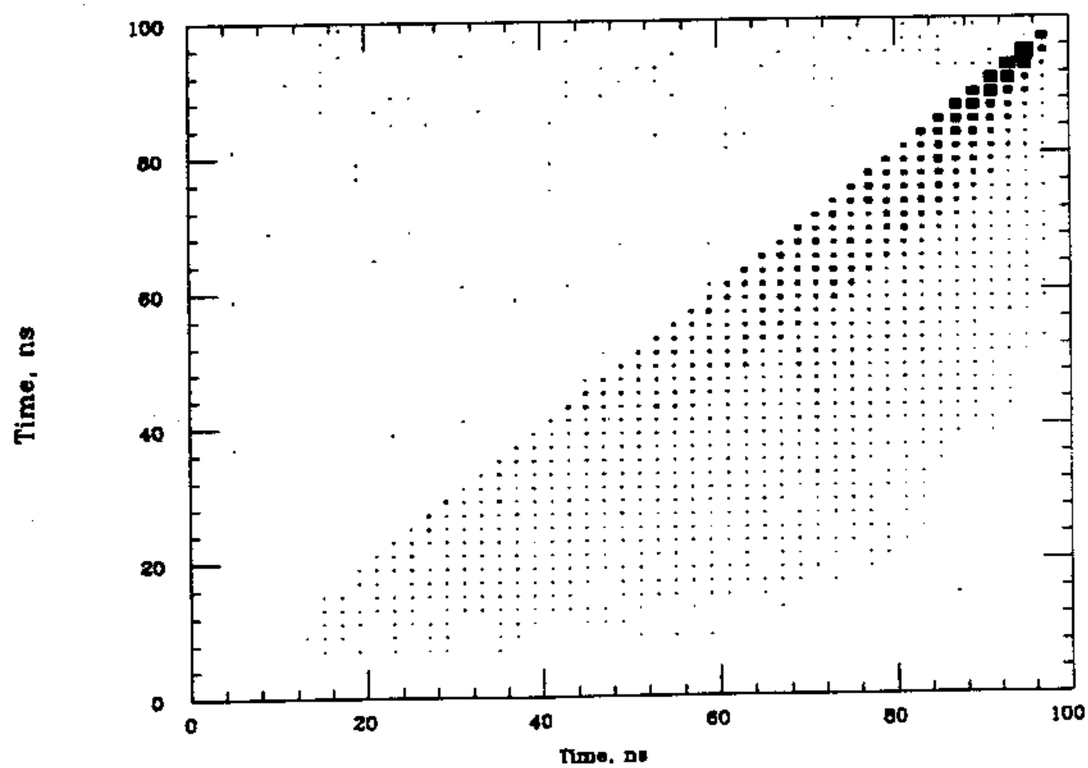

Figure 5: Second time difference versus zenith angle.

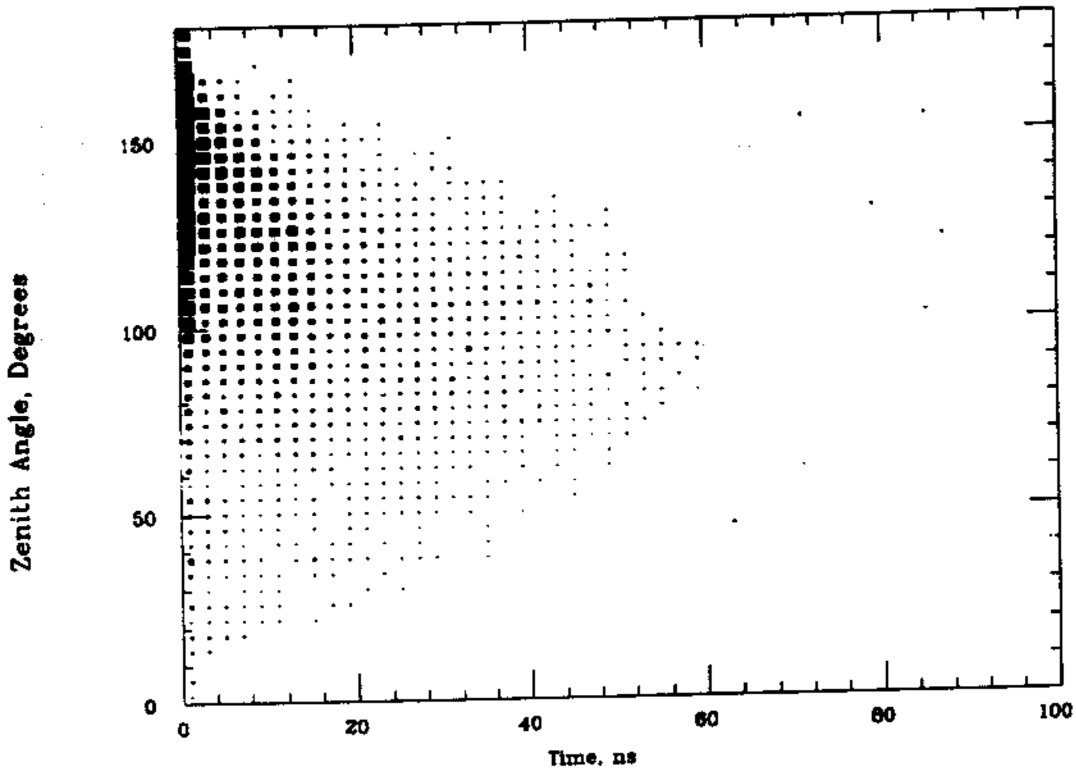

27

11/06/90 muons jgl

Figure 6: Smallest second time difference per event.

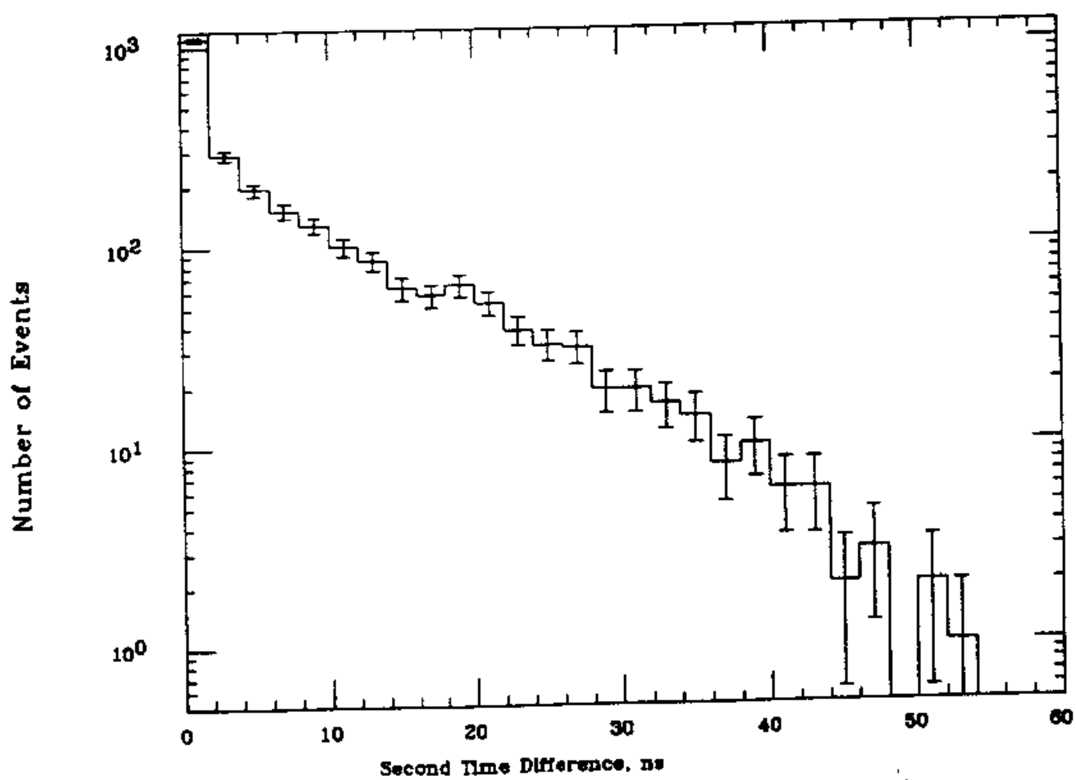

28

09/08/90 muons jgl

Figure 7: Effective area vs zenith angle, >5H

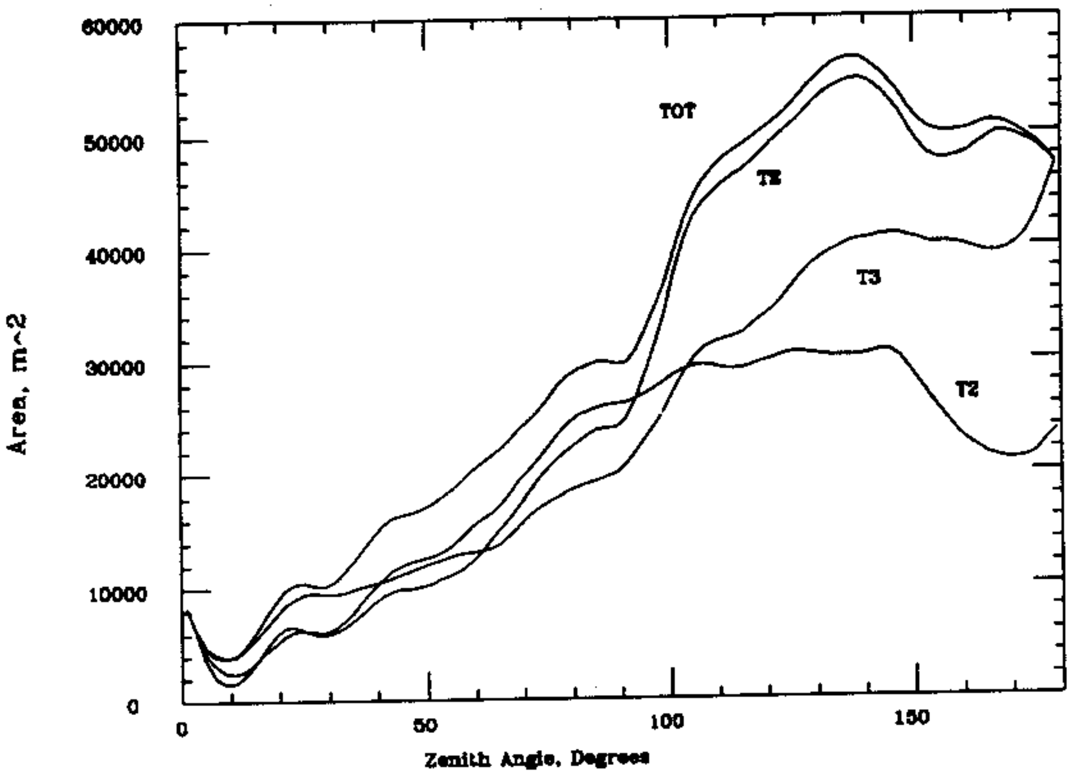

29

Figure 8: Effective area vs zenith angle, >11H, >2S.

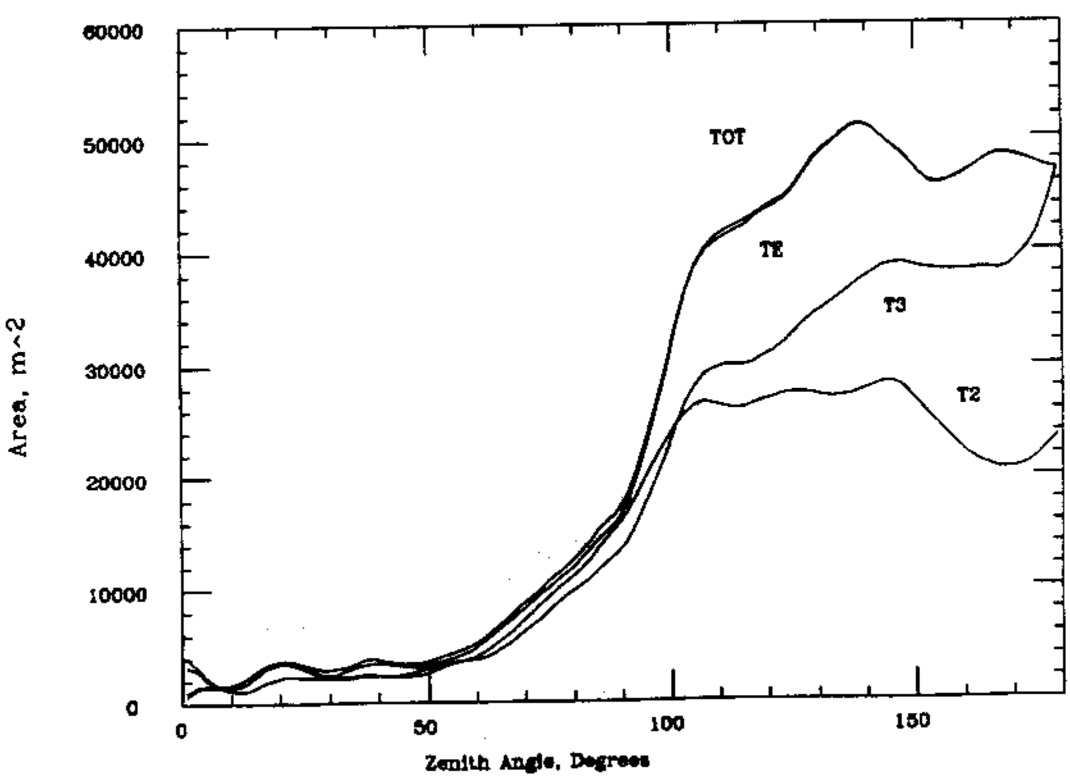

30

Figure 9: Photoelectrons versus slant distance.

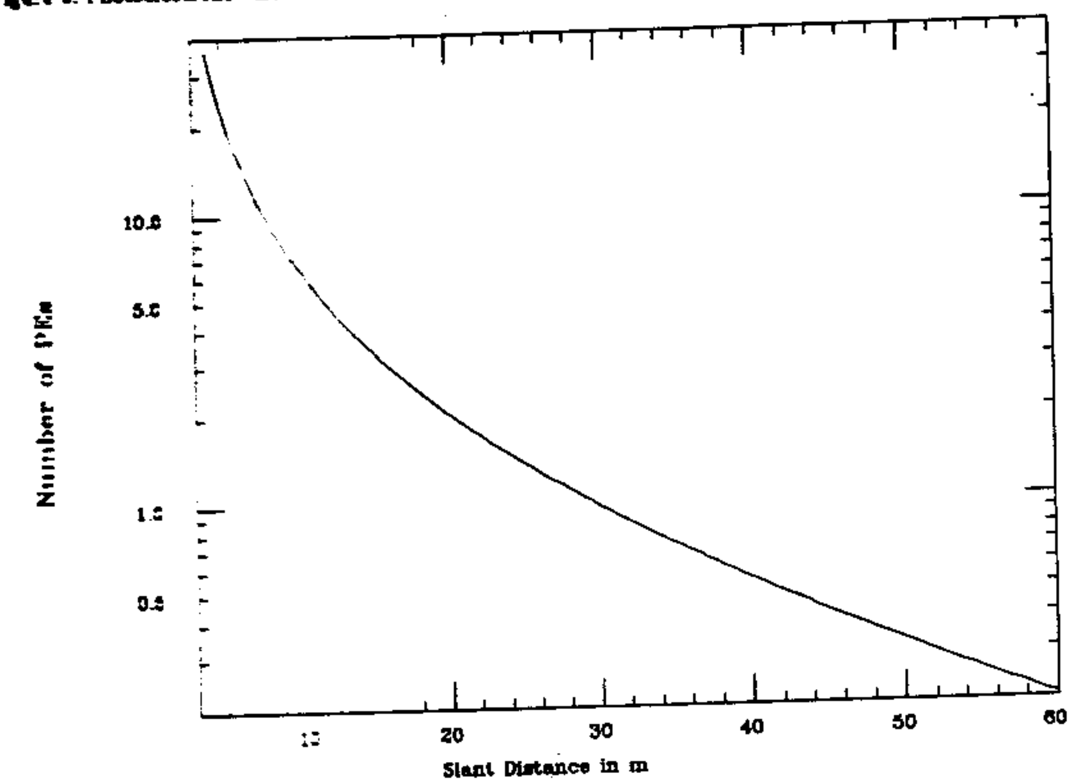

06/08/90 muons jgl

Figure 10: Number of cascade caused hits in fittable eve

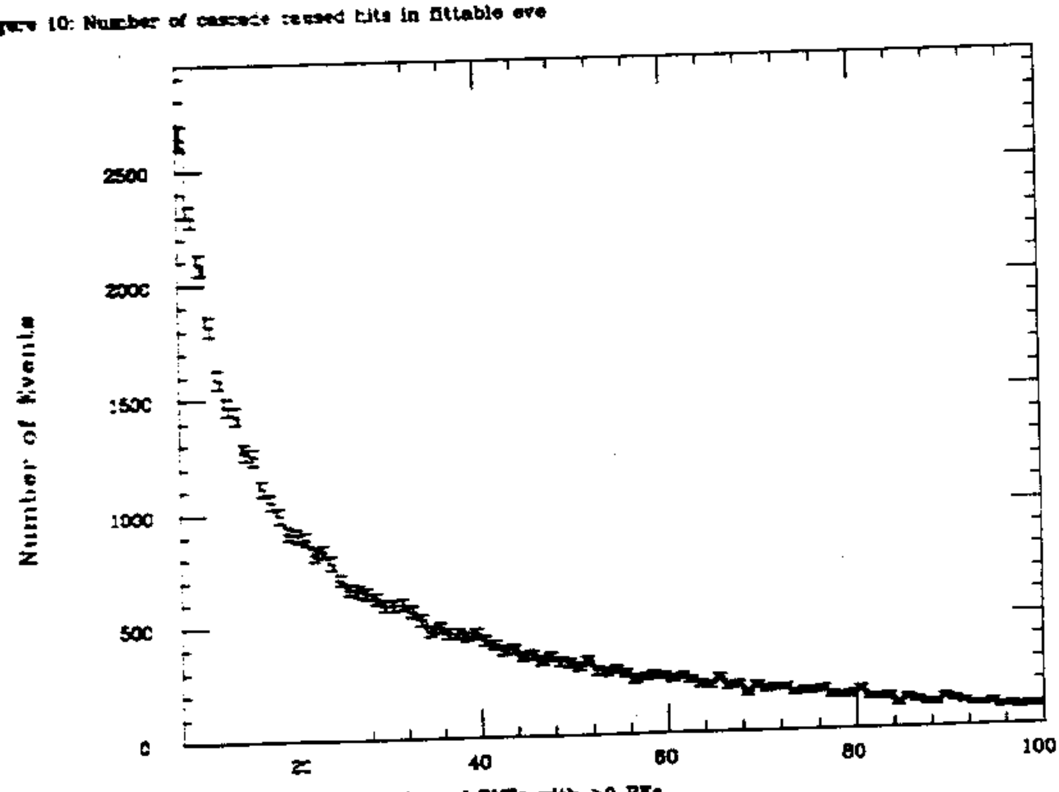

cascades jgl

Figure 11: Cascade caused PEs in fittable events.

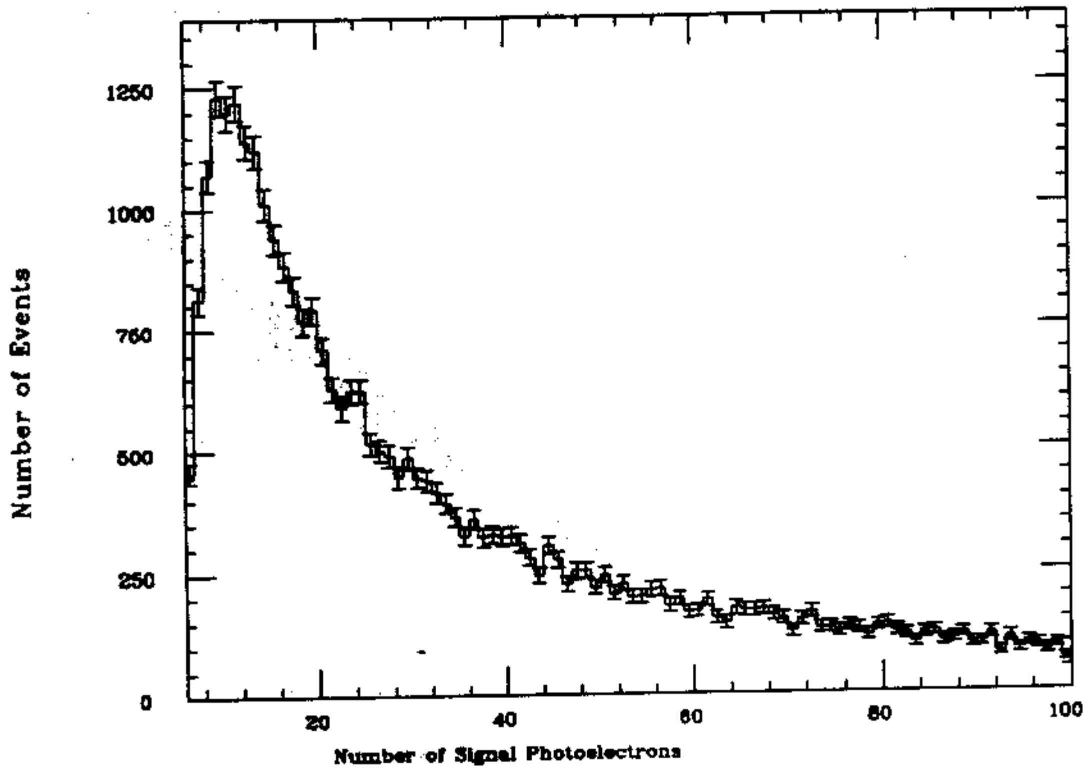

Figure 12: Pair time diffs versus zenith angle.

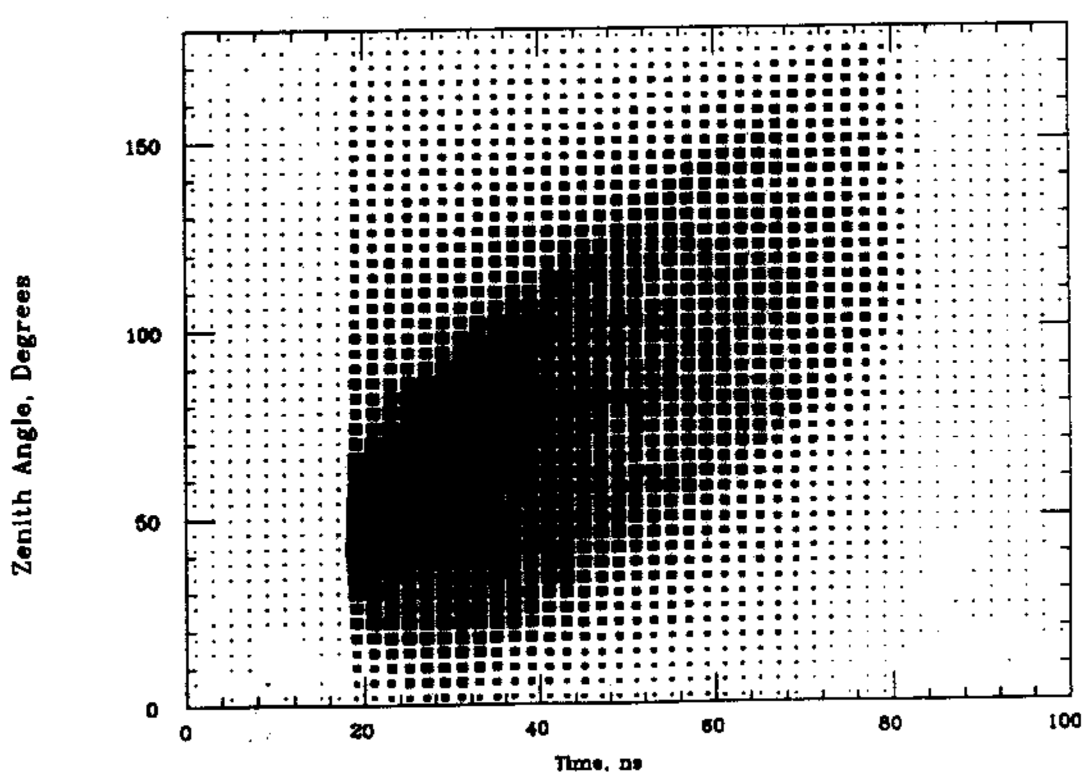

Figure 13: Neighbor pair time differences.

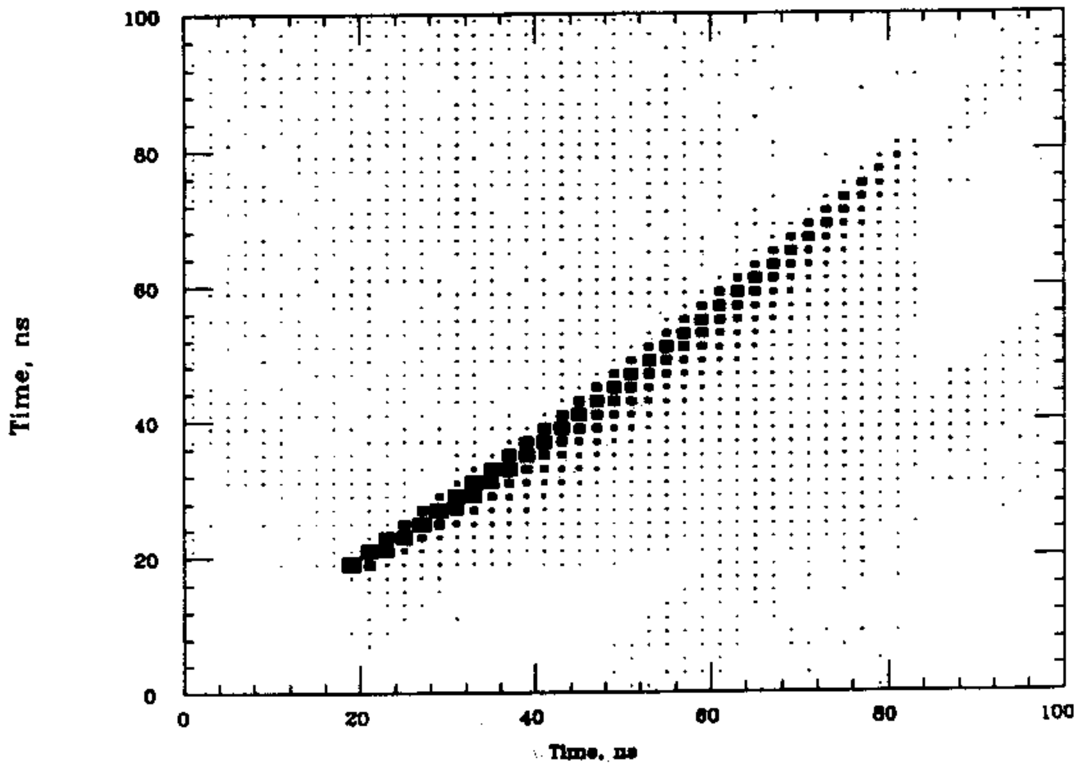

Figure 14: Second time difference versus zentih angle

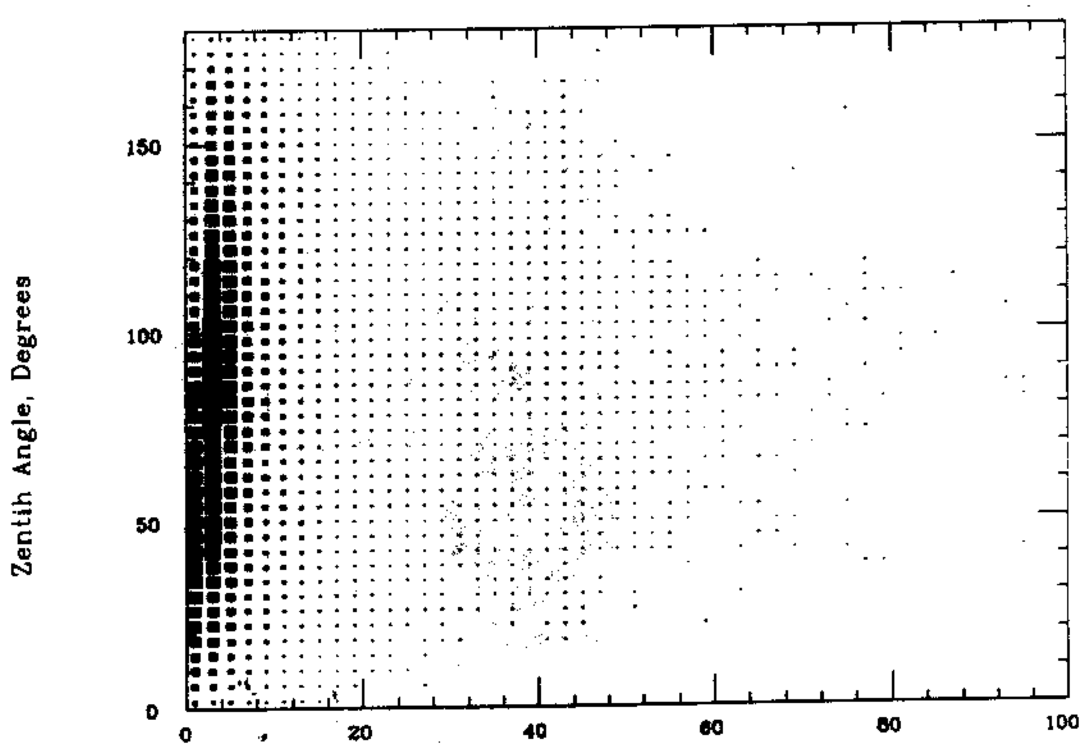

Figure 15: Smallest second time difference per event.

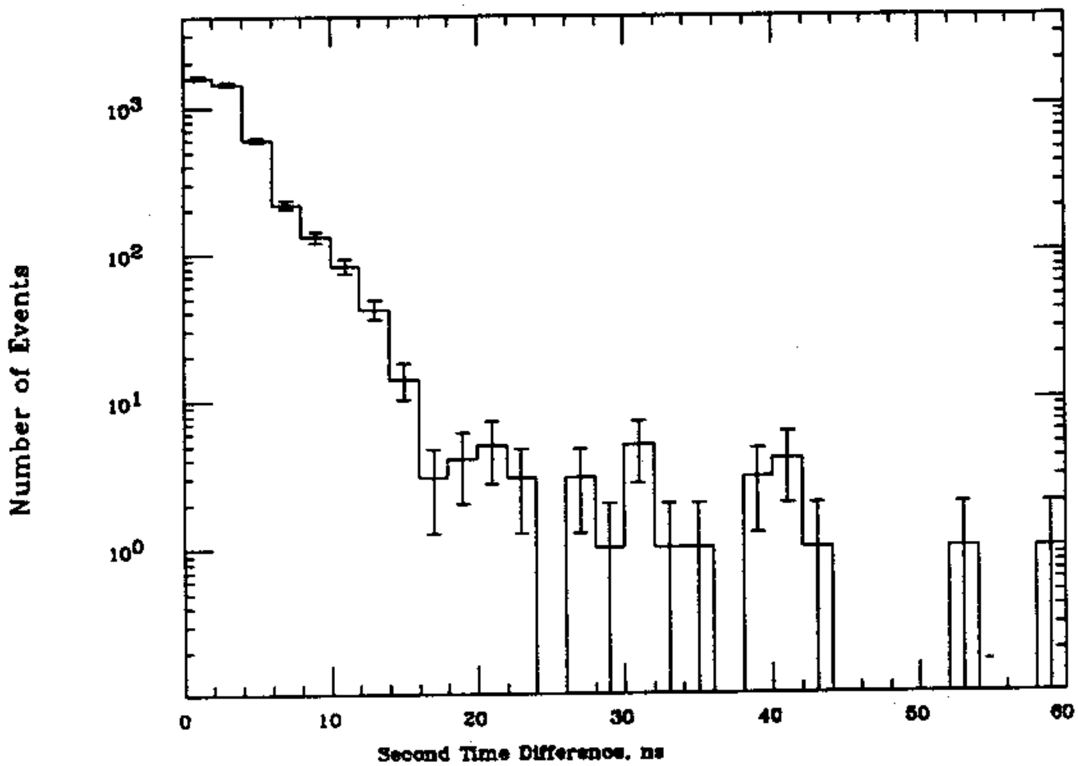

14/06/90 cascades lgl

Figure 16: Effective volume versus cascade energy.

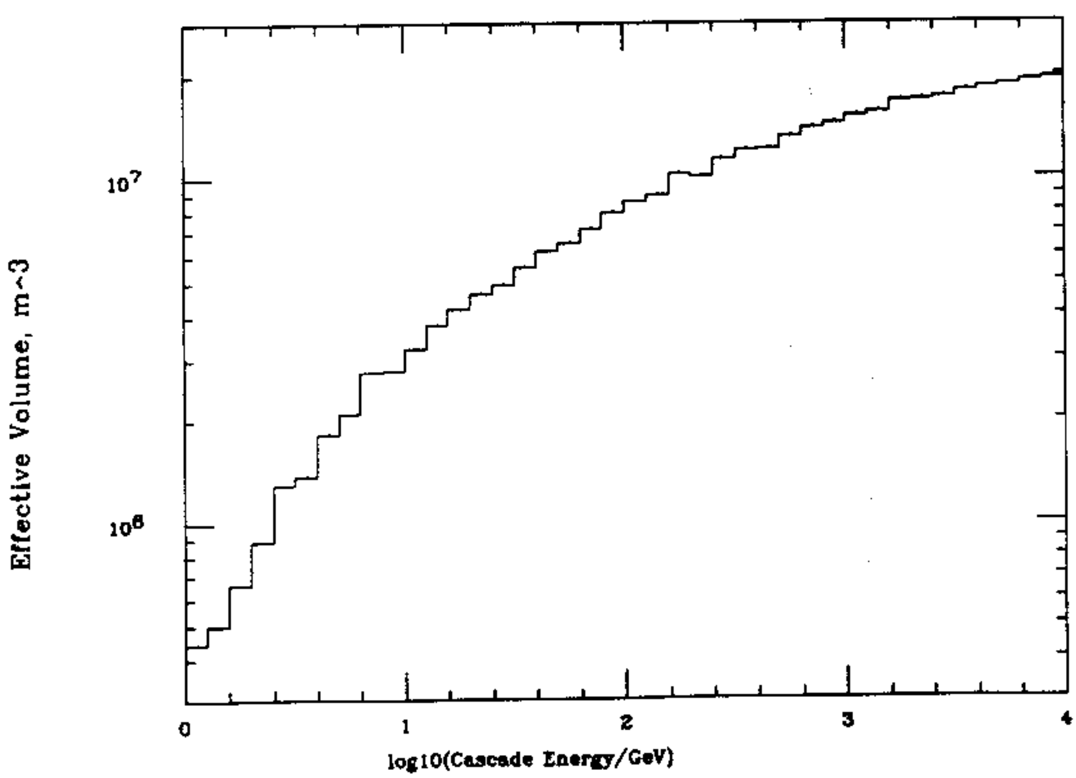

14/08/90 cascades lgl

Figure 17: PEs vs angle, 10 Gev cascade at 20m face on.

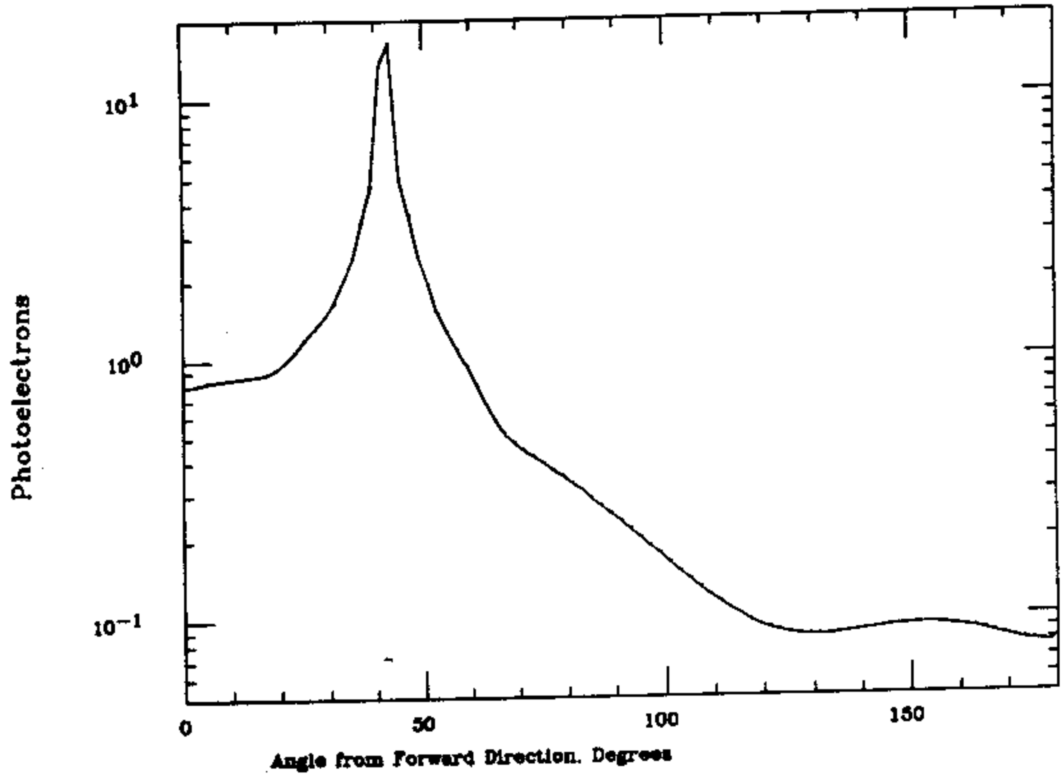

39

14/06/90 cascades jgl

Figure 18: Trigger codes: 1=fittable, 2=T2, 4=T3, 8=TE

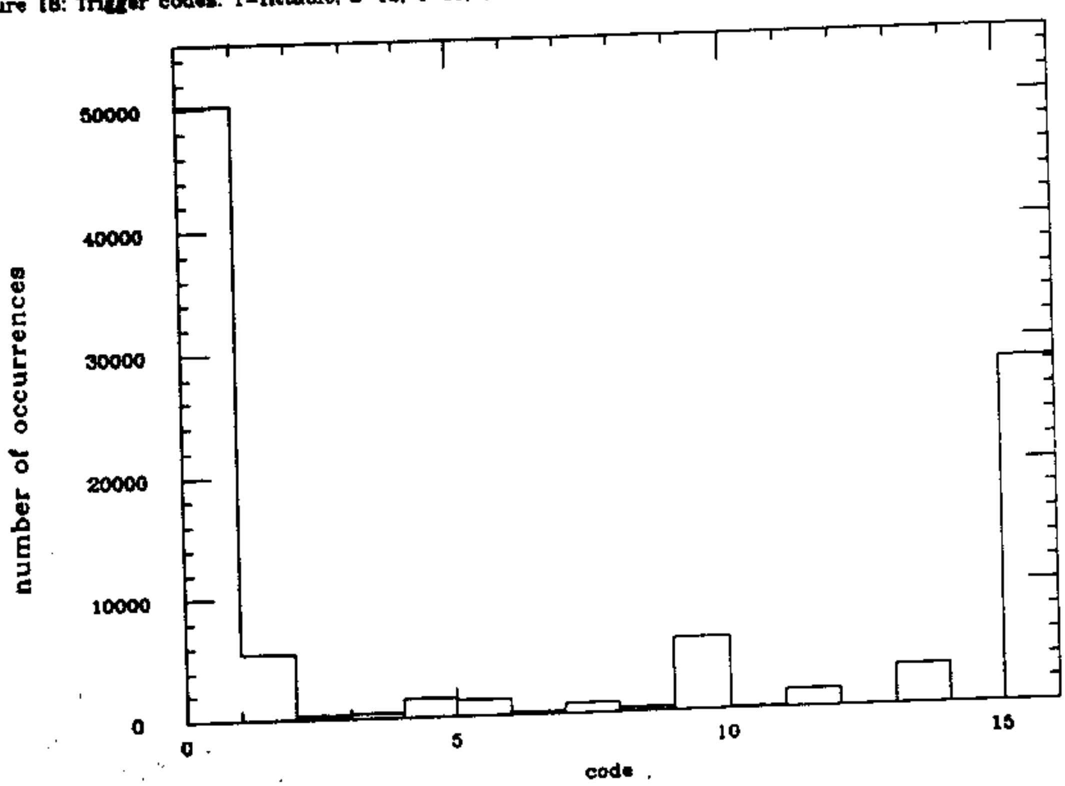

40

cascades jgl

NOTES ON THE SBC FAST DATA PATH

by Phil Ekstrom,
University of Washington

This is my understanding of what got adopted by the Seattle trigger meeting and what that implies.

1) Primary data from each string will be sent down one color of the string's fast data fiber. The second color will be used for redundant data transmission.

The high cost (in various coin) of undersea fiber optic connectors requires that the two colors be carried on a single fiber from SBC to JB. However it is conceivable to separate and recombine the colors in the JB so as to send them to shore via adjacent fibers. Dick Davisson asserts that it could be done with very low loss using passive dichroic couplers. This would allow continued unimpaired operation in the face of one or more failed fibers to shore. The possibility was discussed, but not decided so far as I have heard.

2) SBC messages are of two kinds: optical event messages and rollover messages. Optical events are stacked in the digitizer along with the counter contents (time in units of 2 nanoseconds) as they occur. Rollover events are stacked every microsecond at counter rollover along with the current values of the data input bits. Messages are read out into the link as time permits, first in, first out.

Simultaneously occurring events are transmitted in ascending order of address. This implies that rollover, which is address zero, is sent first before any other events at time zero.

SBC messages on the fast data link share a common overall format; each has three fields, but the interpretation of the fields differs.

	Optical Event	Rollover event
Address field (5 bits)	nonzero number of event source.	zero, marks rollover type.
Time/Data (9 bits)	Time of event in units of 2 NSec	Data field
Up/Down (1 Bit)	State of line after transition	10th counter bit after rollover

Arrangement of fields within the transmitted word is arbitrary and a matter for agreement between the teams working on digitizer and trigger hardware. I have therefore avoided drawing any particular one. They may also agree to some way of communicating FIFO overflow, perhaps by dedicating address 31 to it.

Two nanosecond resolution and 1 microsecond rollover imply the 9 bit time/data field width. Putting all the data from a string down one color channel implies the 5 bit address field width to identify 24 OMs plus 3 calibrators. Total width is 15 bits.

3) A 5B/6B expansion for fiber transmission makes the transmitted message 18 bits long. Sent at 256 MBits/sec the message rate is 14.22 / microsecond and the message period is 70.3 nanoseconds. Subtracting one rollover message per microsecond yields an event throughput of just over 13 events per microsecond.

There are various advantages to making the transmitted bit rate simply related to the digitizer clock rate, for instance an exact binary submultiple. One such advantage is discussed in item 9.

4) Design event burst rate is 12 hits/2 NSec, 48 hits/200 Nsec. The digitizer must provide a fast FIFO buffer to stack the event bursts without loss until they can be sent to shore. This accounts for much of the difficulty of the digitizer, and current designs have gone to a two-stage FIFO with one stage on a separate memory chip.

Even so, the limited density available in GaAs would require segmenting the digitizer into two pieces, each handling half of the event generators and producing half of the messages. The bit specification above appears to assume this division. Each half would need to contribute messages to the single data stream and must either have synchronized counters or generate separate and distinguishable rollover messages.

The greater density available with Si CMOS may offer the option of digitizing events from a whole string on one chip, thereby avoiding the coordination issue. However other considerations may outweigh this advantage. Eric Hazen will decide.

Some aspects of the trigger design will depend on this decision.

5) Since each of the 9 data input bits is sampled and sent every microsecond, each may be treated as a 1 Megabaud serial link. Three of them carry hydrophone signals digitized to 12 bit precision at nominally 80 KHz sampling rate and transmitted as twelve or thirteen-bit synchronous serial bit streams. The 13 bit length will be needed to get sample interval above 12.5 microsec if the LTC1292 serial out digitizer chip is to be used. A fourth stream carries a framing (alias sync) bit to mark the beginning or end of a transmitted word. Exact format (which end first, 12 vs 13, location of framing bit, etc.) is up to UWA who will be dealing with both ends of the link.

Slow Asynchronous streams may be sent by just applying them to one of the data bits. The Neil Brown Unit output and CC data will be sent this way.

This uses six of the nine available bits.

I note that two bits are allocated to "rollover count". I missed the discussion which led to this, so cannot explain it. However note that one bit of rollover count is already being sent via the Up/Down bit, and the discussion below suggests that it will suffice.

This leaves one spare. Again the exact arrangement of bits within the word is arbitrary.

6) Rollover messages serve to increment an extension counter in the trigger. The "tenth" bit of the digitizer counter, sent in the Up/Down bit, should be compared with the lowest bit of the extension counter in the trigger. If a rollover message has been corrupted or otherwise lost, the bits will disagree.

7) Rollover messages divide the event stream into "pages". One page holds the optical event messages received between any two rollover messages. Any pair of events which occurred less than one microsecond apart must result in messages either in the same page or in adjacent pages. Therefore the trigger need only form time differences between events in a two-page block of data.

8) If the "tenth" bit sent with the last rollover message is appended to all times in the first level trigger calculation, then ten-bit twos-complement differences formed between times in the same page or adjacent pages will be correct when interpreted as positive numbers. Overflow cannot occur. This assumes that times belonging to messages received earlier are subtracted from the times of messages received later, not vice versa. Any carries off the end of the register are ignored.

As an example, any two events which yield a difference formed as above with its top four bits zero must have occurred within 128 nanoseconds of each other and in the correct order. Conversely, all events which occurred more than 128 nanoseconds apart will either appear in pages which are not adjacent or else will yield time differences which have at least one bit in their top four nonzero.

Once a pair of differences has passed such a test, we know that the top bit of each is zero. They may then be subtracted from each other with either the same or identical hardware to yield a signed second difference which cannot overflow. That signed number may then be subjected to a two-sided test for magnitude of the second difference.

Example: If such a second difference has its six top bits either all zero or all one, then it lies between -32 and +31 nanoseconds. Less simple-minded tests can be performed to set limits which are not exact powers of two.

In any case, there is no need to carry more than ten bits for any of the first level trigger calculations.

Let me suggest one of two viewpoints to adopt when working all this out yourself to check my conclusions. One way is to regard the initial time values as truly circular -- like numbers on a clock face -- rather than as either positive or signed computer numbers. The other approach is to begin by assuming arbitrarily wide words so that rollover never occurs. Then observe that in the first difference only the lower ten bits are ever nonzero. There is no point in computing high order bits which are always zero, and since you don't need to compute beyond ten bits, there is no need to represent more than ten bits of time.

9) Once interesting regions of data are identified by the first level trigger algorithm, the data can be readied for further processing by appending the local extension counter to the front of the time and subtracting from the now longer word the calibration parameter appropriate to the particular address.

The calibration parameters (I don't call them constants for reasons to be seen below) are determined by observing light pulses from the calibration modules. They depend not only on the various propagation delays for light and other signals, but also on the arbitrary initial contents of the local counter. They are chosen so that when subtracted from the composite word made of transmitted time and local counter contents, they give equal numbers for simultaneous events.

Calibration requires pulses from laser light sources which have a limited life. Therefore it must be a relatively infrequent event, occurring on a scale of some per day rather than per second. If the lock of string clock to the CCC clock reference were to fail, calibration could not be performed often enough to correct for the offset and drift of an ordinary crystal oscillator string clock.

If the transmitted data clock were chosen to be a binary submultiple of the digitizer clock, then a fairly simple circuit could watch both the data clock recovered onshore and a correct onshore reference and generate a value which when added to a calibration constant would form a calibration parameter continuously corrected for clock errors. This mode of backup operation has all of its modest added complexity on shore, and may be preferable to giving each string a more elaborate and extremely stable, e.g., Rubidium, clock at the sea bottom.

10) However long or short the trigger's local extension counter, it will eventually roll over. I suggest that when the "calibrated" time rolls over, a rollover message be sent along with the data so that an additional extension counter of any desired length could be maintained in software during further processing. John Learned can actually have Julian date in nanoseconds if he wants it. Consistent use of rollover messages completely eliminates any loss at counter rollover.

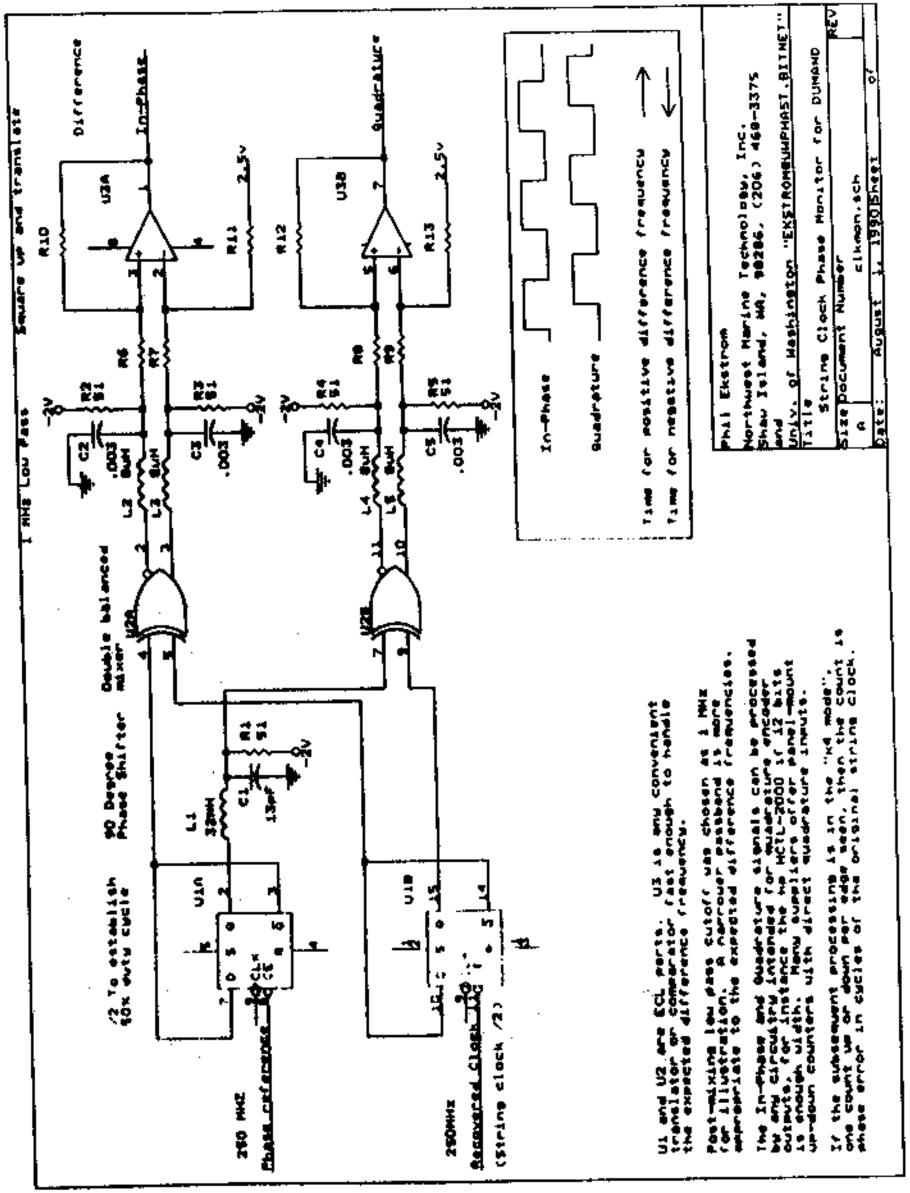

APPENDIX:

Brief Regarding String Clocks

by Phil Ekstrom
8/13/90

At the Seattle meeting I heard concern expressed over the fact that an ordinary crystal clock in the SBC digitizer would not be sufficiently stable to permit continued operation if its lock to the CCC link should fail. The proposal offered there was to put a better, e.g., Rubidium, clock in each String Bottom Controller. This will do the job, but I suggest an alternate way which puts the added hardware on shore where it can be repaired.

The idea is not to control the string clock phase drift, but to monitor and correct for it. Then a crystal clock is stable enough. This can be done transparently with a remarkably small amount of added hardware and code in the trigger processor, plus one precaution taken in the SBC.

In the SBC one must only choose the transmitted baud rate of the fast link to be some simple, preferably binary, submultiple of the string clock: 1/2 implying 250 MHz comes to mind. Then the baud rate clock which the on-shore receiver must recover anyway is a string clock signal available in the trigger whenever the link is working at all.

Next one combines this with the onshore phase reference in a circuit which can be viewed as a quadrature lock-in. The circuit can be very simple, using parts of three ICs. The two difference frequency outputs are quadrature square waves which can be handled with the same circuitry and algorithms used for quadrature angle encoders. In particular, you get to count either up or down as appropriate. It is easy to arrange for one count to be worth one cycle phase error of the string clock. (Actually it would be better for it to be 1/2 cycle or less, which is also easy, to eliminate any contribution to quantizing errors.) Further details depend on Matty's choices here, but those two low-frequency signals could be sent directly to the DSP, where he would maintain an up-down counter of accumulated clock phase error in, say, nanoseconds.

First level trigger processing involving only one string would not need any correction. When adding the calibration parameter to ready the data for correlation between strings, the DSP would also subtract the clock error count. That's all it takes.

See Seattle Proceedings for a circuit.

SBC FAST DIGITIZER SPECIFICATION

by Eric Hazen,
Boston University

I. General Description

A. SBC Inputs

The SBC must digitize OM pulses from 24 OM's and 3 CM's with a least count of 1.95 nS (1/512 MHz) [hereafter referred to as "2nS" and "500MHz" to keep things simple]. Both leading and trailing edges of each pulse are digitized. Any arbitrary signal on the inputs can be digitized, provided that the total number of input transitions doesn't exceed the internal buffer size. 6 auxiliary inputs (for hydrophones, etc) are sampled at 1MHz.

B. Output Data Stream

For each transition, a 10-bit time word, 4-bit channel number, and 1 bit up/down flag are transmitted. Every 1 microsec (when a carry occurs from the 9th time bit), a roll-over word with channel number = 0 is transmitted. The time field of roll-over words is by definition zero, so the 10 time bits and up/down flag are available for auxiliary data. An overflow bit indicates lost data.

C. Physical design

The OM and CM data comes in on optical fibers. The fibers connect to optical receiver boards, which plug into the main fast SBC board. The output of the optical receive cards are differential ECL signals, which feed directly into the digitizer chip.

The fast digitization is accomplished by a two semi-custom ICs, each handling 12 OM inputs and up to 2 CMs. The OM and CM data is fed directly to the chip, along with a 500MHz master clock. The output of each chip feeds directly to a serializer-encoder (perhaps the Gazelle "Hot Rod") chip, and then to a laser driver.

On-chip buffering is provided for 50-64 events, to handle bursts. The worst-case burst expected is one double pulse on each of the 12 OM inputs, within approximately 200nS. This generates 48 transition events which must be buffered.

II. Detailed Digitizer IC Description

A. Front-End (edge detector)

The first task of the digitizer is to synchronize the 14 inputs (12 OM's + 2 CM's) with the 500MHz clock, and to detect rising and falling edges. This is accomplished with several latches and an exclusive-or (XOR) gate on each input channel. Edge detection is done by XORing two successive samples of an input. The XOR output is high when the two samples differ, indicating that the input has changed state. A 14-input OR gate senses when an edge is present on any of the 14 input channels, producing the TRIGGER signal. 14 additional bits record whether the transition on each input is a leading or trailing pulse edge.

The output from this stage is 29 bits:

1	TRIGGER
14	change register
14	direction register

B. Time Stamp

The arrival time of each bit (TRIGGER) is recorded to 2nS accuracy. A 12-bit synchronous counter runs continuously at 500MHz, and the low 10 bits are latched on each TRIGGER. When a carry occurs from the 9th bit (every 1 microsec), a word is latched with a flag set to indicate roll-over. If a trigger is present when roll-over occurs, data is recorded as usual in the change and direction registers, otherwise these fields contain garbage.

C. Auxilary Data

When a clock roll-over occurs (every 1 microsec), 6 external inputs are latched, along with the upper 2 bits of the clock synchronous counter. These 8 bits, plus 2 zeroes, replace the 10 bit time word data.

The data stream at this point consists of:

1	TRIGGER
14	change register
14	direction register
10	time word (or auxiliary input data)
1	roll-over flag

D. FIFO

A FIFO buffer for several words is required in the digitizer for several reasons. First, data can accumulate in the digitizer at a rate of up to 500MHz, and it couldn't possibly be shipped off the chip at that rate. Second, each word may contain several hits, requiring the transmission of several output words off-chip.

A 39 (14+14+10+1) bit word is written into the FIFO on each TRIGGER or roll-over. The FIFO must accommodate input data at up to 500MHz in short bursts, if closely-space transitions occur on several inputs. Data is read out of the FIFO at a maximum rate of 100MHz.

The FIFO may be full when data is to be written. If so, no data is written, but a LOST DATA flag is set in the top word on the FIFO. In addition, a single free location is always reserved for roll-over words at the input to the FIFO, guaranteeing that no roll-over words will be lost.

The FIFO depth required is approximately 64 words (see discussion in section I.)

E. Channel Number Encoder

Each output word from the FIFO is latched into an encoder, which scans the change register for 1 bits. Each time one is found, a word is transmitted off the chip with the current time stamp. A priority encoder converts the 1-of-14 input from the change register into a 4-bit binary channel number. The channel number is used to select the appropriate bit from the direction register for output.

When the roll-over flag is set, a word is transmitted with channel number 0, and auxiliary data in the clock field. If any bits are set in the change register, normal data words are also transmitted, with the clock field set to 0.

The following output word is built:

4	channel number (zero for roll-overs)
10	time word (aux data for roll-overs)
1	direction bit (zero used for roll-overs)
1	lost data (preceding this word)

Roll-over aux data word (replaces time word):

2	roll-over counter
3	Hydrophone (3 12-bit serial streams)
1	framing bit for above
1	EM
1	Command & Control output (serial)
2	spares (set to zero)

F. Output

A 16-bit word is transmitted off-chip. The simplest scheme for doing this is to simply latch the word into an output register, and wait for the laser serializer (Hot Rod) chip to remove it.

III. Conclusions

We have studied the requirements of the SBC digitizer chip in some detail, and believe that the above design will meet them, and can fairly easily be implemented on a single GaAs ASIC. Using Tri-Quint semiconductor's QLSI standard cell family as an example, we have designed and analyzed several of the key (read "hard!") parts of the IC, including the front-end, counters, and FIFO control logic, and encountered no serious problems.

The major remaining design problem is to provide sufficient buffering on the chip to handle bursts. It remains to be seen whether current GaAs or ECL arrays are large enough.

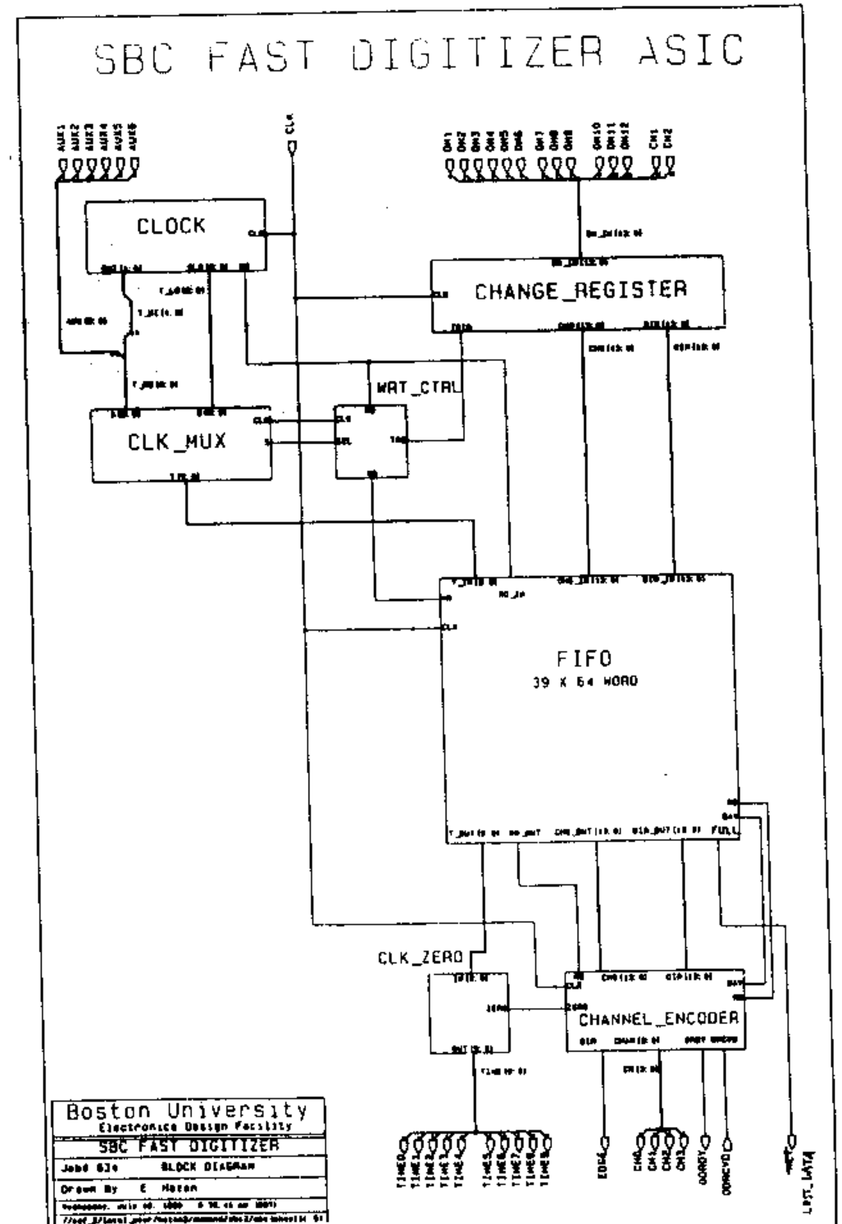

JOM pulse structure

Akira Yamaguchi (Tohoku Univ.)

At Bren meeting (June 1990), we have proposed Japan Optical Module pulse structure having information on "pulse width and charge" of a photo-tube output pulse. This information may be useful for the analysis of multi-muon events etc.

T. Hayashino (Tohoku) has developed the Monte calro simulation for multi-muon events based on the interaction of high energy cosmic ray (proton and iron nuclei) with air (see Minutes of Bern meeting). He showed Frejus data on multi-muon flux was explained well by very high energy cosmic ray of proton 50% and Fe 50% incidents. When we apply this MC simulation to DUMAND site, we can summarize the simulation results as follows:

(1) High μ's multiplicity events ($n_\mu > 10$) were cause only by Fe incidents.

(2) Muons of high μ's event from Fe incident are in spatial spread of ~10m.

(3) On the contrary, muons from proton incident are in small spatial spread of <3m.

OM output pulse will be wider proportional to muon spatial spread. So, the pulse width will give a good information for the analysis of cosmic ray composition. Especially, this physics may be done with 3-string DUMAND because of no need of angular resolution such as 1°. However, we have a question how the multi-muon event is seen at DUMAND array and haw it can be separated from a single-muon event. We need more advanced simulation to study the high muon multiplicity events. T. Hayashino shipped out the multi-muon events on a magnetic tape to J.

Learned, C. Ley, and A. Okada. They will test the effect of 2-pulse structure for the track reconstruction and the analysis of multi-muon physics.

H. Kawamoto (Tohoku) and A. Okada (Tokyo, ICRR) are developing circuits to output the 2-pulse, i. e. both pulse width and charge of a OM output pulse. The pulse structure proposed at Bern meeting is shown in Fig. 1 (a). The first pulse has a width of OM output pulse over threshold. The second pulse is proportional to charge ($\ln Q$) measured using a integrator circuit and starts at 100nsec after the beginning of the first pulse.

After the discussion about the 2-pulse structure at Bern, we concluded to need a notch in order to recognize two pulses as one events by the trigger processor. We have considered to put a notch of 5~10 nsec between two pulses as shown in Fig. 1 (b).

At present, we are trying to test two circuits for notching. One (Fig. 2) is a circuit to make the second pulse from measuring a time from the peak of integrator output pulse to its threshold. The other (Fig. 3) uses a sample / hold IC and output a decay time as the second pulse after sampling and/or holding integrator output pulse. The dead time of these circuits may be several hundred nsec.

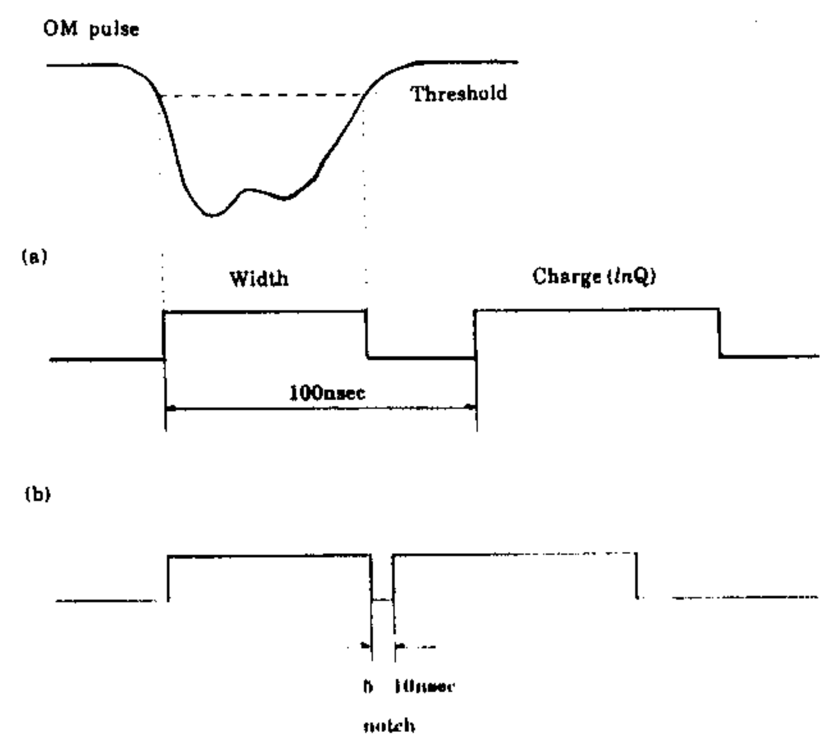

Fig. 1. Two pulse structure.

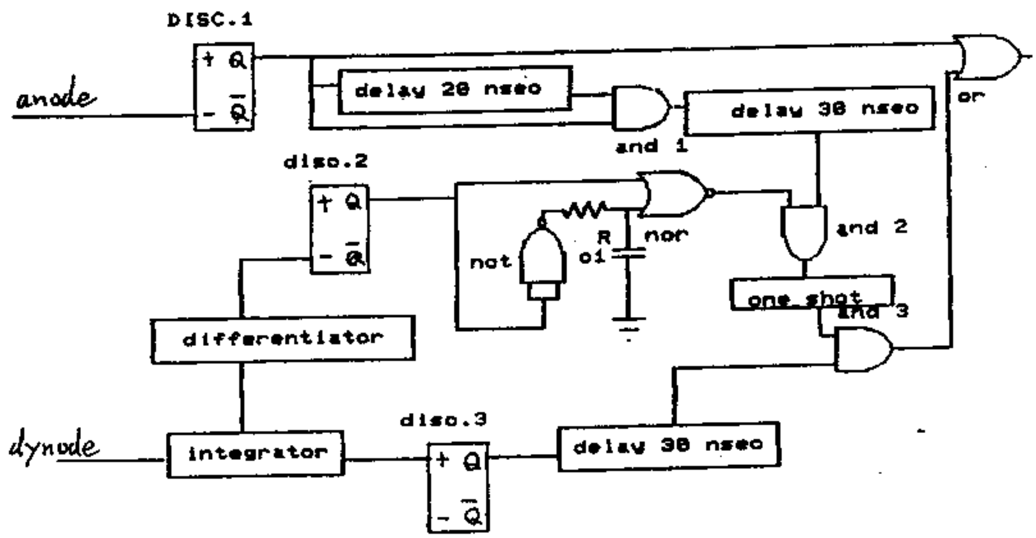

Fig. 2. Circuit-1 : charge pulse is proportional to an integrator pulse width.

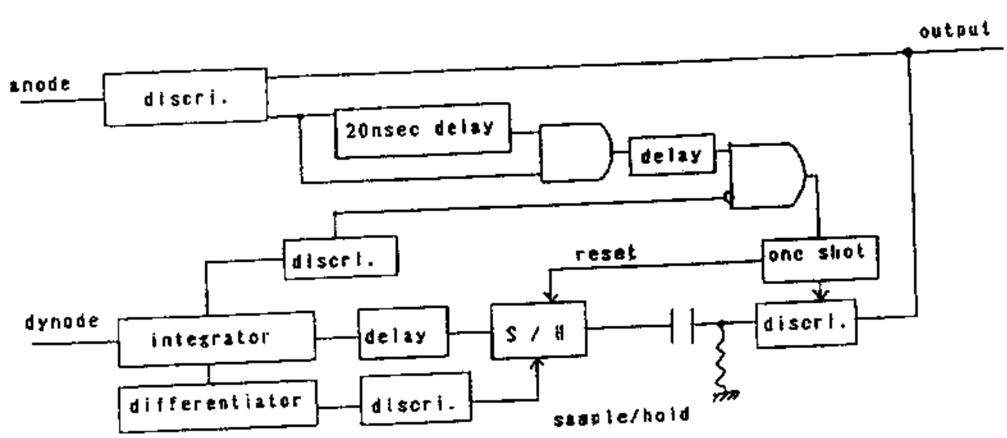

Fig. 3. Circuit-2 : charge pulse is proportional to a decay time measured by sampling/holding an integrator pulse.

FEATURE EXTRACTION AND TRIGGER IN THE DUMAND II DETECTOR.

by Kenneth K. Young,
University of Washington

ABSTRACT.

In order to trigger the DUMAND II experiment for a wide band of event types, it may be necessary to interpret subtle patterns in 3 configuration space , time and pulse height space. With this in mind, we examine the work at accelerator laboratories to carry out this task. The work includes the use of custom hardware, RISC farms, Transputer Arrays, and the application of Commercial Imaging Processors. Each of the technologies have unique strengths and weaknesses. This survey will point out some of these so that we can plan ahead for applications to Dumand II. In general the custom hardware is the fastest but it is also the least flexible so that adaptation to realistic conditions may not be fulfillable.

Dumand Trigger Meeting in Seattle, July, 1990

I. INTRODUCTION.

A typical accelerator experiment will use a large battery of detectors encompassing, tracking devices, energy measurement devices, and timing devices. Interesting events usually leave a signature which is composed of many or all of the devices. A modern trigger system will take many or all of the devices into account before recording the event. This is often necessary to avoid untenable dead time and data record length. In DUMAND II, our basic data are pulses from the optical modules (OM's) only so our task is much simpler. However, the data rate from background is quite high and the signal rate is quite low. We also have to take into account the spatial, time and pulse height coordinates to reconstruct events. The experience at accelerators may be applicable to Dumand II needs.

II. EXPERIENCE AT ACCELERATOR EXPERIMENTS.

II.1 Custom Hardware.

For experiments at FNAL, a system of plug-in boards and a NEVIS BUS system was developed to provide real-time reconstruction of tracks and coordination with calorimeter showers, and muon tracks in experiment E605. A more advanced system was applied to other experiments. The experience in E605 was that it was not very useful. The designed system could not handle the ambient background during the experiment. Since the device was hardware which required long lead times for modification, it was not sufficiently flexible to modify to meet the ambient conditions. In this case, the hardware met the original specifications but it did not have the flexibility to be modified for real conditions. Much later versions of the experiment did use the trigger hardware successfully when there was enough time to make a proper marriage of hardware to real conditions.

The CDF experiment at the Collider at FNAL has constructed specialized hardware to trigger on calorimeter showers. This has worked very well as the design specifications was sufficiently close to the real conditions of the experiment. However, modifications to the trigger to accept more subtle and new signals have proven to be quite difficult and requires much R and D time.

Both hardware systems have not found general outside of the particular experiments which fostered their development. It's probably fair to conclude that these specialized systems require much labour to implement and are not transferrable to other experiments.

II.2 SLD AT SLC

The SLD experiment at SLC has a DAQ system which is rather similar to the Dumand II DAQ. The detectors have local digitization and the digitized information is multiplexed and sent to the counting house via optical cables. See figure 1.

The original intention was to use RISC processors to correct and to trigger on the data stream in real time. As the SLD experiment has come closer to commissioning time, the amount of real-time processing using RISC chips has diminished. Now only the physics corrections to the data is processed in the data stream. The trigger function has been relegated to a later stage. The data stream is first built into an Event using a commercially obtained EVENT BUILDER. The EVENT BUILDER follows the ALEPH design and is produced by a Swiss company, Struck at a cost of $50K. The Event is then processed using conventional 68020 processors in software.

The cautionary tale to draw from SLD that even for a technologically very powerful group like the SLD experiment and with years of lead time, the group had to back down from customized use of RISC chips to the use of a commercial Event Builder and conventional software.

II.3 ZEUS AT HERA

The ZEUS experiment at HERA is a large modern collider experiment using an approximately spherical detector of many kinds of trackers, and calorimeters. The proposed trigger technology in this experiment had novel features which excited my imagination. Typically, for analysis and triggering, the experimenters will unfold a spherical detector into a plane (like world maps). The disadvantage of this is that you lose insights as to the relationships at the boundaries. The typical analysis then writes many special programmes to handle the boundaries and one finds often that the boundary regions are quite troublesome and reuire much more resources than the regions in the centers of the planes. The proposed trigger for ZEUS would utilize a spherical array of Transputers. The Transputers are RISC processor which have the special facility of being connectable to neighbours via its standard multiple ports. In Zeus, conventional commerical technology would be implemented in flexible and straightforward way to handle the 3 dimensional data without the difficulty of boundary regions. Figures 2 and 3 shows the Zeus transputer network.

Transputers are now available in plug in boards for PCs. The GESPAC system has dual transputers on a card which can be assembled into the proper array for your experiment.

The ZEUS physicists compared the T800 tranputer with other processors.

Processor	Speed	Whetstones
Intel 80286/80287	8 MHz	300K
MC 68020/68881	16/12 MHz	755K
IMS T800/30	30 MHz	6000K

Camerini cautions us that the ZEUS experiment has not implemented this hardware in any form even 1.5 years after some of the proponents for this system seemed optimistic that implementation would be quite straight forward.

II.4 LAA EXPERIENCE AT CERN

I heard about the LAA experience from the SPACAL experimenters reporting at SNOWMASS '90. They had applied the DATACUBE parallel pipeline image processor successfully to their data from a calorimeter test run to separate electrons from hadrons. While most e's are quite different from hadrons, some patterns overlap quite significantly. This is important when one requires a high degree of discrimination with the requirment of high efficiency.

The philosophy of the approach by R.K. Bock at CERN is instructive. I quote the abstract from their paper, "FEATURE EXTRACTION IN FUTURE DETECTORS", R. K. Bock et al, NIM A289, p 534 (1990).

Abstract.

"As part of CERN's LAA detector development project, we have studied the possibilities of using modern fine-graine parallel comlputer structures for extracting high-level information (physics features) from fast detectors, in real time. Such information could subsequently be substituted for raw data and thus alleviate bandwidth problems, or could serve for "intellegent" triggering. The main objective is to show how fast today's ditial devices, implemented as custom-made or commercial processors, can execute basic algorithms. The possibilities of integrating them efficiently in the dta flow of experiments, and the extrapolation to be made for technological evolution ove the coming years will be critical parameters when judging realistic applications in the future."

I'll summarize their paper here.

Bock et al considered three commerically available systems:

a)DATACUBE. This is a board system on the VME bus which executes standard image algorthms in hardware at a 10 MHz rate. (this is soon to be 100 MHz) Figure 4 shows the peak finding algorithm configuration used by Bock et al.

b)GAPP, Geometric Arithmetic Parallel Processor, GAPP SCG72 made by NCR, a SIMD machine.

c)ASP, an Associative String Processor. A specific implementation will soon be delivered to Brunel University for track analysis in experiment NA 35.

These systems were evaluated by Bock et al in carrying out specific tasks in pattern recognition. Bock comments that MIPS and MHz parameters are not useful in evaluating systems for YOUR task. The systems were asked to carry out specific tasks within 10 microseconds. These tasks were:

a)peak finding

b)pattern recognition by parametrization

c)calorimeter cluster analysis

d)E_t calculation in calorimeters

e) track finding.

Bock finds:

a) All of the systems were able carry out most of the tasks within 100 microsec. However, technical improvements would soon bring the time to less than 10 microsec.

b) Peak finding in these systems were equivalent to > 10^3 Vaxen. with ASP, peak finding was > 10^6 Vaxen.

c) Global sums were > 10 Vaxen. (Specialized Hardware is much better for this task)

d) DATACUBE is easy to program and interface but is the least powerful.

e) ASP is probably the strongest contender.

SpaCal Experience, see CERN LAA RT/90-01, "ALGORTHMS AND AN IMAGE PROCESSING ARCHITECTURE FOR ON-LINE IDENTIFICATION FROM LATERAL PROFILES IN SPACAL.", Bock et al.

SpaCal is a CERN group who are building and evaluating Spaghetti calorimeters for use in post LEP experiments. The spaghetti calorimeter is composed of scintillating fibers imbedded in a lead matrix. The light is projected on to a plane where it is read-out by PM's. This is a imaging-kind-of calorimeter. The aim is to use the image to distinguish electrons from hadrons in real-time for trigger purposes. The SpaCal group found that implementation of the DATACUBE system provided excellent e-hadron separation.

II.5 NEURAL NETS.

Since patterns of e and hadron showers are "random" patterns they are not charactizeable by linear algorthms which can easily be performed by normal processors. Our visual cortex can easily distinguish patterns and can learn patterns. Attempts to emulate human vision have partially succeeded. David Cutts at Brown University has analyzed e and hadron showers in a calorimeter using feed-forward neural nets and by conventional linear algorthms. He has found that the neural net system is significantly better than the results from linear algorthms. He is considering implementing NN in the D0 experiment for the trigger.

In CDF, Bruce Denby and Myron Campell have received funding from the DOE to make a hardware NN for the purpose of triggering the CDF experiment on low energy electrons. They have studied the NN paradigm using software emulators. This has been quite successful. This group now is trying to implement the NN in hardware using commercially available NN chips. The NN hardware being considered runs in approximately 10's of microseconds.

Both of these projects looks promising but we probably won't see results of the hardware systems for a couple of years.
We can emulate NN for the Dumand II trigger now using standard software packages available commercially or from Cutts or Denby. Denby has a highly documented emulation program which will run on the VAX.

III. CONCLUSIONS.

RISC farm talk is cheap. Powerful groups have gone down this road without producing a viable flexible system.

Commercial producers like DATACUBE may have a product which is directly applicable to our needs.

Neural Nets are fun but no useful hardware examples exist at this time.

The Dumand II DAQ should be sufficiently flexible to take the global trigger systems into its fold. It looks like the DAQ system is acceptable.

List of figures

figure 1.

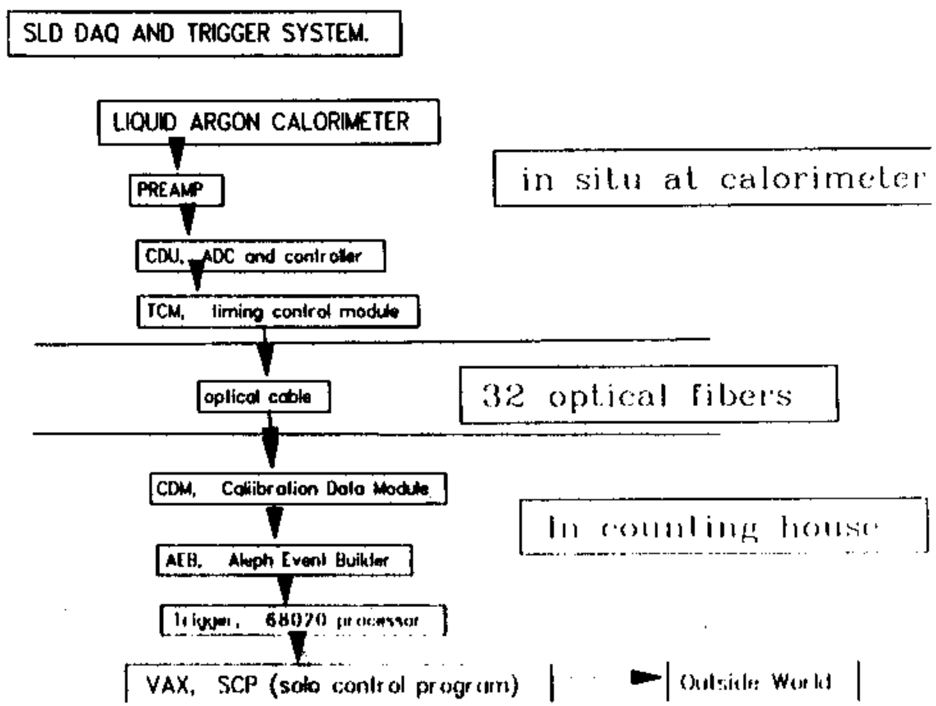

Zeus Figure 2

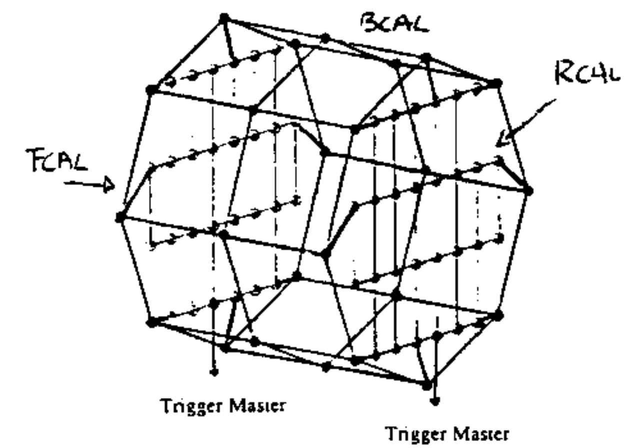

Transputer network for the second level calorimeter trigger. The transputers in the forward, barrel and rear parts of the calorimeter are connected via links. The links between the parts are dashed. The results of the two trigger processors are combined and sent to the global trigger.

TRANSPUTER NETWORK FOR TRIGGERING (OPTIONAL)

Figure 3

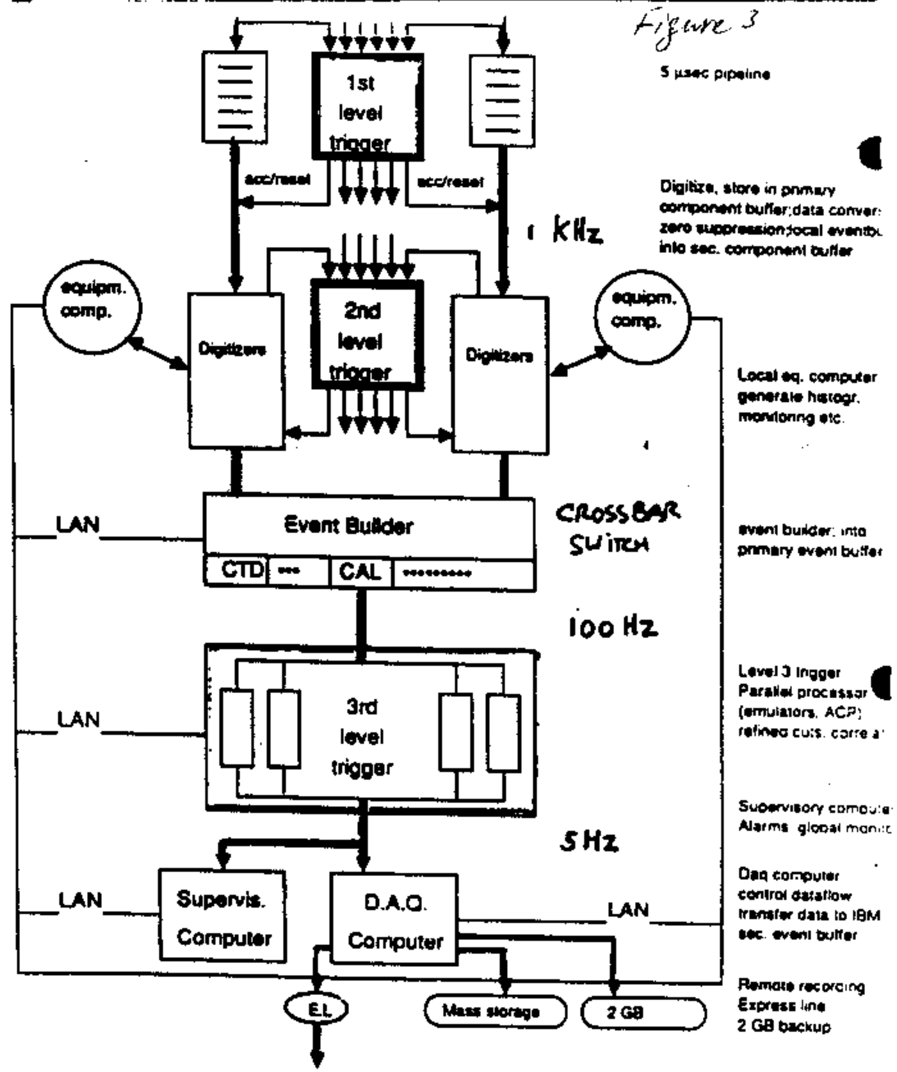

fig. 2. Outline of the ZEUS data-acquisition system

FIGURE 4

An implementation of the "peak-finding" algorithm using standard DataCube modules.

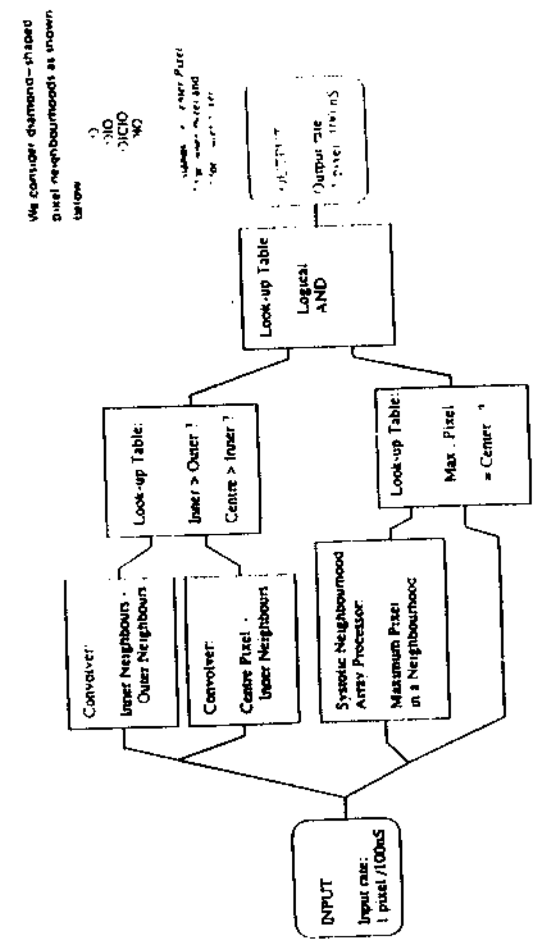

TAXI Chips

by Dick Davisson,
University of Washington

We need some way to convert the data back and forth between the very high clock-rate serial form appropriate for transmission over long distances and the lower clock-rate parallel bus form appropriate for data manipulation.

This is a need we share with others. There are commercially available circuits designed to meet this need. We would be well advised to take advantage of the engineering efforts of others.

One such commercial device of which I am aware is the TAXI chip pair from Advanced Micro Devices (AMD). The pair, comprising a transmitter chip and a receiver chip, work to create in effect a one-deep FIFO with the input register at one end of a high-speed serial data link, and the output register at the other end. The serial data link is not visible from the registers and is completely self-contained. The encoding and decoding of data, and the synchronization of the link are all attended to by the TAXI chips. The chips accept asynchronous data, but to realize the highest data throughput in an asynchronous environment, the TAXIs should be embedded as a layer of a FIFO cascade.

The data rate of the TAXI chips reaches 146 Mbps with the serial clock at 175 MHz. The discrepancy arises from the 5-6 encoding scheme which is used. Their encoding provides some error correction but is mostly useful for error detection. The encoding scheme used seems to be optimized for some task different from ours, meaning that I might have done it differently, but that is a minor point. Overall, it looks like a very useful device.

Not surprisingly, those at the Trigger Meeting who had been involved in the electronics for DUMAND were already aware of the TAXI chips and also of some even faster GaAs chips called Hot Rods. The two types are similar in purpose and operation but they fit their names: the Hot Rod is much faster than the TAXI. The serial clock rate for the TAXI can be anywhere in the range 40--175 MHz. The range for the Hot Rod includes 250 MHz (but I don't know its limits). In general, high clock rates can be put to good use in DUMAND but we must not kill ourselves by overreaching.

8/15/90

Figure 1

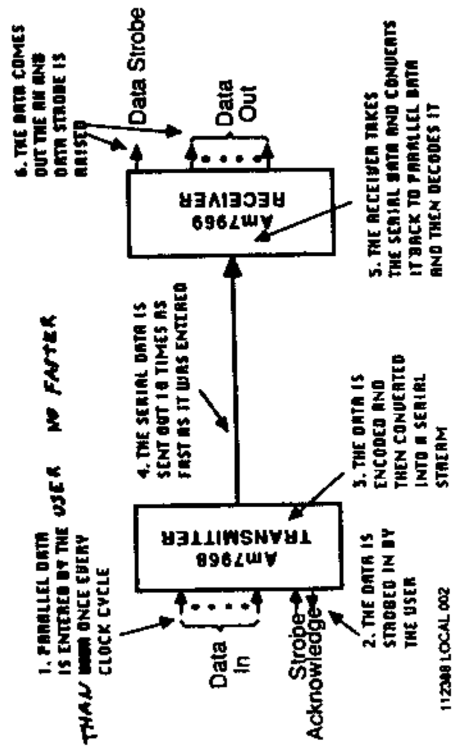

TRIGGER PROCESSOR DESIGN AT PRESENT

by Matthew Jaworski,
University of Wisconsin

July 12, 1990

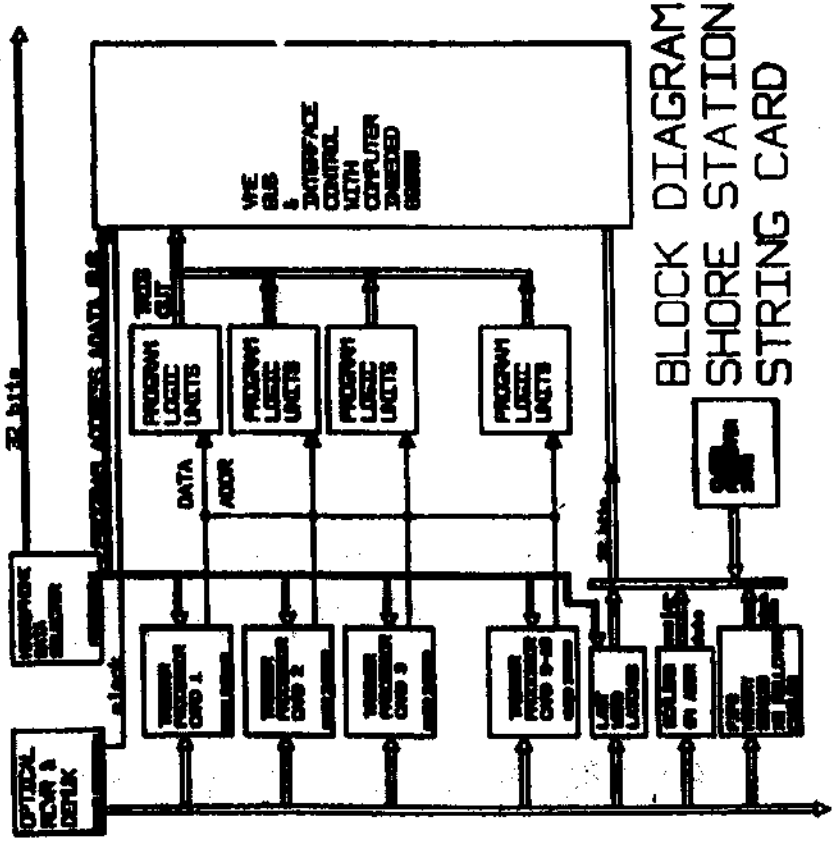

BLOCK DIAGRAM
SHORE STATION
STRING CARD

gazelle

ADVANCE INFORMATION

HOT ROD™

High-Speed Serial Link
Gallium Arsenide

General Description

The HOT ROD Transmitter and Receiver Chipset is a general-purpose interface for very high speed point-to-point serial communication over coaxial or fiber-optic media. HOT ROD is used like a pseudo-parallel register, loading data into one side and outputting it on the other. In this case, however, the "other" side is separated by a transparent serial link.

HOT ROD provides a complete solution for implementing the serial link, including parallel bus interface, using a simple strobe/acknowledge protocol; parallel-to-serial conversion; framing; 4B/5B encoding; error monitoring; clock generation; and clock recovery.

The speed of a HOT ROD system is adjustable from 2.5 Mwords/sec (100 Mbits/sec) at the low end, up to 25 Mwords/sec (1000 Mbits/sec) at the high end. Higher throughput rates can be achieved by cascading HOT ROD systems. HOT ROD's flexible bus interface scheme accepts a TTL-compatible, 40-bit wide word, which could consist, for example, of thirty-two bits of data and eight bits of command or parity information.

HOT ROD Transmitter

The HOT ROD transmitter accepts 40-bit parallel data from the host system to be latched, encoded, and transmitted. When the host is clocked synchronously with HOT ROD, one word is transmitted during each cycle of the 1XCLK so that the 1XCLK or 2XCLK outputs of the transmitter can be used to clock the logic supplying data to HOT ROD. In asynchronous operation, a strobe/acknowledge protocol is used to control data flow to HOT ROD. A rising edge on the STRB input causes the data (D_0-D_{39}) inputs to be latched into the transmitter. In most systems these forty bits will consist of four bytes (32 bits) of data and eight bits of parity or command information. Since all forty bits are treated identically by the chip, the distinction between data, command, and parity bits is solely at the discretion of the user.

Parallel data is latched into the Hot Rod transmitter on the rising edge of STRB. The ACK output rises when this data is transferred from the Input Latch to the Encoder Latch, indicating that the input latch is ready to accept another parallel word.

Features

- Parallel TTL bus interface
  - 40 bits wide
  - All bits are user definable
- Transparent synchronous serial link
  - ECL 100K-compatible
  - AC or DC coupled
  - NRZI 4B/5B encoding/decoding
- High-speed 100-1000 Mbits/sec serial data rate
- Cascadable for multiple Gigabit/sec throughput
- Asynchronous parallel input using STRB/ACK
- On-chip PLL Crystal Oscillator (transmitter)
- On-chip clock recovery PLL (receiver)
  - No external analog circuitry required
- Loopback mode for diagnostic
- Single +5 V supply operation
- Power dissipation
  - 1.3 W max (transmitter)
  - 2.0 W max (receiver)
- Coax or fiber optic compatible
- 58-pin PLCC or LCC

Typical Applications

- CPU to CPU
- CPU to Peripheral
- CPU to Graphics Terminal
- Backplane Extender
- Ribbon Cable Replacement

Point-to-Point Communications

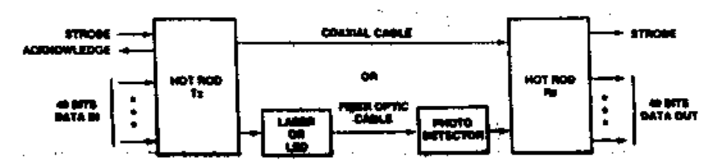

Gazelle Microcircuits, Inc. • 2300 Owen Street • Santa Clara, CA 95054 • (408) 982-0900 • FAX (408) 982-022

* CONSTRUCTION
- FIRST PROTOTYPE WR
- SM FOR SPEED & RELIAB

* TRIGGER LOGIC

HARD WIRED Σ of Δ, Δ of Δ

PLD
XILINX
DSP

IN INCREASING ORDER OF FLEXIBILITY

TRIGGER PROCESSING CD
9U

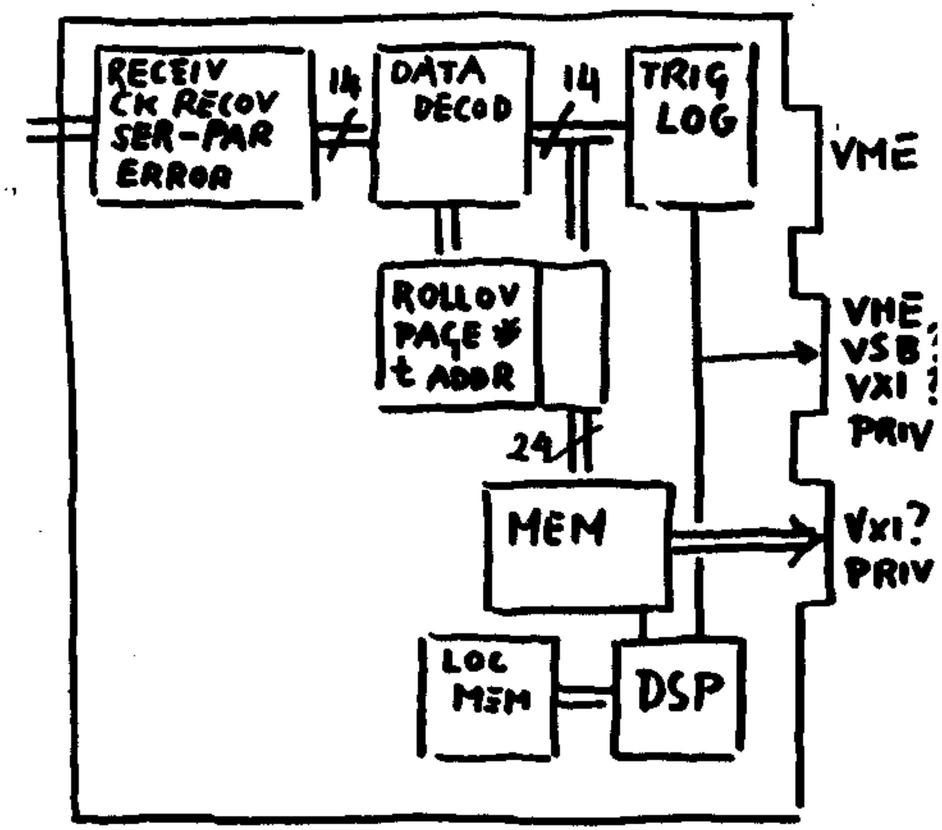

STRING BOTTOM CONTROLLER

by Jeff Bosel,
University of Hawaii

July 12, 1990

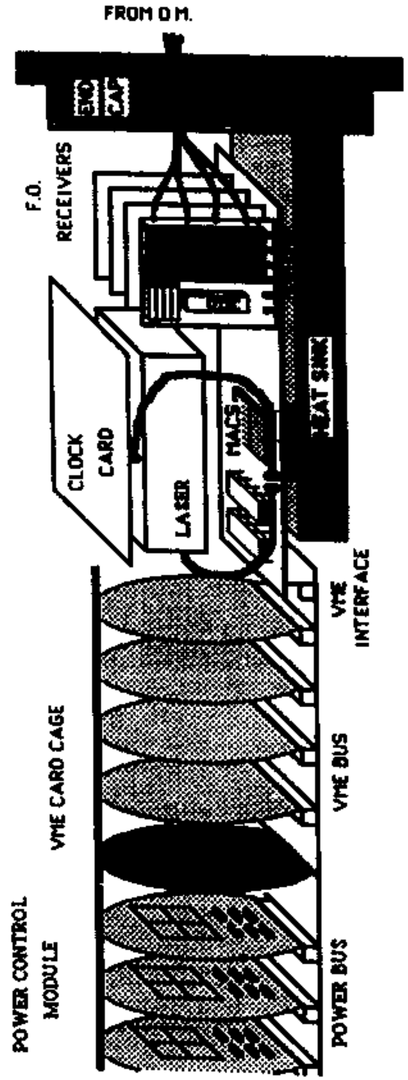

DUMAND SBC

SUMMARY OF DECISIONS
FOR PROCEEDINGS

Critical Decisions for Dumand II
Taken at Trigger Subgroup Meeting, Seattle, 12-13 July, 1990

We have decided in favour of the fast data stream that has been outlined by Phil Ekstrom (see contribution on page 41). This means that the hydrophones will be digitized and entered into the data stream at regular intervals at the rollover word. At the rollover, the NBU data will also be recorded. The SBC will receive the analogue hydrophone data and the digital NBU data. At the JB, the EMC will pass on the digitized NBU, hydrophone data. The TV data will be superimposed either by digitized signals or by a different color FM signal.

1. Rollover 1 microsec, 10 bits time data

2. 2 ns least count for the OM digitization. The subgroup decided that 2 nsec is necessary, but 1 nsec would be too difficult and expensive for the time available.

3. GaAs or Si(CMOS). Both are OK. Eric Hazen will decide.

4. Multihit specifications for OMs: 48 hits/200 ns and 12 hits/2 nsec (Note that 'hit' means input transition, so that one PMT pulse may generate up to 4 electronics hits, if we use the Yamaguchi double-pulse scheme).

o Eric Hazen: will determine what is possible by 10/90
o J. Learned will simulate

5. 2 color allocations to optical fibers: the 2 colors will be used to provide redundant data paths, since the Ekstrom scheme provides adequate rates for data transfer.

6. Rate Limiter will be located at OM and be programmable.

7. SBC data stream will be 16 bit words:

```
Optical Data     |        time      addr
                          9 bits    |1|1| 5 |
                                      start/stop
                                        lost data (overflow)

Rollover         |        data      addr
                          9 bits    |1|1| 5 |
                                      not used
                                        lost data (overflow)
```

Hydrophones	3
NBU	1
C2	1
framing	1
rollover count	2
spare	2

Regarding the Environmental Module at the JB we have decided:

1. G96 bus with 68000 CMOS microprocessor (Gespac MPL4080).

2. Digitize hydrophones with 12 bits.

3. JB-EM: Fast Data:

o Hydrophones. Use "hot rod" chips.
o Video. Use digital or FM. Choice will be up to UWA.

4. UH to build C2, C3 MUX/DeMUX.

5. UH to supply package plans for card cage for EMC.

6. NBU to be mounted at 100m from BOTTOM, not at top of string.

7. In OM strings, all of EM will be in NBU which will talk directly to the SBC.

8. The hydrophones will have local preamps and will talk directly by copper to the SBC. At the SBC, the analogue hydrophone signals will be digitized and the data will be inserted into the data stream at the rollover. (UWA will be responsible for the design of this sub-subsystem.)

These decisions were distributed by E-mail and FAX for prompt review by the Collaboration, and should now be considered firm with one exception: the 2 ns least count for OM data (instead of 1 ns) was found to require further study and a final decision should be taken at the Sendai Collaboration Meeting in October, 1990.