SSCL-348

SSCL-348

Superconducting Super Collider Laboratory

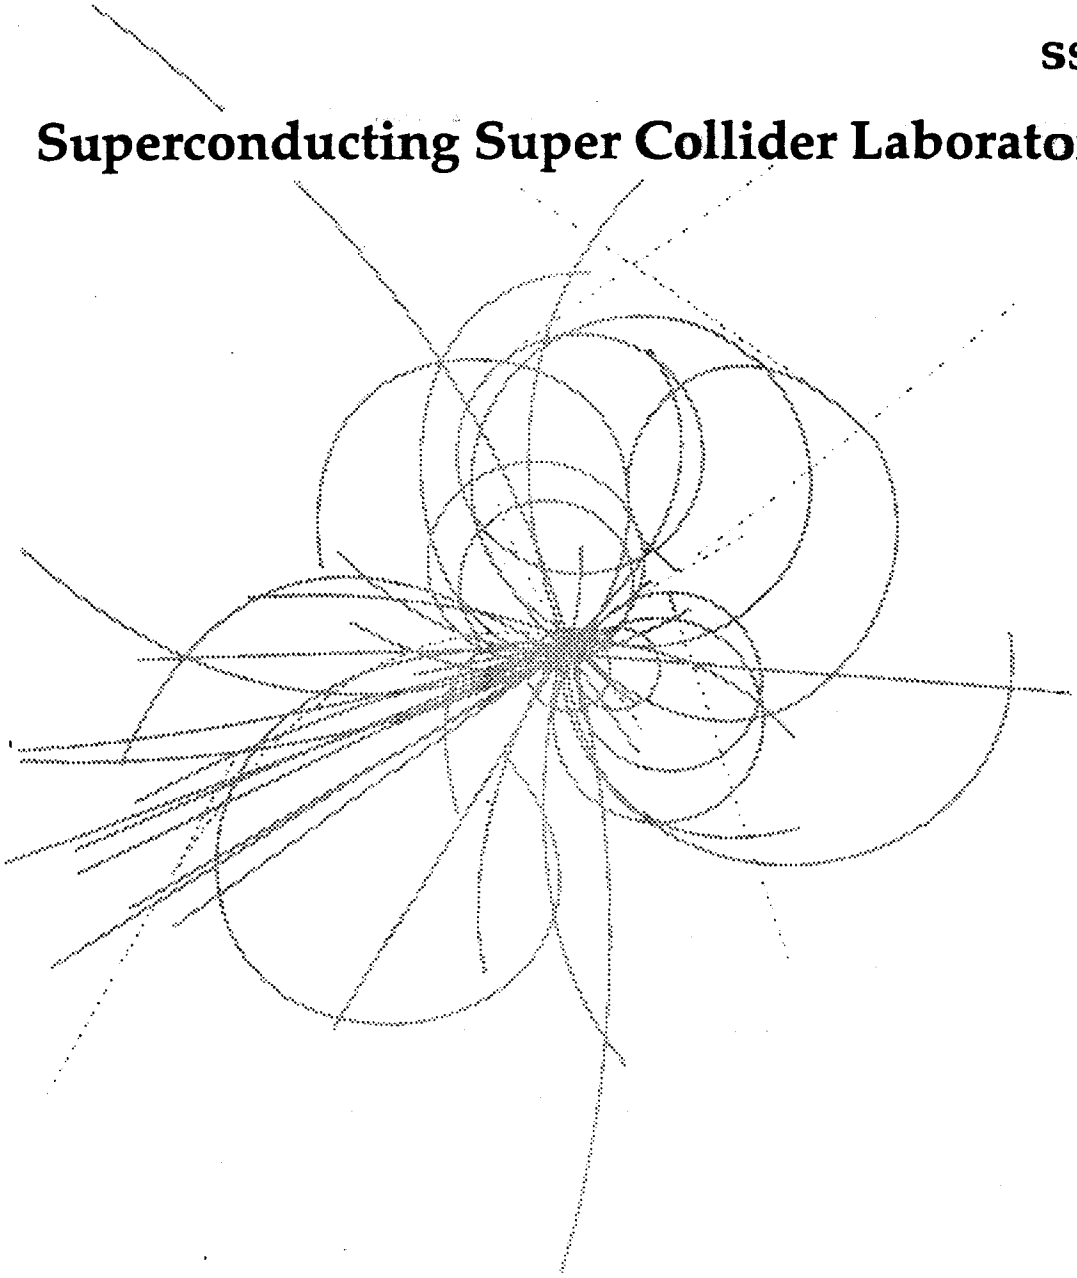

Readout of Overlapping Events

J. Dorenbosch, M. Botlo, A. Fry,
and C. Milner

January 1991

SSCL–348

READOUT OF OVERLAPPING EVENTS

Jheroen Dorenbosch, Michie Botlo, Alan Fry and Cas Milner

Superconducting Super Collider Laboratory*
2550 Beckleymeade Avenue
Dallas, TX 75237

January 1991

* Operated by the Universities Research Association, Inc., for the U.S. Department of Energy under Contract No. DE-AC02-89ER40486.

1.0 INTRODUCTION

This article describes the search for a limited set of solutions to the problem of overlapping events in detectors. First we describe how non-overlapping events are treated. Next, we describe many options for handling overlapping events and try to clarify the consequences. Finally, we try to distill coherent sets of rules.

2.0 NON-OVERLAPPING EVENTS

The level 1 trigger chooses the events to be collected by sending an L1Accept signal to the front-end chips. The signal arrives at the chips a fixed time after the detector input signals from the corresponding beam crossing. The input signals are stored internally for a fixed delay time, and the L1Accept arrives at the chip just before the corresponding data reach the end of the delay. For many detector subsystems, data related to an event are spread out over several beam crossings. Here, the front-end chip must measure all input signals in a certain time window after an L1Accept (Figure 1).

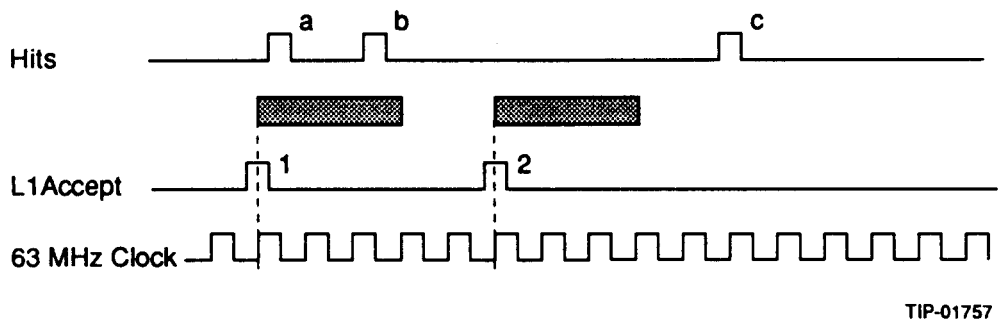

Figure 1. Two non-overlapping events in a straw tube channel. It is assumed the maximum drift time in the straws is less than 48 ns. Hits a and b belong to event 1; hit c is not associated with an event.

Unfortunately, it is unclear whether the level 1 trigger can unambiguously associate all events with the correct crossing. One can hope that the trigger time will be known to within one crossing. It is also unclear whether the level 1 trigger will know the timing uncertainty for each individual event. If the uncertainty on individual events is unknown or is the same for all triggers, one must measure a well-defined additional number of beam crossings for each event. If timing uncertainty happens for only a small fraction of the triggers, one could dynamically control the number of associated crossings. From event to event, a varying number of crossings is measured. For slow systems, where many crossings are measured, this flexibility is unnecessary.

A level 1 Identity number, L1Id, is added to the data by the front-end chips to confirm that event fragments belong to the same trigger. It is the value of an eight-bit counter that is incremented for consecutive events. There is only one counter, common to all channels of a front-end chip. The L1Ids from different front-end chips are compared somewhere downstream in the data acquisition system. Note that the level 2 trigger rejects most events. The L1Ids of the remaining events no longer increment by one. It is also not excluded that consecutive events have the same L1Id.

Another value used for checking is the BunchPhase, which is inserted in the data of each event. It is the value of an eight-bit counter, strobed at the time of the L1Accept. The counter is incremented by the 63 MHz clock. BunchPhase is also used for the time measurement in TDCs, where it is inserted with each hit. Subtraction of the L1Accept BunchPhase from the hit BunchPhase yields the number of cycle times to be added to the hit time measurement.

Several ways to control the measurement time span have been proposed (Figure 2):

a) L1Accept is an 8 ns pulse, and an internal register in the front-end chip is programmed with the (constant) number of clock cycles to be measured. The L1Id is incremented on the rising edge of L1Accept.

b) L1Accept is a level that is high for all clock cycles to be measured. The L1Id is incremented on the rising edge of L1Accept. One has dynamic control of the number of collected crossings.

c) L1Accept (8 ns) is given for all crossings to be measured while L1Reject (8 ns) is given for all crossings to be skipped. This scheme has several advantages:

- Non-consecutive beam crossings can be chosen for measurement, as might be preferred in slow calorimeter readout.
- One does not need an internal register to specify the length of the level 1 delay. However, internal registers and/or additional signals must be used to define which of the chosen crossings belong to the same event and when to increment the L1Id.

d) L1Accept is a pulse that indicates when the measurement time starts. A second pulse, L1End, marks the last crossing to be measured. The L1Id is incremented on the rising edge of L1Accept.

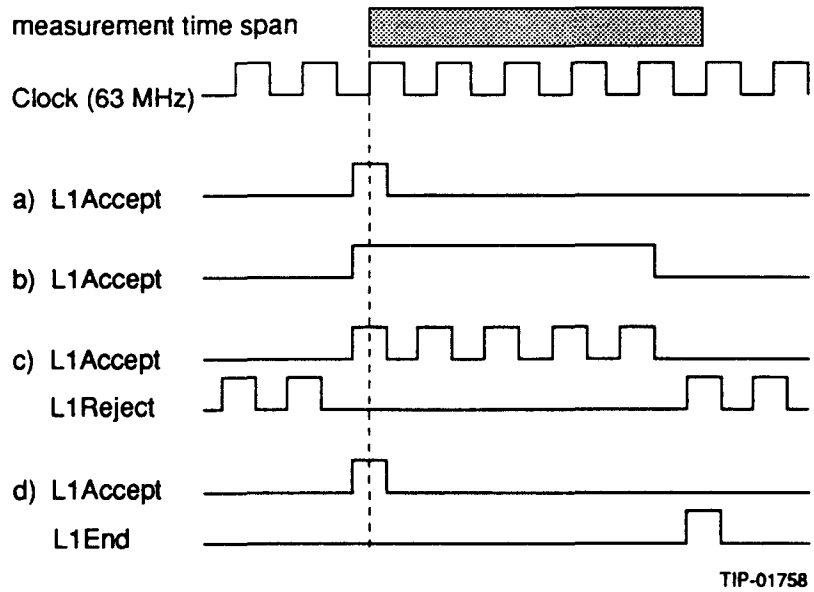

Figure 2. Different ways to control the measurement time span.

3.0 OVERLAPPING EVENTS

When the measurement span of a detector is larger than one clock cycle, it is possible to have several L1Accepts during the measurement. When this happens, data belong to more than one event (Figure 3). A rapid succession of triggers can make a long string of overlapping events (Figure 4).

Event overlaps are related to the level 1 trigger rate instead of to the data rate. They happen even when there are no data present.

The measurement time span is detector-dependent. Hence, different subsystems have different overlap probability. If an overlap happens in the slow muon subsystem, most probably there is no overlap in the silicon tracker. As overlaps are trigger-related, they happen simultaneously in all channels of a detector subsystem.

Note that overlaps are different from pile-up. Pile-up is a local phenomenon, independent of trigger rate. To deal with pile-up one can take consecutive data samples for each event. Event overlap cannot be relied on to solve the pile-up problem.

When data of several events are combined into a common data structure, these events are said to be "married."

4.0 OVERLAP PROBABILITY

The occurrence of event overlaps depends on the trigger rate and the measurement time in a roughly linear fashion. Most affected are muon chambers

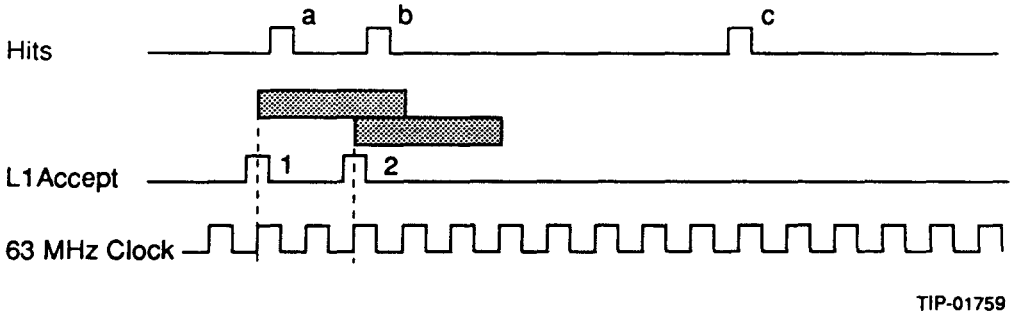

Figure 3. Two overlapping events. Hits a and b belong to event 1; hit b also belongs to event 2; hit c is not associated with any event.

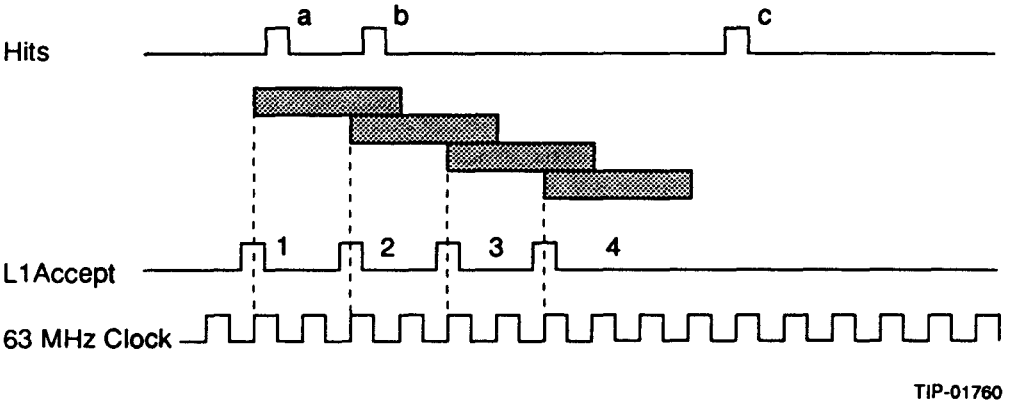

Figure 4. Sequence of overlapping events.

with maximum drift times between 500 ns and 1 μs, and slow calorimeters with measurement times of about 250 ns. At an L1Accept rate of 10^5 Hz, a detector with a response time of 1 μs has a 10% chance to get a second L1Accept during a measurement. Hence 20% of the events will overlap. For response times of 250 ns and 32 ns, the probabilities become small: 2.5% and 0.3%, respectively. The probability of obtaining a string of three overlapping events is about 1% for a measurement time of 1 μs.

One can also estimate the fraction of events accepted by level 2 which will show overlaps. For a subsystem with a measurement time of 1 μs, 20% of the accepted events will overlap with an earlier or later rejected event. Only 0.2% of the accepted events overlap with another accepted event, assuming a level 2 accept rate of 10^3 Hz. At 10^4 Hz this becomes 2%.

5.0 OVERLAP CONTROL

For a fast detector subsystem with a response time of a few beam crossings, one can avoid overlaps completely by imposing a minimal spacing between L1Accepts. For example, the cost of a dead time of 4 clock cycles after each L1Accept is small: less than 1% at 10^5 Hz.

If it is certain that the level 1 trigger rate will always be below 100 kHz, one might impose a dead time of 1 μs per event. The associated dead time would be less than 10%, which might be acceptable if the readout would be significantly simplified. If this dead time is unacceptable, or the trigger rates are higher, one can limit the overlaps in slow detectors by implementing more rules in the level 1 gating logic. Such rules will limit the maximum number of simultaneous events and the longest possible sequence of overlapping events, making the design of front-end chips easier. One possible set of rules is:

- No more than 1 L1Accept per 80 ns.
- No more than 2 L1Accepts per 260 ns.
- No more than 3 L1Accepts per 1000 ns.
- No more than 4 L1Accepts per 4000 ns.
- If any event in a set of overlapping events is accepted by level 2, other events in that set are rejected.

These rules are implemented by the gating logic and do not imply any additional complication for the front-end chip design. The above time values correspond to assumed resolution times for straw tubes (80 ns), liquid calorimeters (250 ns), and muon chambers (1 μs). The values would be set to the resolution times of the actual devices. The 4 μs limit is there to avoid long strings of muon events. For a detector with a response time of 250 ns, the above rules give a maximum of 2 simultaneous events and a maximum string of 3 events; for 1 μs, the maximums are 3 and 4 events, respectively (Figure 5). The total cost in dead time is less than 1% at a level 1 rate of 10^5 Hz, and 11% at 5×105 Hz.

The last rule simplifies the way events accepted by level 2 are handled. There will be no overlapping accepted events. If one wants to correct for the small 0.1-to-1% loss, some extra bookkeeping will be needed in the gating logic. Remember that the L1Accepts and Rejects are common to all subsystems. In the context of this rule, an overlapping event is one that has an overlap in at least one subsystem.

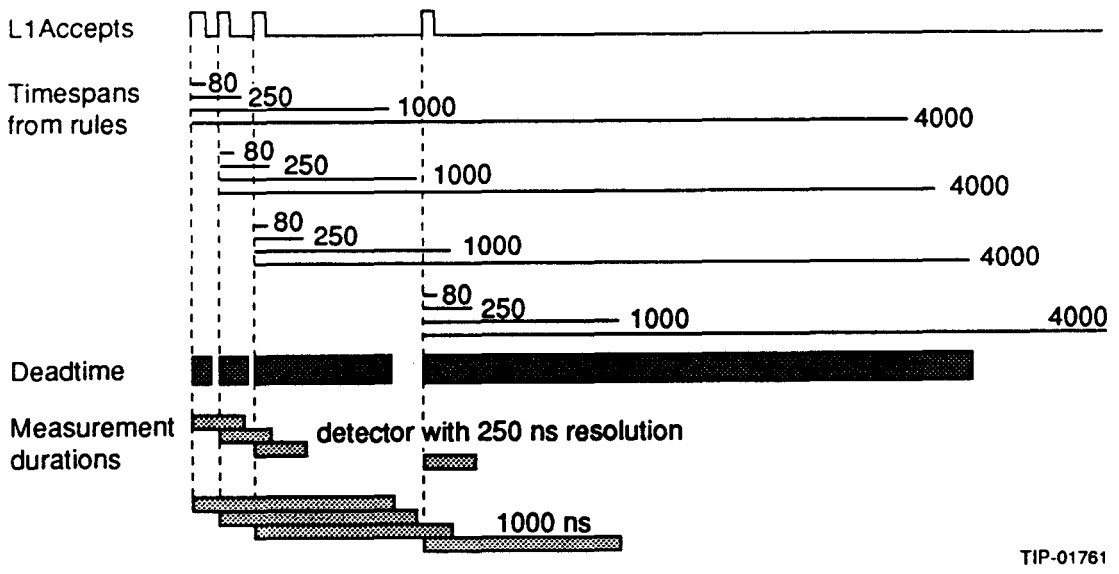

Figure 5. Limiting overlapping events. Each L1Accept defines periods when the number of triggers is limited to 1, 2, 3 and 4 triggers. This limits the maximum number of simultaneously overlapping events and the longest possible string of overlaps.

6.0 CONTROL SIGNALS FOR OVERLAPPING EVENTS

The various L1Accept schemes mentioned above (Figure 2) are changed in different ways to accommodate overlaps. One must first decide whether to decouple overlaps in the different subsystems. Two options are available:

- If an overlap occurs in a slow detector subsystem, an artificial overlap is created in the faster ones.
- Overlaps are created only in those systems where they are unavoidable.

The first approach creates needless complications. Events that could easily be separated become mixed up, and a large amount of extra data may have to be measured. Additional fast control signals might be required. An advantage is that the control signals can be equal for all subsystems. Also, the level 1 Identities stay rigorously equal between subsystems. The second option is preferable.

As seen in Figure 6, different L1Accept schemes have a number of options in dealing with overlaps:

a) L1Accept pulse: The full burden of bookkeeping is put on the front-end chip, which must establish that an overlap occurred, then take

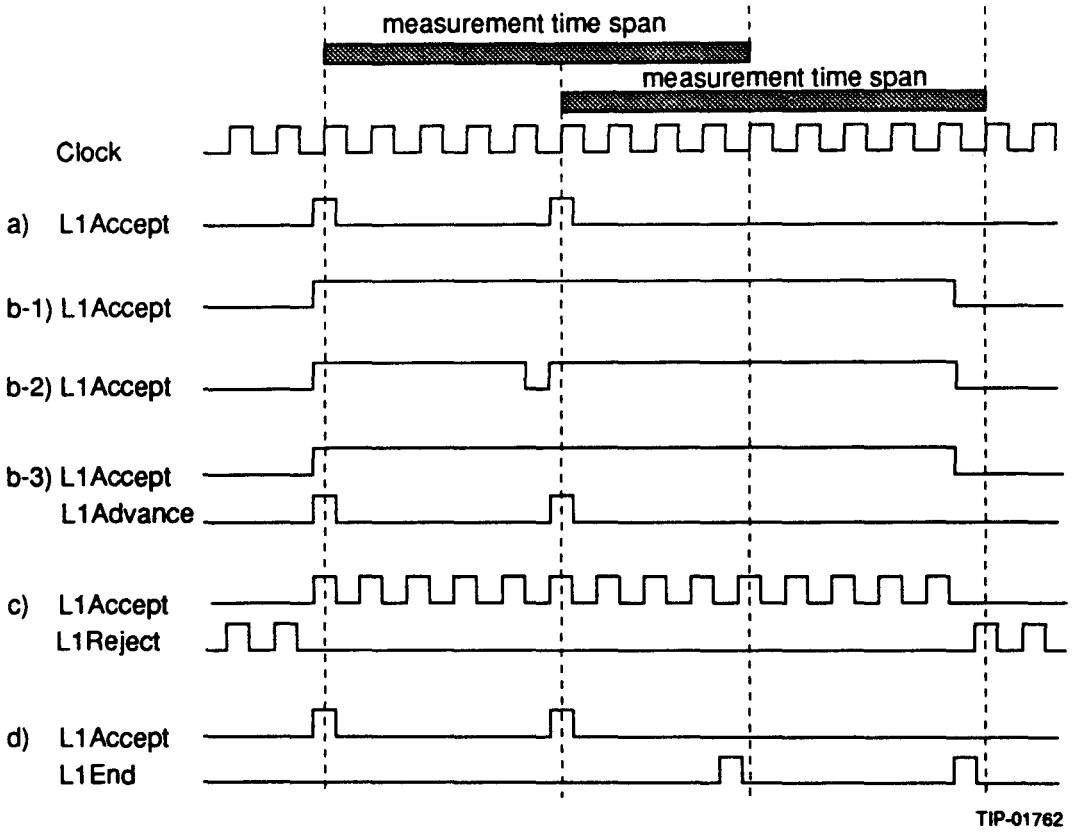

Figure 6. Different options for control signals in case of overlapping events.

the proper action. Events can be separated or married by the front-end chips.

b) L1Accept level:

- $b-1$) The level is held high for the full duration of the overlap. Events must be married. The L1Ids are different between fast and slow subsystems. Extra bookkeeping is required.
- $b-2$) For each overlapping trigger the L1Accept level goes low for a half-cycle. In this case the front-end chip knows when the triggers occurred and has enough data to separate the events. L1Ids can be incremented on each leading edge. Events can be separated or married by the front-end chips.
- $b-3$) Like $b-1$, except another fast signal, L1Advance, will increment the L1Id. In case of overlaps, the dynamic control of the number of measured crossings is lost. However, this is less relevant for the subsystems where overlaps occur, as they already capture multiple crossings for each trigger.

c) L1Accept/L2Reject pulses: It is still possible to measure non-consecutive beam crossings, but separation becomes difficult. An additional signal, like L1Advance, must be used to define which of the selected crossings belong to the same event and when to increment the L1Id.

d) L1Accept/L1End pulses: These signals allow dynamic measurement control and separation without the use of extra signals.

For schemes *a*, $b-2$, $b-3$, and *d*, the L1Id can be incremented for overlapping events. If married events are output by the front-end chip, certain questions must be addressed. Which L1Id is attached to the data block? Are all L1Ids included in the data?

For scheme $b-1$ all data must be married into a single block, and L1Id is incremented only once. L1Ids will differ between subsystems.

7.0 SEPARATING OVERLAPS

Overlapping events can be addressed by separation or by duplication. Separation transforms overlapping events into several normal single events.

Duplication provides multiple copies of the complete overlapping data structure. The data block contains information not related to the event, and the L1Id and BunchPhase may be wrong. Additional information is needed to determine the relative timing of L1Accept and data.

The authors believe that event data must be fully separated before they are written to permanent storage. In fact, they should be separated before they are passed to the physics algorithms running in the level 3 farm. Overlaps should be completely hidden from the user.

To separate overlapping events the following information is needed:

- The data for the total measurement time, tagged by BunchPhase.
- The L1Accept time for overlapping events (BunchPhase).
- The level 2 decision for each event (L2Accept/Reject) if separation is done after the level 2 trigger.

In the delay line of the front-end chips there is only one copy of each data sample. To separate overlapping events, data samples must be duplicated. This can be done in several places:

a) Inside the front-end chip:

1) Inside the front-end chip at L1Accept time. Multiple copies are stored separately when they belong to more than one

event, or complete events are formed inside one large buffer. This is difficult if data are stored in analog fashion.

2) After L1Accept, before L2Accept/Reject. This can be done if accepted data are digitized asynchronously after L1Accept.

3) Inside the front-end chip at L2Accept/Reject time. Here only accepted events need to be separated.

If the front-end chips must provide data to the level 2 trigger as separated events, one must use option 1 or 2.

b) After the front-end chip. Data for overlapping events are merged by the front-end chip into single-event data blocks ("married" events):

1) Separation is done before event building, in the data collection network or in the event builder hardware. In this case data blocks must contain all information needed for separation (Figure 7). As data blocks are assembled inside the front-end chip, it must have a complete set of input signals. This rules out the L1Accept schemes where the front-end chips cannot see the timing of individual L1Accepts.

2) Separation after complete events are built. Information needed for separation is supplied by the gating logic and is included in the event. Note that here it is difficult to preprocess overlapping events in the data collection network.

c) Another option is not to separate overlapping events. If one of the married events is accepted by level 2, the complete data block is output. This implies that all data processing software must be able to recognize and handle the married events and that the information needed to separate overlaps must always remain in the data. This is not a desirable option.

Remember that overlaps are specific for certain subsystems. For subsystems without overlaps, no special action is needed.

8.0 IMPLICATIONS FOR LEVEL 2 PROCESSING

If married event data are sent to the level 2 processors, they must be treated in a special way. For example, overlapping muon drift chamber data must be processed twice, using different event times. This is because hits on a track will line up only when the correct event time is used. Out-of-time tracks, from overlaps or pile-up, will have a bad Chi squared. The event times (level 1 trigger times) must be passed by the gating logic to the level 2 processors. This is easy

L1Id [0]	*BunchPhase*	*L2Id [0]*	*Accept/Reject*
L1Id [1]	*BunchPhase*	*L2Id [1]*	*Accept/Reject*
L1Id [2]	*BunchPhase*	*L2Id [2]*	*Accept/Reject*

Data item	*(BunchPhase, channel number, etc.)*.....
Data item	*(BunchPhase)*.....
Data item	*(BunchPhase)*.....
Data item	*(BunchPhase)*.....
Data item	*(BunchPhase)*.....

Figure 7. Structure of self-contained data block.

if there are only a few processors, but if processing is done in a distributed way this could be cumbersome. Married calorimeter data must be processed twice as well, to extract the correct energy integral from the sampled output pulses.

9.0 LEVEL 2 ACCEPTS AND REJECTS

If events are married, there are different ways to communicate the level 2 decisions to the front-end chips:

- Level 2 decisions for overlapping events are combined. A subsystem with potential overlaps has an interface between the gating logic and the front-end chips that combines the L2Accepts and L2Rejects. The combinatorial rule is simple: if any event inside the overlap is accepted, an L2Accept is given; otherwise, the married event is rejected. Level 2 decisions must be combined for those schemes where the front-end chip does not know that overlapping triggers occurred. Combined level 2 decisions can also be used in most cases where events are married. Here the number of L1Accepts is no longer the same as L2Accepts + L2Rejects.
- The front-end chip implements the L2Accept/Reject rule for the married events. It receives an L2Accept or L2Reject for each event. If events are separated by the front-end chip, this method must be used.

10.0 COHERENT SETS OF OPTIONS

Given all the above, one can try to construct coherent methods to address the overlap problem. In order of preference, the following options are available:

- Completely avoid overlaps for all subsystems by creating enough deadtime after each level 1 trigger.
- Avoid overlaps for all subsystems except the muon subsystem.
- Avoid overlaps in fast subsystems only.

11.0 COMPLETELY AVOID OVERLAPS

Overlaps can be completely avoided by imposing after each L1Accept a dead time equal to the response time of the slowest subsystem. Readout efficiency curves are shown in Figure 8 for response times of 1 μs and 500 ns.[1]

In Figure 9 we show efficiency curves for a response time of 1 μs, with and without a level 2 trigger system that has a dead time of about 10% at 10^5 Hz. At 10^5 Hz, the partial dead times combine to about 15%. The parameters used for level 2 are a FIFO depth of 5 events and an average processing time of 9 μs.[2]

Depending on the requirements for the data acquisition system, avoiding overlaps altogether might be an option, especially if the maximum drift time in the muon chambers is limited.

12.0 AVOID OVERLAPS FOR ALL SUBSYSTEMS EXCEPT THE MUON SUBSYSTEM

Here one creates after each level 1 trigger dead time equal to the response time of the slowest subsystem but one. Assuming a response time of 250 ns and a muon response time of 1 μs, the following rules are proposed:

- No more than 1 L1Accept per 260 ns.
- No more overlaps, except in muons.
- No more than 2 L1Accepts per 2000 ns.
- No more than two overlaps in muons; longest string is 2 events long.
- If any event in a set of overlapping events is accepted by level 2, all other events in that set are rejected.
- No overlapping accepted events.

The efficiency curves for these rules are shown in Figure 10.

Below are three options for the L1Accepts, L2Accepts and L2Rejects. They require different amounts of complication in the front-end chips, but offer increasing flexibility to run the system. Option III is preferred.

I) L1Accept is a level that stays high for the duration of the overlap. The L1Id is incremented on the rising edge. L2Accepts and L2Rejects for overlapping events are combined before they are sent to the front-end chips.

[1] Efficiencies are obtained using the MODSIM discrete event simulation language.
[2] Level 2 processing times were generated using an exponential distribution.

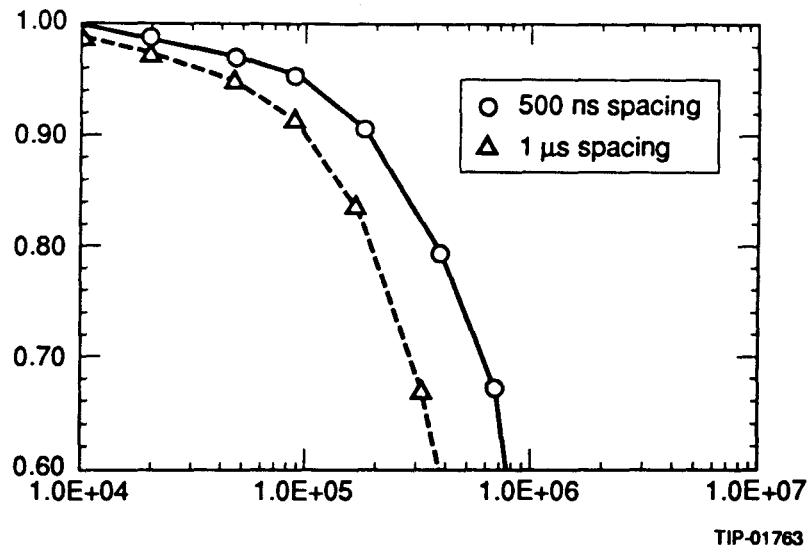

Figure 8. Efficiency versus level 1 trigger rate, for trigger spacings of 500 ns and 1 μs. Efficiencies are shown as a function of the actual first level trigger rate. The actual trigger rate is efficiency $\times$ (true level 1 frequency).

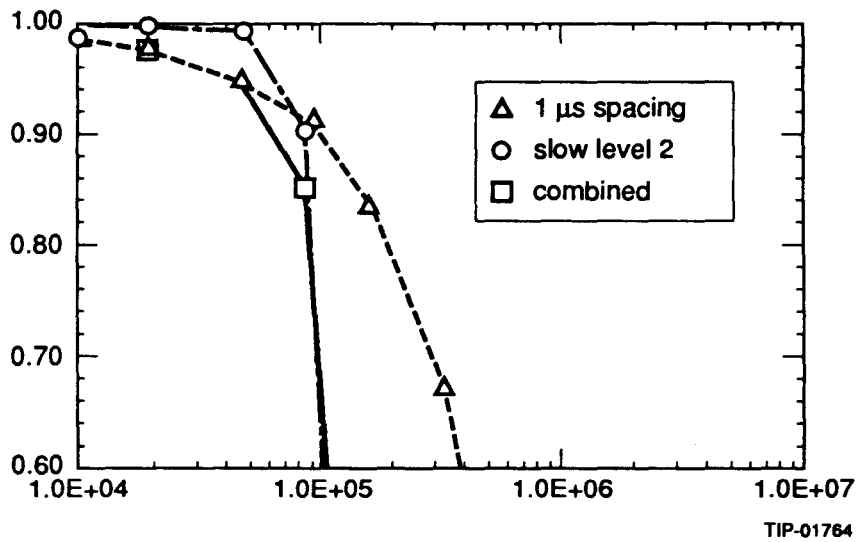

Figure 9. Efficiency curves for a trigger spacing of 1 μs, a level 2 trigger system that has about the same dead time at 10^5 Hz and the combined efficiency.

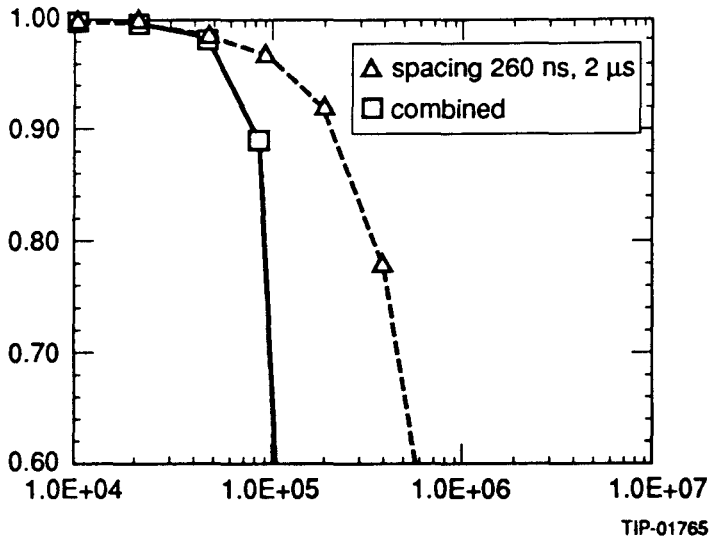

Figure 10. Efficiency curve for a set of rules that spaces triggers by 260 ns and does not allow more than 2 triggers in 2 μs. The efficiency is also shown for the combination of these rules with a level 2 trigger system that has a dead time of 10% at 10^5 Hz.

- Dynamic control of measured samples is possible in systems without overlaps.
- Overlapping events must be married.
- L1Ids are different in different subsystems.
- No extra fast control lines are needed between the gating and the front-end chips.
- Level 2 control signals are different in slow subsystems.

II) L1Accept is a pulse. The L1Id is incremented on the rising edge of L1Accept. An internal front-end chip register sets the number of crossings to be measured for each event. All L2Accepts and L2Rejects are sent to the front-end chips. They are not combined.

- Dynamic control of measured samples is not possible.
- For overlapping events, the front-end chips know when the triggers occurred. Internally it must do some bookkeeping to deal with overlaps.
- No extra fast control lines are needed between the gating and the front-end chips.
- Common level 2 control for all subsystems.

III) L1Accept is a level. For each overlapping trigger the L1Accept level goes low for a half-cycle. The L1Id is incremented on each rising edge of L1Accept. All L2Accepts and L2Rejects are sent to the front-end chips.

- Dynamic control of measured samples is possible in systems without overlaps.
- For overlapping events, the front-end chips know when the triggers occurred and have enough data to separate the events. This approach gives the greatest flexibility.
- No extra fast control lines are needed between the gating and the front-end chips.
- Common level 2 control for all subsystems.

13.0 AVOID OVERLAPS IN FAST SUBSYSTEMS ONLY

The data-taking efficiency can be further improved by allowing overlaps in a slow calorimeter as well and by changing the L1Accept spacing rules to:

- No more than 1 L1Accept per 80 ns.
- No more than 2 L1Accepts per 260 ns.
- No more than 3 L1Accepts per 1000 ns.
- No more than 4 L1Accepts per 4000 ns.

The result is the efficiencies shown in Figure 11. It appears that the level 1 spacing rules are somewhat over-ambitious. The same set of options is available for the L1Accept signal as described above.

Figure 12 compares the efficiencies for the three sets of rules for the level 1 trigger spacing. One needs a very efficient level 2 trigger and readout system to profit from the enhanced performance of the more complicated spacing rules. This is shown in Figure 13, where the three sets of spacing rules are combined with the same level 2 trigger.

14.0 CONCLUSIONS

When an inefficiency of 10% to 15% at 10^5 Hz is allowable, one should avoid overlaps completely by imposing a sufficiently long deadtime after each level 1 trigger. This requires that the maximum drift time in the muon system be below 1 μs.

If a higher efficiency is required, overlapping events are obtained. Complex rules and control signals are needed. Apparently it is sufficient to restrict the occurrence of overlaps to the muon system only. Good efficiencies can be obtained

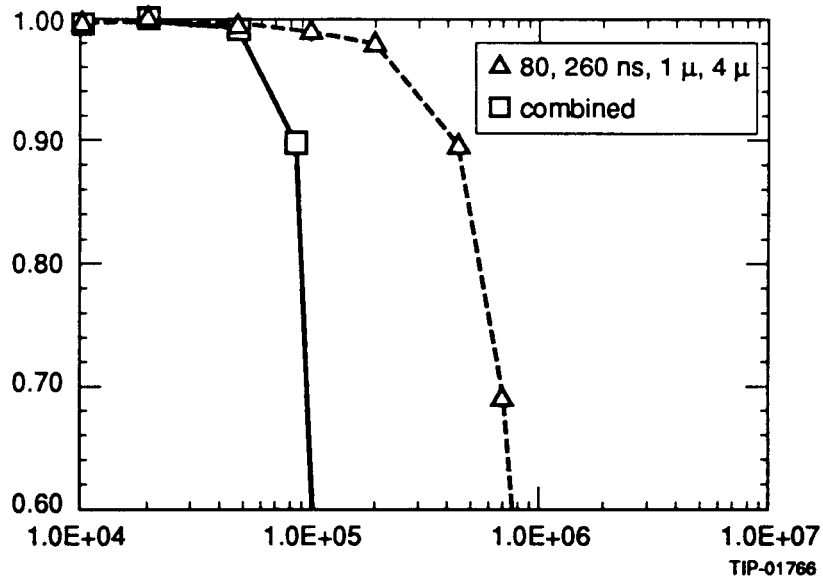

Figure 11. Efficiency curve for a set of rules that does not allow more than 1 trigger in 80 ns, 2 triggers in 260 ns, 3 triggers in 1 μs and 4 triggers in 4 μs. The efficiency is also shown for the combination of these rules with a level 2 trigger system that has a dead time of 10% at 10^5 Hz.

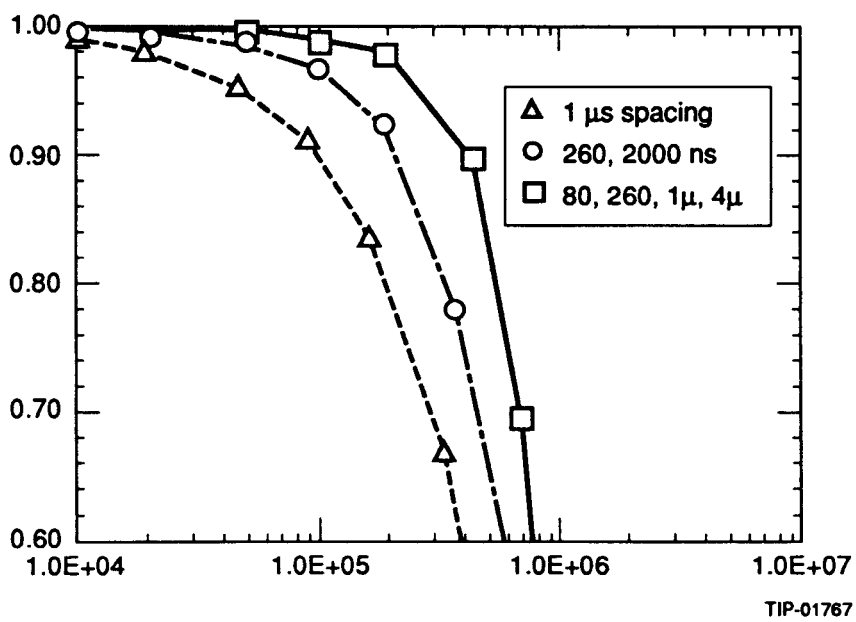

Figure 12. Efficiency curves for three sets of level 1 spacing rules.

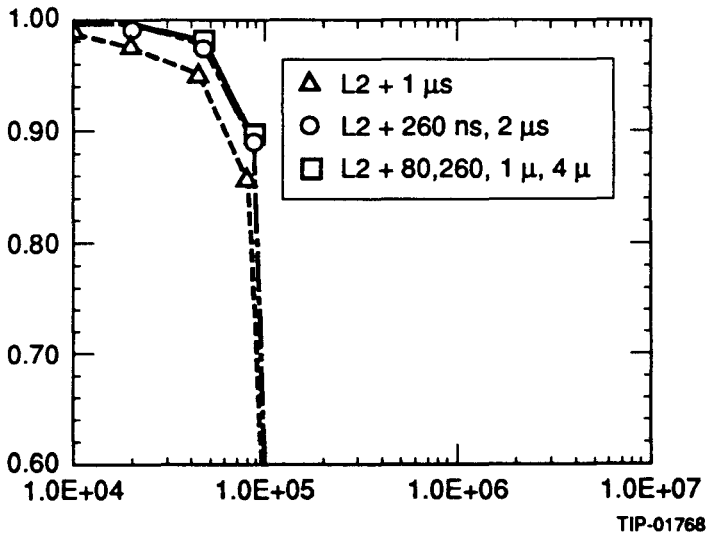

Figure 13. Efficiency curves for three sets of level 1 spacing rules combined with the same level two trigger.

while limiting the number of events in an overlap to no more than two. There is no need to have overlaps in any of the other systems, even in a liquid ionization calorimeter with a measurement time span of 250 ns.

One must choose how to communicate the trigger signals to the front-end chips. In the scheme with the best tradeoff between flexibility and complexity, L1Accept is a level. L1Accept is high for the full overlapping measurement time spans, and goes low for a half-cycle at every new overlapping trigger. This way dynamic control of the measurement time is possible, while the L1Ids will advance in unison in the whole detector. A further simplification can be achieved by imposing a rule on acceptance by level 2 of events in an overlap: if one event in an overlapping set is accepted by level 2, all the other events in the same set are rejected.

In this scheme it is possible to separate the overlapping event data in the front-end chips, or somewhere downstream in the data acquisition system. This can be decided later, when more is known about the complexity of the muon front-end chips.

However, to profit from a very efficient and complex front-end design, a very careful optimization of the total data acquisition system is mandatory.