P1.12

PORTING NESDIS POLAR SOUNDINGS FROM A CRAY TO AN IBM SP

Richard G. Kelley, Jr.*
Computer Sciences Corporation

The Information Processing Division (IPD) of NESDIS' Office of Satellite Data Processing and Distribution makes vertical soundings from Polar-Orbiting Operational Environmental Satellites. As technology evolved, IPD responded by moving software to an IBM Scalable Processor from a Cray J90. This migration of some 650 modules of code was non-trivial. It involved other platforms and extensive validation efforts. The purpose of this poster is to present aspects of the migration to those who may be anticipating similar endeavors.

Before the code was moved, staff evaluated the software. This brought familiarity and uncovered areas of concern, especially those which would require large time investments. The evaluation produced a document containing an analysis and recommendations. This identified how the code ran on the Cray and also recommended not porting some systems because they were antiquated and would be cost-prohibitive to port.

IPD runs satellite processing for three operational satellites, NOAA-15, 16 and 17. These systems typically ingest 14 orbits per satellite daily, with extensive computational requirements. On the Cray each orbit consumed approximately 20 cpu minutes. The new IBM SP system only takes 3 minutes for the same process. (See figure 1.)

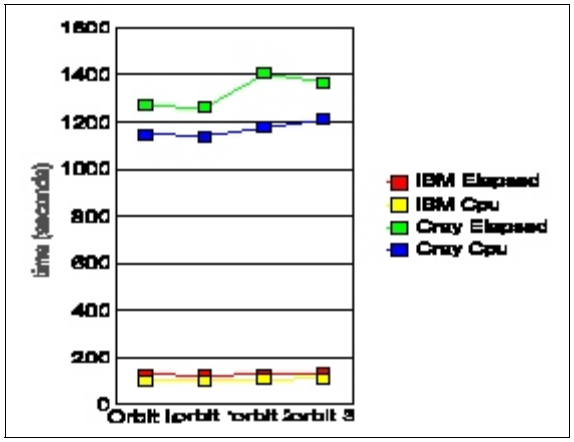

Figure 1 Cray vs SP Timing Runs

*Corresponding author address: Richard Kelley, E/SP1, Rm 3045, FB4, NOAA/NESDIS/IPD, 4401 Suitland Rd, Suitland, MD 20746; email Richard.Kelley@noaa.gov

Besides the work of moving the code, validation was a challenge. Output from the Cray was deemed to be the "truth". However, the Cray was slow and processes would complete on the IBM before the Cray, updating intermediate files. IBM to Cray file comparisons were crucial to validation, so the IBM was forced to slow down to mimic the Cray. Another problem was examining intermediate files for comparison. These binaries had never been looked at before the IBM was enabled. Special tools were developed to dump their contents and then results were compared. An existing, graphics-based tool was used to examine final output files. This tool was especially useful at comparing global scale output. It led to rapid detection of system differences. Another tool was statistics based. This generated weekly, tabular information of environmental parameters such as retrieved temperature, virtual temperature, water vapor, and layer thicknesses. For each of these parameters, comparisons were made of sample size, means, biases, standard deviations, and root mean square errors. This was used by scientific staff to determine the validity of the information produced by the IBM. See table 1 for a typical comparison.

Cray - SP
9 - 15 June 2003
30 - 90 North
Retrieved Temperatures

Pressure (mb)	Sample Size	Rack Mean	Sat Mean	Bias	Std Dev	RMS
0.10	0.00	0.00	0.00	0.00	0.00	0.00
0.20	0.00	0.00	0.00	0.00	0.00	0.00
100.00	5.00	-0.04	-0.05	-0.03	-0.03	-0.03
150.00	5.00	-0.05	0.02	0.05	0.03	0.04
200.00	5.00	-0.00	-0.02	-0.02	-0.02	-0.02
250.00	5.00	0.00	-0.07	-0.06	0.04	0.05
400.00	5.00	0.04	0.08	0.06	-0.02	-0.03
500.00	5.00	0.03	0.08	0.07	-0.03	-0.04
700.00	5.00	0.04	0.09	0.06	0.03	-0.02
850.00	5.00	0.03	-0.05	-0.06	0.00	0.00
1000.00	5.00	0.04	-0.08	-0.03	0.06	0.05

Table 1 Cray minus SP values of retrieved temperatures

Another important aspect was documentation, which existed, but was fragmented and not user-friendly. A single, comprehensive document was generated.

In summary, moving the soundings application to newer technology required advanced planning, including familiarizing porting staff with the intricacies of the software. There were challenges overcome, primarily resulting from machine differences. Validation and acceptance by science experts was key to the successful outcome. Of the lessons learned, the keenest one was the need to develop specialized software to examine binary files. Another was synchronizing the two machines to produce meaningful comparisons.