Remarks on Supervisor/RoI Builder

- Demo B
  - Measured Performance
  - Software Upgrades
- Demo C
  - Current Performance
  - Extrapolations
- Modelling Activities
- Plans for 1998

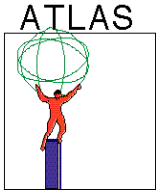

Supervisor/RoI Builder

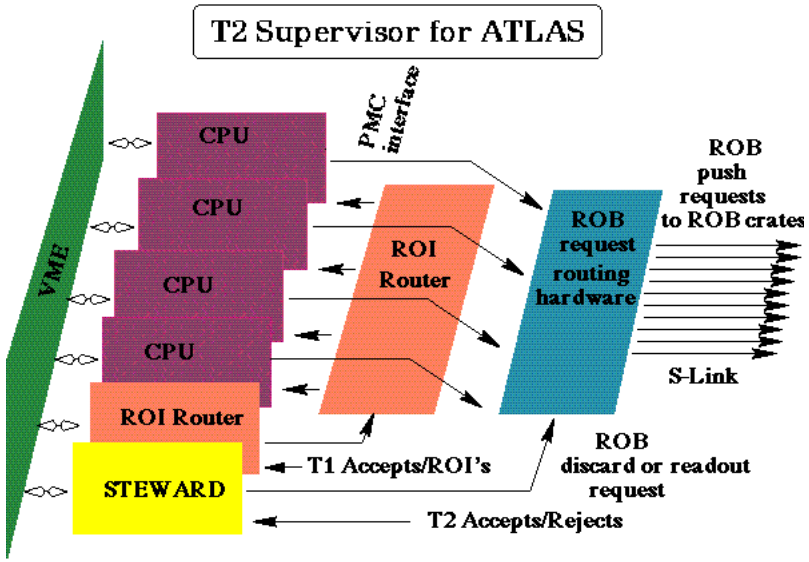

TASKS

STEWARD: Manages (frees and reviews usage of) local and global processors

ROI Router: dispatches ROI information to CPU's

CPU's: determine appropriate free local/global processors to use and determine buffers to request data from and assemble request packet

ROB Router: routes buffer requests to appropriate crates

ATLAS

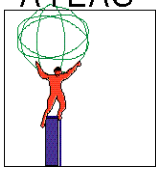

Demo B - SCI

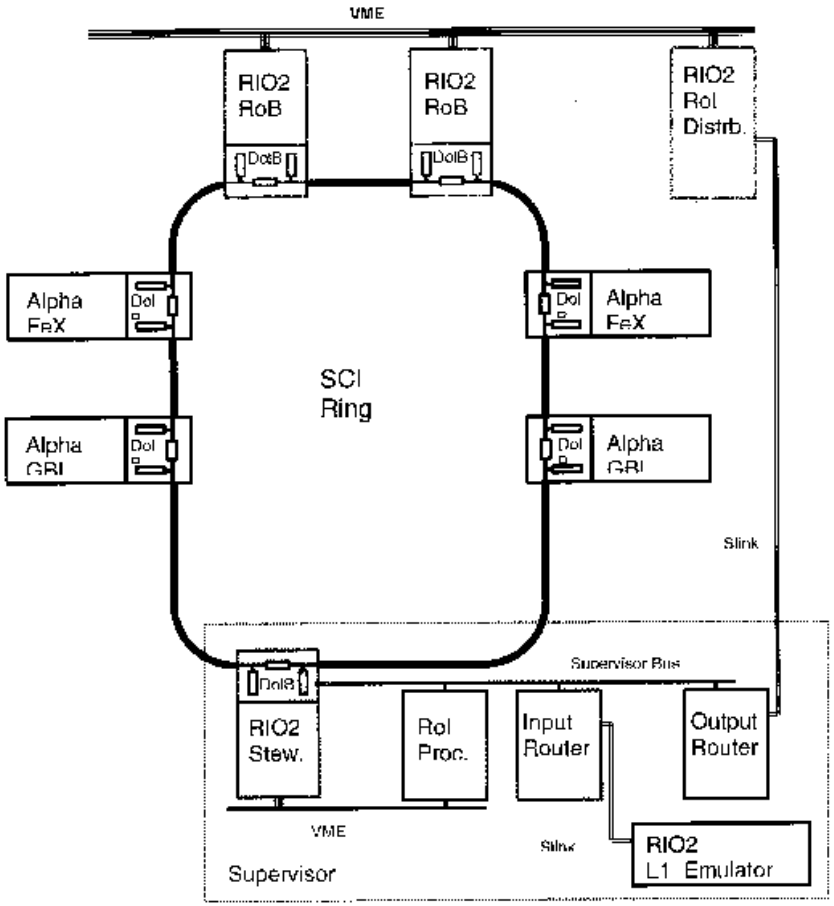

ATLAS

Demo B-SCI
Supervisor Results - May ‘97

- Conditions
  - Simple loop back program: received RoIR over S-Link, transormed records to GOoutR and sent back to supervisor via SCI
  - Free running mode - LVL1 data read from memory by RoI processor
  - One RoI per event
  - Loop back program read blocks of data from S-Link 64 words at a time creating pipeline of ~10 events.
- Results - Time per event: 115 μsec--> 8.7 kHz

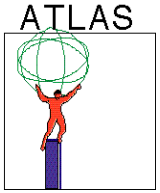

Software Improvements - JLS Aug. ‘97
All times in microseconds

	SCI	VME
Version 2.3 Images	115	169
Version 2.4 Images	44	53
Polled VME mapped no Prefetch	42	35
No histograms	35	31

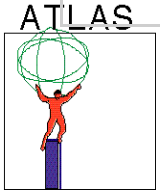

New Version of Programs

- Combine functions of Steward and RoI Processor in a single CPU - simplifies timing measurements
- Event grouping eliminated as much as possible. Record sizes used that provide a steady stream of RoIR's with no long time intervals between groups.
  - clarifies latency measurements
  - may reduce overall rate
  - should shorten latency
- Supervisor/RoI Builder scaling is manifest

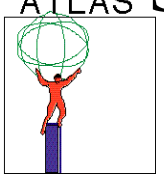

“Free running” timing measurements
All times in microseconds

	SCI	VME
Event time, with histograms	54	49
Event time, no histograms	37	36
Pack RoI Fragments	12	12
Build RoIR	5	5
Send RoIR	9	9
Process GoutR	6	13
Send T2DR (16 events)	12	12

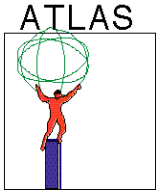

Provisional Conclusions - Demo B

- With 37 μs as time to process an event a rate of ~27 kHz should be achievable.
- Actual measurements show that 21 kHz is maximum sustainable rate.
- Further tests are scheduled for November, 1997.

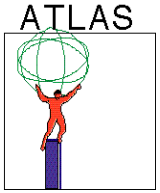

Sequential Event Selection for ATLAS

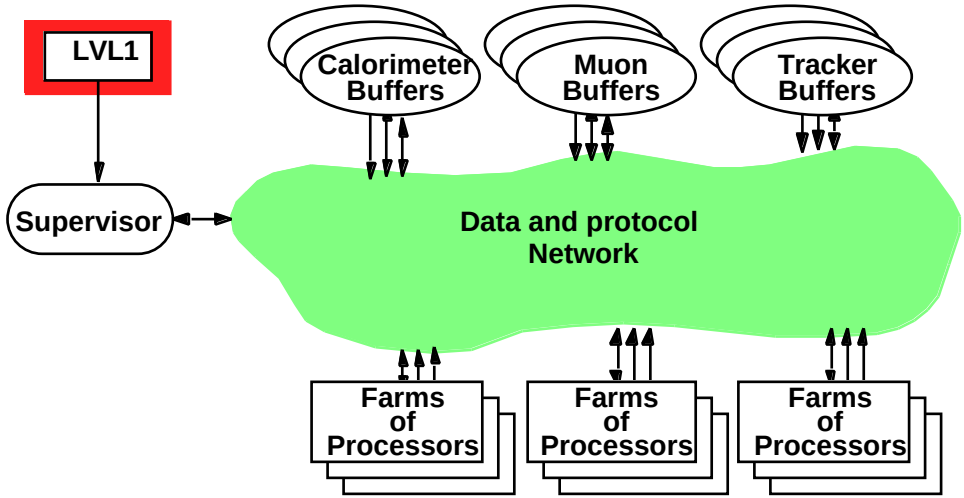

- **Trigger strategy: sequential selection**
  -> Matches physics requirements
- **Trigger architecture: unified protocol / data network**
  -> More flexible, can unify trigger and DAQ networks
- **Data collection protocol: processors request data from sources as needed**
  -> Reduces demands on the communication network

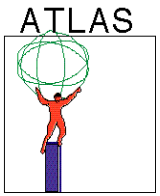

Demonstrator System Configuration

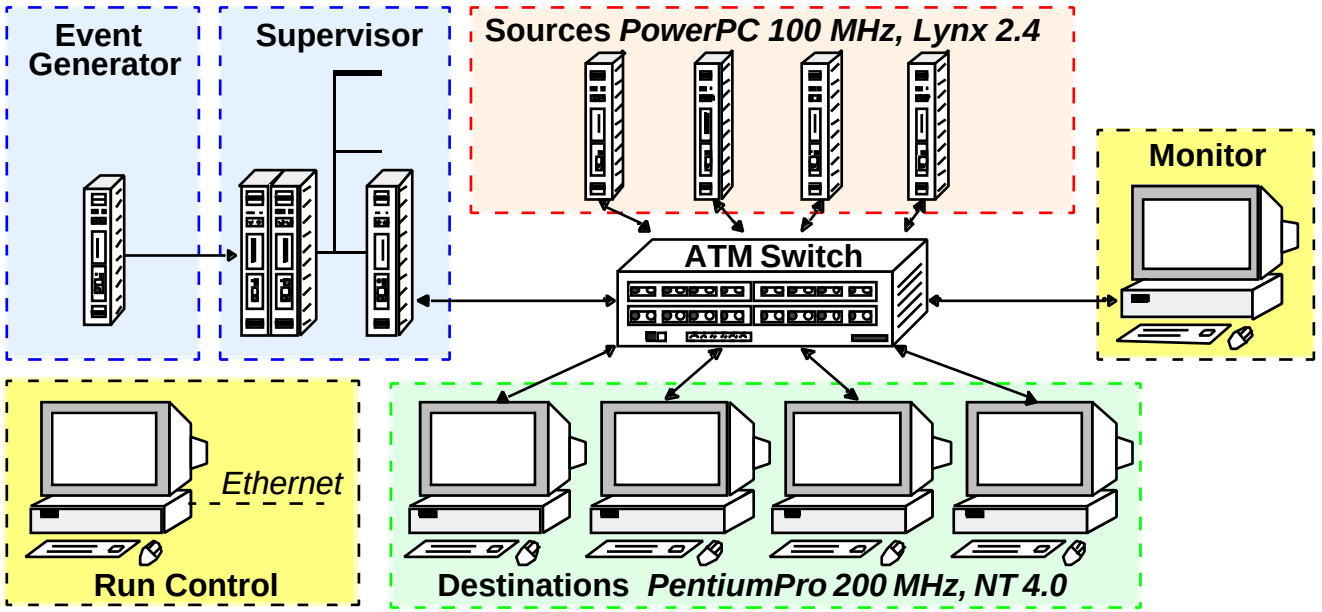

ATLAS

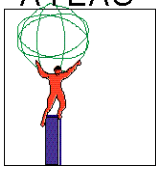

· **The demonstrator implements the concepts of the Sequential Selection Scheme**

Provisional Conclusions - Demo C

- With a single Power PC 100 MHz processor an event rate of ~8 kHz can be sustained
- In emulation mode this rate reaches ~16 kHz
- A 200 MHz Pentium Pro in emulation mode achieves ~22 kHz
- “We think” that with several, faster RoI CPU’s along with various optimizations 100 kHz can be achieved

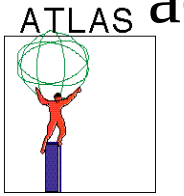

Plans for Next Year

- Further tests with Demo B - November ‘97 at CERN
- Modelling work using measured numbers for DemoB and Demo C - Jos Vermeulen
- No major changes planned for hardware for Supervisor/RoI Builder
- MSU/ANL will concentrate on the LVL1/LVL2 interface