# SEP computer update

*Rick Ottolini*

## HIGHLIGHTS

We have had major changes in hardware and computational practices in the past half year at SEP. Highlights include:

**Routine use of super-computers.**

The Convex has liberated the average researcher from many of the formerly tedious details of programming— no more shoehorning datasets into limited memory resources or learning special vector languages. Applications that used to take months to perfect on the old Vax are implemented now in days. For example, Biondo Biondi's shot profile DMO results in this SEP report are compared to Dave Hale's thesis results. Biondo replicated Dave's results starting from scratch in a matter of days. More detailed examples are given in the case history appendix.

**Networking is an essential resource.**

Our three CPUs (Convex, Sun, Microvax) are etherneted together and connected to hundreds of other computers in the School of Earth Sciences, Stanford, and nation-wide. No longer can each individual computer provide all hardware or software services. Network nodes provide various special functions such as file storage, number-crunching, and graphics output. Also we communicate via mail and computer bulletin boards with colleagues all over the country.

**Computer independent application software ...**

arises from the need to both duplicate and specialize software resources across the network. We have achieved this to a large degree in seismic signal processing applications and graphics. Promotion of the Unix operating system and Ethernet standards among Earth Science computers greatly aids machine independence.

## HARDWARE AND SOFTWARE

Table 1 summarizes the characteristics of the three computers now at SEP. Table 2 describes the School of Earth Sciences network computers larger than an IBM RT. Table 3 describes major application packages on SEP computers.

**Table 1: SEP Computers**

| COMPUTER | **CONVEX** | **SUN III/160** | **MICROVAX II** |
|---|---|---|---|
| CORE | 48 MB | 16 MB | 11 MB |
| DISK | 3960 MB | 142 MB | 71 MB |
| FLOATING POINT | 10-40 MFLOPS | .3 MFLOPS | .5 MFLOPS |
| SCALAR | 4 MIPS | 2 MIPS | 1 MIPS |
| I/O BUS | three Multibuses | VME bus | Q bus |
| SERIAL PORTS | 26 tty, 6 modem | 2 tty | none |
| TAPE | one 9-track | cartridge | cartridge |
| NETWORK | Ethernet | Ethernet | Ethernet |
| RASTER GRAPHICS | Raster Tech | Color console | Monochrome console, two AEDs |
| HARDCOPY | Imagen laser printer | | Printronix line printer |
| OPERATING SYSTEM | UNIX 4.2 | UNIX 4.2 | UNIX 4.2 (Ultrix) |
| LANGUAGES | C, FORTRAN, RATFOR | C, FORTRAN, RATFOR, PASCAL | C, FORTRAN, RATFOR, PASCAL, C++, BASIC, LISP |
| SERVICES | number crunching, terminals, modems, file storage | high quality graphics | number crunching, peripheral support |
| LIST PRICE | $750K | $55K* | $41K |

*Color graphics sub-system included.

**Table 2: School of Earth Science Computers**

| | Core-MB | Disk-MB | Net Name |
|---|---|---|---|
| GEOPHYSICS | | | |
| Reflection Seismology | | | |
| Convex | 48 | 3960 | hanauma |
| Sun III | 16 | 142 | taal |
| Microvax II | 11 | 71 | mazama |
| Seismology/Departmental | | | |
| VAX 11/750 | 4 | 988 | erebus |
| Refraction Seismology | | | |
| Microvax II | 3 | 511 | (installed 2Q86) |
| Borehole Seismology | | | |
| Masscomp | 4 | 550 | loihi |
| Paleomagnetism | | | |
| IBM RT | 3 | 80 | haruna |
| APPLIED EARTH SCIENCES | | | |
| Geostatistics | | | |
| Gould 9860 | 16 | 2700 | summit |
| PETROLEUM ENGINEERING | | | |
| Departmental | | | |
| VAX 11/750 | 6 | 952 | ararat |
| IBM RT | 3 | 80 | (installed 2Q86) |
| Sohio Project | | | |
| Apollo 660 | 4 | 186 | thera |
| GEOLOGY | | | |
| Geology/AES/Departmental | | | |
| VAX 11/750 | 5 | 888 | denali |
| Geochemistry | | | |
| IBM RT | 3 | 80 | (installed 2Q86) |

| **Table 3: Major SEP Applications Software** | | Source |
|---|---|---|
| Seplib | *UNIX based data cube processing kernal* | SEP |
| Segy | *seismic processing library using SEGY headers* | Utah, SEP |
| Sioseis | *Scripps seismic processing package for teaching* | Scripps |
| SOLID | *seismic modeling* | GDC |
| Seis83 | *seismic modeling (ray tracing)* | Cerveny |
| Movie | *fast interactive display of seismic raster cubes* | SEP |
| Vplot | *device independent vector graphics* | SEP |
| Tiplot | *print seismograms on a laser printer* | SEP |
| X windows, SUNVIEW | *bitmap window graphics* | MIT, Sun |
| TeX | *typesetting including tables and math* | Stanford |
| Troff | *UNIX typesetting including tables and math* | AT&T |
| Macsyma | *symbolic algebra solver* | Symbolics |
| Linpack, Eispack | *numerical libraries* | ACM |
| Matlab, FPSlib, Veclib | *array processing libraries* | Stanford, Convex |
| Emycin, MRS | *expert system shells* | Stanford |

## FRONTIER ISSUES

### Standard Processing Package

SEP and other reflection seismic researchers at Stanford continue to be interested in obtaining a commercial seismic data processing package for routine production processing. Such would provide production efficiencies, fill in current software gaps, and avoid needless duplicative SEP software development. The main requirements for a package are affordability and UNIX compatibility. It would also be nice if the package could exploit the strengths of Convex UNIX—large core and disk memories and inter-program communication by pipes.

## Super-computing

The Convex has given us hands-on access to a mini-super-computer of the vectorizing type (Cray-like). The UNIX operating system and vectorizing FORTRAN compiler make using the super-computer as easy as a Vax. Some SEP students have studied the vector unit directly via assembly language. However, hand-tuned code usually performs less than a factor of two over compiled code.

We continues to look at other super-computers, including those of the parallel type. It is not expected that one would come into routine use by SEP unless it provides as transparent an interface as does Convex. How well an algorithm vectorizes or parallelizes is a factor in evaluating new seismic data processing methods.

## Evolution of Movie Software

The survival of the high speed seismic raster graphics package, locally called the 'Movie' program, was an important consideration in changing SEP computers. Not only has this software survived, but continues to proliferate to new hardware and acquire new features. Movies are created at the same rate as seismic hardcopy plotting— about 30 times daily or thirty thousand movies made so far.

We use the three workstations listed in table 1. None has been overall as satisfactory as the old Vax 780 - AED combination, but each setup has advantages. The Sun is an order of magnitude faster than old VAX but lacks hardware scroll and zoom. The Convex - Raster Tech is the slowest, but Convex number crunching permits the realtime display of interactive seismic processing. The Microvax - AED preserves the old VAX speed, but it is inconvenient to copy datasets across the network to it or the Sun. All setups have the advantage of larger core memory ($>$ 8 MB) and less loaded CPU than the old Vax.

Table 4 summarizes the raster transfer speeds of Movie implementations on various UNIX computers with raster graphics devices. The basic test was to continuously transfer a raster rectangle from the host computer to the graphics device. To run the Movie software, a C-language compiler and high speed raster graphics interface is essential. Hardware scroll (pan), zoom (magnification), and color table mapping are also desirable.

**Table 4. Movie Workstation Speeds**

| *HOST CPU* | *Graphics Terminal* | *I/O Channel* | *Pixels/Second* | *Bottleneck* |
|---|---|---|---|---|
| SUN III-160 | SUN III-160 | Memory-mapped | 1,300,000 | |
| VAX STATION II | QVSS (monochrome) | Memory-mapped | 430,000 | Q-bus |
| VAX 11/780 | AED 1024 | DMA | 300,000 | Unibus |
| VAX STATION II | AED 1024 | DMA | 300,000 | Q-bus |
| CONVEX | Raster Tech | DMA | 200,000 | DMA software |
| VAX STATION II | QDSS (color) | DMA | 100,000 | DMA software |
| APOLLO 660* | APOLLO 660 | Memory-mapped | 60,000 | graphics hardware |
| IRIS 1400* | IRIS | not available | 50,000 | graphics hardware |
| ANY | ANY | Ethernet | 40,000 | protocol software |

*New models have been introduced since these tests were made.

Work in progress on high speed raster seismic graphics includes:

(1) A single hardware independent raster graphics code.

(2) Interactive display of realtime seismic data processing such as 2-D power spectra, statics, stacking, and migration made possible by the speed of the Convex.

(3) Visualization of 4-dimensional data relationships. We find the current 3-D paradigm of a data cube with each side a filmloop to be inadequate for some datasets. The 4-D model replaces the filmloop with a rectangular array of panels (spreadsheet-like). Each row or column can be a face of a display cube or the display cube itself.

(4) Bitmap graphics display parameter controller (Macintosh-like).

### A.I. Software

The symbolic expert system called Macsyma has been an extremely useful aid to SEP researchers. However, expert system *shells* (see table 3; reported on in SEP-37) have attracted little attention at SEP. As previously reported, I constructed examples for selecting and using signal processing software. This is interesting from a computer science or engineering development point of view, but does not necessarily provide a more powerful set of tools for assisting the sesimic data processing researcher.

### Networks and Workstations

Like other groups at Stanford and industry, we are experimenting with combining large computers into networks with powerful workstations. We once considered devoting the Convex as an array processer and file storage host and using office-based microvax class machines (i.e. Sun, Apollo, IBM RT) for routine administrative computing. This is the model for the electrical engineering Convex elsewhere at Stanford. However, our average seismic or graphics application file I/O requirements greatly strain the capacity of the Ethernet computer connection. In addition, because peripheral makers don't provide interfaces to major I/O buses or networks, we can't globally distribute the data storage or I/O peripherals in the most accessible way. Therefore, we use the Convex as the routine administrative computer as well an array processor and file host and have assigned the workstations to specialist roles. We expect that workstations may come into greater use as networks become more transparent (due to industry wide adoption of Sun's Network File System standard) and workstations become more common at SEP.

## MINI SUPER-COMPUTER CASE HISTORIES

### Convex speed in velocity analysis (Jon)

A velocity spectrum program described elsewhere in this report was used extensively in a decon program described elsewhere in this report. I was amazed and delighted by the Convex speed in this application.

Velocity analysis is a process that has much overhead. But it is a process that is often used repetitively and the overhead need only be done once if enough RAM is available. Usually the repetition is over "midpoint." In the deconvolution application in this report the repetition is over "regressor" or "lag." I'll use the word "panel" to refer to a plane of constant midpoint or constant regressor. Looking at a velocity program you see the cost of generating the first panel is about 5 times the cost of subsequent panels. Compound this with the vectorizability of the panel repetition loop and you are looking at speedups that are amazing. Of course these speedups cannot be realized in a computer or an array processor that does not have sufficient RAM to hold many panels.

I was used to a single NMO taking about a minute on the VAX so I was amazed and delighted that a 40 point velocity spectrum repeated 43 times took only about three minutes on the Convex C-1. This looks like a speedup of 500=40*43/3 over the VAX. Our best speedups in SEP-44 were only about 60, multiplying this by the factor of 5 for saved overhead gives 300. The discrepancy between 300 and 500 probably results because I never benchmarked the same program on the VAX or tried to isolate the CPU time versus the computation time. Anyway, it should be obvious that I would not have considered such an application on the VAX.

Of course you are likely to be looking at better tuned programs than I was running on the VAX, so in absolute terms the parameters for the three minute job are

750 time points
64 offsets
40 velocities
43 panels

My only effort at optimization was to arrange that the vectorization index (the panel index) ran down sequential locations in memory. In other words, historically a program has a dimension line like

```
REAL VELAN ( TIME, VELOCITY, PANEL )
```

but the dimension statement in my program was

```
REAL VELAN ( PANEL, TIME, VELOCITY )
```

According to the Convex company's description, this gave me a factor of about 2.5 in speed, but I did not check it.