STAR Note 195:
An Estimate of STAR Computing Needs

D. Olson, M. Bloomer

LBL

8−Feb−1995

Introduction

This document describes an estimate of the computing resources that will be needed in order to carry out the data processing required for physics analysis of STAR data. First a model for the processing is described in Section 2. Sections 3 and 4 are an analysis of what is known about current data processing in EOS and STAR, respectively. Finally a projection to what STAR's real needs will be is described in Section 5.

Computing Model

For the purpose of estimating the computing resources needed for STAR the following computing model is used (see Figure 1). The sequence of data flow is:

1. The STAR DAQ writes raw data to tapes that get stored in the off−line tape warehouse.
2. The raw data tapes get staged for Pass1 processing (event reconstruction). The reconstructed events get written to the mass storage system.
3. The Pass2 processing reads the reconstructed events and produces DST's with identified particles suitable for physics analysis. These DST's are still all events.
4. The Pass3 processing reads a selection of the Pass2 DST's and generates smaller samples tailored to analysis, with filtered output, of Micro−DST's. These Micro−DST's are the main source of data for physics analysis. The Micro−DST data is written back to the mass storage system and, if necessary, written to tapes for shipping off−site.
5. Physics analysis on−site at BNL reads Micro−DST's from the mass storage system and stores results back to the mass storage system.
6. Physics analysis off−site (with insufficient network access) reads Micro−DST's from tape and generates results which also get stored in the mass storage system. Note that physics analysis off−site with good network access is considered the same as physics analysis on−site (5.).

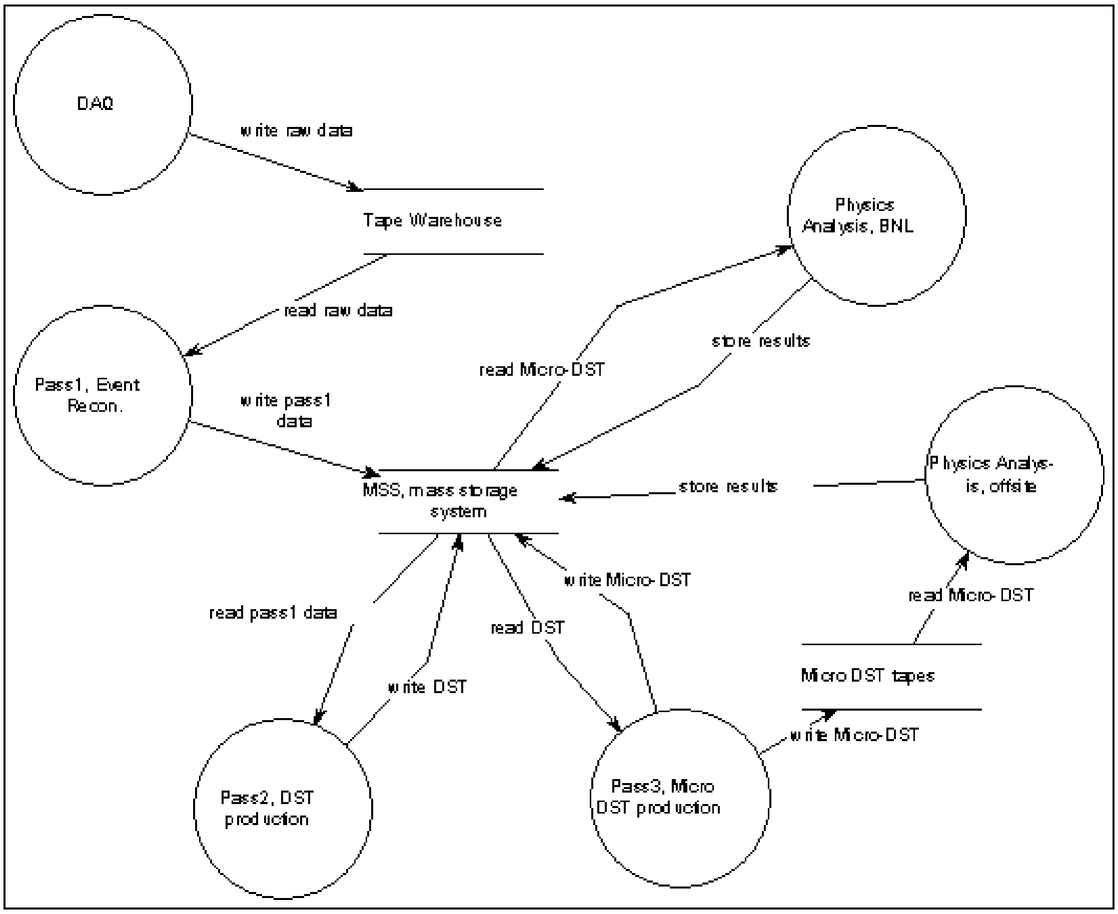

Figure 1

The elements of this figure are:

- DAQ − data acquisition writing raw data
- Tape warehouse − room where all off−line tapes are stored. Tapes from warehouse must be handled by an operator to be used for processing.
- MSS − the on−line mass storage system, something like Unitree.
- Pass1 − basic event reconstruction through track finding but not yet the best particle identification, applied to all events
- Pass2 − applying best calibrations to pass1 tracking to do particle ID and generate particles for physics analysis (DST production), applied to all events
- Pass3 − select events from DST's with filters on both the type of events and which subset of each event is saved in the micro−DST. This selection is generally run separately for each type physics analysis and with differing selection and filtering criteria.
- Micro−DST tapes − data tapes for physics analysis sent to remote sites, used by physicists without sufficient network access to the HSM at BNL.
- Physics analysis − processing micro−DST's and ntuples in order to generate physics histograms.

The data flow characteristics in this model are:

Estimates based upon EOS data analysis

The EOS data processing should be a fairly reliable estimate for what STAR will require. The EOS TPC is a very similar pad readout geometry, similar electronics, and comparable track density to what we expect for STAR.

The data shown below is from processing Au+Au at 1.2 A GeV. The cpu times are for a 55Mhz Hypersparc processor (65 SPECfp, 11 MFLOP).

Pass1

The first histogram below shows the amount of cpu time used for hit finding vs. multiplicity per event. It is clear that this stage is rather insensitive to the multiplicity.

The second histogram shows the relative cpu time used by the various modules for event reconstruction analysis used for EOS. The dominating part is the track reconstruction, which includes particle ID and V0 reconstruction.

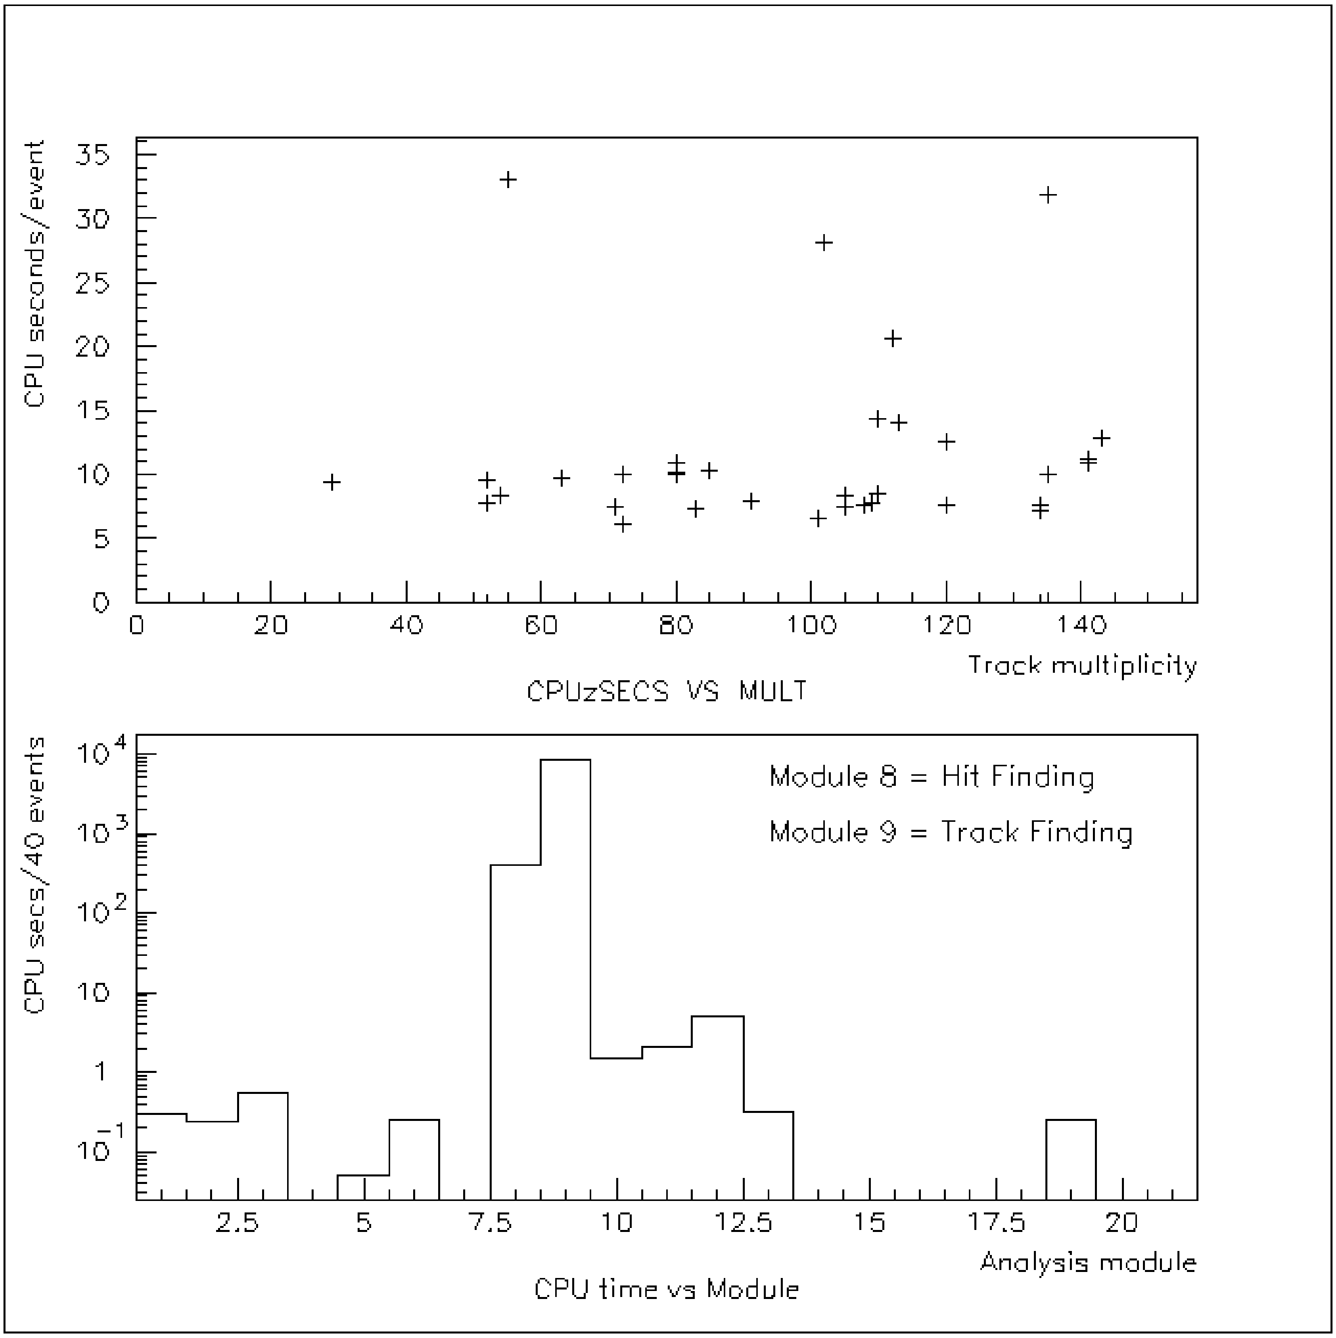

Figure 2

Figure 3.

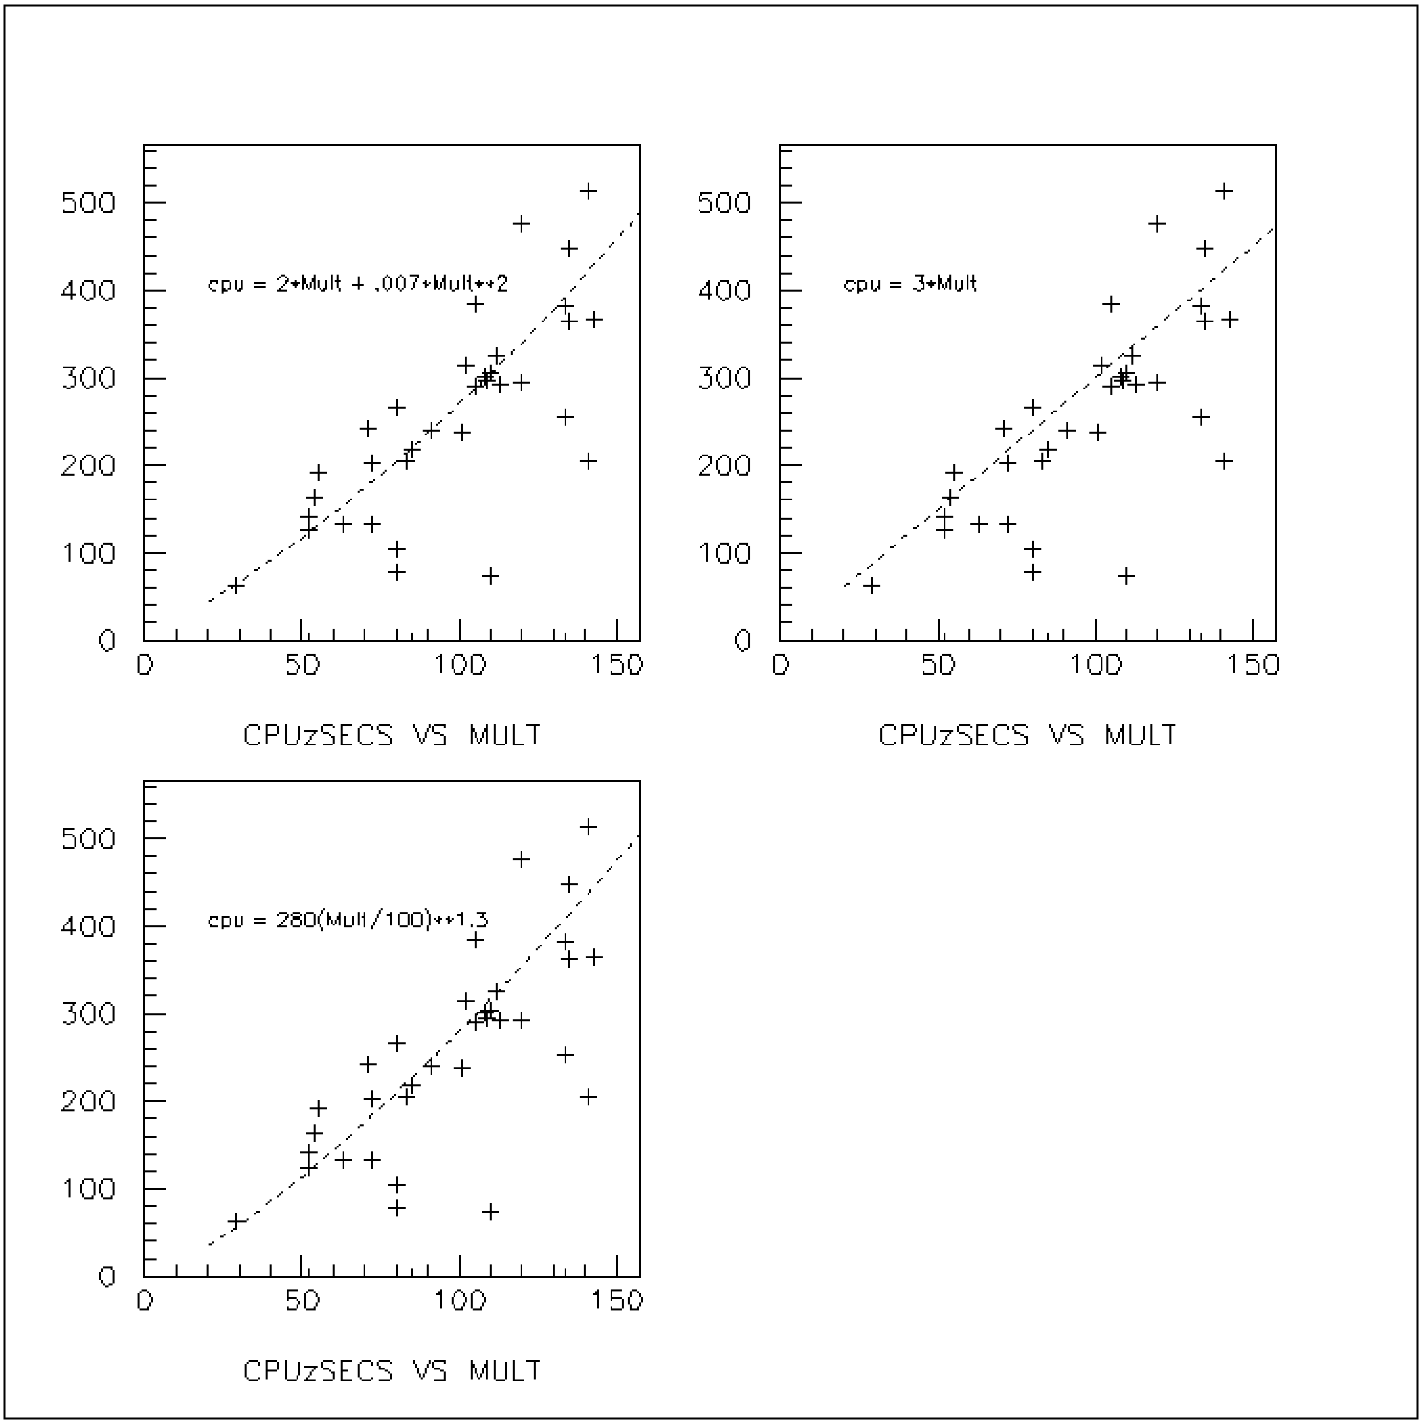

Figure 3 shows a sample of CPU time vs. reconstructed track multiplicity for EOS TPC data processing. Also shown are three different parameterizations of the time to reconstruct an event.

The experience of EOS indicates that as the tracking evolves from the initial track finding to what is required to fix broken tracks, refitting with the vertex for best momentum resolution and other maturation improvements, one can expect the tracking time to increase about a factor of 10.

Current STAR processing

At present the STAR event reconstruction is at a rather immature stage. The tracking software is reconstructing tracks but all the problems of optimization for matching broken tracks, correcting for "real data" distortions and impurities, and properly reconstructing the decayed particles is not included. However, the cpu times at this stage are:

Task	cpu seconds (R3000, 70SPECfp, 12MFLOP)
TPC cluster recon	2540 s
SVT cluster recon	~850 s (estimate, not measurement)
TPC tracking	300 s (rejecting outliers and refitting)
SVT tracking	~480 s (multiple passes assumed)
SVT−TPC matching	150 s
TPC pid	30 s

This comes to a total of 4350 seconds (52 GFLOP−seconds).

Extrapolation to "real" STAR data processing

From EOS

Extrapolating the three parameterizations of EOS tracking to 3000 reconstructed tracks for STAR yields GFLOP−secs of:

- 100 GFLOP−sec/event, linear (33 NT MFLOP−sec/event)
- 256 GFLOP−sec/event, power law (3080 (NT /100)1.3 MFOP−sec/event)
- 760 GFLOP−sec/event, polynomial (22 NT + NT 2 /13 MFLOP−sec/event)

Of these three methods of extrapolation we expect that the power law method is the best. A linear dependence is too low since there are many aspects of tracking and physics analysis algorithms that loop over combinations of tracks. The polynomial method is probably too large since for most algorithms one does not need to loop over all combinations of tracks in the whole event but rather all tracks in some region of phase space.

From current STAR

The only extrapolation one can do with the current STAR data processing numbers is to state that as the code matures and becomes both more accurate and reliable that the cpu time required will increase. Gauging from the EOS experience, the increase in time may be a factor of 5 to 10.

Operational considerations and other factors

There are numerous other factors which contribute to the total processing requirements. The most significant of these are:

- **simulations –** simulating events and detectors response at the raw data level is more time consuming than reconstructing the real data, however, the number of events that need to be simulated at this level are probably considerably fewer than all the real events. Allowing an additional 50% of cpu cycles for simulations should be sufficient.
- **Pass 2, 3 –** the additional passes through the data after event reconstruction add significantly to the

bandwidth requirements of the data processing but have considerably smaller cpu cycle requirements.
- **annual duty cycle** – the accelerator is expected to run 34 weeks/year. If the combined accelerator – STAR detector duty factor is 80% and STAR records 1 event/second, we will acquire 16M events/year, or an annual average of 0.5 event/second.
- **data processing efficiency** – the efficiency with which the data processing facility can be utilized will be less than 100% for many reasons, we will assume an 80% utilization.

Total cpu requirement

By combining all the factors above and choosing the 256 GFLOP–sec/event, which agrees with current EOS extrapolation and projected increase of present STAR processing, we get:

(256 GFLOP–sec/event) * (1.5 simulations) * (1/0.8 utilization of computing) * (0.5 event/sec annually)

which yields a computer facility capacity of 240 GFLOP for STAR data processing.