ParCAT

What is ParCAT?

- Command-line utility
  - Utilizes MPI, pNetCDF, written in C
  - Provides scriptable and embeddable interface
  - Buildable and runnable on diverse architectures – laptops to supercomputers
- Provides parallel pre-processing of large datasets
  - Spatio-temporal averages (seasonal, monthly, yearly, arbitrary), frequency distribution, and differencing of two datasets (plus difference of averages and average of differences)
  - Takes advantage of multi-core HPC offerings to improve parallelism and lessen IO issues
- Reads netCDF input files, produces netCDF output files

Example 1

- CLM dataset, 15 year, monthly data. 349 variables, ½ degree resolution. ~71gb, 180 files (no CMIP5 postprocessing)
- Compute overall average for each variable across the time series using Titan
- Each node opens one or more files; each core processes one or more variables. Cores consolidate results and nodes create one or more output files.

Example 1 Timing Results

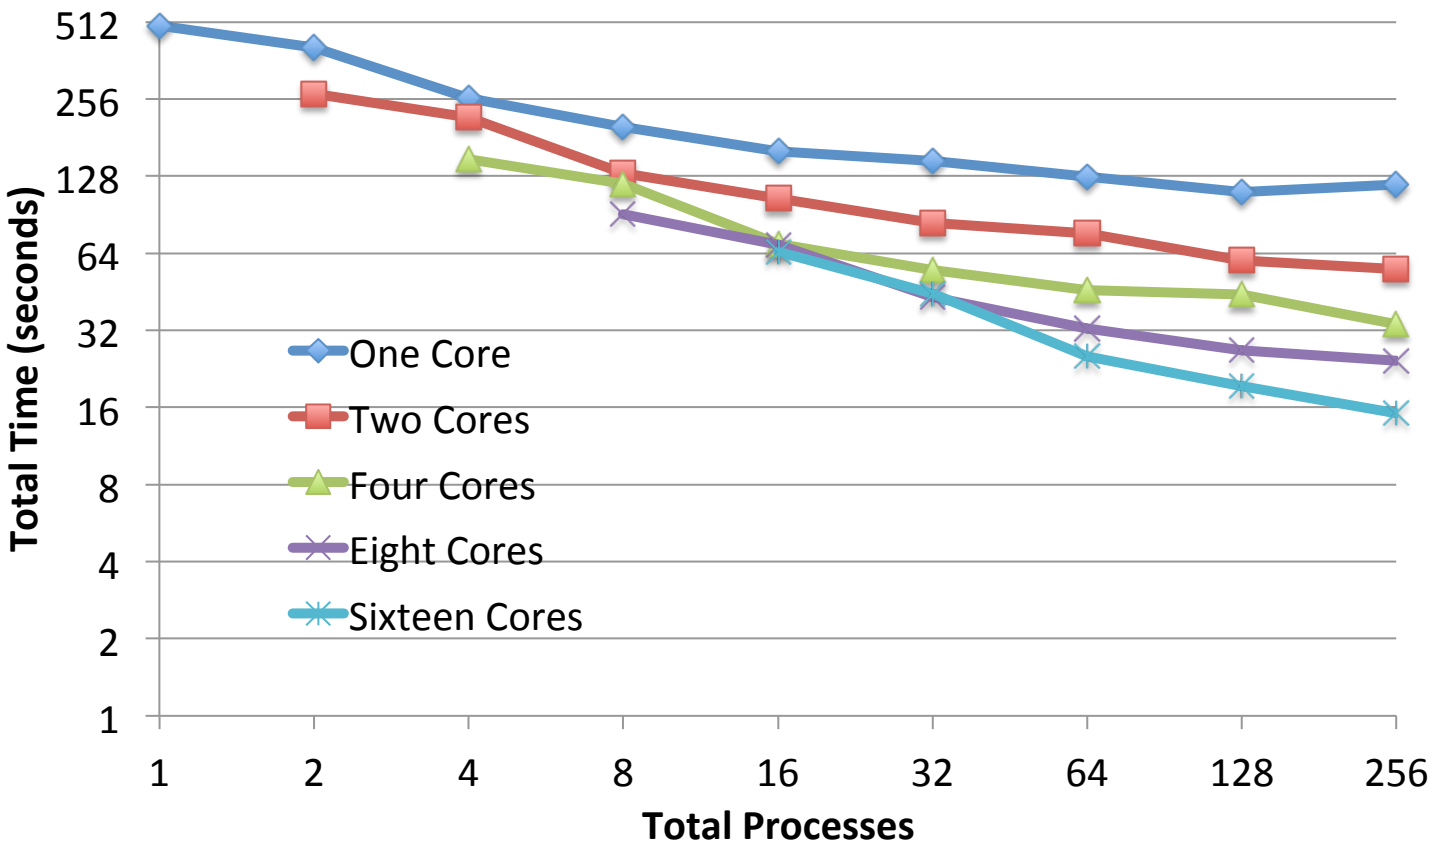

- Time went from ~475 seconds to ~15 seconds using only 256 processes on Titan
- The improvement should be even better for larger datasets

Example 1 – Loading into UV-CDAT
Average GPP over the Entire Data Set

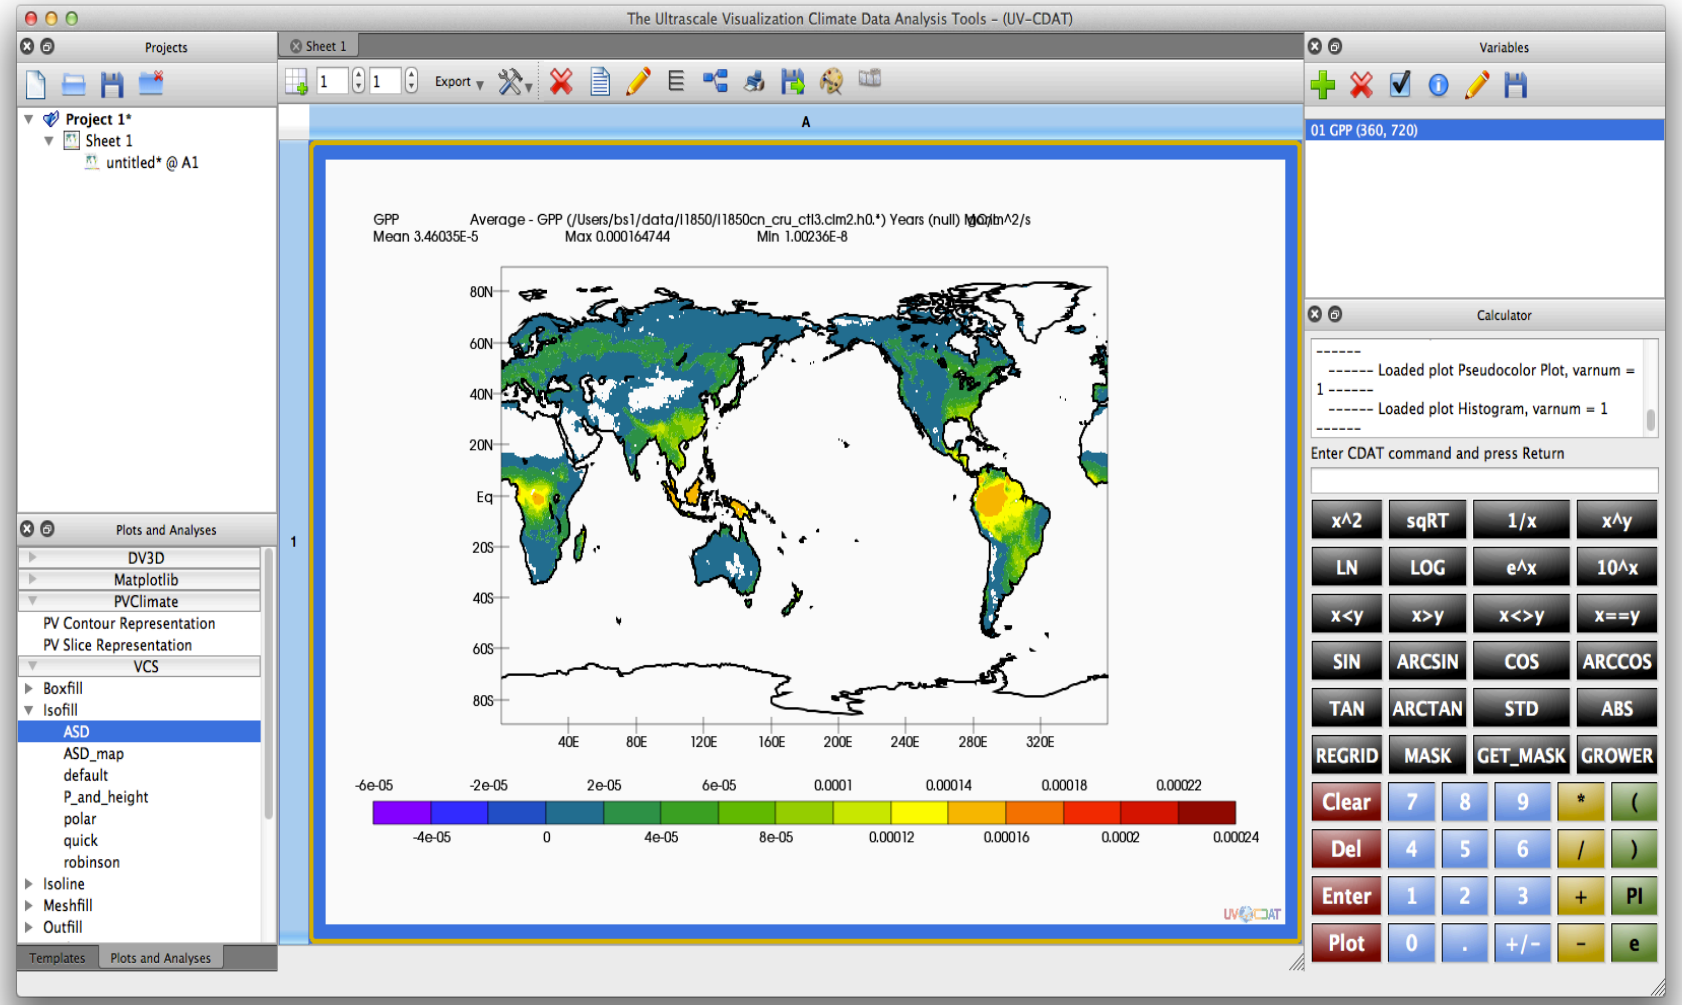

Example 2

- CAM dataset, 52 years, monthly data. 28 variables, 1 degree resolution. ~18gb, 28 files (postprocessed into CMIP5 standards)
- Compute overall average of each variable on Titan
- Each node opens one or more files, each core processes one or more time steps. Results are combined into one or more output files

Example 2 Timing Results

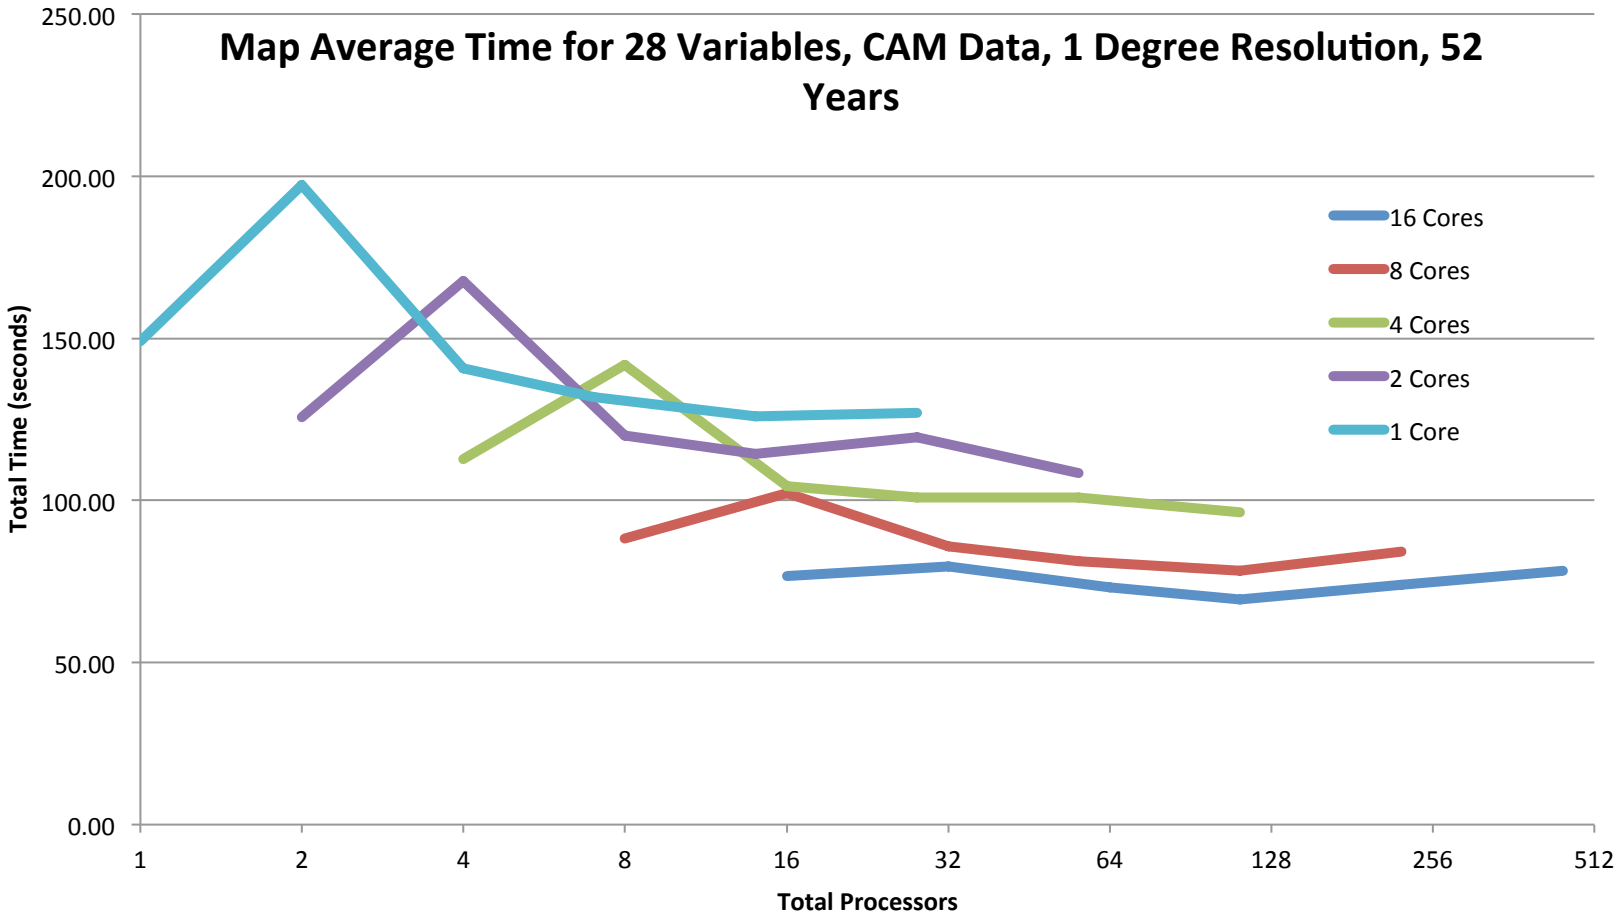

- Time went from ~150 seconds to ~50 seconds. More improvement should be possible with more variables or timesteps