

Recent Queuing Simulations of the L2STTpp

S.L. Linn – Florida State
10 Apr 02

Abstract

Recent timing measurements of TFC I/O and processing times were incorporated into the STT queuing simulation. Latencies of up to $250\mu\text{s}$ were found with no dead time incurred for luminosity conditions four times higher than expected in run 2b.

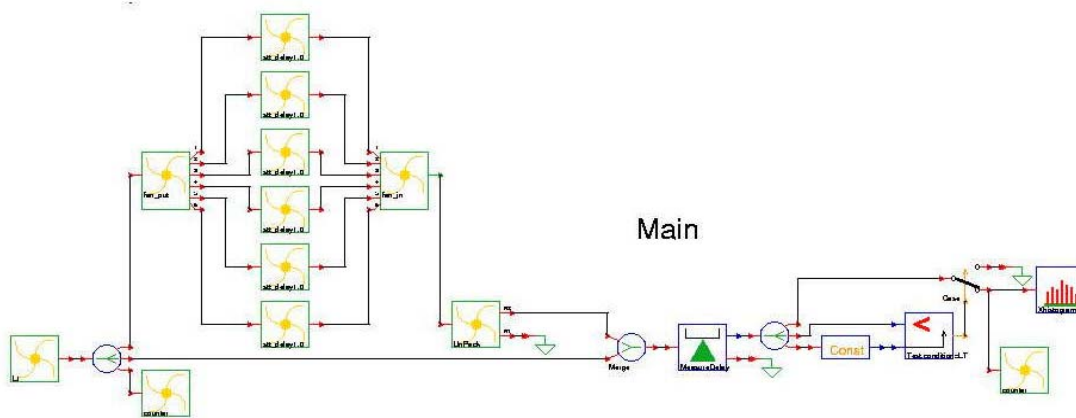


Figure 1. Schematic of the STT queuing simulation showing the six independent sectors.

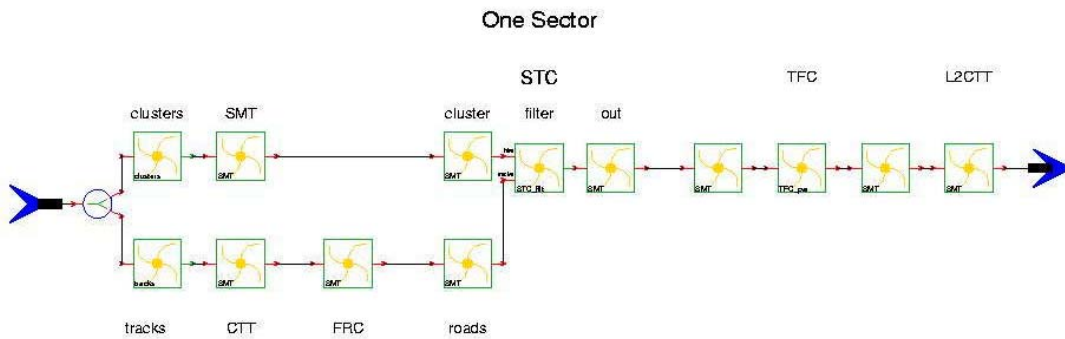


Figure 2. Schematic of one sector of the STT queuing simulation.

1. Queuing Model

The STT is modeled as six independent sectors as shown in Fig.1. Each sector is modeled as a chain of processors with variable delays and buffers (queues) as shown in the Fig.2. Each sector randomly selects the number of tracks/sector and the number of clusters/STC channel. The track distribution shown in Fig.3 is given by a double exponential and the clusters distribution is given by a double Gaussian. Events are generated with a Poisson distribution with a mean time between events of $100\mu\text{s}$ and a minimum time between events of $9\mu\text{s}$. Dead time inefficiency is computed as the ratio of input to output events. Events can also be lost when buffer limits are exceeded or processing time exceeds $500\mu\text{s}$.

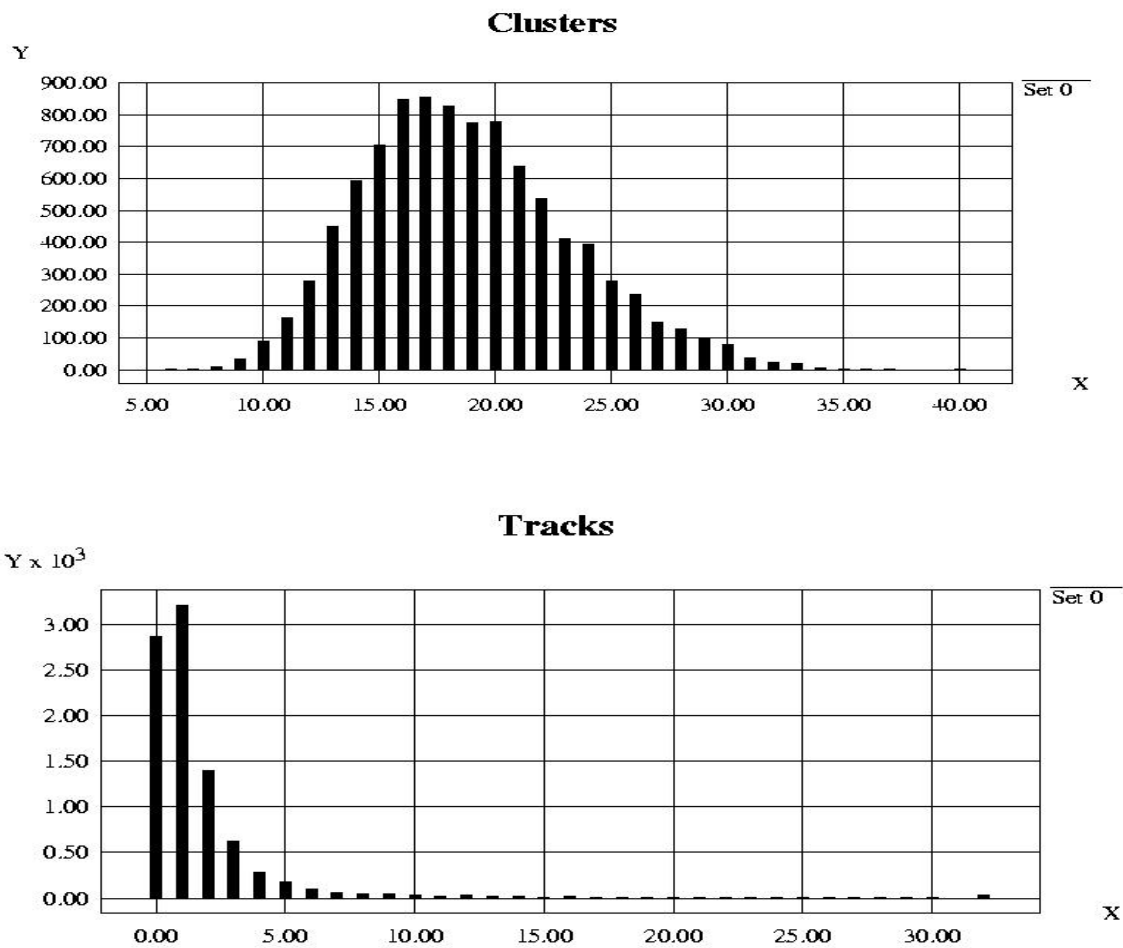


Figure 3. Top-Clusters distribution modeled as a double Gaussian distribution. Bottom-Track distribution modeled as a hyper-exponential distribution (mean=1).

2. TFC Model

Due to limitations of the queuing simulation, it is not possible to model individual tracks within the context of the full STT simulation, which uses the number of tracks and clusters as attributes of an event. For this reason a separate simulation of the TFC was done to determine a parametric function to describe the TFC delay as a function of the number of clusters and tracks. The part of the TFC between the input and output DPM's buffers was simulated separately, so that individual track processing times could be measured. Two quasi-independent banks of four DSP's were modeled to include XBUS transfer times and processing. Processing times were taken from recent measurements (Hobbs and Pancake, 4 Apr 02). Odd and even numbered tracks were processed by separate DSP banks respectively.

To simplify the calculation the number of fits (N_F) was assumed to be 1.5 resulting in a processing time for a single road (t_2) given by

$$t_2 (\mu s) = 23.13 + 0.58 N_{HR}$$

If $N_{HR}=10$, the following function can be used to describe the processing as a function of the number of tracks (N_R),

$$t(\mu s) = 27.7 (m+1) + 0.6 (8m - N_R), \text{ where } m = \text{int}[(N_R - 1)/8].$$

The resulting distribution is a step-like function which was approximated by a step function with processing times of : $26\mu s$ ($N_R=1-8$), $54\mu s$ ($N_R=9-16$), $81\mu s$ ($N_R=17-24$), and $109\mu s$ ($N_R=25-32$). Smearing of the processing times due to a variable N_{HR} and N_F was also studied; however, the variation was found to be less significant than variations due to N_R . The function used is shown in Fig 4.

For the full simulation of the STT, the simple step parameterization was used with I/O times per road (t_1, t_2) of

$$\begin{aligned} t_1(\mu s) &= 3.31 + 0.03 N_{HR} \text{ (event buffer to IDPM) ,} \\ t_3(\mu s) &= 1.70 + 0.06 N_{HR} \text{ (ODPM to Hotlink buffer).} \end{aligned}$$

3. Results

For the processing times associated with clustering finding and road processing, the default timing estimates were used. Total processing latency was measured for an average number of tracks/sextant of 1, 2, and 4. In all but the last case, the dead time was zero, and buffers were always less than four events deep. In the last case of four tracks/sextant, a 0.1% dead time was measured, but this is very sensitive to the exact shape of the input distribution, which was always truncated at 32 tracks. The latency distribution for $\langle N_R \rangle \geq 4$ is shown in Fig.5 where maximum latencies of about $250\mu s$ are rarely exceeded. The four peaks correspond to events with less than 8, 16, 24, and 32

tracks, and the tails and fill-in correspond to events that were buffered while waiting for processing.

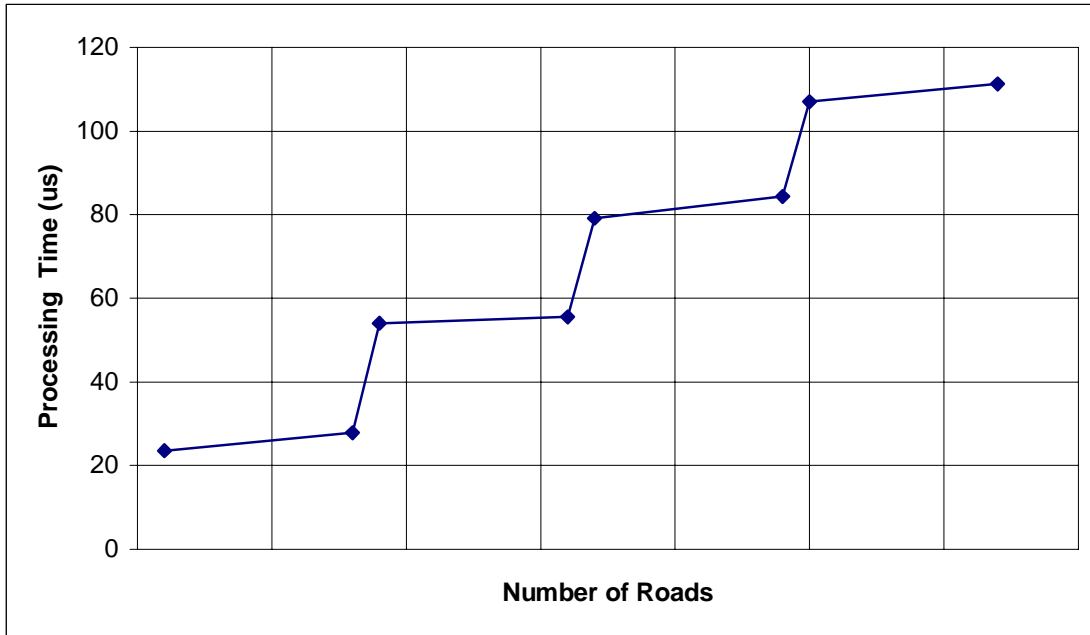


Figure 4. Processing time vs. Number of Roads

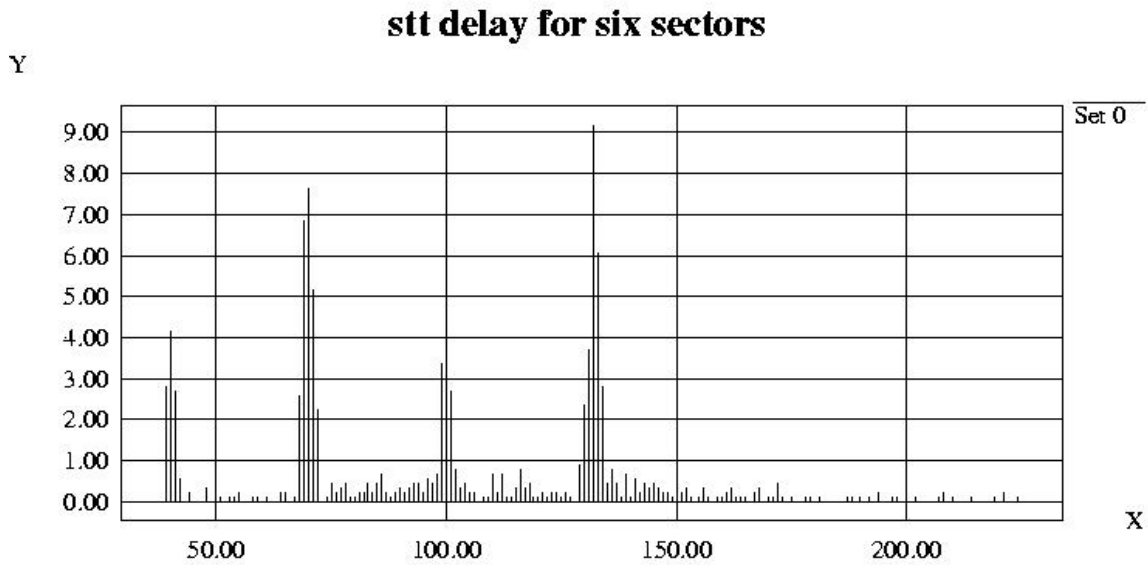


Figure 5. Events vs. total processing time(μ s) .