

DESIGN OF A TRIGGER AND DATA ACQUISITION SYSTEM
FOR THE SLAC/LBL/LLNL B FACTORY*

Don Briggs

*Stanford Linear Accelerator Center,
Stanford University, Stanford, CA 94309, USA*

Abstract

This report compares and contrasts the approaches taken at KEK and SLAC in the design of trigger and data acquisition systems for the high-luminosity environment of a B Factory. The description of the SLAC design presented obtains from work in progress by a working group^[1] of an ongoing workshop.

*Invited talk presented at the KEK Workshop for B Factory,
Tsukuba, Japan, April 15-18, 1991.*

* Work supported by Department of Energy contract DE-AC03-76SF00515.

1. Introduction

The B Factory trigger and data acquisition design studies at KEK and SLAC began with similar specifications but evolved in different directions. The KEK design uses a synchronous passive pipeline approach; whereas the SLAC design pursues an asynchronous active pipeline approach.

2. Comparison of Specifications and Goals

The design luminosities of the KEK and SLAC B Factories are essentially the same, between a few times 10^{33} and 10^{33} . The trigger and data rates presents new challenges to the design of trigger and data acquisition systems. The SLAC and KEK designs target the same main physics channels, namely, the Υ resonances including the low-multiplicity $\tau^+\tau^-$ final states. The expected physics rates are up to 200 Hz. The SLAC design provides reasonable efficiency for 2- γ physics at an additional rate of a few tens of Hertz. The background rates are comparable. The flux of cosmic rays is the same, but the SLAC design presents a far larger acceptance to cosmics early in the trigger process.

3. Comparison of Synchronous and Asynchronous Trigger Systems

Synchronous trigger systems use system-wide control signals to carry timing information (e.g., strobe, gate, clear, or reset signals). Asynchronous trigger systems decouple control from timing information, and require a more complicated form of hand-shaking, or data transfer protocol. We distinguish the response of the two forms in environments of high trigger rates. In a synchronous trigger, the leading impact of high rates of background on the physics derives from the trigger acceptance. It is time-dependent, and either zero or unity across the system. In an asynchronous trigger, the channel-by-channel occupancy dominates. See Figure 1.

An asynchronous trigger can be designed to be more robust against high background rates and to degrade gracefully. See Figure 2. We cannot guarantee the background rates with any certainty.

At SLAC, we think that the conservative approach is to provide an asynchronous trigger, which we think is more robust.

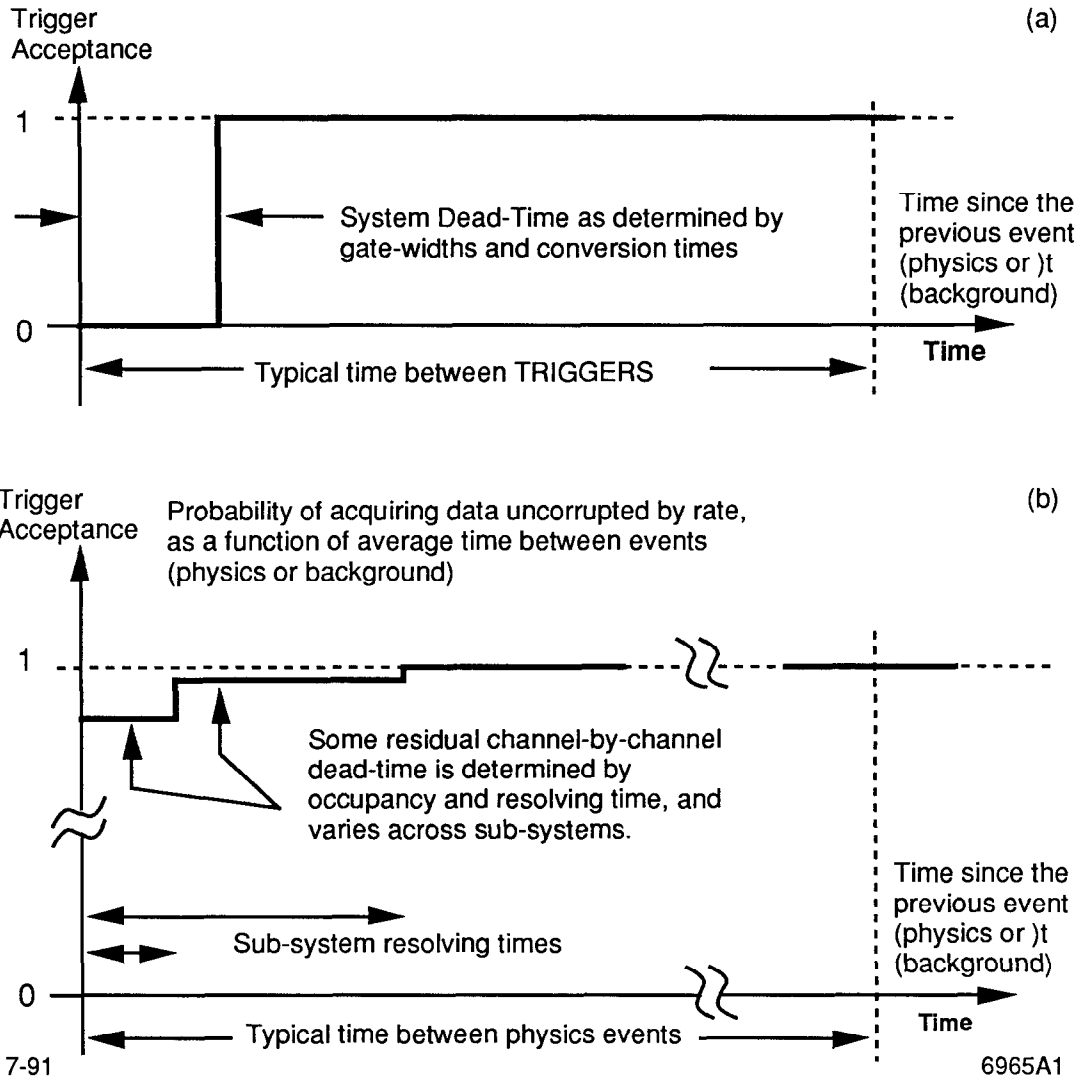


Figure 1. (a) Trigger acceptance as a function of the time since the previous trigger in a synchronous system. In addition, the data may acquired may be corrupted by rate. (b) The probability of acquiring a physics event uncorrupted by rate effects, as a function of the time since the previous event (physics or background) in an asynchronous system.

4. Review of the SLAC Detector Components

Moving from the Interaction Point outward, the SLAC detector design includes a Silicon Vertex Detector (SVD), a drift chamber, a particle identification system, a calorimeter, a conventional coil, and a muon system. The SLAC detector design does

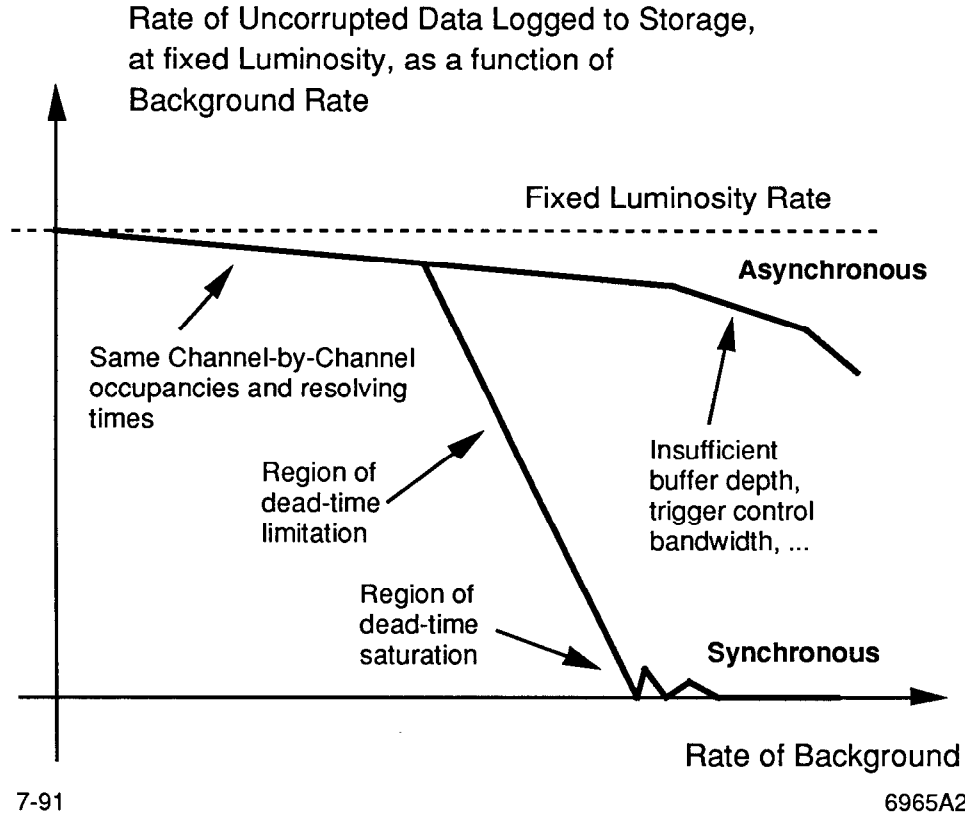


Figure 2. The rates at which data are logged to tape or disk, for fixed luminosity, as the background rate is increased, comparing a synchronous and an asynchronous system. At some point, a synchronous system becomes dead-time limited, then saturated. An asynchronous design can survive a higher background.

not yet include an intermediate tracker. Discussion on such a device proceeds on the physics merits, rather than on considerations of trigger implementation. The Drift Chamber will use a Low-Z gas, probably Helium with some CO₂ and iso-Butane. Early discussions focussed on jet-cell designs, but a small-cell design now dominates. We expect that we can use FADCs and avoid TDCs.

The SLAC detector design presently excludes Time-of-Flight for particle identification. A Time-of-Flight system is very advantageous for a trigger because it defines a narrow time window for all charged particles in its acceptance. We must deny participation in the trigger process to the forms of particle identification considered so far by the Particle Identification Working Group in the SLAC Workshop, namely,

CRID and Aerogel.

The calorimeter comprises towers of Cesium Iodide crystals instrumented with vacuum phototetrodes or Silicon photodiodes, followed by charge-sensitive amplifiers. The calorimeter instrumentation must support relatively modest occupancy, so we expect that more than one channel of peak-and-hold amplifier per CsI tower can provide asynchronous multihit capability. We expect a conventional coil to provide the magnetic field. We expect to instrument the flux return for muon identification and tracking.

5. Approach Used in the Trigger Design at SLAC

Two performance requirements drive the design. The first is that the performance of the trigger must not limit the effective luminosity of the accelerator. Project management sets as a goal minimum deadtime. The second is that the trigger must be fully efficient and redundant. In implementation, we strive for

- graceful degradation in case backgrounds are greater (or different) than expected;
- orthogonality of charged and neutral trigger at the lowest level; and
- pipelined operation at the lowest level.

We require, even in extremely adverse conditions, to take physics data at reduced rates and to study backgrounds, i.e., that the system degrade gracefully. A system that “freezes up” above some critical background rate we find unacceptable. Orthogonal triggers at the lowest level allow direct measurement of efficiencies and systematic effects of trigger acceptance. Pipelined operation at the lowest level decouples the throughput rate from the latency. Longer latency (the time between receiving the input and forming the output) allows more powerful discrimination and a lower data rate to the next stage. Asynchronous pipelined operation eliminates system-wide data capture operations (clear, gate, reset, readout) as a bottleneck to throughput. A prominent specification is that the trigger be completely testable—that it support single-step diagnostic operation, and every internal register must support diagnostic read-write.

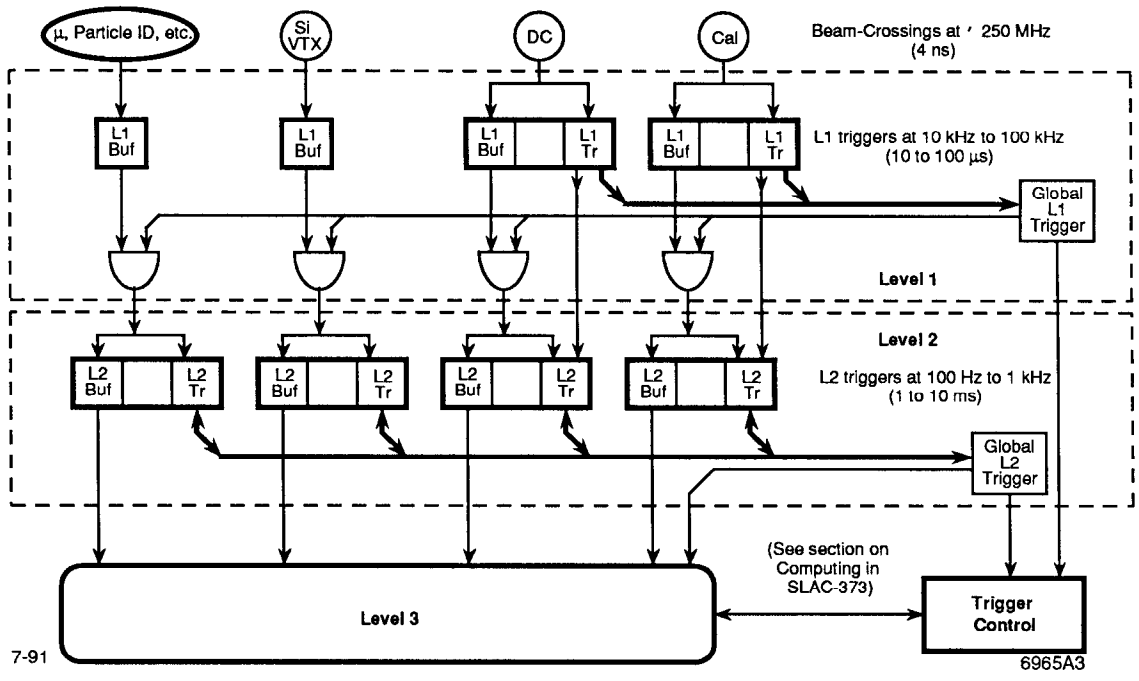


Figure 3. SLAC trigger and data acquisition architecture.

6. Architectural Overview

Figure 3 presents an overview of the SLAC design. Each subsystem is a source of raw data at the top. At the bottom, a farm of workstations is the destination of data. Level 1 of the trigger represents decision-making without resort to microprocessors, and no communication between subsystems. Level 2 represents the part of the trigger that might well use microprocessors and spans subsystems.

In Level 1, the drift chamber and calorimeter trigger independently. The drift chamber identifies events with a set of tracks satisfying the following criteria: two or more tracks to half its outer radius or farther from the Level 1 fiducial volume; and one or more tracks to its outer radius from the Level 1 fiducial volume. We call this the "1.5 track" trigger. The Level 1 fiducial volume is a cylindrical tracking volume about the interaction point, coaxial with the drift chamber, with an outer radius about one-tenth the inner radius of the drift chamber, and the same extent in length (length 3 m, radius ≈ 2 cm, probably larger than the beampipe). That is to say, we expect that tracks can be recognized and reconstructed to the precision commensurate

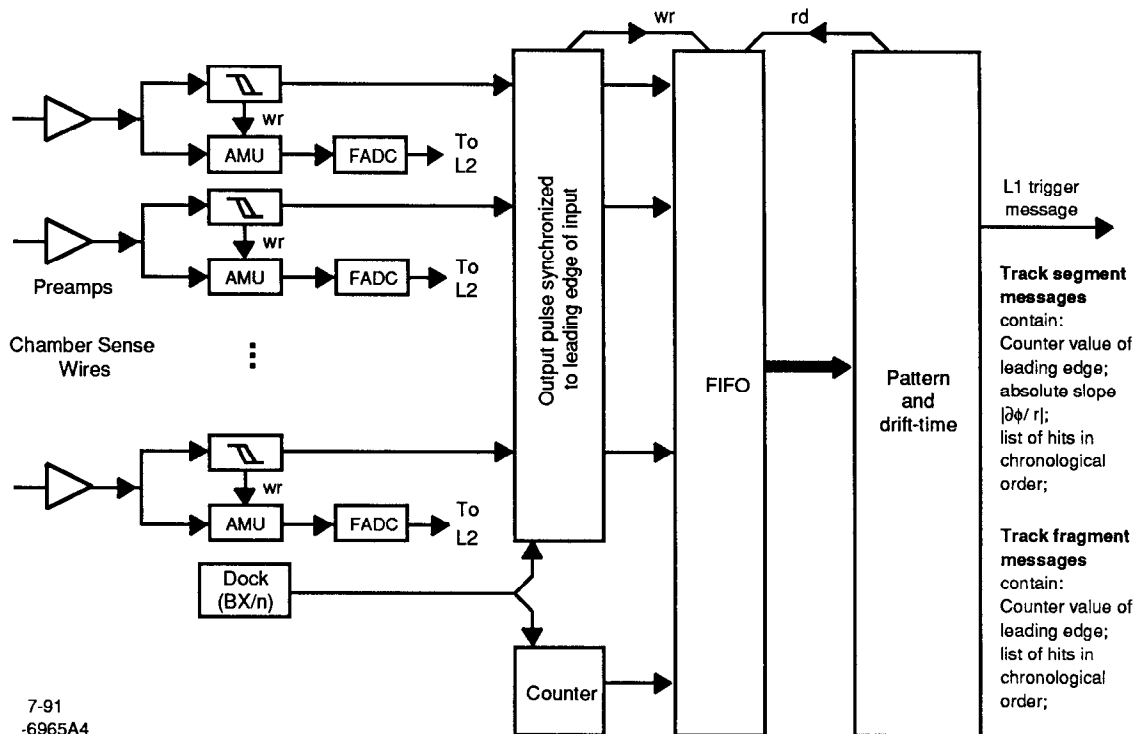


Figure 4. Drift chamber jet cells know naturally how to pipeline the trigger decision.

with that volume without resort to microprocessors. The calorimeter identifies events with total energy above a threshold of a few GeV or a pattern of energy deposition consistent with two separated minimum-ionizing particles. Events identified by the drift chamber or calorimeter at Level 1 go to Level 2. Level 1 provides orthogonality of the charged and neutral triggers.

7. Level 1 in the Drift Chamber (DCL1)

In initial discussions, a jet-cell design for the drift chamber was the leading candidate, but a small-cell design was not ruled out. The Trigger Group saw jet-cells as superior based on the reduced channel count (2000 channels of FADCs vs. 7000 channels of ADCs and 7000 channels of TDCs) and the difficulty of pipelining TDCs asynchronously. Figure 4 presents the front end of the jet-cell design.

The Tracking Group recommended the small-cell geometry, citing lower wire tension and thinner endplates, better systematics, longer lifetimes, and lower per-channel

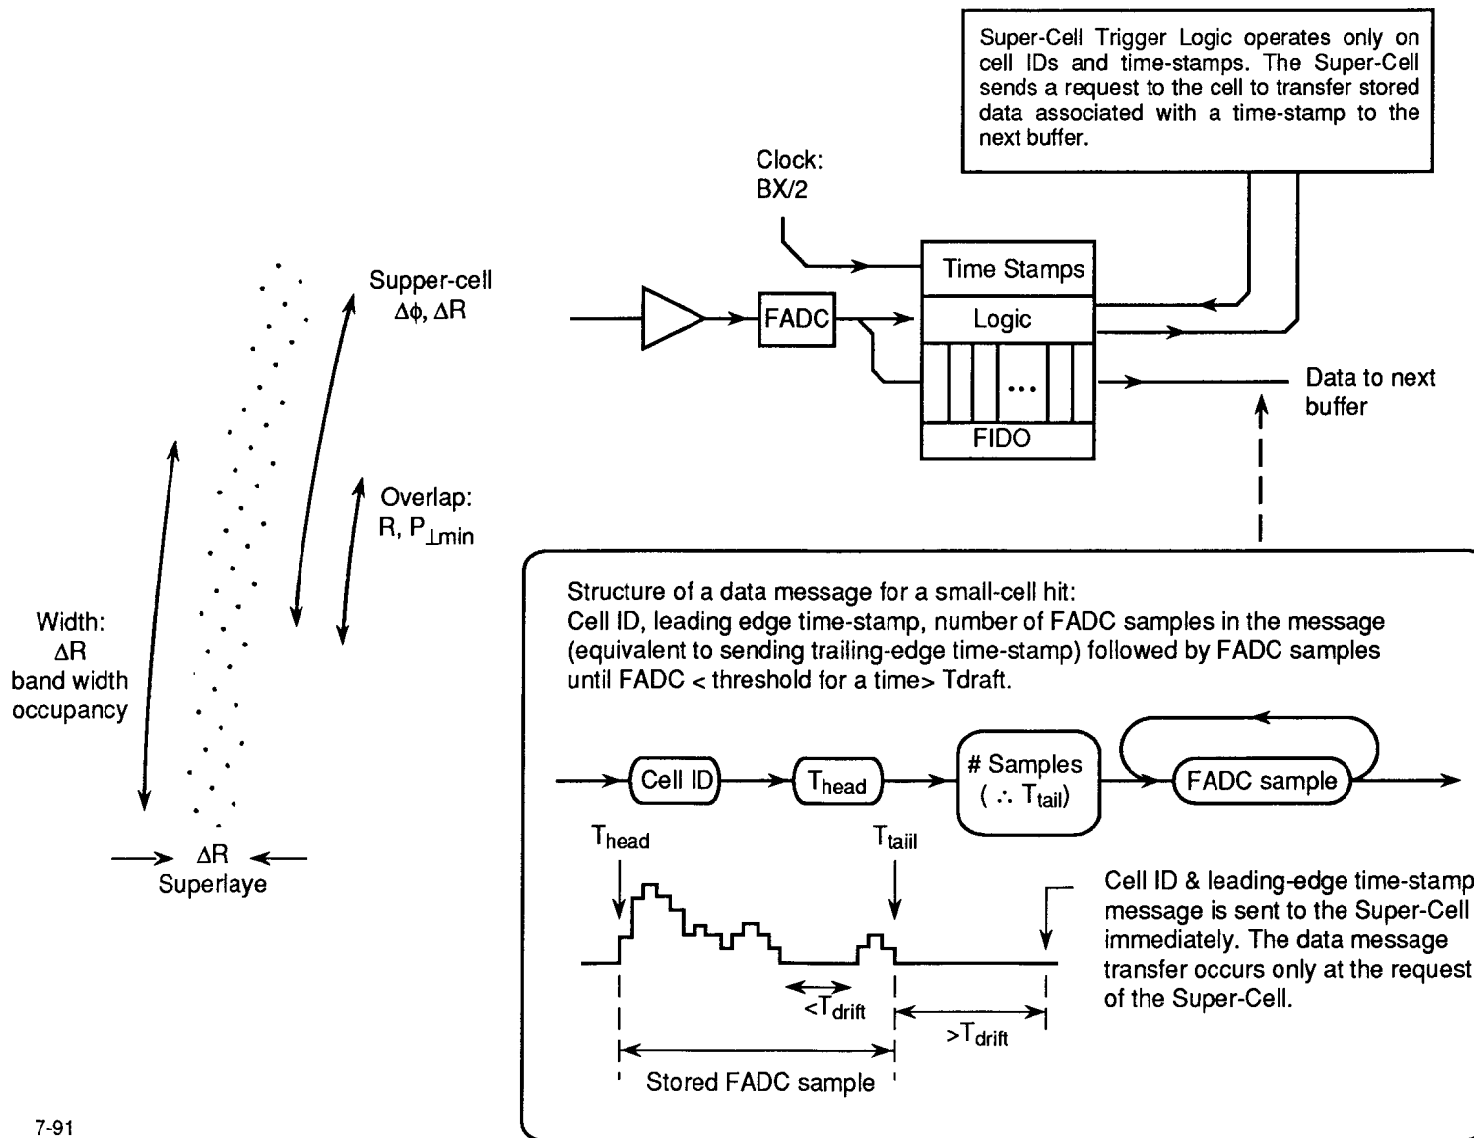




Figure 5. Small-cell logic.

Track fragment message:



Track segment message:



7-91
6965A6

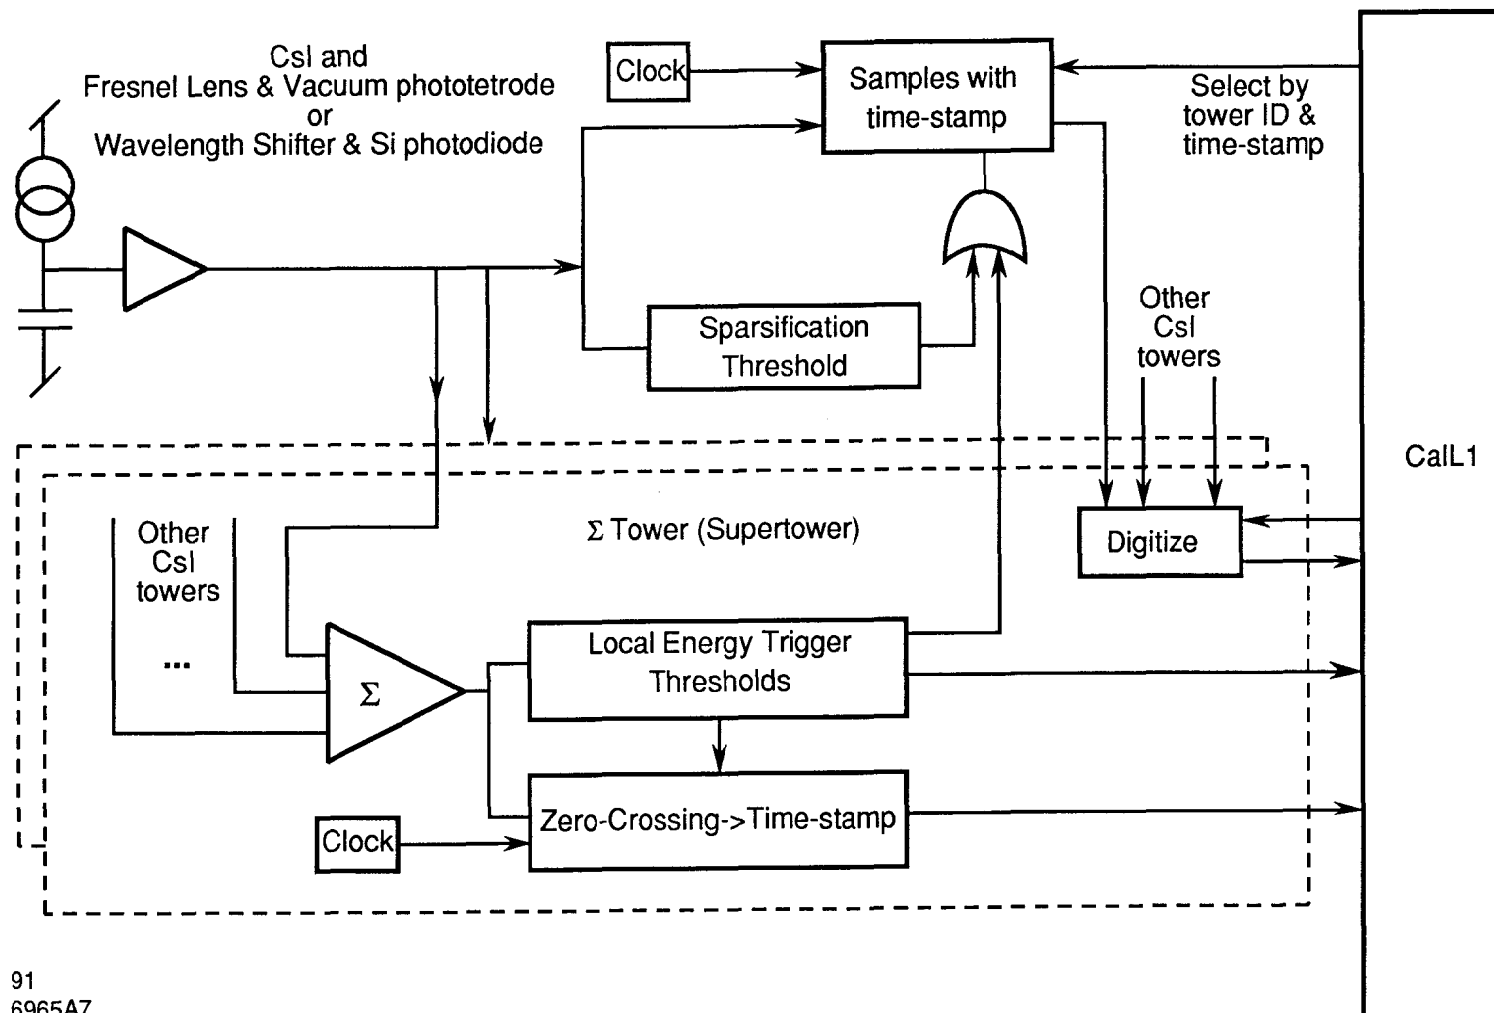
Figure 6. Syntax diagram for messages from Super-Cell.

occupancy. The groups arrived at a compromise—we shall attempt a design for a small-cell chamber using 7000 channels of FADCs and associated logic. Figure 5 presents the compromise small-cell front-end design. The small cells are grouped into Super-Cells. Figure 6 presents a syntax diagram for the output from the Super-Cell. The DCL1 node receives messages from Super-Cells, identifies candidate tracks, and builds a preliminary track list. If the track list satisfies the “1.5 track” criterion, the list and its supporting data messages move to Level 2.

8. Level 1 in the Calorimeter (CalL1)

Figure 7 shows the front end of the Calorimeter trigger and data acquisition. The lowest-level components of the Calorimeter Level 1 trigger comprise groups of about 5×5 contiguous Cesium Iodide towers, or Supertowers. The number of participating elements is reduced from about 10^4 to about 400. The signal-to-noise ratio for a minimum-ionizing particle in a Supertower will still be about 80 : 1. To guarantee efficiency at the boundaries of Supertowers, we deploy two overlapping layers of Supertowers, as in Figure 8, requiring 800 participating elements. Each tower reports to one Supertower in each layer. Both layers report to the CalL1 node.

The CalL1 node spans both layers of Supertowers and identifies events with total energy above a threshold energy; or a pattern of energy deposition consistent with two



91
6965A7

Figure 7. Block diagram of calorimeter trigger and data acquisition at Level 1. CsI towers are grouped into Supertowers.

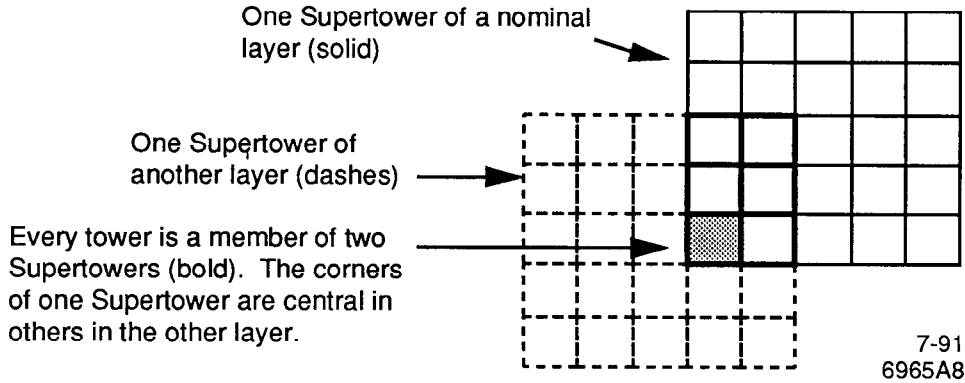


Figure 8. Two overlapping tilings of Supertowers provide efficiency and redundancy.

separated particles of minimum-ionizing energy or greater. The threshold energy will be determined by accelerator backgrounds, and we expect it to be a few GeV. CalL1 accepts about 3 kHz of cosmic rays. This is quite different from the KEK design which provides strong rejection of cosmic rays very early. The sampling envisaged here is sample-and-hold with a slow, high-precision ADC for digitization. A feature of the Supertower is that it provides a time-stamp which is useful even for minimum-ionizing particles.^[2] Let $w(t)$ be the peak-normalized impulse response of the Supertower. Then Λ is the effective shaping time (of order $1 \mu\text{s}$ or less in CsI, with PMT readout).

$$\Lambda^2 = \frac{\int_0^\infty w^2(t) dt}{\int_0^\infty \dot{w}^2(t) dt} \quad (1)$$

Let θ be the signal-to-noise ratio of a Supertower in response to some energy deposition. Then the lower limit of the time resolution σ_t of one Supertower for a measurement with energy deposition θ is

$$\sigma_t = \frac{\Lambda}{\theta} \quad (2)$$

Values of Λ , depending on doping of the CsI and the choice of vacuum phototubes or Silicon photodiodes, range from 300 ns to a few microseconds. Evaluated for minimum-ionizing response, $\theta = 80$, $\Lambda = 1 \mu\text{s}$, $\sigma_t = 12.5 \text{ ns}$, or 3 beam-crossings. A time-stamp with resolution approaching this limit is very useful. The effective

time resolution for CalL1 improves with greater energy deposition, beginning with the other Supertower required for the two minimum-ionizing criterion.

9. Level 2

The Level 2 trigger interrogates the SVD (and probably the muon system) and imposes a much smaller fiducial volume. Level 2 takes advantage of timing information in the drift chamber and the resolution of the SVD for more precise tracking. The Level 2 fiducial volume is again a cylindrical tracking volume about the interaction point, about 5 cm in length and 1 cm in radius. Level 2 resolves low-momentum looping tracking and identifies events that (1) have two charged tracks emerging from its fiducial volume, or (2) satisfy the total energy criterion. We have investigated the notion of a calorimetric veto in Level 2 to reject cosmics, and have determined that such a veto is deleterious to the tau physics.

Level 2 manages the prescaling. We choose a set of event types that we can identify at Level 2. Level 2 suppresses each event type by its prescaling rate using a random sampling method. Prescaled event types include Bhabha events and the beam-gas background to the $2\text{-}\gamma$ physics. There are two reasons to identify Bhabhas at Level 2. The rate of Bhabhas into the detector acceptance, as high as 200 Hz, is large enough to provide a useful diagnostic to accelerator operations as an online measure of luminosity. Also, should the online data processing throughput be marginal, we can prescale Bhabhas. Prescaling would require tight cuts, whereas a measure of luminosity would use loose cuts.

The dominant background to the $2\text{-}\gamma$ physics comes from beam-gas interactions. We can identify those events that satisfy DCL1 but fail Level 2 on the fiducial volume in z . A sample of this background is useful for making corrections.

As with SVD designs for SSC, the SLAC B Factory SVD will be “smart” pixels or strips—each strip or pixel will generate a channel-level event when its analog voltage exceeds its threshold. Each channel event stores the analog peak locally and causes a record to be written at the chip level, comprising a time-stamp and a strip address or a row/column address. The time-stamp may be content-addressable. Each strip

or pixel exceeding threshold becomes insensitive for a few hundred nanoseconds, but the expected channel occupancy is low.

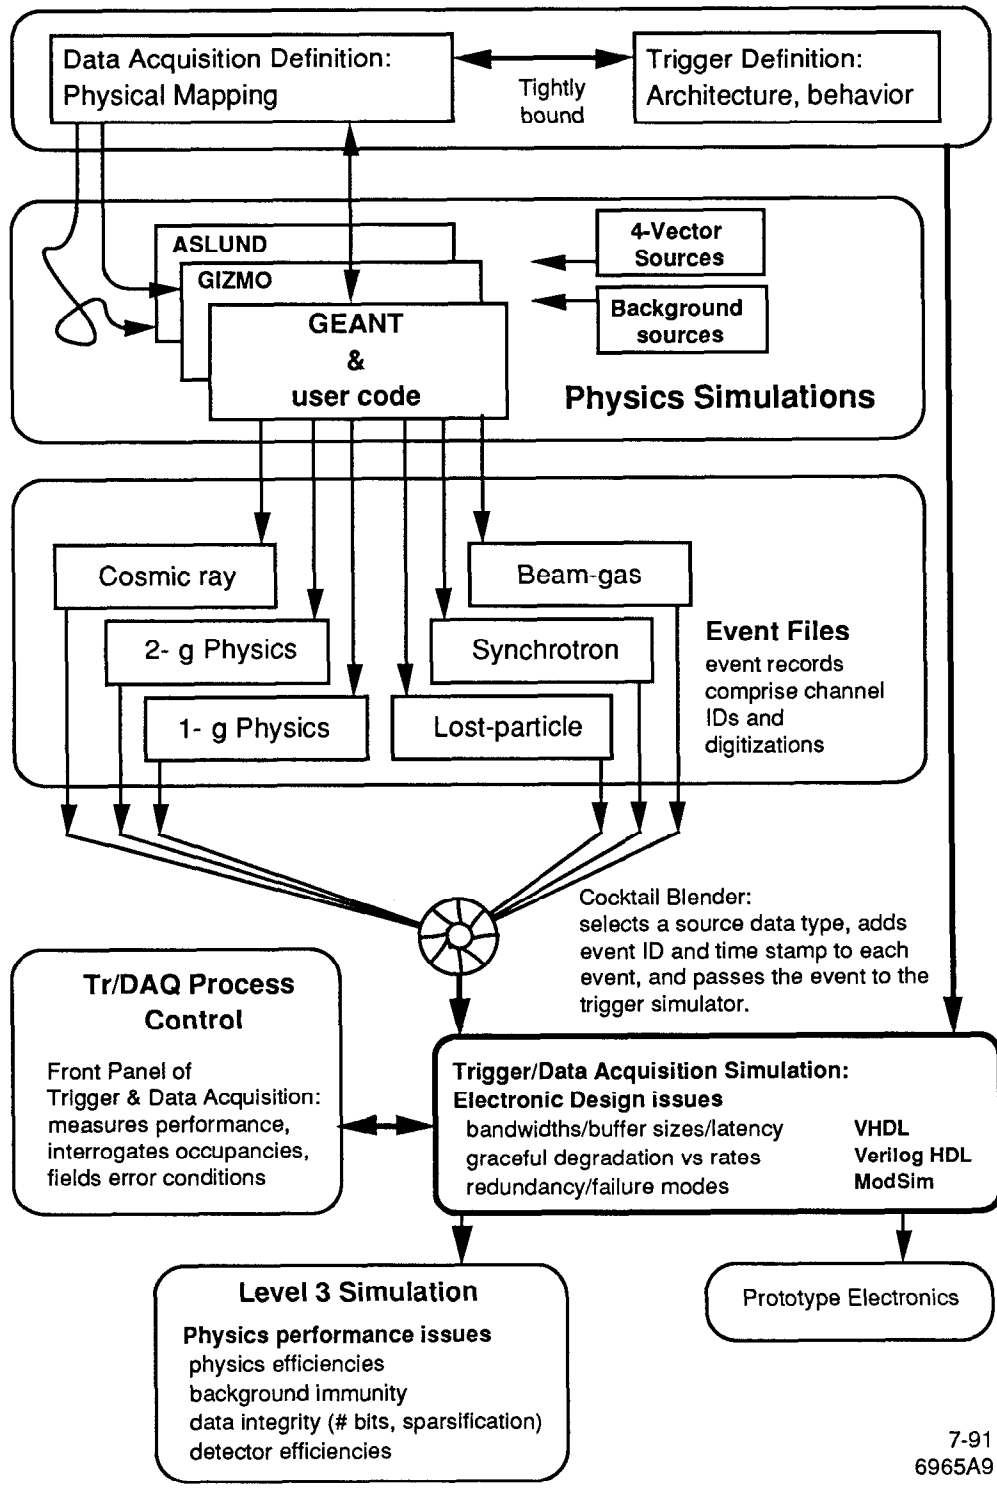
10. Backgrounds

The main backgrounds are lost particles, synchrotron radiation, beam-gas elastic scatters, and cosmics. The lost particle background arises from beam particles that interact with the residual gas atoms in the beam pipe far upstream, losing enough energy that they are lost from containment and hit a mask or magnet near the detector.^[3]

Synchrotron radiation amounts to 100 kW near the detector. For the synchrotron radiation, the detector acceptance was calculated including γ s scattering through mask tips, bend sources far upstream, and γ s backscattered from masks downstream. The rate expected is less than 10^{-2} per beam crossing into the drift chamber. This rate grows by a factor of about 2.5 if the innermost quadrupole magnets are misaligned by 1 mm. The Super-Cells of DCL1 reject this background of isolated hits in the drift chamber. CalL1 rejects this background because its power spectrum lies below its thresholds.

The lost particle background rate was calculated assuming 1 nTorr within ± 30 m of the IP, and 5 nTorr elsewhere. Sources within ± 185 m of the IP were included, using Decay Turtle and EGS. The soft γ s are rejected by CalL1. The charged particles have low p_t , and very few contribute to the “0.5” track at DCL1. Level 2 rejects them based on the fiducial volume. The calculations suggest acceptably low rates, but a full Monte Carlo is required.

We expect beam-gas elastic scatters to pass DCL1 at a few hundreds of Hertz, and Level 2 at a few tens of Hertz. We expect about 3 kHz of cosmic rays into the Calorimeter. A few hundreds of Hertz cross the fiducial volume of DCL1, and a few tens of Hertz cross the fiducial volume of Level 2. Some ten percent of cosmics produce a delta ray into the drift chamber, but very few of them pass the DCL1 fiducial volume cut. As for cosmics that graze the Calorimeter and pass CalL1 based on deposited energy, Level 2 can reject these readily if it has resort to the muon system.



7-91
6965A9

Figure 9. Further work includes more detailed definition and simulation.

11. Further Work

A complete and consistent design at a low level of detail will provide a behavioral specification for prototype electronics to be used in beam tests and cosmic ray tests. Clearly, this design requires sophisticated and rigorous simulation. See Figure 9. The inputs to the simulation are Monte Carlo events of the physics channels and the backgrounds. By increasing the simulated rates of backgrounds well above those expected, we challenge the design as we further determine the details. The goals of the simulation are to determine necessary buffer sizes, bandwidths of control and data paths, and latencies of the decision-making elements; to identify bottlenecks and failure modes; to verify graceful degradation; to compare busses (e.g., VME, VXI, Fastbus); and finally, to serve as a performance specification for production engineering. We expect to pursue this effort with available CAE software.

12. Comparison of Designs


Figure 10 compares the KEK and SLAC designs in overview. Clearly, the KEK design is less complex, less expensive, and more familiar. It takes advantage of Time-of-Flight information, which is prompt and powerful. Its robustness derives from the constraints that span detector subsystems. The SLAC design is more of a design departure. The systematics in the KEK design are more difficult to understand if the backgrounds are high. The transition from a two-track requirement to three tracks depends on the backgrounds. The SLAC design is robust based on its redundancy. The orthogonality of the charged and neutral triggers in the SLAC design affords a direct measure of the efficiencies. The SLAC design should tolerate high backgrounds better. The SLAC detector design started with no Time-of-Flight and no inner tracker, forcing the trigger design to be more complex, especially in the light of the tau and $2\text{-}\gamma$ physics measurement goals. The KEK design included these elements from the beginning, considering the trigger problem. Both designs provide flexibility.

13. Conclusions

The KEK and SLAC designs for B Factory Trigger and Data Acquisition are different because the detector designs and physics goals are different. The design of

KEK


- Step 0: TOF majority yields time-stamp
- Step 1: Impose Calorimeter
- Step 2: PDC & CDC starts LookUp Table
- Step 3: Digitizer



SLAC

Level 1: DC looks for tracks which get at least half-way to outer layer (half-track) with coarse tracking/large fiducial volume, and sends tracks to lever2.

Level 2
IF:
a) There are more than "1.5 tracks" in DC; OR
b) Calorimeter trigger; Or
c) Overlap of Cal & DC (prescale for diagnostics);
THEN
Identify and Prescale Bhabhas OR
Interrogate all other subsystems for data associated with time-stamp.
If SVD confirms smaller fiducial volume or if Calorimetric energy is sufficient, pass all information to Level 3.



Level 1: Cal looks for at least 2 Σ towers $\geq 0.8 E_{\text{imp}}$ or total energy \geq few GeV, and sends hits to level 2. Should provide loose time-stamp, even for minimums.

Figure 10. Comparison of designs in broad overview.

an asynchronous pipelined trigger presents an initial challenge of complexity, but the effort pays off by decoupling the throughput from the latency. The more powerful discrimination afforded by longer latency at Level 1 provides a lower input rate for Level 2. The KEK design is elegant, simpler, and well-matched to the KEK detector and physics goals.

14. Acknowledgements

The SLAC design for a trigger and data acquisition system is due to the working group of the ongoing workshop, led by Bill Wisniewski. Other members include: Tim Bolton, Chris Hearty, Walt Innes, Chang-Kee Jung, Andy Lankford, Gerard Oxoby, Peter Rowson, and Phil Rubin and Mike Shaevitz. The estimations of background sources and rates from the accelerator are due to the group headed by Hobey DeStaebler, including Dave Coupal, Chris Hearty, and Mike Sullivan.

References

1. "Proceedings of the Workshop on Physics and Detector Issues for a High-Luminosity Asymmetric B Factory at SLAC," SLAC-373.
2. This discussion derives from the literature on Radar range-finding. See, for instance, *Radar Handbook*, Merrill I. Skolnik, McGraw-Hill, 1970.
3. "An Asymmetric B Factory Based on PEP," SLAC-372. See Section 4.2, "Estimation of Detector Backgrounds."