# MODELLING MICROPROCESSOR FARMS FOR SSC DATA ACQUISITION*

Dave Cutts

Brown University

Providence, Rhode Island

Cathy van Ingen

Fermilab

Batavia, Illinois

Data acquisition, involving both data collection and accompanying filtering, is a crucial part of an SSC experiment. The extreme demands on the acquisition system for I/O bandwidth and real-time computation capability that makes irrecoverable event selections requires that the system operation be well designed and fully understood. Modelling is a valuable and perhaps necessary tool for a successful design of such complex systems. As a demonstration of the modelling process we have made a few studies of simple systems consisting of a large number of microprocessors arranged in a "farm", a system in which each event is routed from the digitization/readout crates directly to a single microprocessor, one of a large parallel array, which performs the complete software event filter on that event. Subsequent events are transferred in turn to other processor nodes which are idle. The results of our crude simulations agree with the expected performance of such farms, and suggest that realistic modelling of SSC acquisitions will be important.

We have employed an IBM software package called RESQ ( for RESearch Queueing ) to study the operation of a microprocessor farm. The basic simulation imagines events each in turn being transferred via a single readout system to an available processor in the farm. After a variable analysis time the event is either rejected or readout (as to a host or recording device), and the processor is made available for another event. Figure 1 gives a schematic of this simple model. The symbols associated with the event buffers and the CPU farm indicate "passive queues", for which "tokens" (in this case buffers or CPUs) can be allocated and released. On the other hand the symbols for the readin (data cables to dual-port memories, say), the filter algorithm within a CPU, and the readout from the CPU refer to "active queues" with which are associated variable "service times".

The simulations depend on a number of parameters which are readily varied, including

1. number of CPUs in the farm
2. mean analysis time per event
3. event frequency
4. event length
5. event rejection probability
6. input bandwidth, given as number of parallel 40 MByte data cables

With particular values of these parameters a number of runs of the RESQ package were made to explore the gross features of acquisition farms suitable for the SSC environment.

The result of a study of acquisition farm performance as a function of the input bandwidth is illustrated in Figure 2. These simulations used a 1000 processor farm, each unit requiring 1 second on average for a filter analysis, and rejecting events with a .999 probability. The input rate was 500 events per second each of mean length 1 MByte. Plotted in Figure 2 are both the mean length of the input queue (number of events in buffers waiting to be readin) and the percentage utilization of the input system. Note that as the input bandwidth is decreased input utilization steadily rises but the input queue length stays low until at about 50 percent utilization. Beyond this point, however, the readin becomes quickly saturated. This result from our simple model agrees with the natural requirement that

(input bandwidth) > (event size) * (input rate)

As shown in Figure 2, the study suggests further, that "excess" bandwidth is important, presumably because of the variable event size and input frequency. Indeed, for this system one would like a total input bandwidth of 10**9 rather than the 5*10**8 given by the above rule.

Another study of a SSC processor farm assumed processors 5 times as powerful, so that the mean analysis time per event was 0.2 seconds, but with an input event trigger rate of 1000 Hz. As before, the average event length was 1 Mbyte, and the rejection

probability was .999. Rather than study the variation of input bandwidth, we fixed the input capacity at 1.4*10**9 (35 parallel 40 Mbyte cables) and varied the number of CPUs in the farm. As shown in Figure 3, the percentage utilization of the farm steadily rises as the number of CPUs is reduced. However, above about 75 percent utilization the system is frequently overloaded, as shown by the rapidly rising mean queue length per CPU. This feature illustrates the obvious rule that

(number of CPUs) > (events/sec. input) * (analysis time)

or for our particular system, that there be more than 200 CPUs. Again, the simulation highlights the need for capacity greater than the minimum, to keep deadtime small (which presumably occurs due to fluctuations such as in event analysis times).

These models of SSC processor farms are very simplistic, and certainly do not provide detailed design information. They do illustrate gross features of such farms that perhaps are self-evident in retrospect but are very important and worth exposing. These models and the simulations were thrown together and run in a few hours; certainly one could do a much more detailed and relevant design study with these or similiar tools. The crucial role played by the data acquisition system in SSC experiments would suggest that such studies are very important.

* Work supported in part by U.S. Department of Energy.

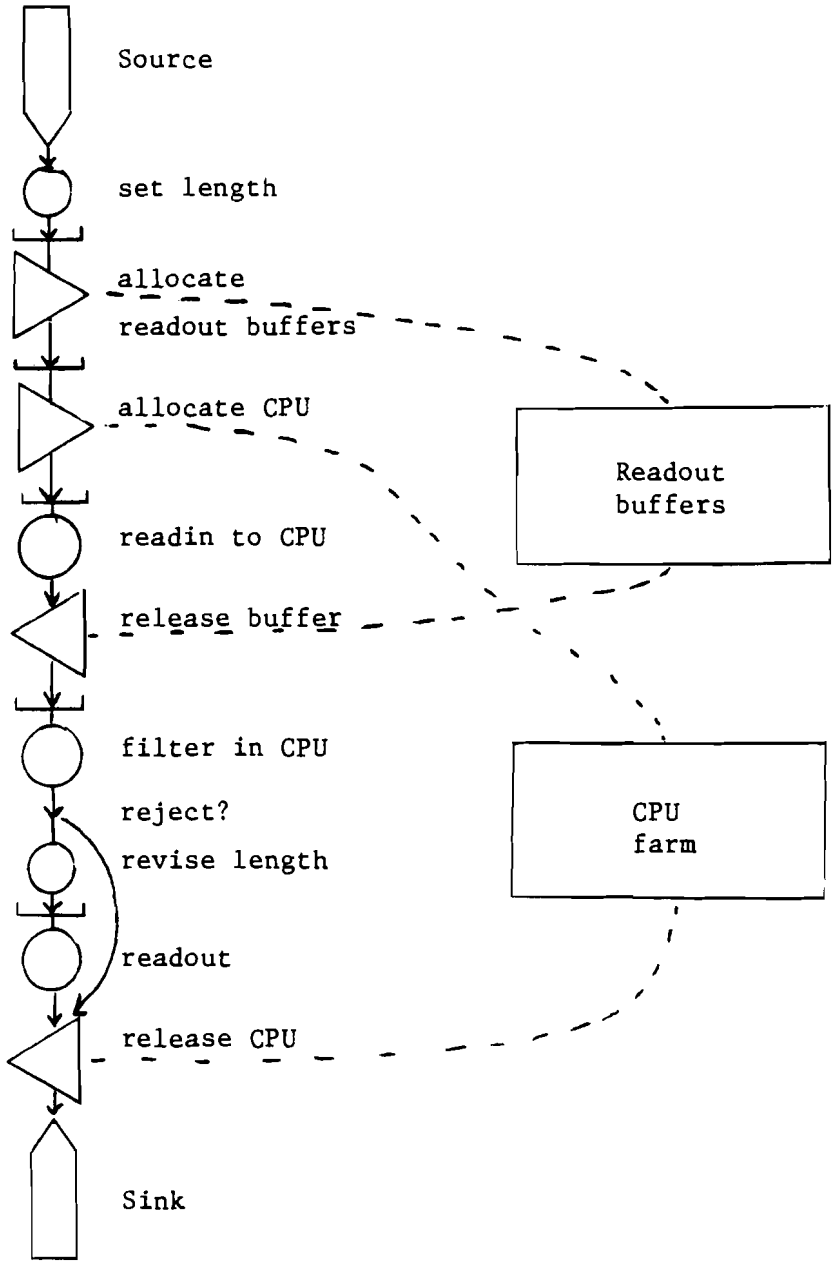


Figure 1. Model for a simulation of SSC farm

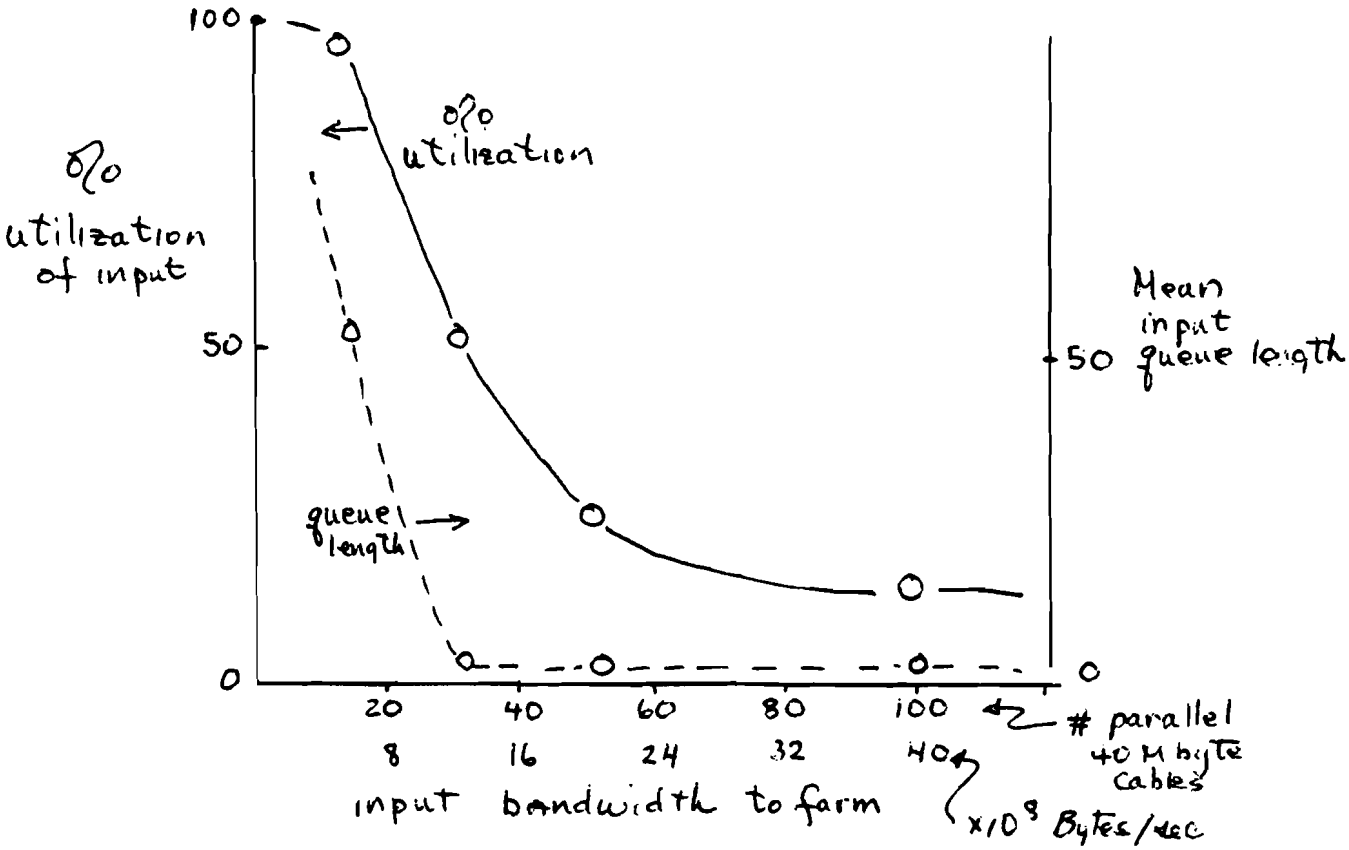


Figure 2. Performance as a function of input bandwidth.

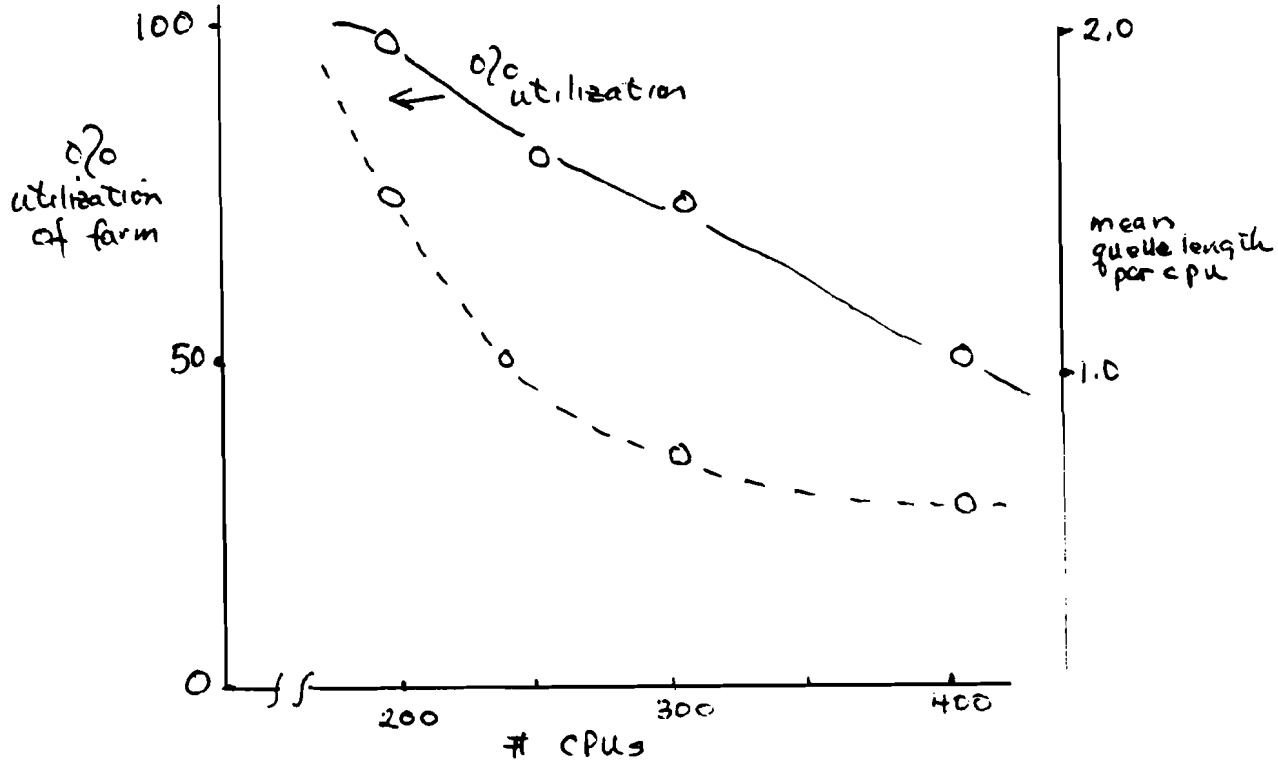


Figure 3. Performance as a function of farm size.