CDF/DOC/TRIGGER
Version 4.0
May 27, 1997

Specifications and Design of the Level 1 Calorimeter Trigger PreFRED Modules for the Run II Upgrade.

Alexis Amadon, Peter Wilson
University of Chicago

1 Introduction

L1 (Level 1) Calorimeter PreFRED consists of 2 sub-systems or modules whose functions are[1] :

1. to sum up the energy deposited in all 24 calorimeter wedges and apply the ΣE_t and $\not{E}_t$ trigger thresholds (SUMET module)
2. to summarize the tower over threshold triggers (TOWTRG module)

In this note, we specify and present the design of these 2 independent modules.

2 SUMET

2.1 Specifications

The SUMET Module receives 2 10-bit transverse energy sum words ($E_t(\phi)$) from each of the 12 CRATESUM boards. Each word is associated with a ϕ-wedge and has a 0-512 GeV range with a 0.5 GeV least count. A total of 240 differential signals are thus to be brought in with 12 cables.

The output to FRED consists of 4 bits signaling the outcome of 4 programmable thresholds applied to ΣE_t or $\not{E}_t$ (typically 2 thresholds for $\not{E}_t$ and 2 thresholds for ΣE_t, reconfigurable to 3 or 4 thresholds for $\not{E}_t$ and 1 or 0 threshold for ΣE_t). The trigger bits are to be sent out to the J3 backplane of the crate, so the FRED module can quickly access them. The thresholds will have to span a certain range and granularity. A minimum number of bits are thus required for ΣE_t and $\not{E}_t$ before the thresholds are applied. The dynamic range proposed here for $\not{E}_t$ and ΣE_t are 0-256 GeV and 0-2048 GeV, respectively, with a 1 GeV least count for both.

The output to FRED should be sent in phase with that of the other PreFRED Modules. The latency for the data throughput should not exceed 3 CDF clock cycles, excluding the output alignment waiting time.

ΣE_t, $\Sigma E_x = \Sigma E_t \cos\phi$ and $\Sigma E_y = \Sigma E_t \sin\phi$ are to be sent to L2 in case of a L1 Accept (L1A), and to DAQ along with a bunch counter value, the 4 trigger bits and $\not{E}_T^2$ in case of a subsequent L2 Accept (L2A). So these words will first have to feed a pipeline until a L1A or L1R (Reject) signal is received from TSI for the event of interest. On a L1A, as for the other modules, TSI transfers a common 2-bit buffer address which tells which of 4 buffers the data should be directed to on the board. On a L1R, the data is read out of the pipeline without being transferred anywhere. Subsequently, if L2 rejects the event, no signal is sent, and the DAQ buffer can be overwritten on a subsequent L1A. If L2 accepts the event, the contents of that buffer are read out over VME using block transfers by the DAQ.

Should one or more $E_t(\phi)$ input words be saturated, the SUMET output words should also be saturated.

2.2 Design

2.2.1 A common data path for the three summing processes

Initially, it had been guessed[1] that SUMET be split into 2 physical modules as the load of operations involved (ΣE_t, $\Sigma E_x = \Sigma E_t \cos\phi$, $\Sigma E_y = \Sigma E_t \sin\phi$, $E_T^2 = (\Sigma E_x)^2 + (\Sigma E_y)^2 \ldots$) seemed to be too large to be carried out on a single board. However a rethinking of the problem led to the following picture. Instead of processing the ΣE_t, ΣE_x, and ΣE_y computations in parallel, which would require the duplication of a lot of adders, they could be processed sequentially on the same data path, under the control of a sequencer whose frequency would be 3 times that of the system clock. Then the functional steps for ΣE_t, ΣE_x, or ΣE_y would simply be pipelined on the same board.

2.2.2 The re-use of the L1 DIRAC Receiver Aux-Cards

The problem of having to feed in 240 differential signals into a single board then arised. To resolve this issue and re-use the input connecting system developed for the DIRAC board[2], the CRATESUM data will come in two stages within a CDF clock cycle: since each CRATESUM has to send E_t for two adjacent detector wedges, it seems natural to have it done in two sequential steps of 66 ns, on the same input data lines. A total of 12 even wedge E_t's are thus sent in the first phase, whereas the 12 odd wedge E_t's are sent in the second phase (or vice-versa). This reduces the number of incoming wires by a factor of two: it enables to use 12 20-pin connectors instead of 12 high-density 40-pin ones, therefore allowing the use of the Aux-cards developed for DIRAC to transform the differential input data signals into individual bits directly usable by the SUMET logic.

2.2.3 Speed performance requirements

In this scheme, a logic step has then to be repeated 6 times per event (ΣE_{x1}, ΣE_{y1}, ΣE_{t1}, ΣE_{x2}, ΣE_{y2}, ΣE_{t2}, where the 1 and 2 indices refer to the sets of even and odd wedges).

Thus a 22 ns clock cycle operation is required. This 45.45 MHz frequency is to be provided by the sequencer introduced earlier. That sequencer also has to yield 132-ns clocks shifted by 22 ns with respect to one another, that allow to distinguish which of the 6 CDF-clock cycle phases is being processed.

2.2.4 FRED alignment pipeline and L1 FIFO

As for the output to FRED, in order for the data to be aligned with that of the other PreFRED modules, a variable length pipeline is foreseen in which the trigger bits are stored while waiting for the proper delay to come out. That waiting time consists of a coarse delay in units of CDF clock cycles as well as an adjustable fine delay (in steps of 22 ns) that will allow to align the PreFRED outputs within 22 ns of a CDF clock cycle. Once the Global L1 system is integrated and the alignment delays between modules are perfectly known, the coarse length of the pipeline and the phase-selecting fine delay will be freezed to constant values.

The L1 FIFO will store the bunch counter value, ΣE_t, ΣE_x, ΣE_y, $\not{E}_T^2$, the four trigger bits and an error bit (see below) while waiting for a L1A or L1R. The implementation of the L1 FIFO and buffer management is similar to the one developed for DIRAC and CRATESUM, so as to keep the U. of C. production consistent.

2.2.5 Processing steps

The logic functions to be processed by the SUMET board are the following (see functional diagram in Figure 1):

1. The 12 $E_t(\phi)$ words are weighted by an unsigned coefficient ($|\cos\phi_1|$, $|\sin\phi_1|$, 1, $|\cos\phi_2|$, $|\sin\phi_2|$ depending on the cycle phase)

2. They are then summed two by two in cascade until a half detector sum has resulted; in order to properly compute ΣE_x and ΣE_y, a sign in the cascading sums is introduced depending on the cycle phase and the wedges being added

3. After a 66 ns wait for the second half sum to be formed, the half sums are in turn added, yielding ΣE_x, ΣE_y, and ΣE_t in that order.

4. Each one of the sum results is then formatted to fit a certain resolution (see Table 1). Their absolute value is computed before being squared. $(\Sigma E_x)^2$ then waits 22 ns for $(\Sigma E_y)^2$ to be available as they are added to form $\not{E}_T^2$. ΣE_t, ΣE_x, ΣE_y and $\not{E}_T^2$ are then ready to leave the processing chip to feed the L1 FIFO. These outputs are aligned and latched out on the desired 22 ns phase.

5. Meanwhile, ΣE_t and $\not{E}_T^2$ are compared with their thresholds so that four trigger bits are yielded. These bits then feed a pipeline which will lead them to FRED with the proper PreFRED alignment delay. At the same time these bits are sent to FRED, they also feed a L1 FIFO along with the bunch counter value.

During the adding process, an overflow bit is kept track of at every stage. It is ORed with all other overflow bits involved in either the ΣE_t or $\not{E}_T^2$ calculation. In the end, in case of a ΣE_t or $\not{E}_t$ overflow (occuring in particular if one of the input words is saturated), the ΣE_t

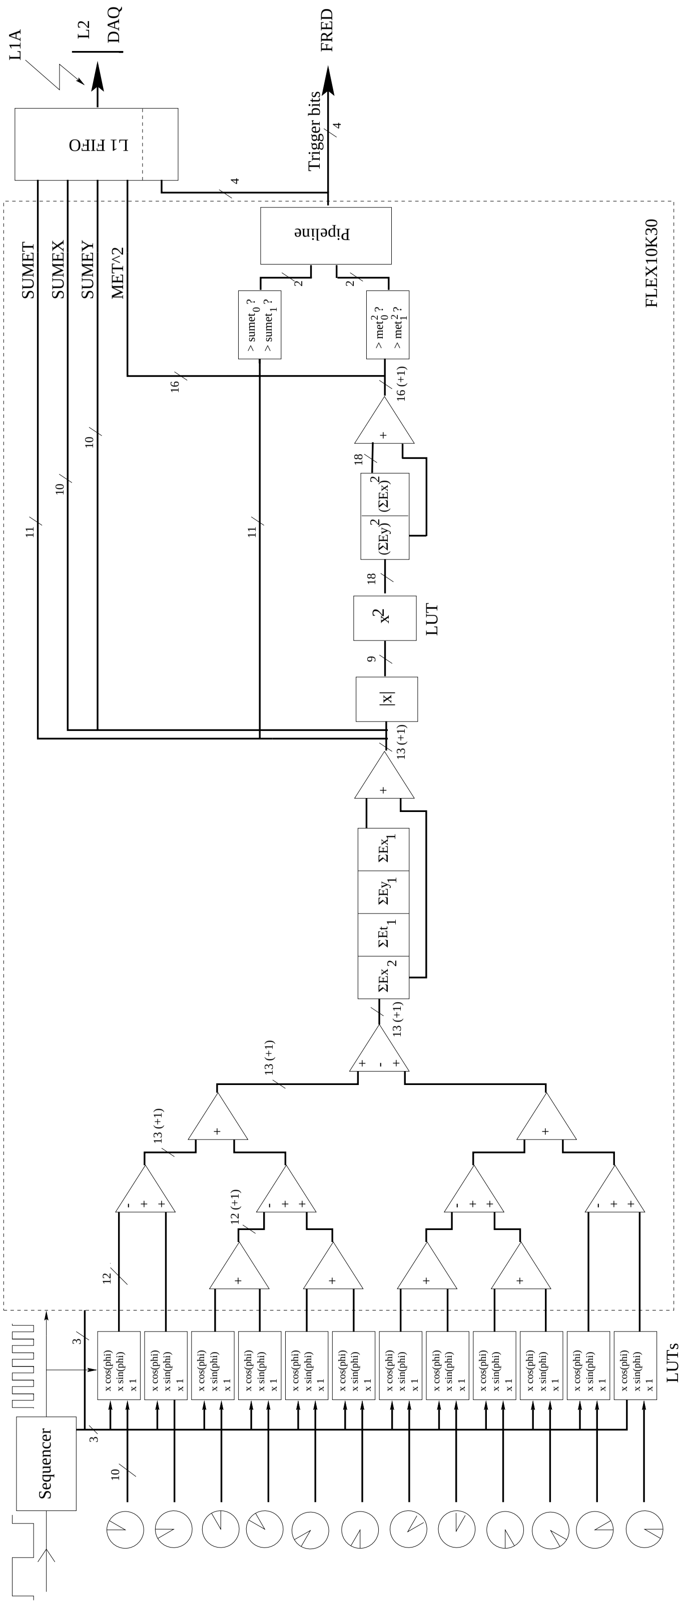

Figure 1: Functional diagram of the SUMET module: the board receives data at an incoming clock cycle of 66 ns; that frequency is multiplied by a factor of three by the sequencer chip to allow processing of ΣE_x, ΣE_y, and ΣE_t on the same data path. The sectors in the circles represent the transverse detector sections to which the incoming data belong (a sector covers an even and an odd wedge). The numbers along the paths represent the number of bits transmitted. The (+1)'s indicate that an overflow bit is kept track of.

Table 1: Format of the output data. LSB = Least Significant Bit. * Signed words are actually absolute values appended with a sign bit at the MSB location; they are not two's complements.

Output word	# of bits	type	LSB	Full scale
ΣE_t	11	unsigned	1 GeV	2048 GeV
ΣE_x	10	signed *	0.5 GeV	256 GeV
ΣE_y	10	signed *	0.5 GeV	256 GeV
$\not{E}_T^2$	16	unsigned	1 GeV^2	(256 GeV)^2
error	1	set to 1 if ΣE_t turns negative		
trigbit	4	each bit set to 1 if above threshold		

or $\not{E}_t$ words are ORed with their respective overflow bits so the results come out saturated themselves. In either case, the trigger bits are set to 1, so as to accept the saturating event.

If at any time in the process, the ΣE_t word (before being formatted to 11 bits) has its bit sign turned to 1, it means an error occurred in the summing sequence. That ΣE_t sign bit is then identified as an error bit and sent out along with the data to the L1 FIFO. It also directly feeds the Control module so as to alert the TSI via VME.

2.2.6 Connection to L2

Upon a L1A, ΣE_t, ΣE_x, and ΣE_y will be sent out from the L1 FIFO to L2 via a front panel differential connector. These words will be the only data signals passing through the front panel of the board. In addition, the 2-bit DAQ buffer address associated with the event and received from the TSI will be sent in parallel with them, so as to tag the data for L2. This total of 33 bits can be accomodated by a single 68-pin differential Honda connector (see pin assignment for that connector in Table 2).

2.2.7 Connection to the J3 backplane

As mentioned earlier, every 66 ns, the Aux-card will receive 120 input signals from the CRATESUMs, feeding the SUMET module through the J3 backplane. Now it is foreseen that the backplane will basically be the same (or very similar) as the one used for the L1 Calorimetry crate, which has 24 point-to-point connections between each DIRAC and their associated CRATESUM. In a similar way, the PreFRED modules will use subsets of these connections to send their trigger bits to FRED. In the SUMET case, only 4 output connections will be used. The input/output pin assignments on the backplane are the same for every PreFRED module (see Table 3).

2.3 Implementation

The new generation of Programmable Logic Devices with Embedded Memory allows a great deal of flexibility, integration, and registered speed performance, which are the qualities looked upon for this application. By including a maximum number of logic functions within such a device, we effectively reduce the number of parts needed for the board, increase the speed performance and enable future re-design of the logic by simply reprogramming or reconfiguring the device. Such an implementation has been achieved at the simulation stage with the Altera 356-pin FLEX10K30 BGA chip: all of the logic framed with a dash line on

Table 2: Pin assignments for the SUMET front panel connector (output to L2): bit #: data. SUMET = ΣE_t, SUMEX = ΣE_x, SUMEY = ΣE_y, L1BA = DAQ buffer address. For each one of those words, bit 0 represents the least significant bit. SUMEX9 and SUMEY9 are sign bits.

	Column	
Row	A	B
1	1: SUMET0-	35: SUMEX6-
2	2: SUMET0+	36: SUMEX6+
3	3: SUMET1-	37: SUMEX7-
4	4: SUMET1+	38: SUMEX7+
5	5: SUMET2-	39: SUMEX8-
6	6: SUMET2+	40: SUMEX8+
7	7: SUMET3-	41: SUMEX9-
8	8: SUMET3+	42: SUMEX9+
9	9: SUMET4-	43: SUMEY0-
10	10: SUMET4+	44: SUMEY0+
11	11: SUMET5-	45: SUMEY1-
12	12: SUMET5+	46: SUMEY1+
13	13: SUMET6-	47: SUMEY2-
14	14: SUMET6+	48: SUMEY2+
15	15: SUMET7-	49: SUMEY3-
16	16: SUMET7+	50: SUMEY3+
17	17: SUMET8-	51: SUMEY4-
18	18: SUMET8+	52: SUMEY4+
19	19: SUMET9-	53: SUMEY5-
20	20: SUMET9+	54: SUMEY5+
21	21: SUMET10-	55: SUMEY6-
22	22: SUMET10+	56: SUMEY6+
23	23: SUMEY0-	57: SUMEY7-
24	24: SUMEY0+	58: SUMEY7+
25	25: SUMEY1-	59: SUMEY8-
26	26: SUMEY1+	60: SUMEY8+
27	27: SUMEY2-	61: SUMEY9-
28	28: SUMEY2+	62: SUMEY9+
29	29: SUMEX3-	63: GND
30	30: SUMEX3+	64: GND
31	31: SUMEX4-	65: L1BA0-
32	32: SUMEX4+	66: L1BA0+
33	33: SUMEX5-	67: L1BA1-
34	34: SUMEX5+	68: L1BA1+

Table 3: Pin assignments for the PreFRED-J3 Backplane connection. TRIGBIT = output trigger bits to FRED, CSIN = input from CRATESUMs. In the case of the SUMET board, only the four least significant TRIGBITs are used, and the CSIN pins transmit the wedge transverse energies. In the case of the TOWTRG module, CSIN brings in trigger bits. The index associated with CSIN represents the channel number, which follows the wedge numbering scheme such that channel 0 is associated to wedges 0 and 1, channel 1 is associated to wedges 2 and 3, etc. The numbers in parentheses are the bit numbers with 0 as least significant. AUX_SUM_EN is connected to ground on the board so as to forbid the summing of channels enabled for DIRAC. The Receiver Aux-card makes no connection to the TRIGBIT pins.

	Row						
Pin #	z	a	b	c	d	e	f
1	GND	GND	TRIGBIT(0)	TRIGBIT(1)	TRIGBIT(2)	TRIGBIT(3)	GND
2	GND	N/C	TRIGBIT(4)	TRIGBIT(5)	GND	TRIGBIT(6)	GND
3	GND	N/C	TRIGBIT(7)	GND	TRIGBIT(8)	TRIGBIT(9)	GND
4	GND	N/C	TRIGBIT(10)	TRIGBIT(11)	GND	TRIGBIT(12)	GND
5	GND	GND	TRIGBIT(13)	GND	TRIGBIT(14)	TRIGBIT(15)	GND
6	GND	N/C	CSIN0(0)	CSIN0(1)	CSIN0(2)	CSIN0(3)	GND
7	GND	N/C	CSIN0(4)	CSIN0(5)	GND	CSIN0(6)	GND
8	GND	N/C	CSIN0(7)	CSIN0(8)	CSIN0(9)	CSIN1(0)	GND
9	GND	GND	CSIN1(1)	GND	CSIN1(2)	CSIN1(3)	GND
10	GND	N/C	CSIN1(4)	CSIN1(5)	CSIN1(6)	CSIN1(7)	GND
11	GND	N/C	CSIN1(8)	CSIN1(9)	GND	CSIN2(0)	GND
12	GND	N/C	CSIN2(1)	CSIN2(2)	CSIN2(3)	CSIN2(4)	GND
13	GND	GND	CSIN2(5)	GND	CSIN2(6)	CSIN2(7)	GND
14	GND	N/C	CSIN2(8)	CSIN2(9)	CSIN3(0)	CSIN3(1)	GND
15	GND	N/C	CSIN3(2)	CSIN3(3)	GND	CSIN3(4)	GND
16	GND	N/C	CSIN3(5)	CSIN3(6)	CSIN3(7)	CSIN3(8)	GND
17	GND	GND	CSIN3(9)	GND	CSIN4(0)	CSIN4(1)	GND
18	GND	N/C	CSIN4(2)	CSIN4(3)	CSIN4(4)	CSIN4(5)	GND
19	GND	N/C	CSIN4(6)	CSIN4(7)	GND	CSIN4(8)	GND
20	GND	N/C	CSIN4(9)	CSIN5(0)	CSIN5(1)	CSIN5(2)	GND
21	GND	GND	CSIN5(3)	GND	CSIN5(4)	CSIN5(5)	GND
22	GND	N/C	CSIN5(6)	CSIN5(7)	CSIN5(8)	CSIN5(9)	GND
23	GND	N/C	CSIN6(0)	CSIN6(1)	GND	CSIN6(2)	GND
24	GND	N/C	CSIN6(3)	CSIN6(4)	CSIN6(5)	CSIN6(6)	GND
25	GND	GND	CSIN6(7)	GND	CSIN6(8)	CSIN6(9)	GND
26	GND	N/C	CSIN7(0)	CSIN7(1)	CSIN7(2)	CSIN7(3)	GND
27	GND	N/C	CSIN7(4)	CSIN7(5)	GND	CSIN7(6)	GND
28	GND	N/C	CSIN7(7)	CSIN7(8)	CSIN7(9)	CSIN8(0)	GND
29	GND	GND	CSIN8(1)	GND	CSIN8(2)	CSIN8(3)	GND
30	GND	N/C	CSIN8(4)	CSIN8(5)	CSIN8(6)	CSIN8(7)	GND
31	GND	N/C	CSIN8(8)	CSIN8(9)	GND	CSIN9(0)	GND
32	GND	N/C	CSIN9(1)	CSIN9(2)	CSIN9(3)	CSIN9(4)	GND
33	GND	GND	CSIN9(5)	GND	CSIN9(6)	CSIN9(7)	GND
Connector Key (equivalent to 3 pins)							
34	GND	N/C	CSIN9(8)	CSIN9(9)	CSIN10(0)	CSIN10(1)	GND
35	GND	N/C	CSIN10(2)	CSIN10(3)	GND	CSIN10(4)	GND
36	GND	N/C	CSIN10(5)	CSIN10(6)	CSIN10(7)	CSIN10(8)	GND
37	GND	GND	CSIN10(9)	GND	CSIN11(0)	CSIN11(1)	GND
38	GND	N/C	CSIN11(2)	CSIN11(3)	CSIN11(4)	CSIN11(5)	GND
39	GND	N/C	CSIN11(6)	CSIN11(7)	GND	CSIN11(8)	GND
40	GND	N/C	CSIN11(9)	AUX_SUM_EN	TRIGBIT(16)	TRIGBIT(17)	GND
41	GND	GND	TRIGBIT(18)	AUX_SPARE	TRIGBIT(19)	N/C	GND
42	GND	N/C	N/C	N/C	GND	N/C	GND
43	GND	N/C	N/C	N/C	N/C	N/C	GND
44	GND	N/C	N/C	N/C	N/C	N/C	GND

Figure 1 has been fitted by the Altera development software and runs properly with a 22 ns clock cycle. Nevertheless, the beginning ϕ-dependent weighting function cannot easily be implemented in such a PLD due to the large combinatorial multiplication process.

2.3.1 ϕ-dependent weighting

The $E_t(\phi)$ initial weighting can be performed via look-up tables, i.e. the readout of static RAMs whose input address is the 10-bit $E_t(\phi)$ word itself in association with 3 cycle phase bits selecting alternatively one of the 6 weighting coefficients. These 3 bits come from the sequencer such that the first 2 factors to be used for a new event are $|\cos\phi|$ and $|\sin\phi|$, since more time is needed at the end for the $\not{E}_t$ computation. To preserve a certain resolution after weighting by these factors, the SRAM output data word needs to be more than 10-bits long: 12 output bits are thus foreseen with a least count of 0.125 GeV for the ΣE_x and ΣE_y cycle phases.[1] At this stage, no sign is affected to the data yet. For the ΣE_t cycle phase, the energy scale is kept with a 0.5 GeV least count, which, with 12 bits available, offers a full scale of 2048 GeV. If an input channel comes saturated (wedge $E_t > 512$ GeV), the memory output bits are all set to 1's, regardless of the CDF-clock-cycle phase.

So 12 SRAMs are required with a 13-bit wide addressing and at least 12 bits for the output (8K$\times$12); moreover these memories should withstand a 22 ns clock cycle. Synchronous 16K$\times$16 fast SRAMs are offered on the market with read cycle times of less than 15 ns (Motorola MCM62996). These chips should satisfy our needs, pending board simulation and prototype tests.

2.3.2 Adding, squaring, threshold comparison: the PLD integration

After the ϕ-dependent weighting, all subsequent logic functions enumerated earlier can be implemented in the Altera FLEX10K30BC356-3 PLD chip: these include the cascade adders, the absolute value and squaring functions, the threshold comparators and the alignment pipeline. They were effectively coded using the Altera High-Level Development Language (a language similar to VHDL, only simpler) and its parameterized "megafunctions". From that source code, the Altera Development Environment fits all required logics within the single FLEX10k30 chip. The 144 input bits coming from the SRAMs as well as the 52 output bits are well accomodated by the 246 I/O pins available on that device.

Each symbol of the functional diagram specifies a logic step to be executed at every 22 ns clock rise. Before each step, the data are registered for a better control of the phasing issues. In particular, the sign of some cascade adders (add or subtract) has to be switched by the clock depending on the ΣE_x or ΣE_y cycle phase. All but the output of the first stage of cascade adders require an extra sign bit, bringing the number of bits per word feeding the next stage of adders to 13. In addition, for each summing, an overflow bit is also registered and ORed together with all other overflow bits from the same event (this is done separately for the ΣE_t and $\not{E}_t$ processes). The waiting for the second half of the detector data to come through involves 4-word-deep shift-registers, both for the half-sums and the overflow bits. Eventually, the result of the overall sum is truncated to fit the output specifications. Of the 12 bits effectively available for ΣE_t, only the 11 most significant bits are kept for a 2048 GeV full scale and a 1 GeV least count. For $|\Sigma E_x|$ or $|\Sigma E_y|$, the MSB of the summing result

[1] A simple Monte Carlo simulation shows the $\not{E}_t$ numerical resolution deteriorates by less than 0.4 GeV in the scheme presented here.

is considered as overflow whereas the 2 LSB are dropped, yielding a 9-bit word with a 256 GeV full scale and a 0.5 GeV least count. This word in turn feeds the address port of a 512×18 embedded memory which acts as a look-up table to produce the square of ΣE_x or ΣE_y. Then $(\Sigma E_x)^2$ and $(\Sigma E_y)^2$ are aligned and added, yielding the $\not{E}_T^2$ result as well as a last overflow bit to be ORed with the $\not{E}_t$ overflow sequence. In the process, the two LSB are dropped from $\not{E}_T^2$ so that the least count for the result becomes 1 GeV^2.

No reduction of the $\not{E}_T^2$ 16-bit word to a smaller 8-bit $\not{E}_t$ is foreseen, as it would require too big a look-up table to fit in the PLD and consequently involve the loss of clock cycles needed to operate it. Thus for the last compare-to-threshold stage, simple 16-bit-wide comparators are used on the $\not{E}_T^2$ path. Similar comparators deal with ΣE_t. Since the system has to be reconfigurable to accomodate as many as 3 or 4 $\not{E}_T^2$ thresholds (and 1 or 0 ΣE_t thresholds), 4×16 bits are reserved in the PLD for the threshold allocation. For the ΣE_t thresholds, only 11 out of 16 bits are used. [2] In a last stage, the 4 resulting comparison trigger bits are ORed with the $\not{E}_t$ and ΣE_t final overflow bits, so as to a priori accept events which caused the saturation of any energy scale encountered during the summing process. Only then are those four bits simultaneously written into the PreFRED alignment pipeline; this has been implemented with a 3-bit programmable depth allowing up to 8 stacked 4-bit words, corresponding to a maximum alignment delay of a microsecond. Within a CDF clock cycle, as mentioned in the design section, the output data can be latched with a 22 ns resolution by selecting 1 of 6 shifted 132 ns clocks with a 3-bit programmable fine delay.

The FLEX10K30 PLD is programmed with volatile memories within it. After the power is turned off, it needs to be reprogrammed or "reconfigured". It is planned to have a dedicated Flash Memory (FM) on the SUMET board which will automatically reconfigure the PLD as soon as the power is back on. That FM is large enough to contain several configuration settings, so that different programs can be selected for the PLD: this would be one way of allowing three of four thresholds for $\not{E}_t$ as opposed to only two. Also there could be a configuration where the $\not{E}_t$ and ΣE_t overflow bits would be ignored by the 4 output trigger bits (i.e. not ORed with them), in order to disable unwanted data acquisition caused by some known hot channel. However in that case, unless the configuration program is rewritten to eliminate that particular channel from the adding process, $\not{E}_t$, ΣE_t, and their consequent trigger bits would have no validity.

2.3.3 Timing diagram and latency

The next figure shows a timing diagram of the various signals going through the FLEX10K30 PLD, as simulated with the Altera development package[3]. There is no rest-time between the processing of separate events: the whole sequence of events is pipelined and each event is hatched in 6 22-ns data bunches: $E_{x1}, E_{y1}, E_{t_1}, E_{x2}, E_{y2}, E_{t2}$ in that order. The 12 "etin" lines represent the incoming 12-bit signals coming from the SRAMs. In this particular instance, time $t_0 = 605$ ns indicates the arrival of the first data bunch (E_{x1}) into the PLD chip. All incoming data for the first event extend to $t_0 + 132$ ns. The "clk" signal is fed in by the sequencer, in phase opposite to the arrival of data. "et0" is merely the output port of the input registers associated with the incoming E_{t0}. The "clken" are the cycle phase taggers,

[2]The thresholds themselves are to be down-loaded before the beginning of a run. This is done accross a dedicated 16-bit port connected to a 4-word-deep shift-register located within the PLD, and enabled only for that initialization time.

[3]No input values appearing on the diagram intend to reproduce real data here (in particular there is no consistency between the cosine- and sine-weighted values).

i.e. 132-ns clocks coming from the sequencer. From them, the signs "sg" needed in operating the additions/subtractions of ΣE_x or ΣE_y are obtained logically with simple AND/OR gates. "sum" is the signed 13-bit ΣE_x, ΣE_y, or ΣE_t. That "sum" value is truncated when going through the absolute value function, which yields the 9-bit "abssum". ΣE_t is made available for output as the truncated 11-bit "sumet" while ΣE_x ("sumex") and ΣE_y ("sumey") are extracted from their absolute values and re-concatanated with their sign bits. $\not{E}_T^2$ ("metsq") also comes out from the pipeline shorly thereafter, so that all those data are gathered into a 132-ns clocked register, whose timing can be adjusted by 22 ns increments ("fine-delay") to synchronize the data departure from the chip with the first L1 FIFO write request. In the example shown on the diagram, fine-delay = 2 causes all output to be delayed by 66 ns after their time of first availability (2×22 ns + 22 ns prior to that delay for setting up the variable delay process). "dataout[10..0],[20..11],[30..21],[46..31],47" are respectively ΣE_t, ΣE_x, ΣE_y, $\not{E}_T^2$, and the error bit as they leave the chip. The 4 trigger bits can be issued out at the same time, before the specified 396 ns latency requirement: "out[51..48]" represents the trigger bits as they become available. Or in the more general case where an alignment delay is needed to present those trigger bits to FRED, the "align-delay" 3-bit word is used in units of CDF clock cycles. In this instance, align-delay = 1 causes the trigger bits to be issued 1 CDF clock cycle after the other output, as the "tofred" word.

3 TOWTRG

3.1 Specifications

The TOWTRG Module receives 20-bit tower-trigger summary words from the 12 CRATE-SUM boards[3] and condenses them into a single 20-bit detector summary word to be sent out to FRED, and DAQ. The 20-bit input tower-trigger summary words from the CRATE-SUMs will include a priori 10 single object trigger bits (e.g. 1 electron with $E_t > 20$ GeV, 1 photon with $E_t > 25$ GeV...) and 3 di-object pairs of bits (e.g. at least 2 electrons with $E_t > 10GeV$...) which tell whether 0, 1, 2 or more objects with the di-object trigger were found in the associated wedge pair. In the case where more than 2 objects fulfill a di-object trigger requirement, the result of the 2-bit adder is set to 3. The output tower-trigger summary word should reflect the same pattern. For the di-object triggers, out of the two resulting bits, only the most significant bit will tell FRED whether at least two objects of the same nature were found with the corresponding threshold. The lower significant bit will also be used by FRED for combined di-object triggers like e-μ. Thus in this scheme, 16 bits out of 20 are used by the system, which leaves 4 bits for additional information if required in a future development (tri-object trigger?).

The timing requirement is that all steps leading to the TOWTRG summary should be performed within one clock cycle (132 ns). However, the output must be sent in phase with the other PreFRED modules.

The main issue resides in the connectivity of all input signals to a single TOWTRG board: the 20 bit words from the 12 CRATESUM boards arrive on differential signal cables.

3.2 Design and Implementation

The logic design of the board is straightforward since the tower trigger summary will be performed in the same fashion as in CRATESUM with the DIRAC data: the single object trigger bits will have to be ORed whereas the di-object bits will simply be added two by two. Then an overflow in the adding process means more than 3 objects pass the di-object threshold, in which case the 2-bit result is set to 3.

This logic is simple enough that it can easily be fit in the same type of PLD (FLEX10K30) as that used for SUMET. Moreover, the number of input bits ($2\times$ 120) is identical to that of the SUMET board; the number of output bits (20) is lower than for SUMET (32). This enables the same board design for the 2 modules, as well as the use of the same kind of Aux cards (DIRAC's) to receive the incoming data. Only the PLD configuration would have to differ. Thus, as for SUMET, the data would be received in 66 ns sequences from the CRATESUMs: each 20-bit summary word would have to be cut in two and sent in two sequential packets across the same kind of cables as that used to feed the SUMET Aux-card. The first sequence should be made of 10 bits that include the 6 di-object trigger bits as the di-object summing process takes longer than the ORing of the single trigger bits. The 2nd sequence would then be made of the remaining 10 single object trigger bits.

The input look-up tables used in SUMET would be dummy in the TOWTRG scheme, merely echoing the incoming data to the PLD. The sequencer generating the 22 ns clock would still be used to quickly process the data along registers between the di-object cascade adders, in order to meet the one-CDF-clock-cycle latency requirement. Eventually, the 20 resulting bits should be sent out of the PLD not only to the L1 FIFO, but also to FRED, after a PreFRED alignment delay implemented by a small FIFO within the PLD (FRED

FIFO), as is done in SUMET. The difference here is the larger number of bits to be sent to FRED, which has to be taken into account in the joint board design.

3.2.1 Connection to L2

No connection to L2 is required for TOWTRG, although the PreFRED Front Panel connectors used by the SUMET module will be available to transfer the 20 TOWTRG output trigger bits if found necessary.

3.2.2 Connection to the J3 backplane

As in the case of SUMET, every 66 ns, the Aux-card will receive 120 input signals from the CRATESUMs, feeding the TOWTRG module through the J3 backplane. As for the output to FRED in the TOWTRG case, 20 point-to-point connections will be used on the backplane. The pin assignments for the TOWTRG connection to the backplane (Table 3) show the locations of the input trigger bits from the CRATESUMs as well as the output trigger summary bits.

4 Conclusion

According to the current design scheme, both the SUMET module and the TOWTRG module will be implemented using an identical IC board with different programs/LUTs loaded in their PLDs/memories. In both instances, the data will come from the CRATESUM boards in two sequences. Then the re-use of the DIRAC-specified Aux-Card seems most appropriate to bring the incoming data to the board. Moreover the output trigger bits to FRED, whether they come from SUMET or TOWTRG, can be implemented with point to point connections using the L1 P3 backplane specified to link the DIRACs to the CRATESUMs.

A 22 ns clock generated by a local sequencer is used to process the data, i.e. compute ΣE_x, ΣE_y and ΣE_t in sequence in the SUMET case. A total of 12 16K$\times$16 SRAM chips are needed for the ϕ-dependent weighting while only one PLD chip is required for the subsequent processing of the entire SUMET/TOWTRG function (FLEX10K30). In addition, the board requires a VME interface, a Controller, a start-up Flash-Memory and the L1 FIFOs and DAQ buffers. Some of these peripheral functions which have already been developed for the DIRAC module only need duplication (VME Interface) or slight modification (L1 FIFOs/DAQ buffers...) for the SUMET/TOWTRG module (see schematics in appendix).

Moreover, the SUMET design is general enough that it can also apply to the other more simple PREFRED modules. It is foreseen to end up with the same boards for all PREFREDs, with just different programs in their PLDs.

References

[1] H.Frisch et al. *Conceptual Design of the L1 Calorimeter Trigger for the Run II Upgrade*, CDF Note 2909.

[2] J.Wahl *DIRAC*, Trigger Note 90.

[3] D. Toback *CrateSum, Upgrade for Run II.*

5 Appendix: some PreFRED schematics

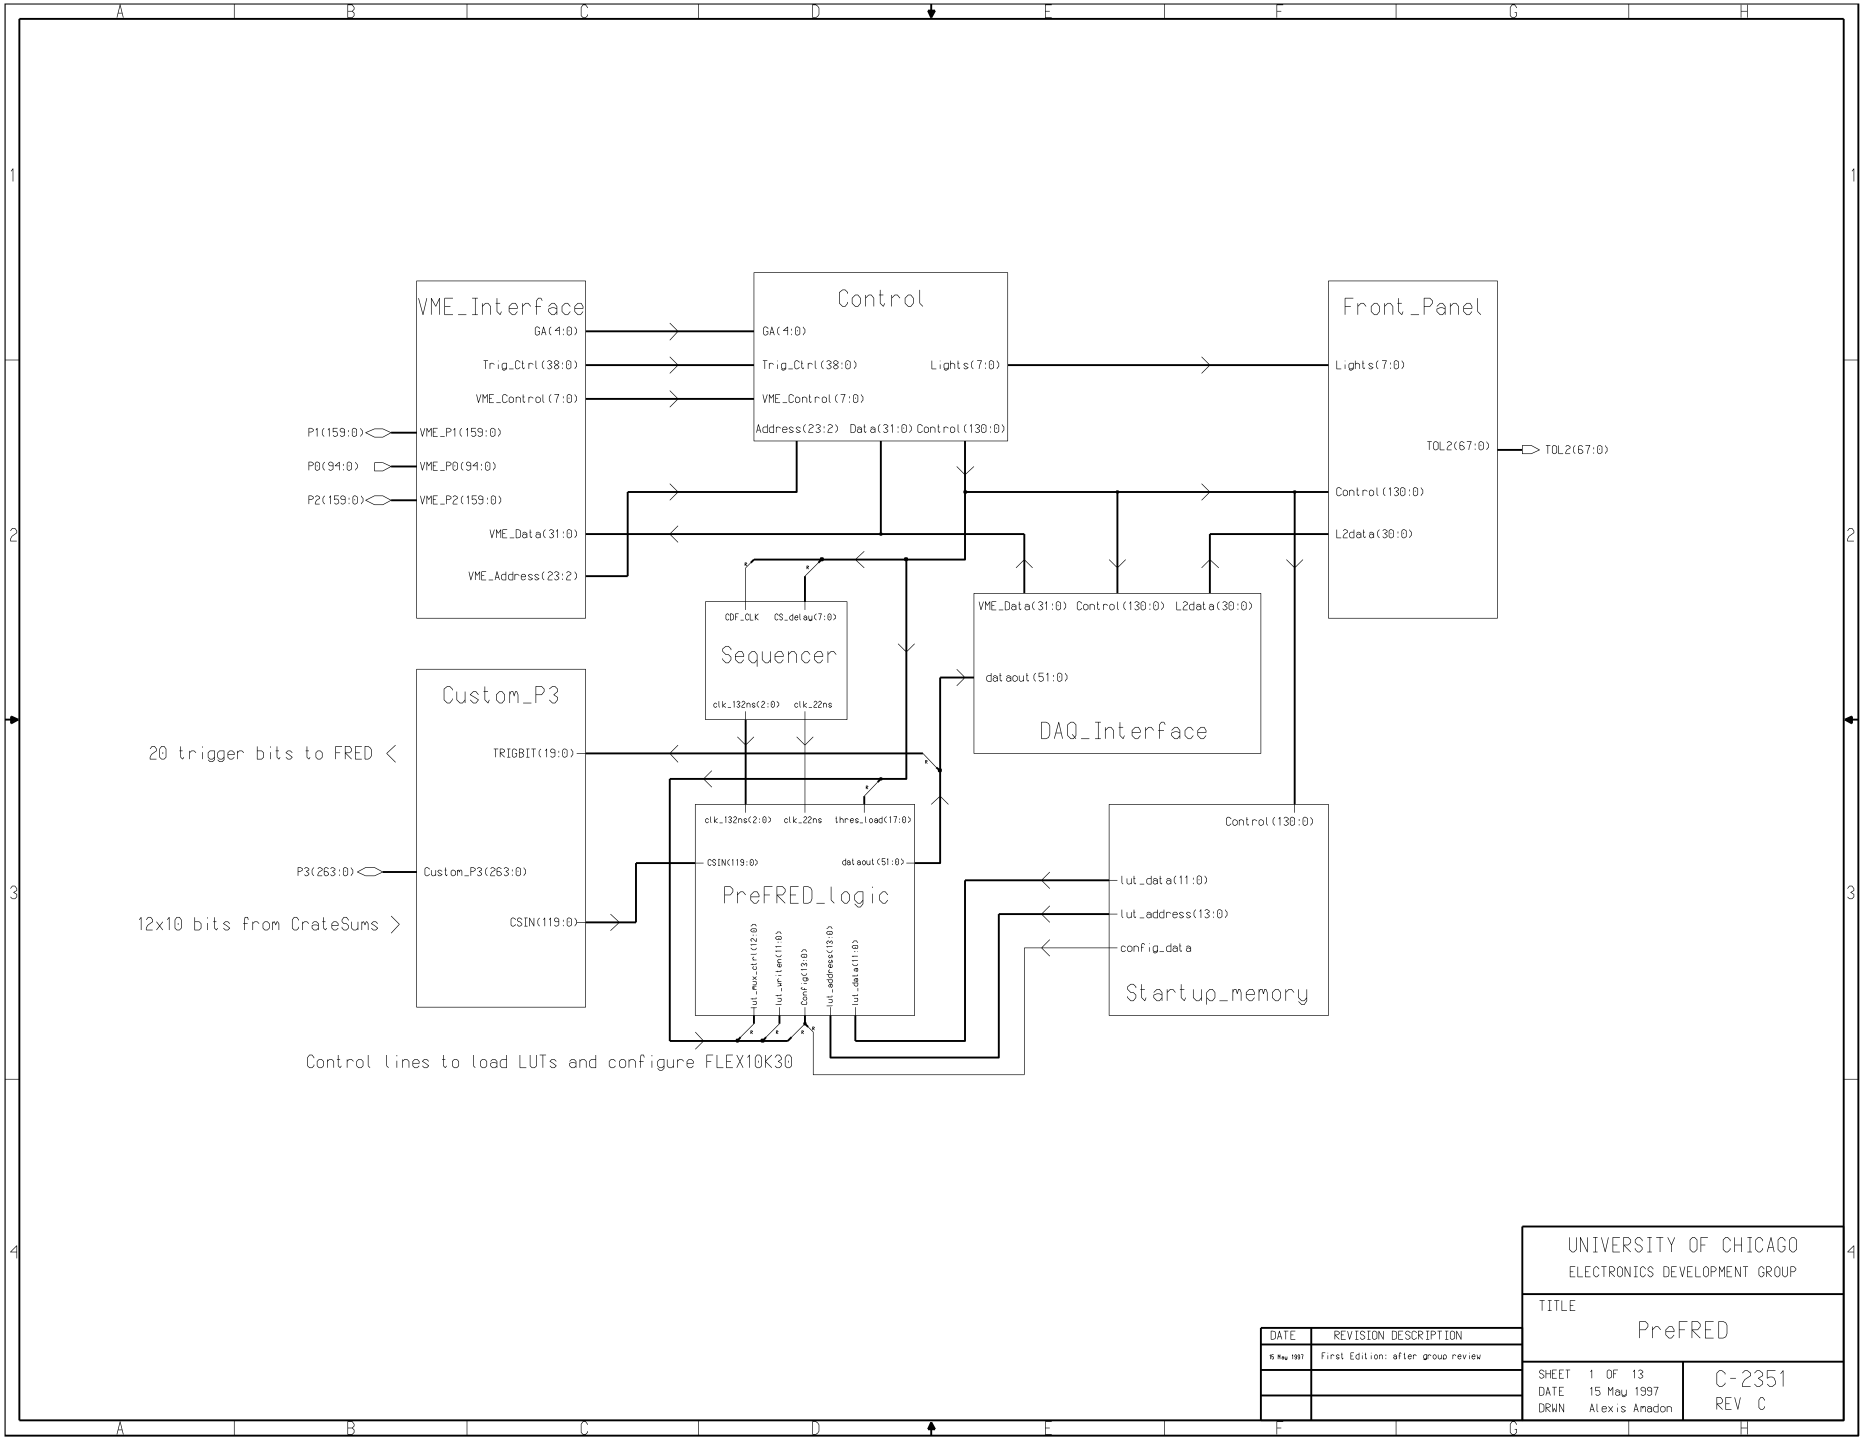

VME_Interface
GA(4:0)
Trig_Ctrl(38:0)
VME_Control(7:0)
P1(159:0)
VME_P1(159:0)
P0(94:0)
VME_P0(94:0)
P2(159:0)
VME_P2(159:0)
VME_Data(31:0)
VME_Address(23:2)
Control
GA(4:0)
Trig_Ctrl(38:0)
VME_Control(7:0)
Lights(7:0)
Address(23:2) Data(31:0) Control(130:0)
Front_Panel
Lights(7:0)
TOL2(67:0)
TOL2(67:0)
Control(130:0)
L2data(30:0)
Sequencer
CDF_CLK CS_delay(7:0)
clk_132ns(2:0) clk_22ns
DAQ_Interface
VME_Data(31:0) Control(130:0) L2data(30:0)
dataout(51:0)
Custom_P3
20 trigger bits to FRED
TRIGBIT(19:0)
P3(263:0)
Custom_P3(263:0)
12x10 bits from CrateSums
CSIN(119:0)
PreFRED_logic
clk_132ns(2:0) clk_22ns thres_load(17:0)
CSIN(119:0)
dataout(51:0)
lut_mux_ctrl(12:0)
lut_writen(11:0)
Config(13:0)
lut_address(13:0)
lut_data(11:0)
Startup_memory
Control(130:0)
lut_data(11:0)
lut_address(13:0)
config_data
Control lines to load LUTs and configure FLEX10K30
UNIVERSITY OF CHICAGO
ELECTRONICS DEVELOPMENT GROUP
TITLE
PreFRED
SHEET 1 OF 13
DATE 15 May 1997
DRWN Alexis Amadon
C-2351
REV C
DATE
REVISION DESCRIPTION
15 May 1997
First Edition: after group review

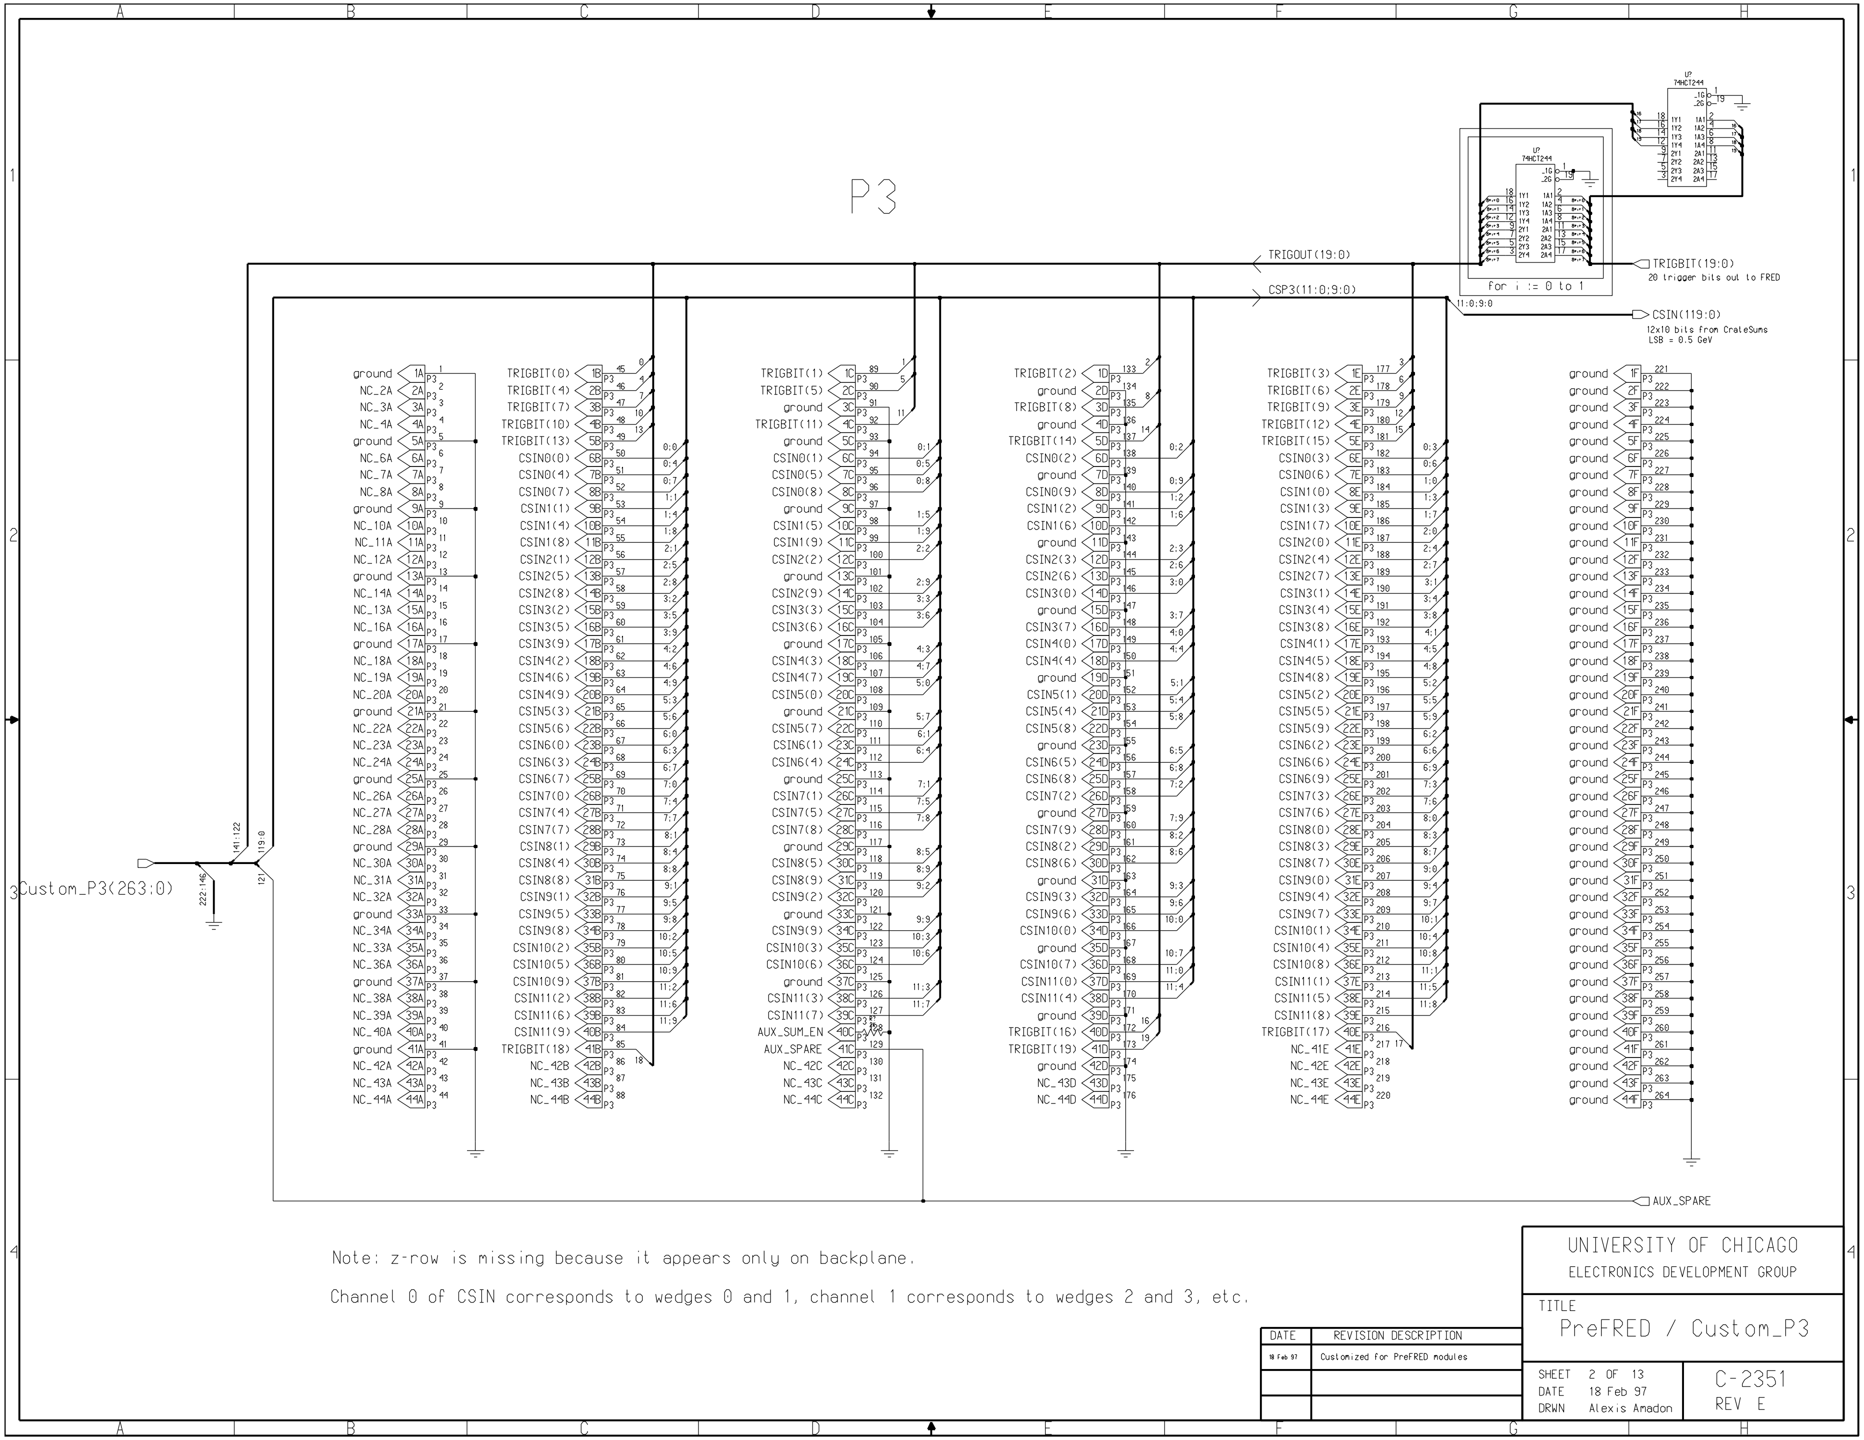
P3
Custom_P3(263:0)
TRIGOUT(19:0)
CSP3(11:0;9:0)
TRIGBIT(19:0)
20 trigger bits out to FRED
CSIN(119:0)
12x10 bits from CrateSums
LSB = 0.5 GeV
For i := 0 to 1
74HCT244
AUX_SPARE
Note: z-row is missing because it appears only on backplane.
Channel 0 of CSIN corresponds to wedges 0 and 1, channel 1 corresponds to wedges 2 and 3, etc.
DATE
REVISION DESCRIPTION
18 Feb 97
Customized for PreFRED modules
UNIVERSITY OF CHICAGO
ELECTRONICS DEVELOPMENT GROUP
TITLE
PreFRED / Custom_P3
SHEET 2 OF 13
DATE 18 Feb 97
DRWN Alexis Amadon
C-2351
REV E

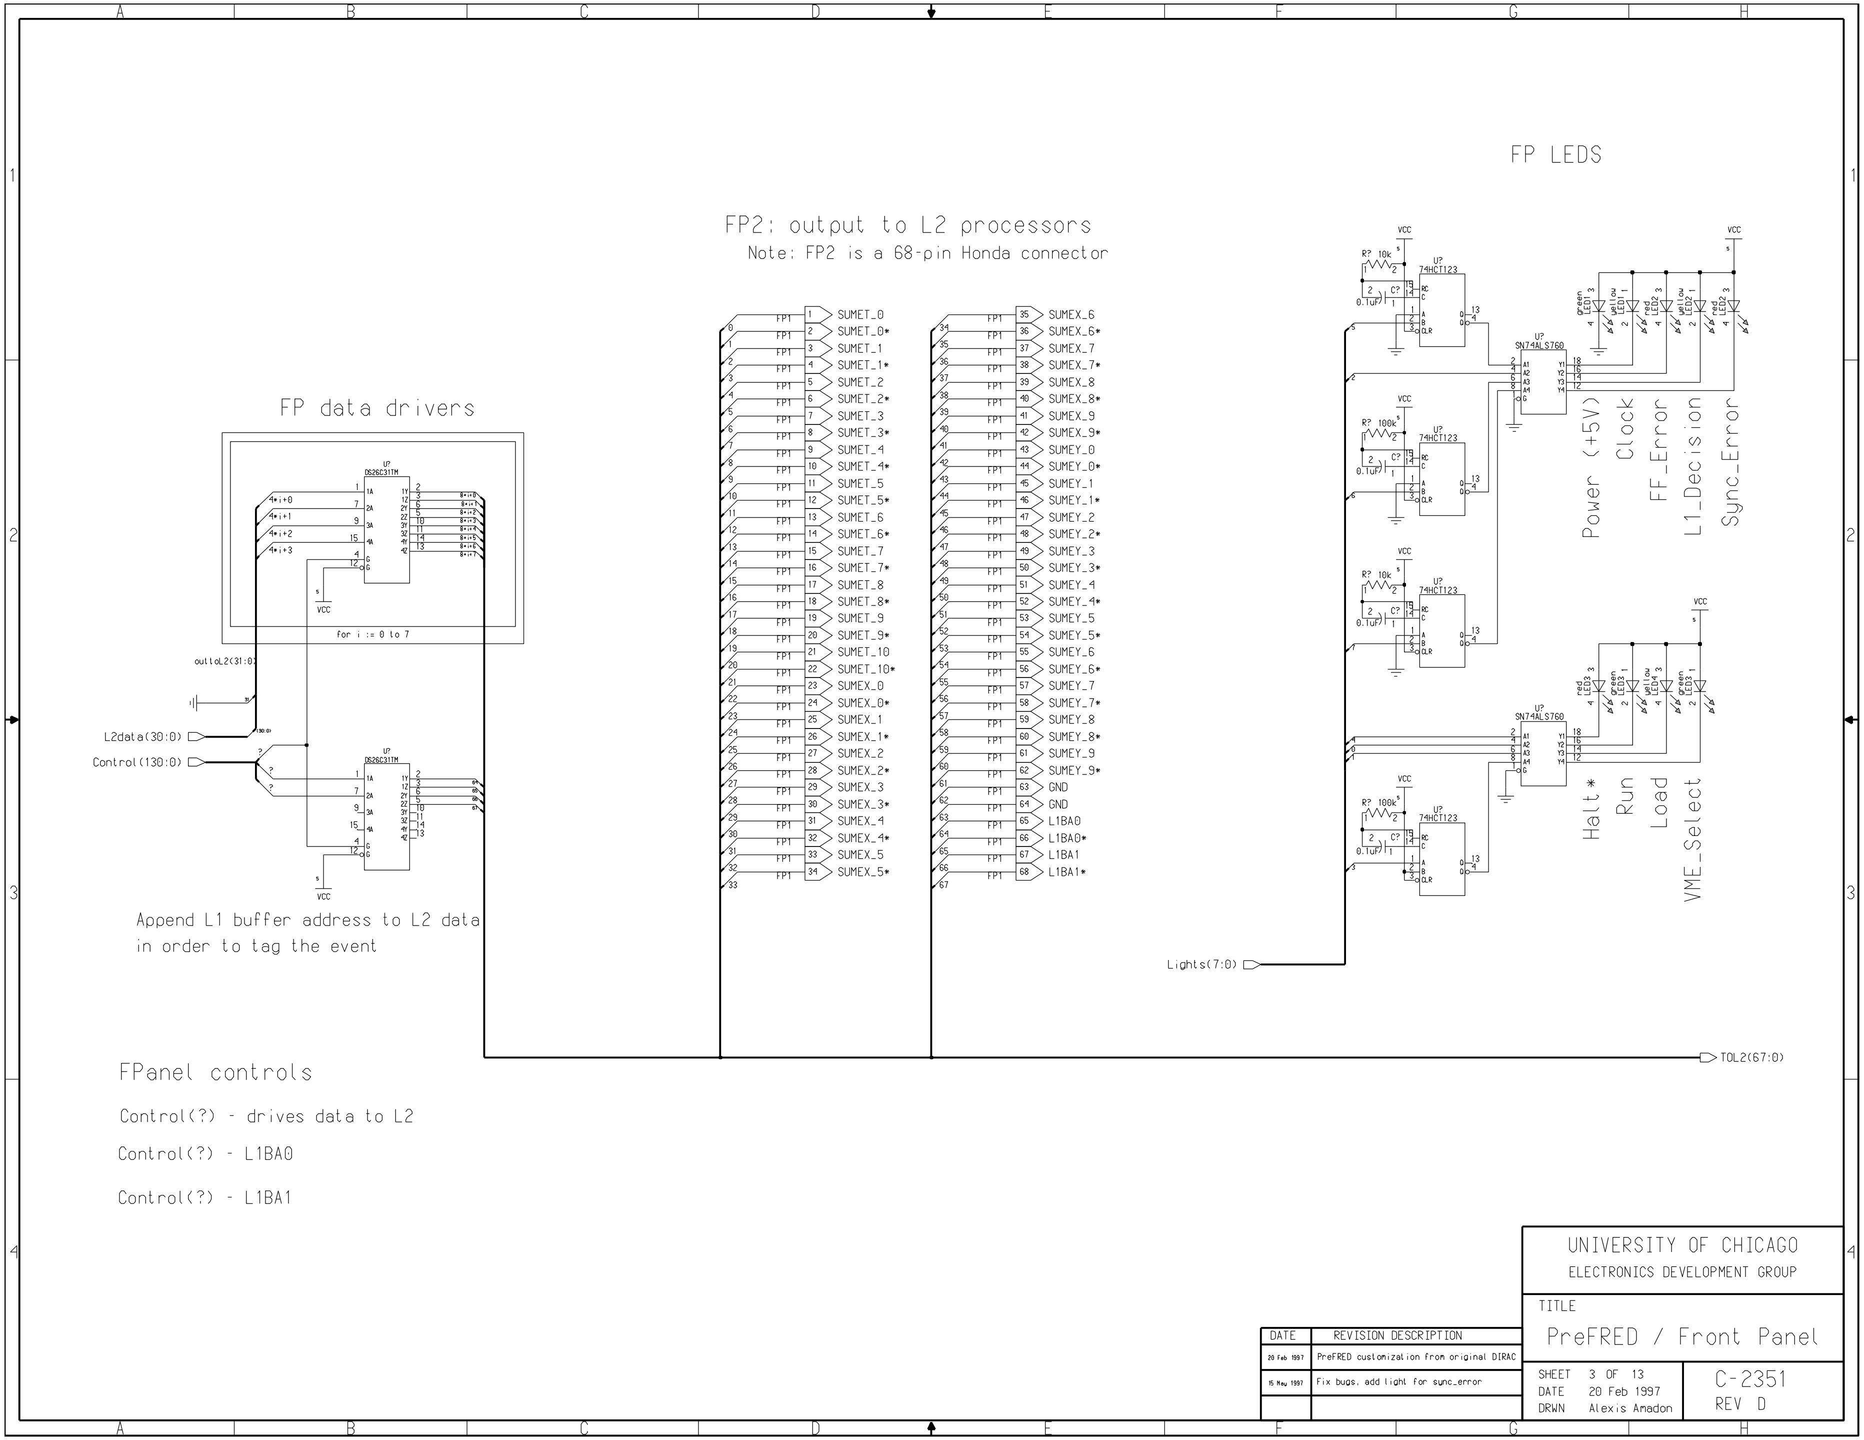

FP data drivers
For i := 0 to 7
DS26C31TM
outtoL2(31:0)
L2data(30:0)
Control(130:0)
Append L1 buffer address to L2 data
in order to tag the event
FPanel controls
Control(?) - drives data to L2
Control(?) - L1BA0
Control(?) - L1BA1
FP2: output to L2 processors
Note: FP2 is a 68-pin Honda connector
TOL2(67:0)
FP LEDS
74HCT123
SN74ALS760
Lights(7:0)
Power (+5V)
Clock
FF_Error
L1_Decision
Sync_Error
Halt*
Run
Load
VME_Select
UNIVERSITY OF CHICAGO
ELECTRONICS DEVELOPMENT GROUP
TITLE
PreFRED / Front Panel
SHEET 3 OF 13
DATE 20 Feb 1997
DRWN Alexis Amadon
C-2351
REV D
DATE
REVISION DESCRIPTION
20 Feb 1997
PreFRED customization from original DIRAC
15 May 1997
Fix bugs, add light for sync_error

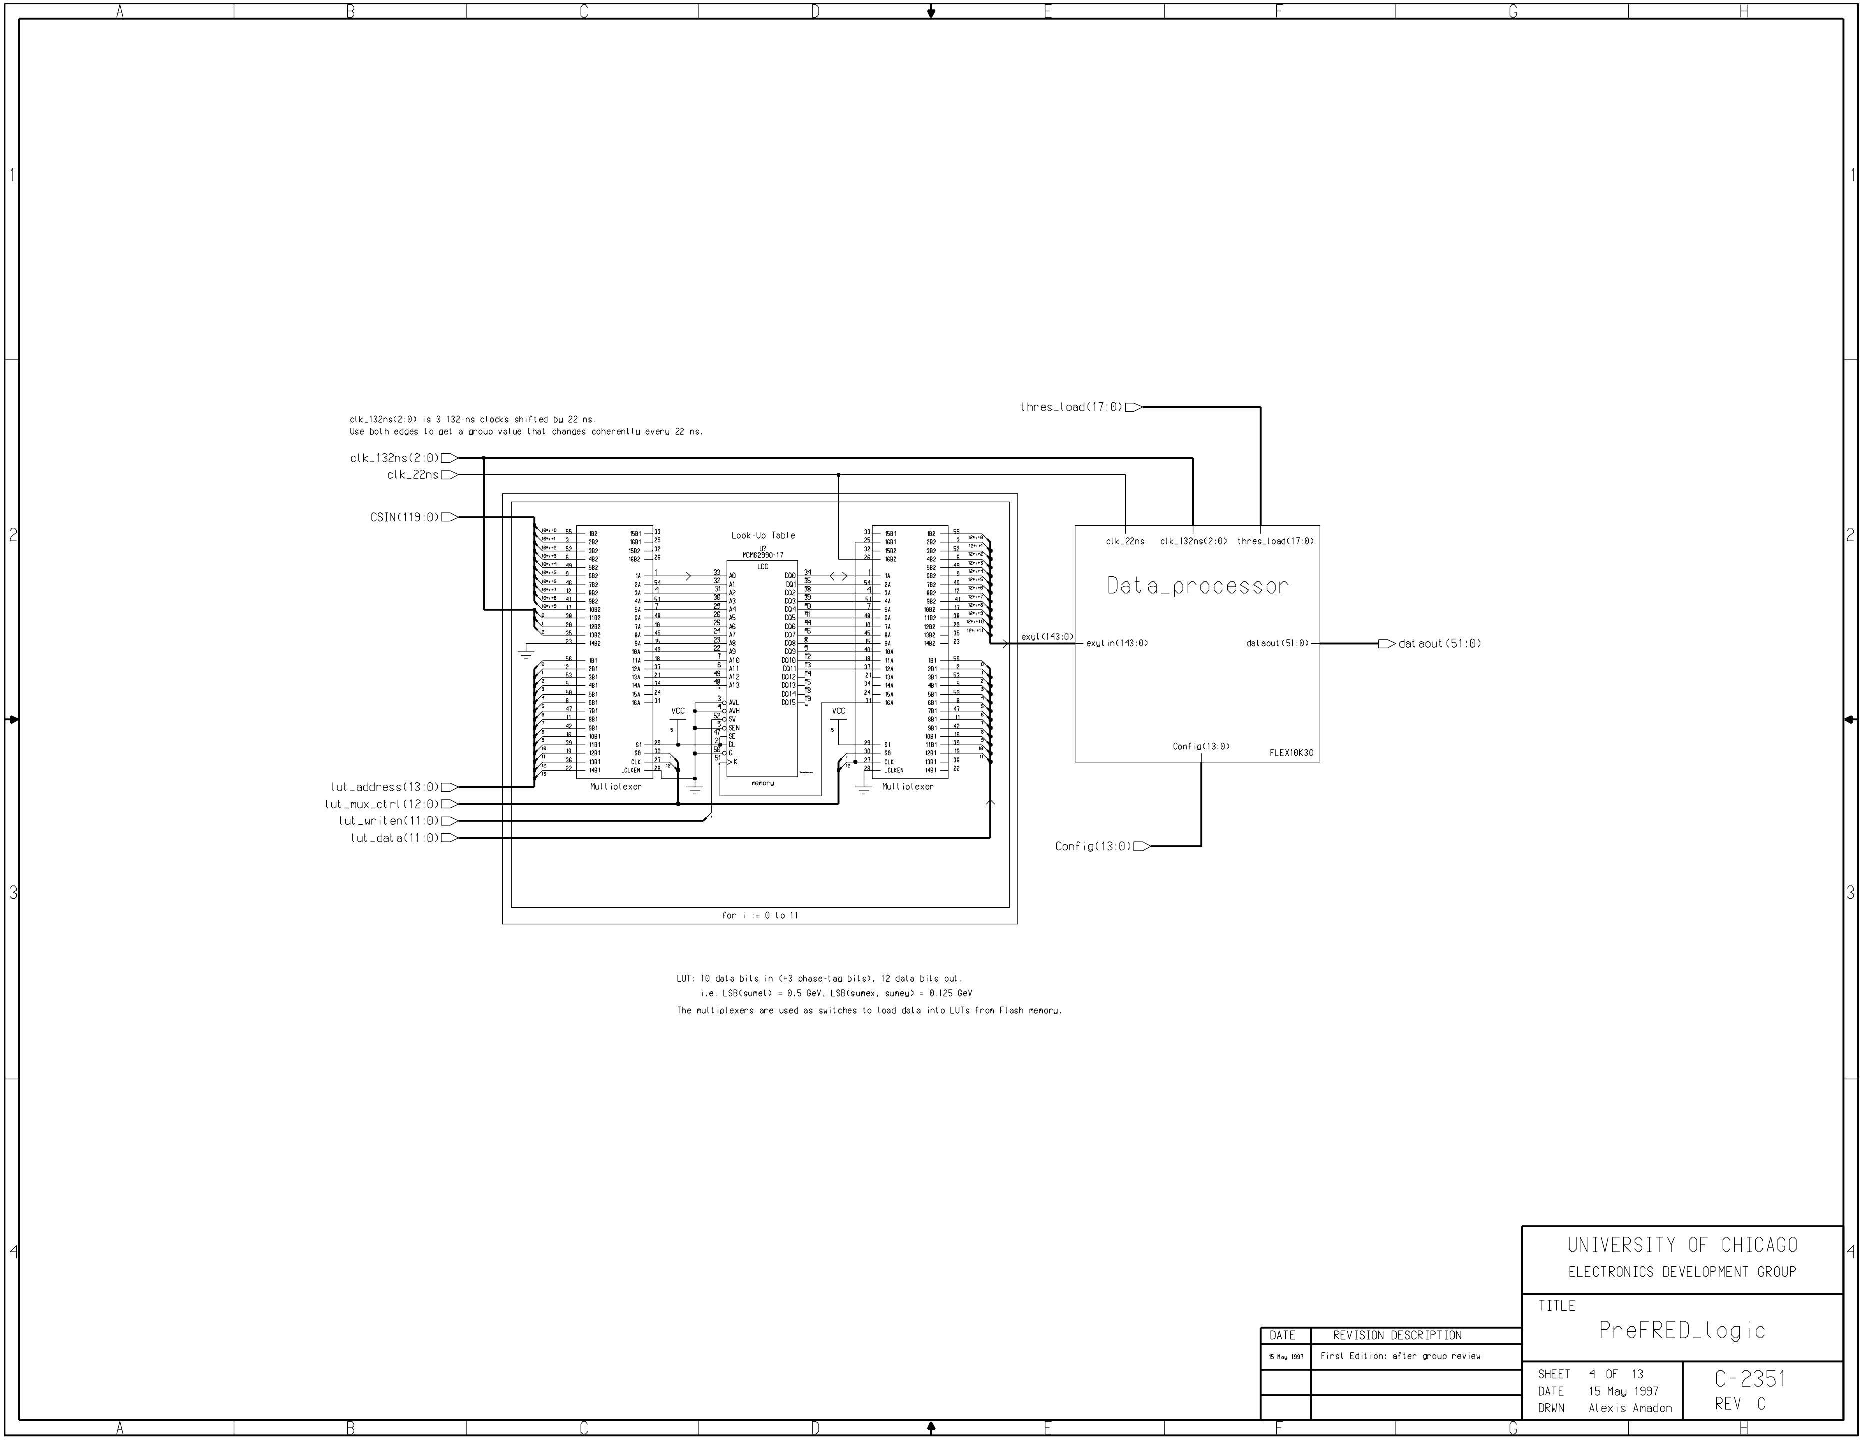
clk_132ns(2:0) is 3 132-ns clocks shifted by 22 ns.
Use both edges to get a group value that changes coherently every 22 ns.
clk_132ns(2:0)
clk_22ns
CSIN(119:0)
lut_address(13:0)
lut_mux_ctrl(12:0)
lut_writen(11:0)
lut_data(11:0)
thres_load(17:0)
Config(13:0)
Look-Up Table
U?
MCM62990-17
LCC
memory
Multiplexer
Multiplexer
VCC
VCC
for i := 0 to 11
exyt(143:0)
Data_processor
clk_22ns
clk_132ns(2:0)
thres_load(17:0)
exytin(143:0)
dataout(51:0)
Config(13:0)
FLEX10K30
dataout(51:0)
LUT: 10 data bits in (+3 phase-lag bits), 12 data bits out.
i.e. LSB(sumet) = 0.5 GeV, LSB(sumex, sumey) = 0.125 GeV
The multiplexers are used as switches to load data into LUTs from Flash memory.
DATE
REVISION DESCRIPTION
15 May 1997
First Edition: after group review
UNIVERSITY OF CHICAGO
ELECTRONICS DEVELOPMENT GROUP
TITLE
PreFRED_logic
SHEET 4 OF 13
DATE 15 May 1997
DRWN Alexis Amadon
C-2351
REV C

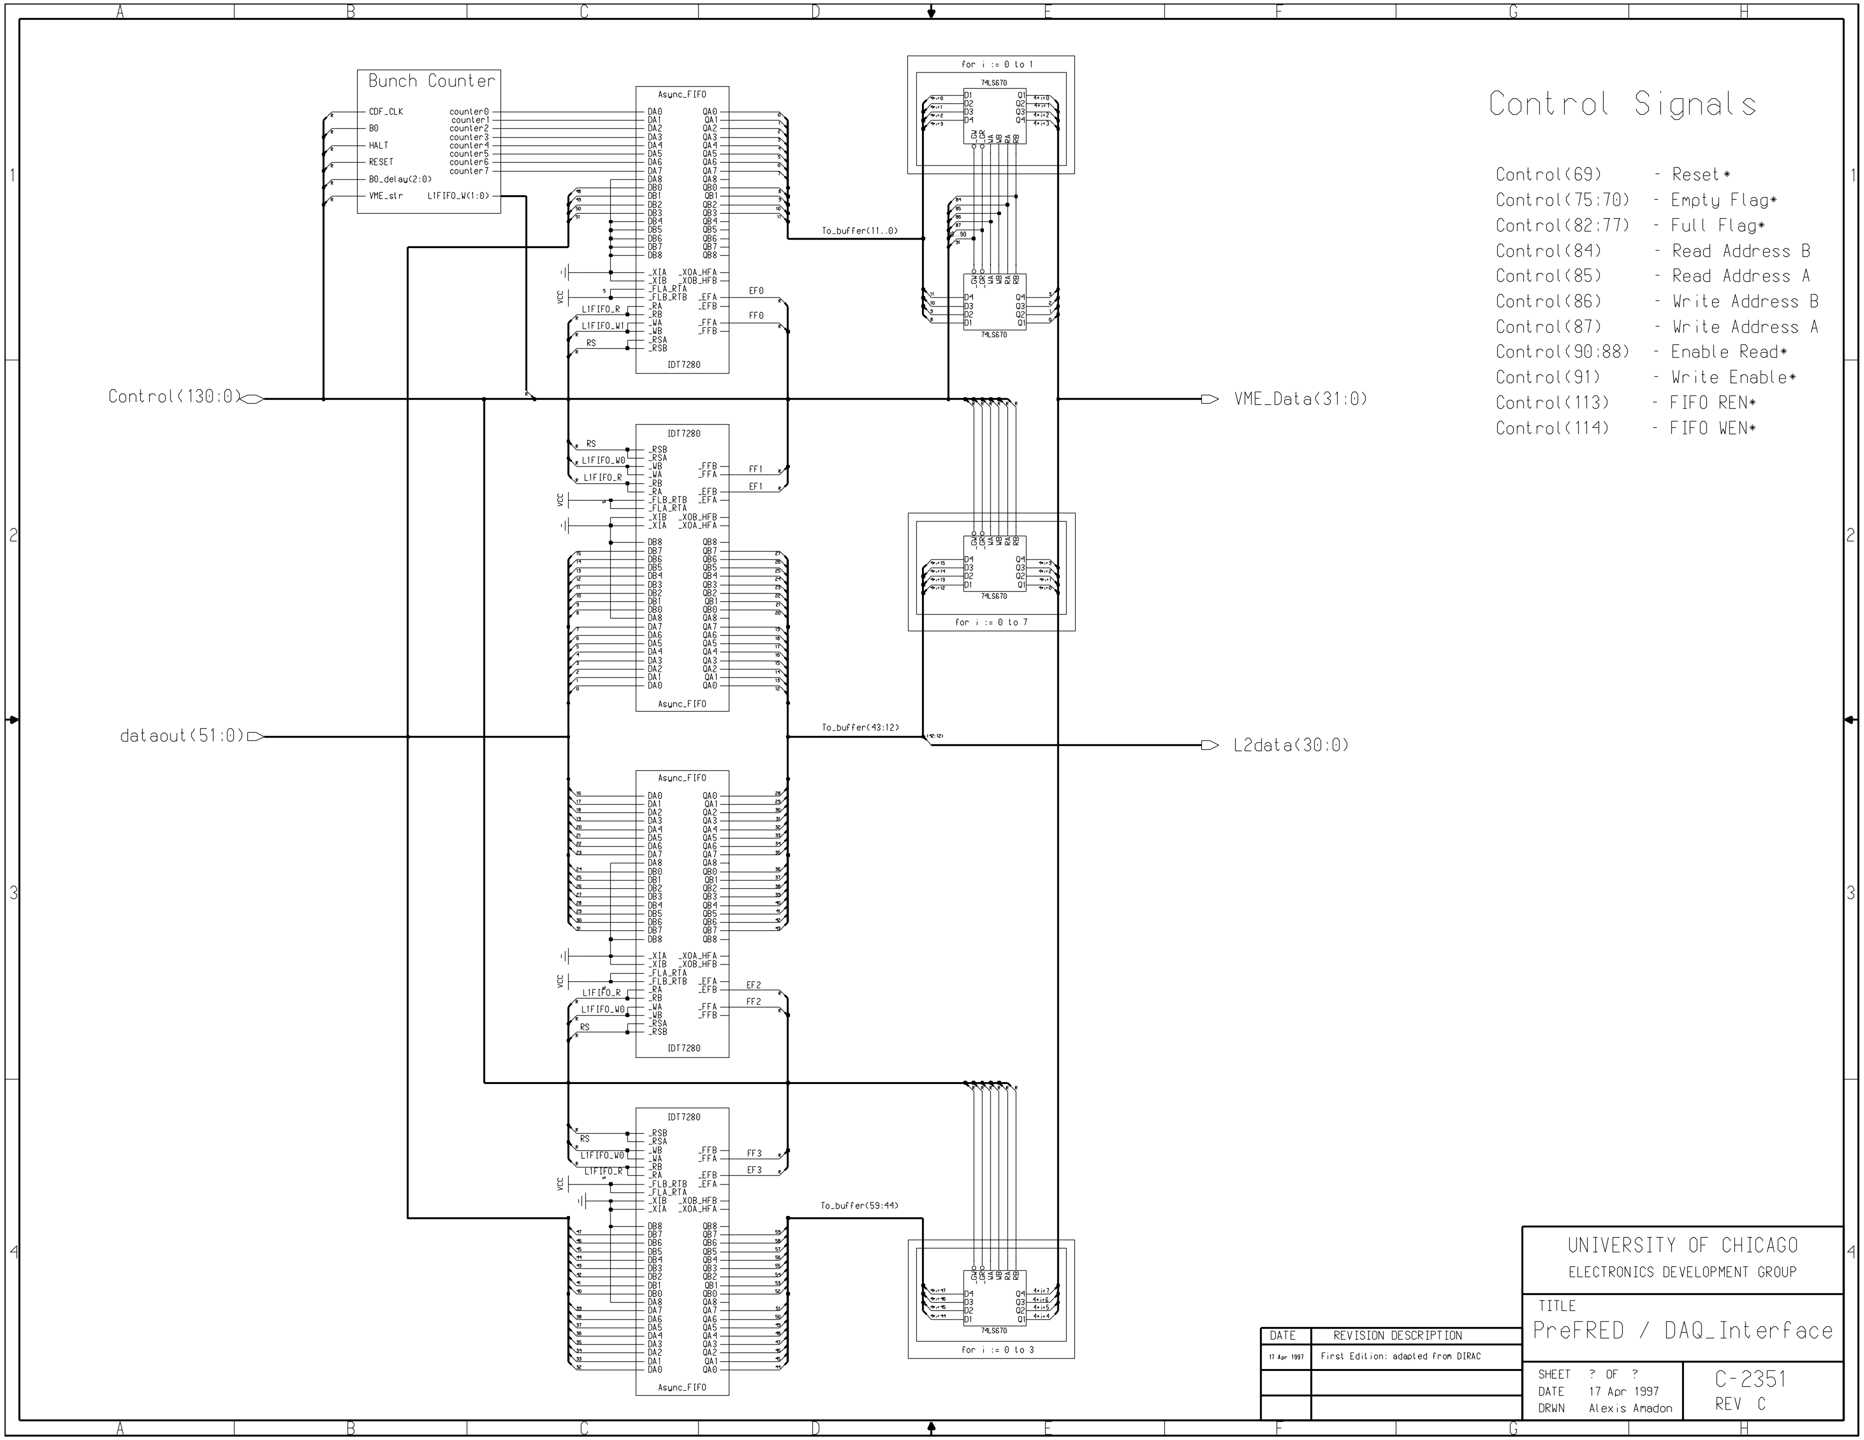
Bunch Counter
Async_FIFO
IDT7280
74LS670
Control(130:0)
dataout(51:0)
VME_Data(31:0)
L2data(30:0)
To_buffer(11..0)
To_buffer(43:12)
To_buffer(59:44)
for i := 0 to 1
for i := 0 to 7
for i := 0 to 3
Control Signals
Control(69) - Reset*
Control(75:70) - Empty Flag*
Control(82:77) - Full Flag*
Control(84) - Read Address B
Control(85) - Read Address A
Control(86) - Write Address B
Control(87) - Write Address A
Control(90:88) - Enable Read*
Control(91) - Write Enable*
Control(113) - FIFO REN*
Control(114) - FIFO WEN*
UNIVERSITY OF CHICAGO
ELECTRONICS DEVELOPMENT GROUP
TITLE
PreFRED / DAQ_Interface
SHEET ? OF ?
DATE 17 Apr 1997
DRWN Alexis Amadon
C-2351
REV C
DATE
REVISION DESCRIPTION
17 Apr 1997
First Edition: adapted from DIRAC

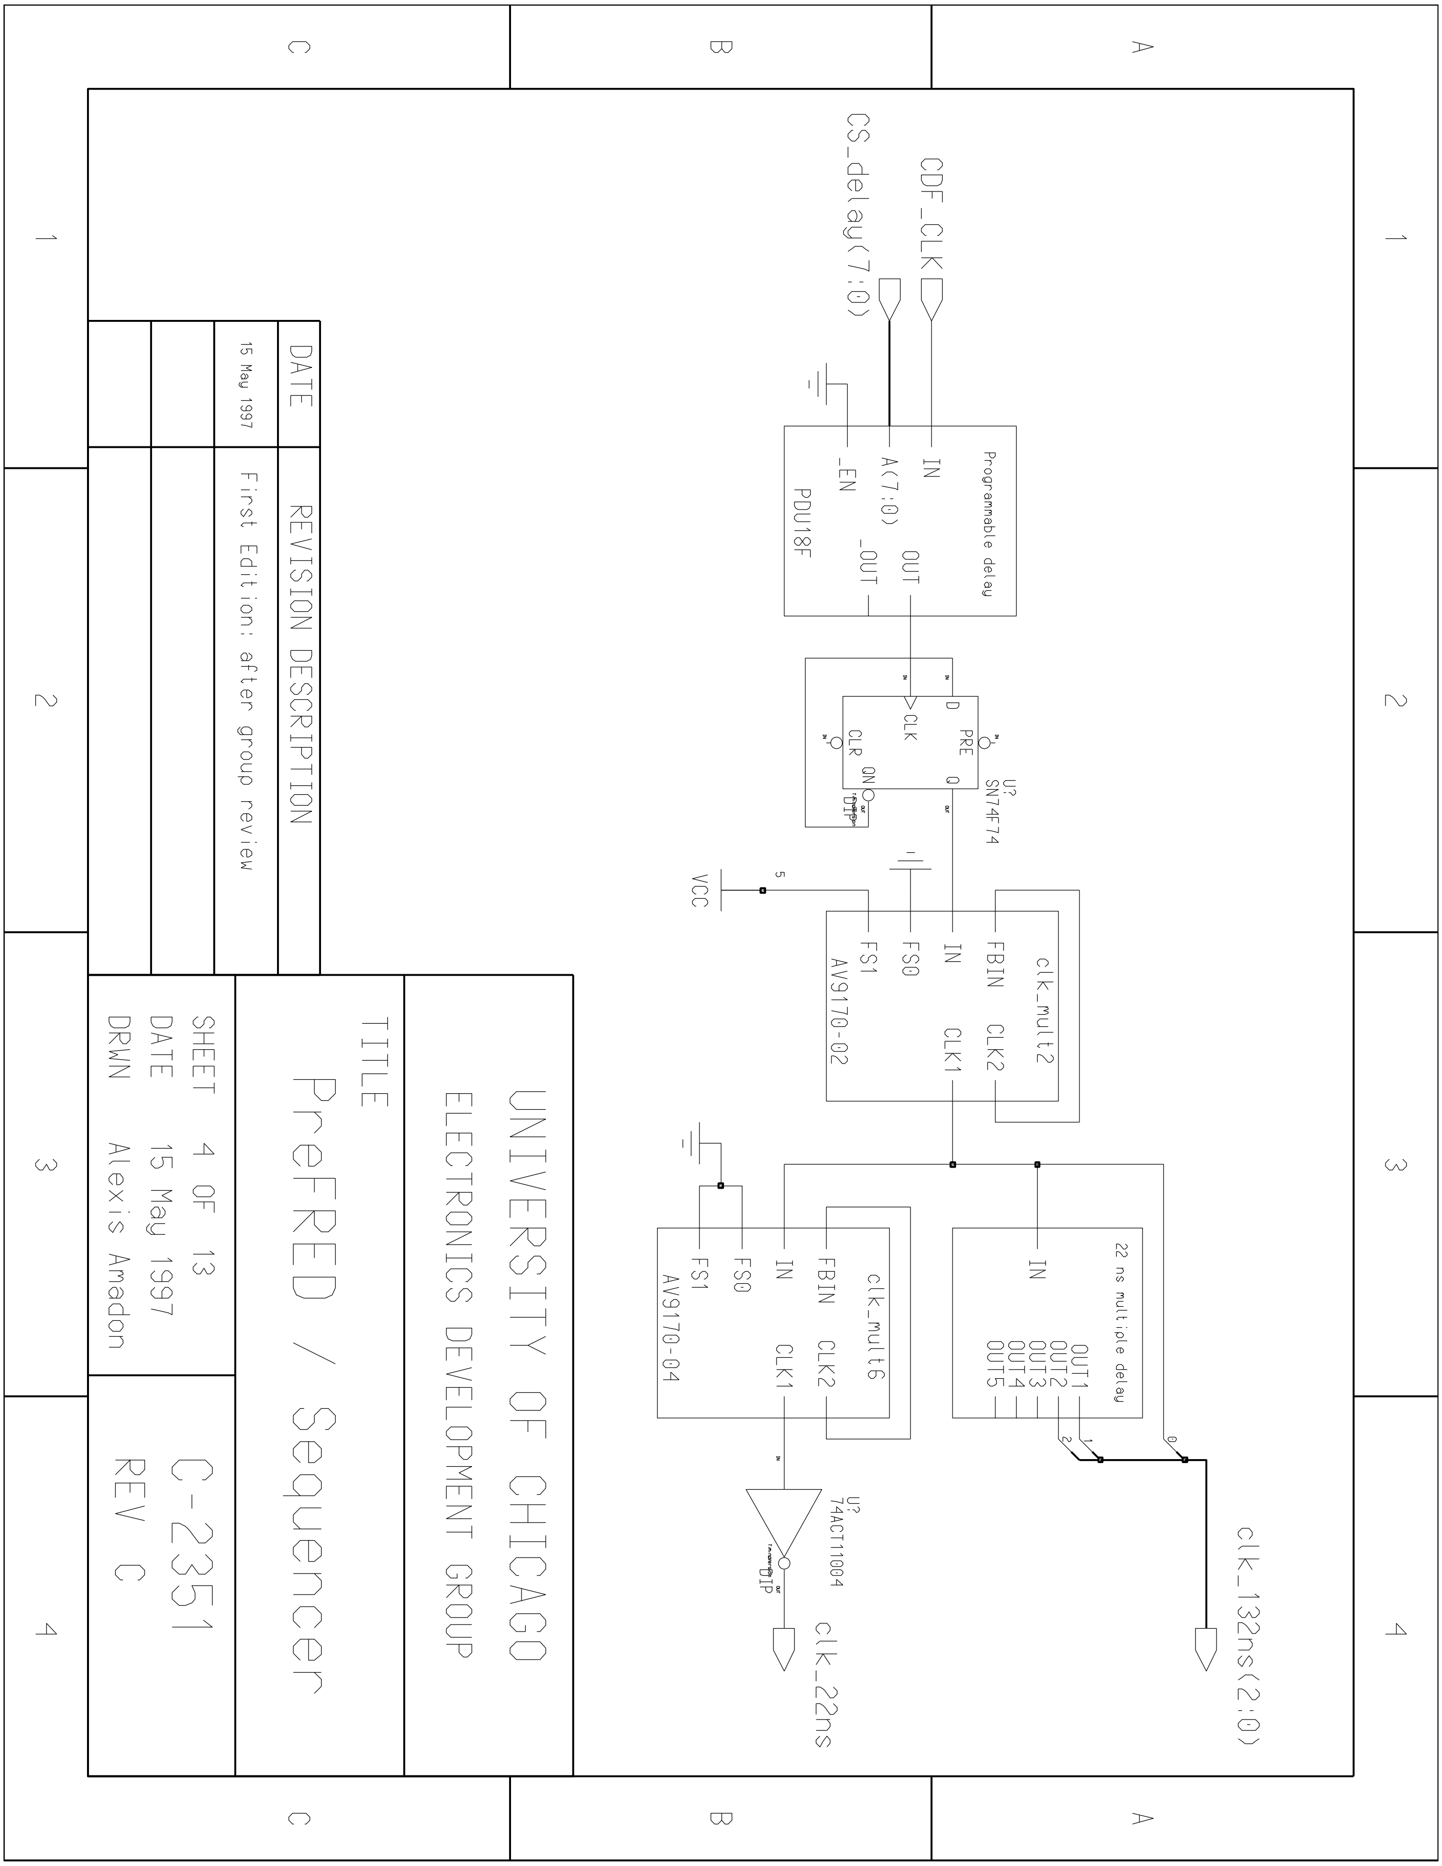

A
B
C
1
2
3
4
CDF_CLK
CS_delay(7:0)
Programmable delay
IN
A(7:0)
_EN
OUT
_OUT
PDU18F
D
CLK
PRE
CLR
Q
QN
U?
SN74F74
VCC
5
clk_mult2
FBIN
IN
FS0
FS1
CLK2
CLK1
AV9170-02
22 ns multiple delay
IN
OUT1
OUT2
OUT3
OUT4
OUT5
clk_mult6
FBIN
IN
FS0
FS1
CLK2
CLK1
AV9170-04
U?
74ACT11004
clk_132ns(2:0)
clk_22ns
DATE
REVISION DESCRIPTION
15 May 1997
First Edition: after group review
UNIVERSITY OF CHICAGO
ELECTRONICS DEVELOPMENT GROUP
TITLE
PreFRED / Sequencer
SHEET 4 OF 13
DATE 15 May 1997
DRWN Alexis Amadon
C-2351
REV C