Status of MAP at Liverpool

T. Bowcock, M. McCubbin, A. Moreton, G. Patel
Liverpool University
C. Parkes
CERN

LHCB Software Week, November 1999

Overview

- What is MAP?
- Status
- Current usage
- Performance
- Future plans

What is MAP?

Monte Carlo Array Processor:

needed for generating the necessary large MC samples

300 processors: 400MHz PII, 128 Mbytes memory, 3 Gbytes disk

LINUX operating system with batch facility

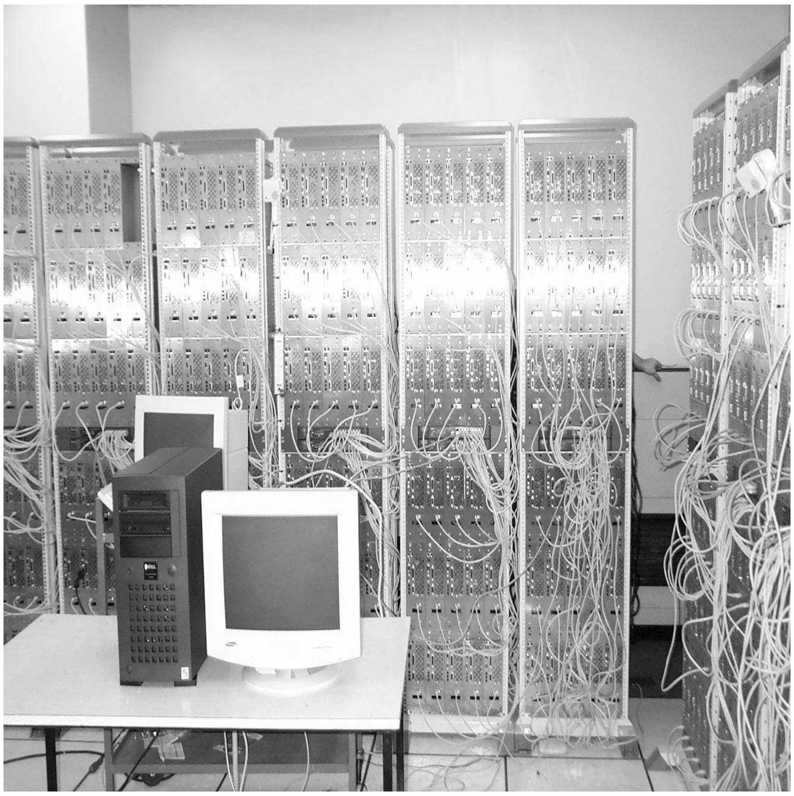

Status

- 150 of 300 processors have been commissioned

  of these, 5 down for various hardware reasons (may be fixable)

- extra air-conditioning Friday 19/11/99:

  commissioning of extra 150 processors imminent

- submit jobs one at a time (no queueing yet):

  command: mqjob jcf outdir

  jcf: job control file with name of directory of files to send to MAP (executable, cards and input files), number of processors to run on, names of files to be returned

  output comes back into directory outdir.output

  check on status of job: mapq

- reasonably reliable (and improving): regularly get output back from all 145 processors

Current usage

- VDET optimisation using SICB v220, using "standard" cards files for a number of "standard" physics channels

  include histos (based on AXACC) to estimate the primary and secondary vertex resolutions and the decay length resolution

  eventually use AXSEL to estimate event selection efficiencies and backgrounds

  different VDET designs via a set of vsrp.cdf files:

  number of stations = 18, 27 (with z positions scaled)

  silicon thickness = 150μm, 300μm and 500μm

  inner radius = 0.8cm, 1.0cm

  don't produce data at the moment, just keep histos ... maybe ntuples in near future?

- soon: a test generation of 5 million inclusive b events (type 360000) for Erica Polycarpo using SICB v219(?), producing ntuples

Performance

- ran SICB v116 $B_d^0 \to \pi^+\pi^-$ generation producing AXACC histograms over recent weekend

  175k events produced with 145/150 of the processors, 1200 events per processor ... no "underlying" events considered

  time distribution of completed jobs:

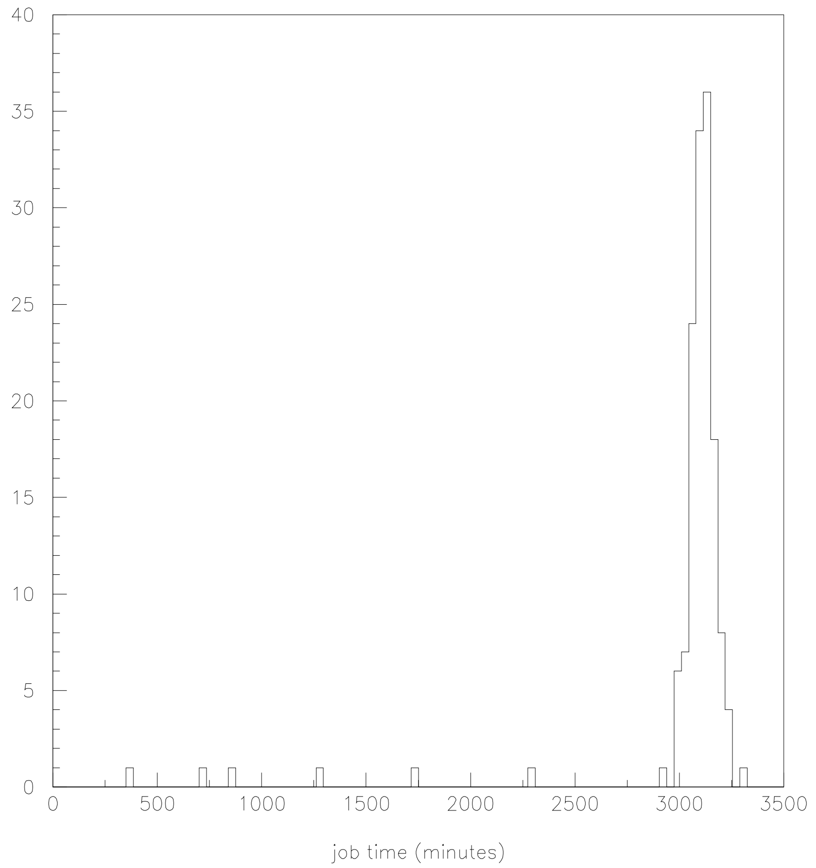

  reason for short jobs unknown, but they introduce inefficiencies

Performance (cont'd)

- main problems at the moment:

  getting working version of SICB v220 with our modifications

  I/O and non-random packet loss transferring files, causes non-negligible job failure rate

  user interface needs developing (queueing, priorities), possibly introducing Web-style PCSF job submission form

Future plans

- 4 Tbyte disk space trial scheme probable, early December for decision on funding, March 2000 for getting disks, June 2000 fully operational
- until enough disk space, not considering running external production and no customers requested

  carry on doing tasks such as VDET optimisation at Liverpool
- iron out the problems and fully implement the 300 processors

  timescale: of the order of 6 months