Realtime Astronomy Signal Processing in PC Clusters

D. Backer, P. Demorest, and R. Ramachandran
University of California, Berkeley

I. Stairs, and R. Ferdman
University of British Columbia

D. Nice
Princeton University

Wide bandwidth signal processors for pulsar and other science at a number of large telescopes around the globe are now based on computing in cpu clusters rather than custom hardware. This approach is driven by a number of features: (1) The processing algorithms are in software where they are easily accessible to change by the general user community for particular scientific applications. (2) The algorithmic software can be ported to clusters with ever more powerful cpus. (3) The data can be represented by many bits to provide accuracy of power measurements for time-variable signals.

A major challenge in this approach has been input of the digitized data into the cluster. The generic architecture involves a set of cpus acting as "data servers" with commercial digital I/O boards running at speeds up to 1 Gbs. The data servers pass data along with minor processing to a high speed switch that feeds data segments to the main cluster of "slave cpus", which run specialized algorithms. Modern data servers with appropriate mother board design can (just) robustly handle a 1 Gbs input and subsequent 1 Gbs output.

Cluster-based processors are now employed at Jodrell Bank (COBRA); Nancay (LBP); Parkes (CPSR); Green Bank Telescope (CGSR,GASP) and Arecibo (ASP). Each of these can handle realtime signal processing of bandwidths ranging from 64 to 128 MHz depending on unit and observing parameters. The principal application for all of these processors is the removal of dispersion in pulsar signals via Fourier filtering and subsequent detection and synchronous averaging modulo the apparent pulse period. Dispersion removal requires forward and reverse FFTs of lengths that often exceed the L2 cache; FFTW library is employed for optimum performance.

Our collaboration is involved with LBP, ASP & GASP, generically Astronomy Signal Processors. The ASPs employ 4 data servers with I/O card from Electronic Design Team, 16 slaves and a master cpu. Athlon (LBP) and Xeon (ASP,GASP) processors are used. Each data server handles 32-MHz of data; with 8b sampling this generates 1 Gbs flow. The Astronomy Real-Time Software (ARTS) package is very modular with extensive use of shared memory for temporary storage of data segments and intermediate data products. This approach allows critical "sampling" of the data for real-time assessment of various stages of processing, much as directional couplers and test-points are used in hardware. ARTS also allows buffering the full data rate to disks for offline processing. A primary application is observations where there is more than one pulsar in the beam. Another is where the slaves just can't keep up with dedispersion processing in real time. With our total of 4 Gbs we require use of slave disks. A potential development will be to port data to VLBI Mk5 units.