jvb7u@fermi.clas.Virginia.EDU (Jon Brinkmann) writes:
> In a nutshell, PAW was designed to be used for reducing experimental
> data. Any type of data, but there is a package designed for High
> Energy Physics. It has vastly superior memory management, graphics and
> least squares capability to any interactive language I've ever
> encountered.

This must not be the same PAW I fight with every day. The PAW I use
has a memory management system that constantly produces damaged ZEBRA
files, and sporadically lets HBOOK and GEANT stomp all over each
others memory zones. Of course, what do you expect from a memory
manager written in Fortran (which makes it nearly impossible to debug,
by the way)? Note that the memory it manages is all declared in one
humongous common block, so there is no real dynamic memory as such.
There's a hard builtin limit to the size of files that it can handle
(not that it tells you this, it just makes the files unusable if you
attempt to write too much data to them). The graphics look clunky and
it's very difficult to customize plots with comments, etc. I won't
even mention the Fortran interpreter that silently gives wrong answers
when you declare the type of external functions, silently gives wrong
answers when there are tabs in the source, etc. etc. etc.

> I guess it has to be, when a single "shot" of an accelerator
> produces hundreds of megabytes of data!

PAW is not suited for looking at that quantity of data. PAW is
designed to work with greatly reduced data samples, on the order of
tens of thousands of tuples. It beats batch jobs and line printer
output, but falls far short of being a "good" tool.

Charles Allen   Internet: cca@physics.purdue.edu
Department of Physics   HEPnet: purdnu::allen, fnal::cca
Purdue University   Bitnet: cca@fnal.bitnet
West Lafayette, IN 47907-1396   talknet: 317/494-9776