More Queueing Simulation of the Silicon Track Trigger DRAFT

S.L. Linn
Florida State University
(November 6, 1999)

Abstract

A new model of the Silicon Track Trigger (STT) was implemented in the Ptolemy simulation package. The system dead time was found to be negligible, so not bottleneck exhists for the shared PCI bus.

I. STT SPECIFICATION

The model simulated follows U. Heintz "Specification of L2STT system for queueing simulations". Figure 1 shows the schematic for one sextant(crate) of the simulation.

A. Parameters

The average values of the system parameters are given by:

- N_t = 2(Hyper-exponential): number of L1CTT tracks per crate(sextant), maximum=32.
- $N_h = 36$(Double Gaussian): hits per 'detector', maximum=90.
- $H = 3.6 \pm 2.8$(Gaussian): hits per cluster (Silvia).

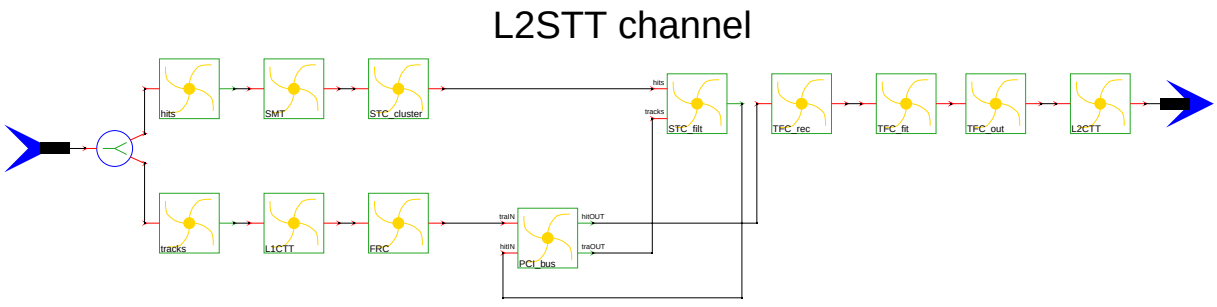

FIG. 1. Schematic of one sextant of the STT.

- $C = 3.7 \pm 1.2$ (Exponential): average number of clusters in a road.
- $W = 2$: number of 32 bit words/track

B. L1 - 0 event buffer

- - Poisson: $\mu = 100(\mu s)$
- - Minimum time between events: $7(\mu s)$

C. L1CTT - 0 event buffer

- - Format: $\Delta = 0\mu s$
- - Transmit to FRC: $\Delta = 2(10 + N_t)/53(\mu s)$

D. FRC - 0 event buffer

- - Receive from L1CTT: $\Delta = 0(\mu s)$
- - Transmit to STC: $\Delta = (10 + N_t)/26(\mu s)$

E. STC(roads) - 0 event buffer

- - Receive from FRC(PCI): $\Delta = (12 + N_t)/33(\mu s)$
- - Lookup and load: $\Delta = 0(\mu s)$

F. SMT - 0 event buffer

- - Digitize: $\Delta = 3(\mu s)$
- - Transmit to STC: $\Delta = 2N_h/53(\mu s)$

G. STC(cluster) - 0 event buffer

- - Recieve from SMT: $\Delta = 0(\mu s)$
- - Clustering: $\Delta = 0.1(\mu s)$

H. STC(filter) - 0 event buffer

- - Process all clusters: $\Delta = N_h H/26(\mu s)$.

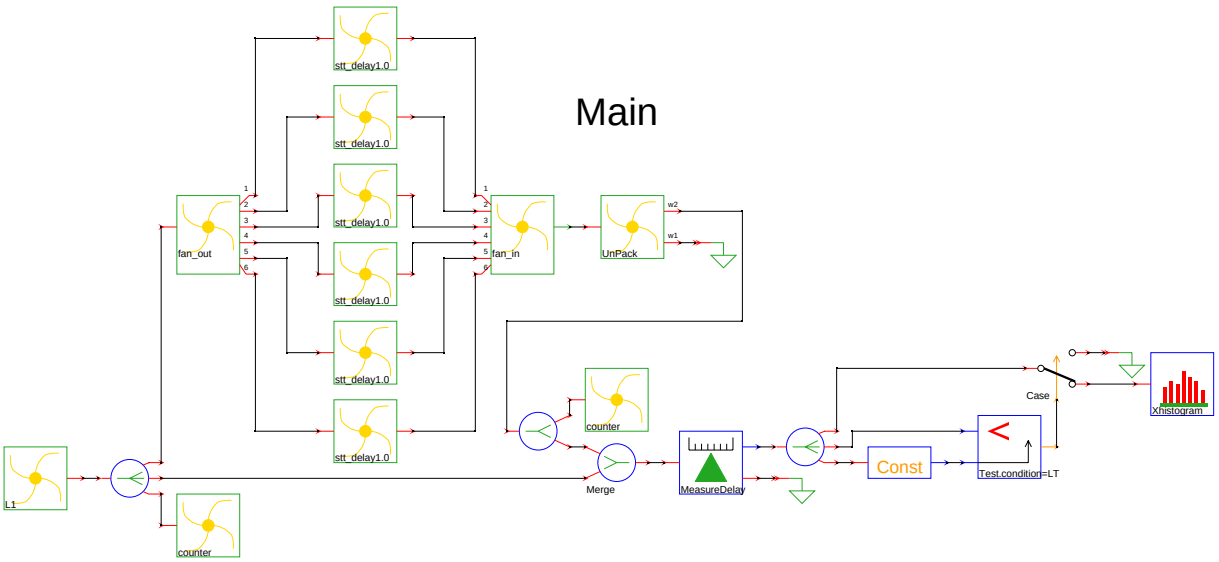

FIG. 2. Schematic of the main STT simulation.

I. STC(output) 1 event buffer

- - Transmit clusters in roads to TFC: $\Delta = N_t T/26(\mu s)$

J. TFC - 16 event buffer

- - Recieve cluster: $\Delta = (2 + 5 + N_t T)/33(\mu s)$
- - Look up constants: $\Delta = 0$
- - Fit tracks: Δ is given by the distribution in Fig. 8.
- - Transfer tracks to L2CTT: $\Delta = (10 + N_t + N_t W)/4(\mu s)$

K. L2CTT - 16 event buffer

- - Recieve track data: $\Delta = 0$
- - Format and sort: $\Delta = 0$
- - Transmit data: $\Delta = 0$

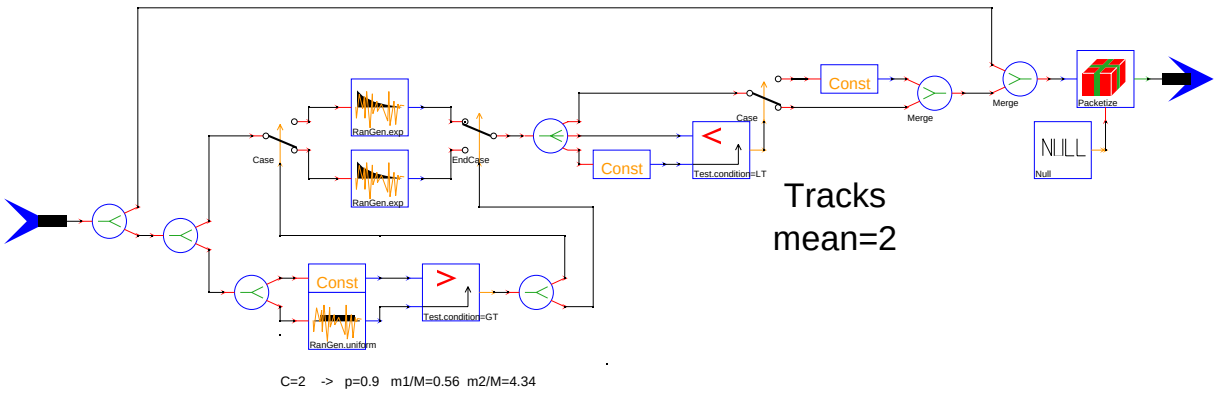

FIG. 3. Schematic of the hyper-exponential track generation galaxy.

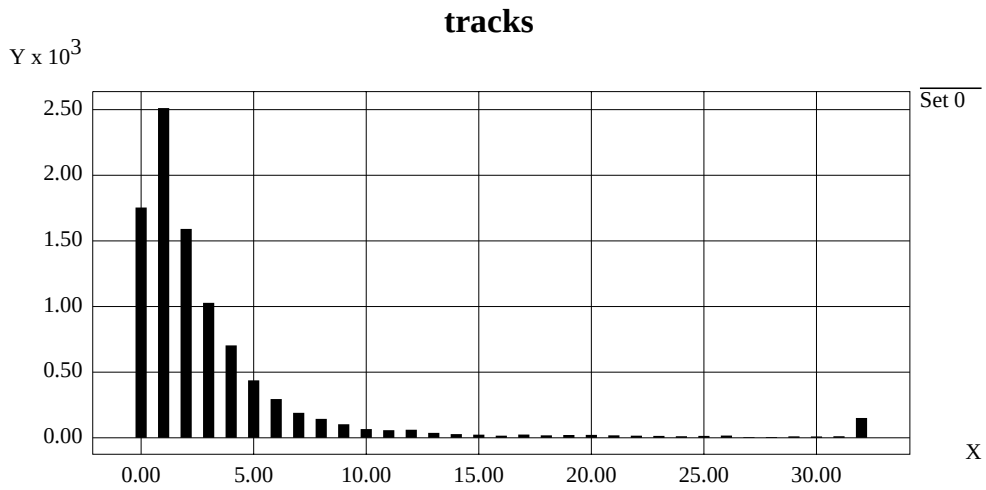

FIG. 4. Track multiplicity distribution

II. THE MODEL

Figure 2 shows the main schematic of the simulation. The STT is modeled as six independent sectors, where delay times for a single event depend on the sextant which takes the longest time to process. Dead times are taken as the ratio of lost to total events, and an event is lost when it encounters a busy processing element with a full queue. Total system latency is measured relative to an undelayed branch in the system. The L1 galaxy generates Poisson distributed events with an average time separation of 100μs, with a minimum time between events of 7μs.

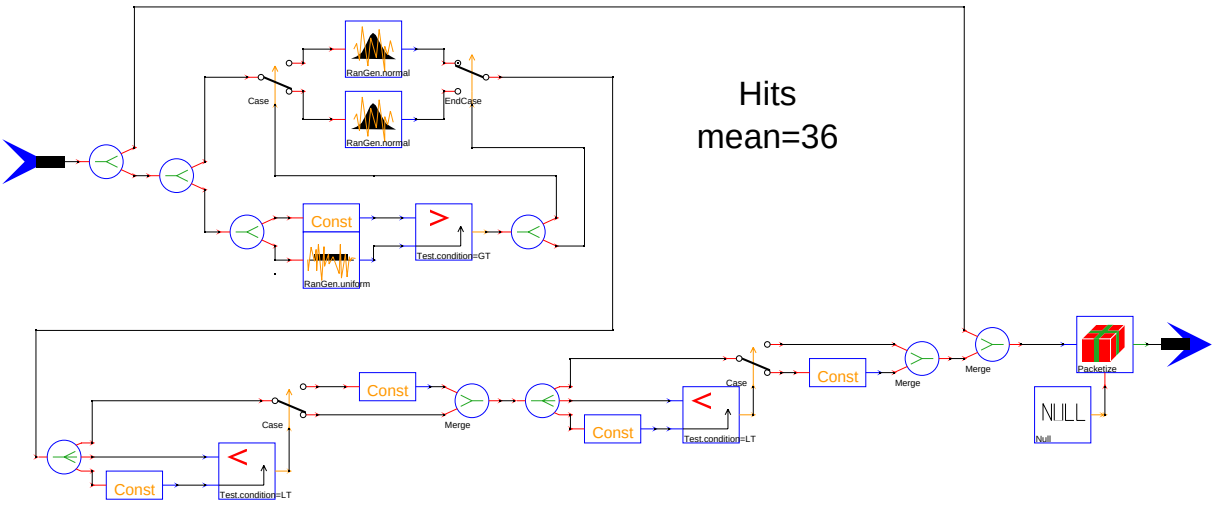

FIG. 5. Schematic of the double-Gaussian hit generation galaxy.

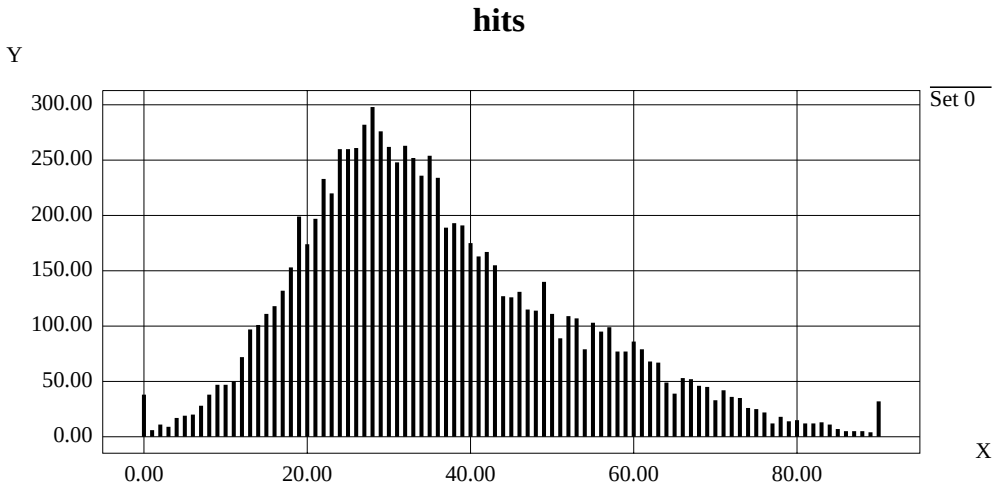

FIG. 6. Hit multiplicity distribution

Each event carries a time stamp and a unique event number. After a L1 trigger, the event is fanned out to the sextants, the number of L1CTT tracks (N_t) is generated by a the hyper-exponential galaxy shown in Fig.3, and the distribution is dhown in Fig.4. The number of SMT hits (N_h) for each sextant is generated a the double-Gaussian galaxy shown in Fig.5, and the resulting distribution is shown in Fig.6. Both N_t and N_h are packetized and carried with the event so that each module can use these quantities for delay calculations which depends on event size. This feature allows the correct treatment of correlations throughout the chain. For most elements, the simulation uses a queue and serve galaxy like the schematic

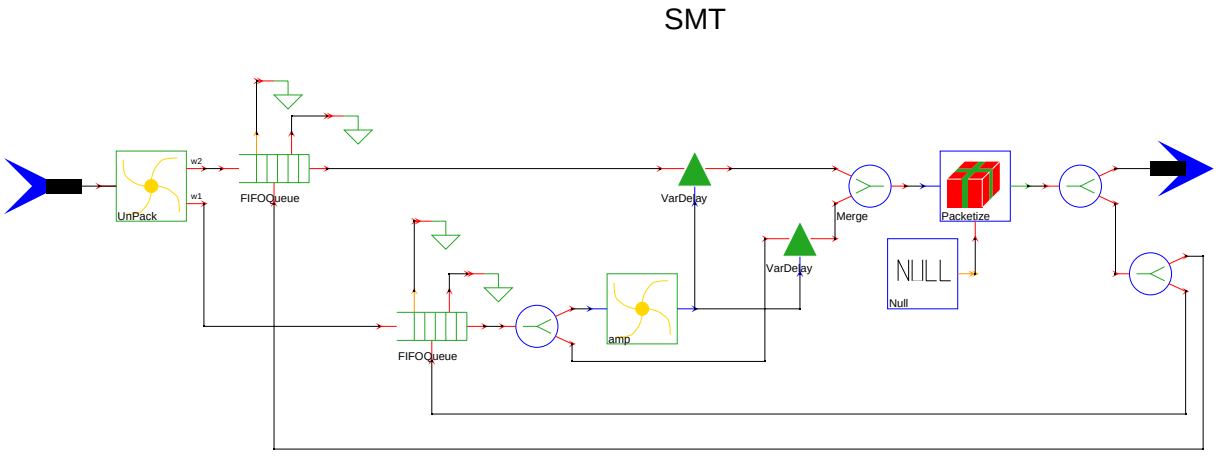

FIG. 7. Schematic of the SMT processor galaxy.

Module	Offset	Gain	Buffer	Max	Old
L1CTT	0.377	0.038	0	1.6	
FRC	0.385	0.038	0	1.6	
STC(roads*)	0.364	0.030	0	1.3	
			0	4.5	0.0
SMT	3.000	0.030	0	6.4	3.0
STC(cluster)	0.100	0.000	0	0.1	-
STC(filter)	0.000	0.011	1	1.0	-
				1.1	2.0
STC(output*)	0.000	0.154	0	4.9	-
TFC(receive)	0.212	0.121	-1	4.1	-
				9.0	8.0
TFC(fit)	-	-	16	-	-
TFC(output)	2.500	0.750	-1	26.5	7.0
L2CTT	0.000	0.000	16	-	

TABLE I. Table of simulation parameters for new and old simulations. Left to right are the module name, offset and gain, buffer size, maximum delay, and old average delay.

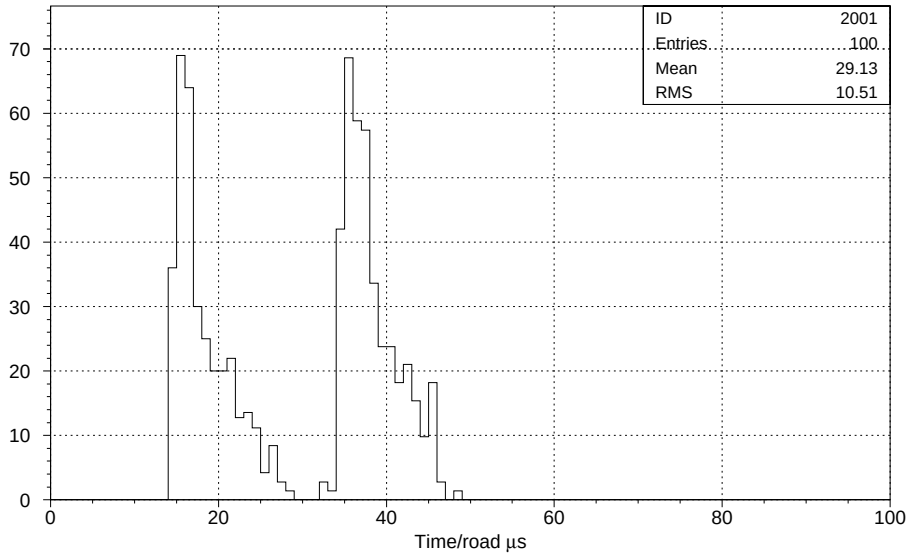

FIG. 8. Fitting time for one TFC processor.

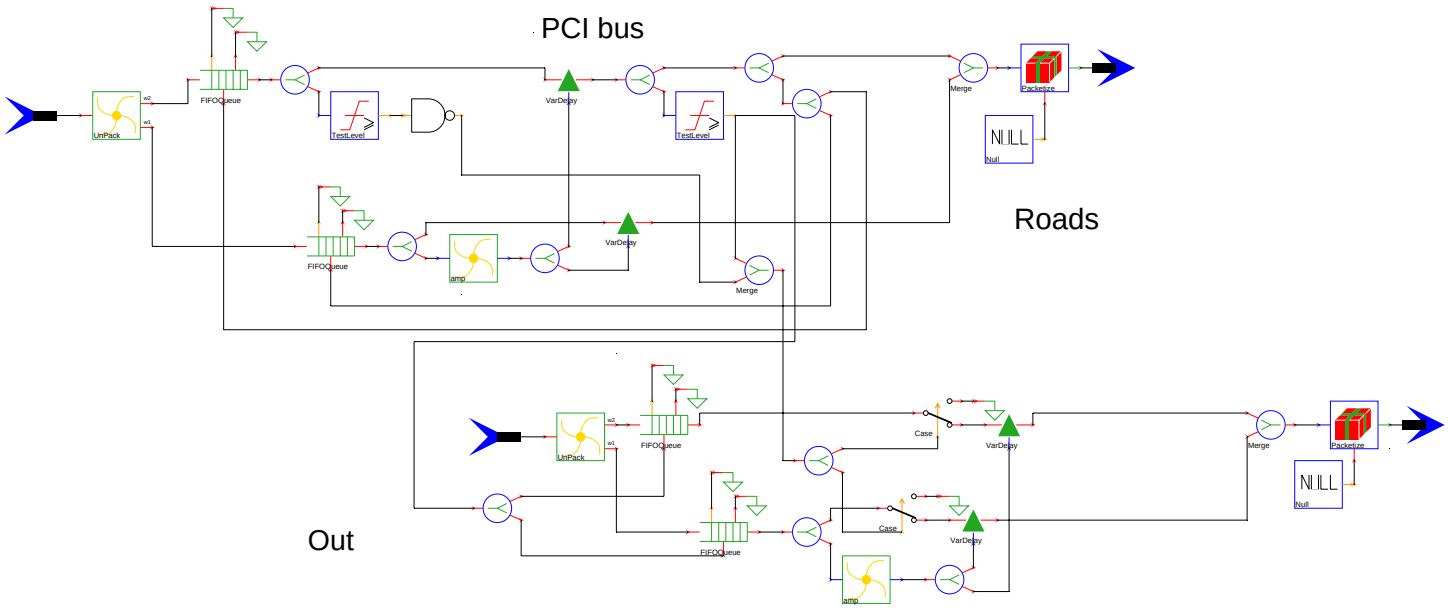

FIG. 9. Schematic of the PCI bus galaxy showing delay chains for the STC roads, STC output, and bus logic.

shown in Fig.7. In this galaxy the queue can be of any size, including zero. The galaxy inside each module contains an 'amplifier' that multiplies N_t or N_h by a 'gain' and adds and 'offset' to give the delay time. Table I gives the module, gain, offset, and delay for the maximum values of N_t or N_h.

A third random number arises from the track fitting time which is taken from the double exponential distribution as shown in Fig.8(Hobbs).

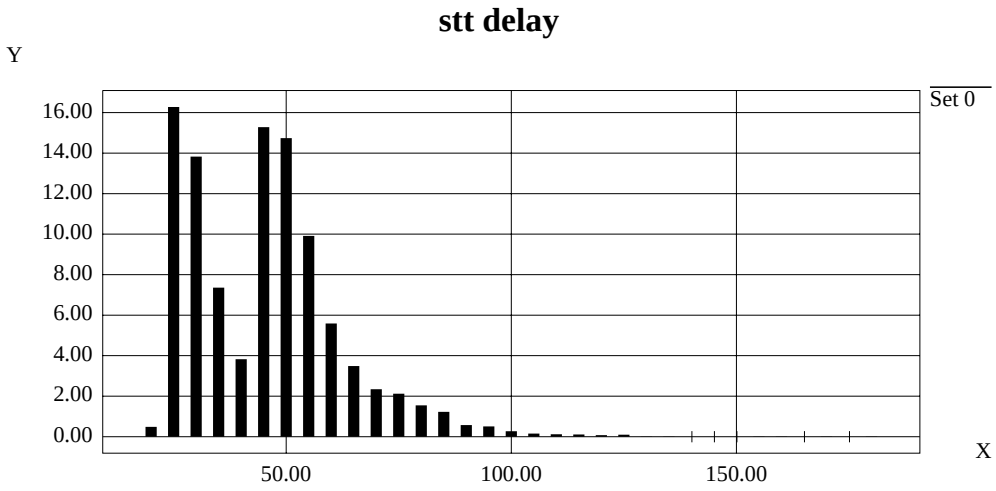

FIG. 10. Latency of one STT sextant.

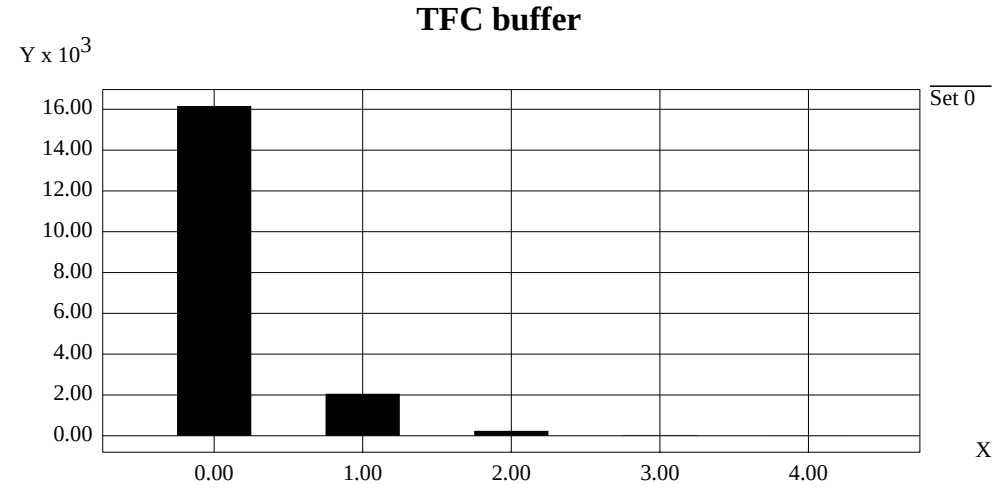

FIG. 11. Occupancy of the TFC buffer.

The transfer of FRC roads to the STC and the transfer of filtered clusters to the TFC occur through a shared PCI bus. A bus protocal was implemented using the schematic shown in Fig.9. The protocal gives priority to the road transfers.

III. RESULTS

The latency for one STT sextant is shown in Fig.10, where the track fit delay structure is reproduced. The full system latency is a 'Worst-of-N' convolution of this distribution with the five other sextants. The total dead time was found to less than one percent for the nominal values. This differs from the previous result for several reasons. First, the 7μs

minimum between L1 accept events is not counted in this dead time because it it common to the full L2 system. Subsequent delays less than 7μs cause no dead time. Finally the track fiting time is always less than 50μs and can be handled by the 16 deep buffer before it. Figure 11 shows that this buffer rarely exceeds one event and is never filled to capacity. Table I also shows the maximum delay for each module for $N_t = 32$ and $N_h = 90$. The current model is grouped so that old and new delay modules are comparable . Even when these extreme values are used, the dead time is less than one percent, and a buffer after the track fit module is not necessary.