Online Computing Systems Hardware Environment

John Haggerty, BNL

January 10, 1997

The PHENIX online computer environment consists of:

- VME crates with at least one CPU per crate
- Sub-Event Buffer computers (SEB)
- Assembly and Trigger Processors (ATP)
- PHENIX Acquisition Computer (PAC)
- X-terminals, PC's, NC's, and other display devices
- A network which connects all this together

VME Processors I

Although their arrangement is not certain, there will be several different kinds of crates which will be dedicated to specific modules. The types are:

- Master timing system (1-2)
- Global Level 1 (1-2)
- Within each subdetector:
  - Timing and Control (1-12)
  - Local Level 1 (as necessary) (1-6)
  - High voltage and ancillary systems (1-12)
  - DCM's (10-40)

The model is that each crate will have at least one VME cpu module on a network which will be responsible for running code needed to initialize, test, and monitor the operation of the modules in its crate. (Note that we assume that it is not necessary for a single VME cpu to be able to address modules in other crates; we do not plan on bus extenders and such.)

VME Processors II

We assume that the VME cpu's will be running VxWorks. (Is VxWorks necessary? Probably *not*, but it is still a reasonable decision.) What is the basis for selecting a PHENIX standard cpu?

- VME speed, addressing modes
- CPU speed
- VxWorks support
- Compilers, other software support (CORBA)
- Network capability (wouldn't it be nice if they were on a 100BASE-T network?)
- Price, continued availability
- EPICS support (MVME 167 only?)

What are some of the modules we'll investigate?

- Motorola (January 16, 10:00)
  - 680x0 based (MVME162 et al.)
  - PowerPC based (MVME 1600, 2600, 3600)
- Heurikon
  - 680x0 based (Nitro240 et al.)
  - i960 based (V960D et al.)
  - MIPS based (baja4700 et al.)
- DEC α (BSP?)
- VMIC 200 MHz Pentium Pro?

See http://www.vita.com/ for more than you want to know, though Board Support Packages from Wind River may be more relevant (http://www.wrs.com).

Time scale for decision: March-June, 1997

Sub-Event Buffers I

We have agreed that these will (at least for the time being) be general purpose computers running a rather complete operating system. Could they be VME? Yes, although it would appear to be difficult to satisfy all our demands. I tend to think of them by sub-detector, but they might be divided further (by granule?).

We require that they:

- Have a high speed PCI interface which must be capable of pushing data from PCI into memory at no less than 10 Mbyte/sec, and the faster, the better.
- Have available high speed network interfaces and software device drivers known to be capable of a good fraction of their maximum speed.
- Have plenty of capacity for increasing memory, and no problems using large sections of memory for data buffering.
- A pleasing combination of software drivers, compilers, debuggers, CORBA support, and other development tools.
- Cost, continued availability

PHENIX Data Acquisition Architecture

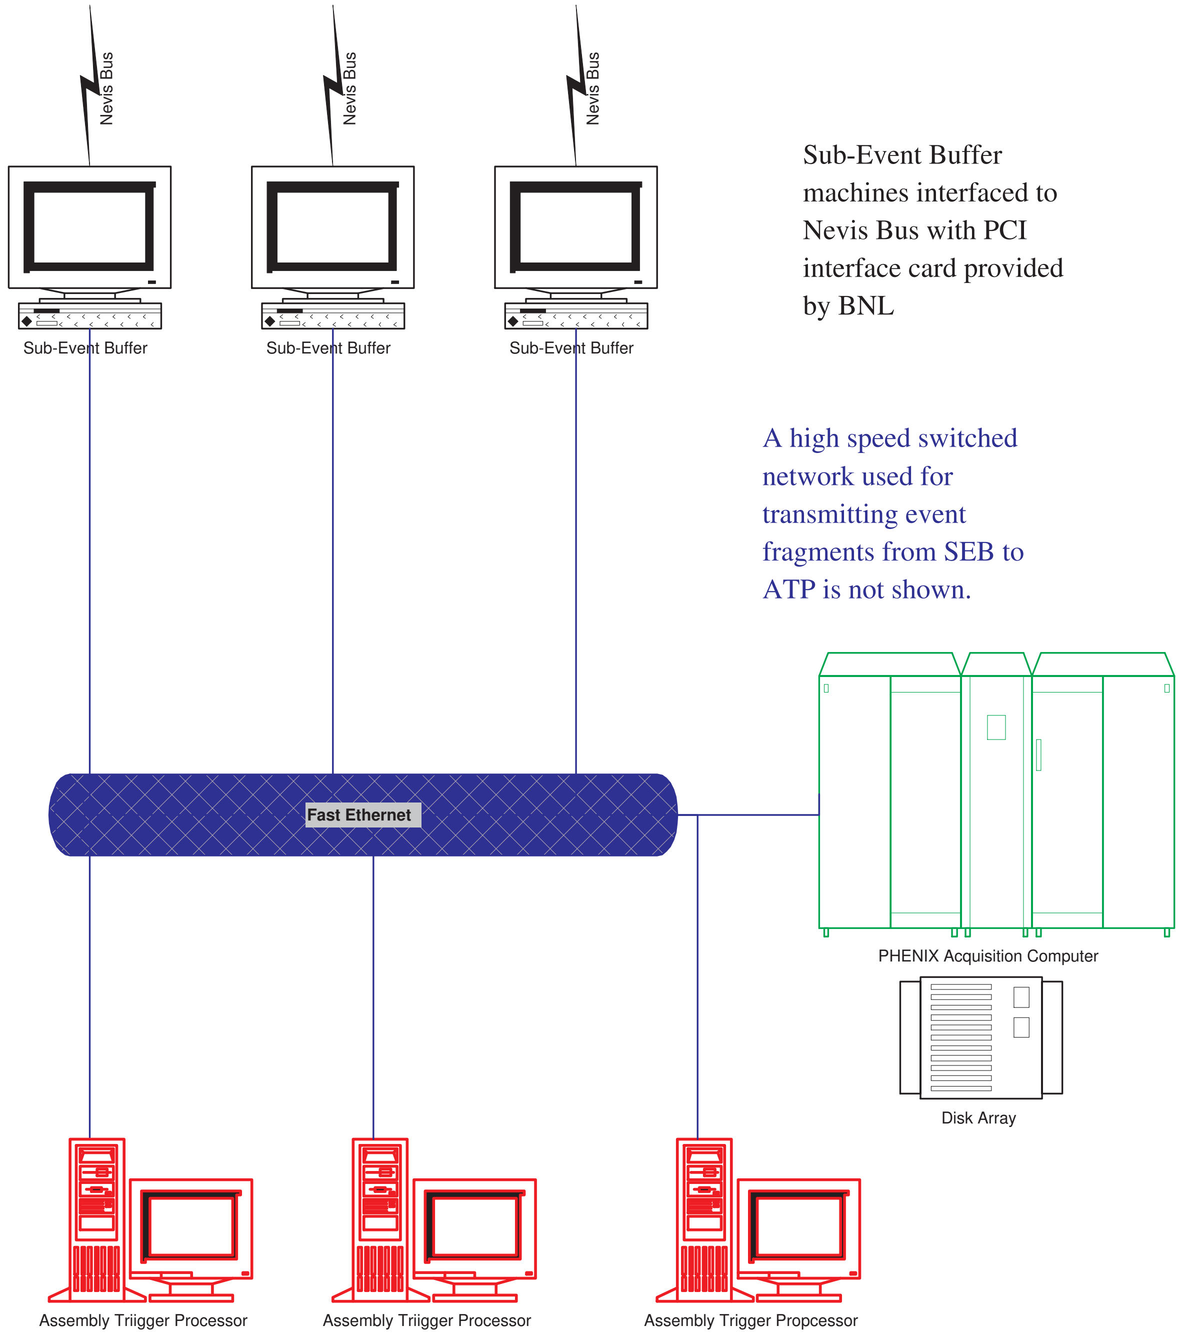

Sub-Event Buffer machines interfaced to Nevis Bus with PCI interface card provided by BNL

A high speed switched network used for transmitting event fragments from SEB to ATP is not shown.

Sub-Event Buffers II

What are some of the possibilities:

- Intel based motherboards running NT, Solaris, or Linux
- DEC α running DEC UNIX or NT
- SGI O2 running IRIX
- Power Mac?
- HP, IBM, other workstation?

A decision will be made in early February, 1997. This accelerated schedule is necessary in order to develop the PCI interface to the Nevis Bus, its software, and be ready for a sector test.

Assembly and Trigger Processors

Data from the Sub-Event Buffers will be pushed across a switched network of some sort into a group of processors now called Assembly and Trigger Processors (ATP). Initially, they can be identical to the SEB's for testing and code development.

Networking

Note that this does *not* refer to the fast switched network which is used for transferring data from SEB to ATP. This is the network(s) which interconnect the VME processors, SEB's, ATP's, PAC, display devices, and provides a connection to the RCC network.

The goals of this network are:

- Fast Ethernet whenever possible (follow Gould and RHIC plans, share equipment when possible)
- Single module failures isolated (no thin wire repeaters)
- Outside traffic bridged off, and PHENIX operational regardless of connection outside counting house (PHENIX DNS in counting house, for example).
- High performance for disks remotely mounted from RCC

PHENIX Acquisition Computer

Although it is antithetical to the notion of a fully distributed system, it is quite likely that PHENIX will want a small number (one?) of very powerful computers to serve as a common interactive host, serve a very large array of disks, and support an extremely high speed network connection to the outside. This machine may be used for online "orchestration," running the user interface, and calculating routing for SEB to ATP.

If this machine is necessary (I can't see how too live without it, but I could be wrong), we will want it in our counting house in 1998 in order to start configuring it, loading it with disks, and using it.

Other Hardware Issues

- Data logging in the counting house (DVD? DLT? DTF?) If we ever do this, it would be nice to
- Disk system–size of file systems, RAID systems
- Network connection to RCC capable of sustaining 20 Mbyte/sec
- Display devices that don't take a huge amount of system management (remote management?) consisting of X-terminals PC's, NC's, Java Stations.
- Miscellany (printers, etc.)