Salvo Seismic Imaging Software on C-Plant

Curtis Ober
Parallel Computational Science
Sandia National Laboratories
Albuquerque, NM 87185-1111
email: ccober@cs.sandia.gov

January 7, 2000

Introduction

Salvo is a 3D prestack-depth-migration software package that utilizes algorithms with better physics and improved mathematical methods to produce higher quality seismic images than traditional methods. Most of the remaining oil/gas in the United States lie in regions of complex geologies where traditional seismic imaging methods are frequently unable to produce high quality images, which increases the risks of drilling dry wells. With well drilling costs of order 10 millions dollars a piece, drilling dry wells can potentially threaten an oil company's existence as well as result in higher oil/gas costs that eventually get passed onto the consumer. Having higher quality seismic images lessens the risk of drilling unsuccessful wells and allows additional oil/gas reserves to be discovered. Prestack depth migration, even without the better physics in Salvo, is a very computational-intensive and data-intensive process challenging today's most powerful computers. Salvo utilizes massively parallel computers, using improved computational and data input techniques as well as algorithmic enhancements, to reduce the cost and time to produce higher quality images.

To reduce the computational costs, the Salvo algorithm solves the paraxial form of the 3D acoustic wave equation. Additionally, operator splitting is applied to the paraxial wave equation which allows efficient tridiagonal solves to be employed. Salvo also uses phase correction filters to compensate for approximations invoked during the theoretical development. To reduce the runtime, two types of parallelism are implemented within Salvo: frequency and spatial parallelism. Frequency parallelism spreads the individual frequencies across processors, while spatial parallelism divides the x-y domain among the processors. These parallelisms can be used separately or in combination. Frequency parallelism needs fairly trivial communication routines (broadcasts and gathers). Spatial parallelism is much more communication intensive and uses MPI's derived data types. To address the I/O issues, some of the compute nodes (usually one is sufficient) can be dedicated to reading and writing data. This creates an asynchronous I/O system within Salvo, which allows the reading and writing of data to be hidden behind the computations.

Porting and Validation

The porting of Salvo was fairly painless, requiring just a few hours to get the correct libraries (*i.e.*, MPI) and compilers, and accurate images have been obtained for frequency parallelism. Salvo is primarily ran with frequency parallelism because of its higher parallel efficiencies. Recently spatial parallelism was tested, and Salvo consistently core dumped. This does not occur on SGI platforms. Since the current C-Plant communications may have bugs, C-Plant developers have suggested waiting for the new communications package before continuing the debugging effort. It is believed that the non-trivial communication used in the spatial parallelism could be causing the core dump. Other applications have seen similar communications-related problems.

The I/O system on C-Plant has been a known problem, however it has improved over the last six months. Results from the frequency-parallelism tests are shown in Figures 1, 2, and 3 for three different dates during the past six months. For a small fixed-sized problem (256 frequencies) with one I/O node, the runtimes are shown in Figure 1. Runtimes with the output written to a NFS mounted disk show very poor scalability. However if the output is not written, runtimes and scalability are much improved.

Figure 2 shows runtimes for a larger fixed-sized problem (1024 frequencies), and results from the ASCI Red machine are shown for comparison. The ideal scaling lines are separated by a factor of ~ 2.7 in runtimes, which is near the theoretical value of 3 for processor speeds. Runtime with output written to a NFS mounted disk have improved. At 16 processors, though, the runtimes plateau. This is because the I/O node requires ~ 200 seconds to write the output data. So no matter how many more nodes are used, the runtime will remain fairly constant. When the output is not written, Salvo scales very well to 128 processors.

For the same problem as used in August 1999, validation tests were rerun in December 1999 and the results are shown in Figure 3. Runtimes with output written to the NFS mounted home directory have continued to improve. The runtime plateau occurs at 64 processors, where the I/O node requires ~ 30 seconds to write the output data. When the output is written to the newly-added RAID disks, no runtime plateau is seen up to 256 processors. The scalability is very good, and it now appears that Salvo can be run without concerns of poor read and write times.

C-Plant Experience

The C-Plant meetings were used for the exchange information between application developers about current problems and successes on C-Plant (as noted above with the communications-related core dump). This exchange not only included the application specific problems but also every day items such as logging-in, and killing processes. These meetings also allowed the C-Plant developers to present information on the current status of the operating system, debugger, and queuing system to a major of the users and allow users to give feedback to the developers.

The Salvo validation tests (along with other applications) provided information on the poor I/O system on C-Plant. The C-Plant development team responded which is shown by the improvements over the past six months. Runs from ASCI Red helped find some errors associated with the compilation of Salvo on C-Plant, and provided a gauge for the relative performance of Salvo on C-Plant. Overall, confidence in completing runs on C-Plant has increased through this effort and C-Plant will be the primary architecture for future parallel computations with Salvo.

Future Work

Dependent on future funding, C-Plant will be used to produce seismic images for various research topics including, vertical cable acquisitions, phase encoding, and reduced-source migrations. The major difficulty with using C-Plant (along with other architectures at SNL) is the bandwidth between C-Plant machines and the user's workstation (or other machines where visualization and analysis occur). Seismic imaging, as mentioned earlier, is highly data intensive, requiring and producing large volumes of data. One survey required 6 weeks of CPU time on the Intel Paragon. In 16 hours of this runtime, enough data was produced to require 24 hours to move the data off the machine. With this poor bandwidth, only a portion of the final image could be saved. With faster machines such as C-Plant, this problem will be more pronounced. Although it is believed that this is a known problem, it continues to be a weak link in computational services.

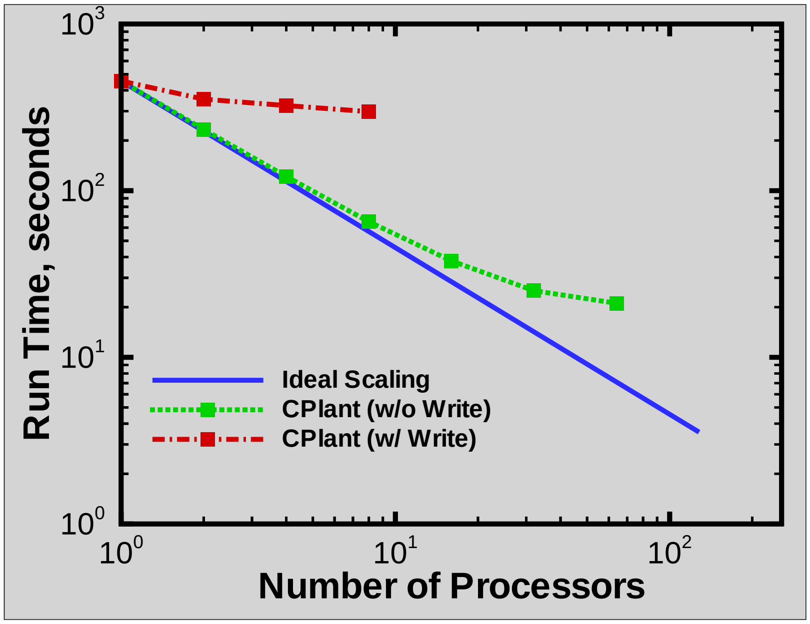

Figure 1: Salvo runtimes for a small fixed-sized problem. Runs completed by June 18, 1999.

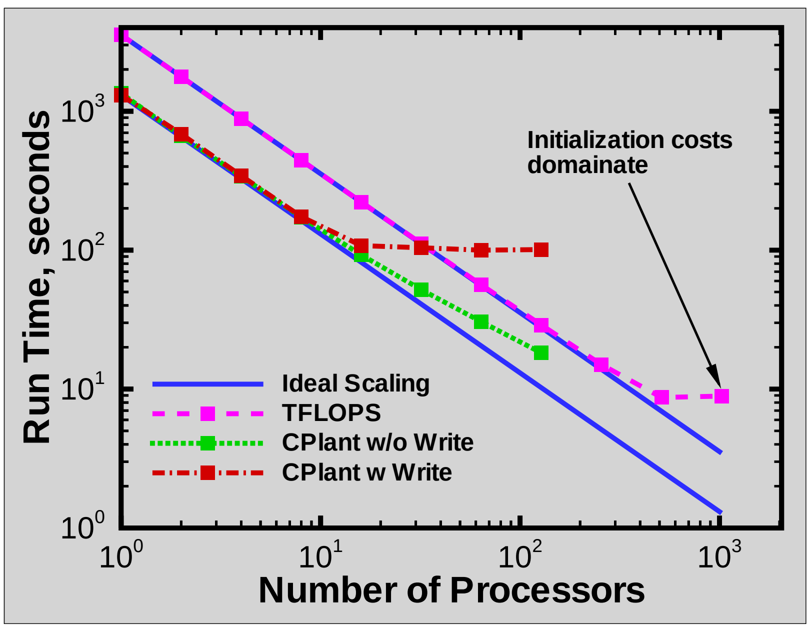

Figure 2: Salvo runtimes for a fixed-sized problem. Runs completed by August 6, 1999.

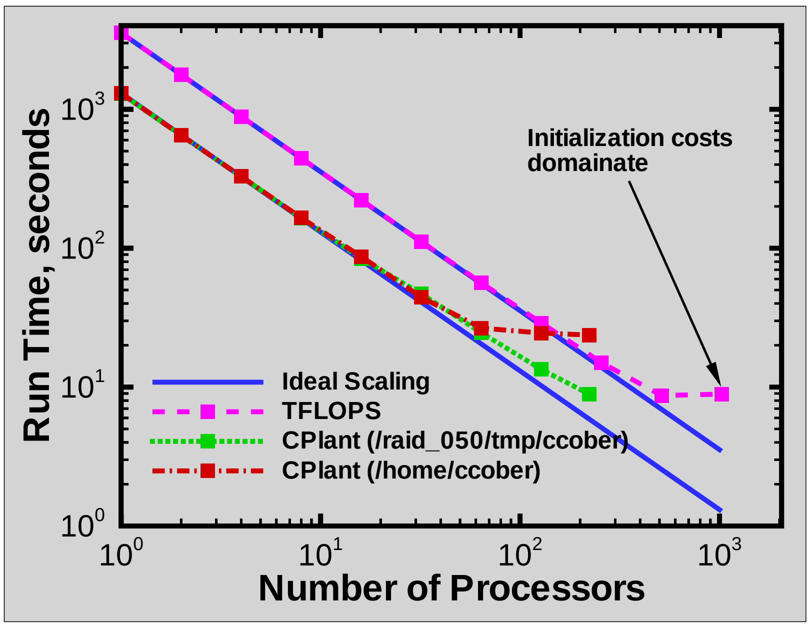

Figure 3: Salvo runtimes for a fixed-sized problem. Runs completed by December 6, 1999.