NEAR REAL-TIME IMAGE RECONSTRUCTION

C. DENKER, G. YANG and H. WANG
Big Bear Solar Observatory, New Jersey Institute of Technology, 40386 North Shore Lane, Big Bear City, CA 92314, U.S.A. (e-mail: cdenker@bbso.njit.edu)

(Received 12 March 2001; accepted 21 March 2001)

Abstract. In recent years, post-facto image-processing algorithms have been developed to achieve diffraction-limited observations of the solar surface. We present a combination of frame selection, speckle-masking imaging, and parallel computing which provides real-time, diffraction-limited, 256× 256 pixel images at a 1-minute cadence. Our approach to achieve diffraction limited observations is complementary to adaptive optics (AO). At the moment, AO is limited by the fact that it corrects wavefront abberations only for a field of view comparable to the isoplanatic patch. This limitation does not apply to speckle-masking imaging. However, speckle-masking imaging relies on short-exposure images which limits its spectroscopic applications. The parallel processing of the data is performed on a Beowulf-class computer which utilizes off-the-shelf, mass-market technologies to provide high computational performance for scientific calculations and applications at low cost. Beowulf computers have a great potential, not only for image reconstruction, but for any kind of complex data reduction. Immediate access to high-level data products and direct visualization of dynamic processes on the Sun are two of the advantages to be gained.

1. Introduction

The power of parallel computing has not yet been exploited in experimental solar physics. Parallel computers in supercomputing centers have been used for numerical calculations in astrophysics, however, they cannot be used for real-time data processing because the observational data cannot be transferred to a supercomputing center in real-time. It is typical for campaign-type observations that only a few data sets are analyzed per year, and only at the end of this process does the scientist find out whether or not something really interesting has been captured. The time lag between the observations and the data analysis is especially too long for rapid response, such as needed for space weather warnings and flare forecasting. Furthermore, the time lag renders the whole scientific enterprise less efficient than it needs to be considering today's computer technologies, and the limitations on data storage also limits the duration of the observing sequence. Parallel processing of solar data will literally provide a new window through which we can observe the Sun in exquisite detail to study the evolution of granulation, filigree, pores, sunspots, prominences, filaments, and flares. The underlying data-processing algorithms are understood, but the complexity is such that only parallel computing enables us to visualize and interpret large data sets efficiently. The development of multi-threaded and parallel numerical algorithms for high-spatial resolution imaging and

Solar Physics **202:** 63–70, 2001.

multidimensional spectroscopy becomes even more important in the context of proposed, next generation, 1- to 4-m class solar telescopes.

In Section 2, we give a brief overview of the image reconstruction algorithm and the observations taken at the Big Bear Solar Observatory (BBSO). Section 3 deals with the hardware configuration while Section 4 focuses on the implementation of the software. In the final section, we discuss some of the preliminary results of near real-time image reconstruction.

2. Observations

Our image reconstruction algorithm is based on the speckle-masking method or triple correlation technique introduced by Weigelt (1977). We use a modified implementation of the code by de Boer (1993). This method determines the modulus and phase of the object's Fourier transform. One advantage is the highly redundant and reliable calculation of the Fourier phases so that we can apply a very efficient noise filter (de Boer, 1996). A more complete bibliography on solar speckle interferometry is given by Denker (1998) and recent results from speckle-masking imaging at BBSO are given by Denker and Wang (1998) and Denker *et al.* (1999).

Speckle-masking imaging requires a sequence of short-exposure images to 'freeze' the wavefront aberrations which makes it possible to separate the object information from the information about the atmospheric turbulence. The first steps in preprocessing the data concern the usual corrections utilizing the average dark and flat-field images. The image displacement within the sequence is then removed with respect to the image with the highest rms-contrast. In the future, this step in the data reduction will be replaced by a correlation tracker which compensates for image motion at a rate of up to 400 Hz. Since speckle interferometry is only valid for a small region, the isoplanatic patch, the 'specklegrams' have to be divided into a mosaic of partially overlapping images. Each isoplanatic patch has a size of approximately $5.1'' \times 5.1''$. The differential image displacement in these stacks of 64×64 pixel images is also removed. We use the spectral ratio technique to derive the Fried-parameter r_0 (von der Lühe, 1984). This method relies on theoretical models of the signal transfer function (Korff, 1973) and of the short-exposure modulation transfer function (Fried, 1966). The amplitudes of the object's Fourier transform are corrected according to the classical method of Labeyrie (1970). To derive the phases of the object's Fourier transform, we use the speckle-masking method. Back-transformation of the modulus and phases of the object's Fourier transform yields a mosaic of partially overlapping speckle reconstructions. Finally, we align these reconstructions very accurately and put them together.

Figure 1 shows a white-light (520 ± 3 nm) speckle reconstruction of a sunspot. The image was obtained with the 65-cm vacuum reflector at BBSO on 27 August 1999. The penumbral filaments, penumbral grains, granulation, and filigree in Figure 1 show details down to the diffraction limit of the telescope which is

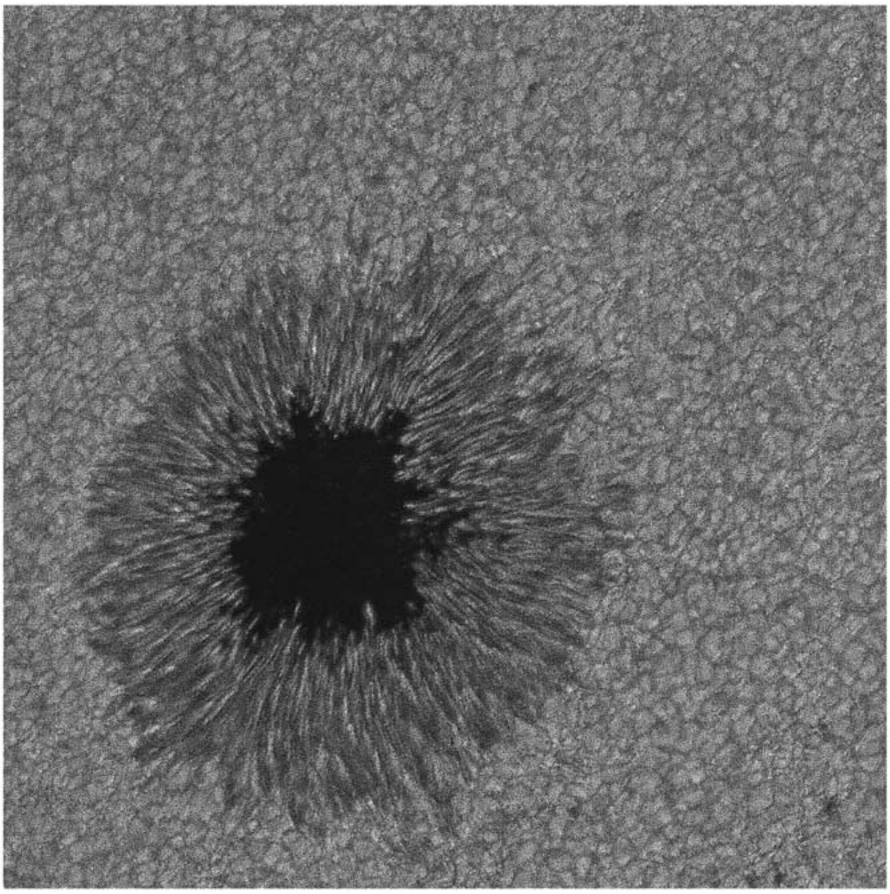

Figure 1. White-light (520 ± 3 nm) speckle reconstruction of a sunspot (NOAA 8673) observed on 27 August 1999 with the 65-cm vacuum reflector at BBSO. The field of view is 76×76 arc sec^2.

$\lambda/D = 0.165$ arc sec at 520 nm or 120 km on the solar surface. We used a 1024 × 1024 pixel, 12-bit, Silicon Mountain Design 1M15 CCD camera to take sequences of 150 short-exposures images (4 ms) at a frame rate of 15 frames s^{-1}. Only the 50 images with the highest rms-contrast are saved and used for the speckle reconstruction. A detailed description of the seeing characteristics at BBSO is given by Goode *et al.* (2000). The time interval between consecutive sequences was approximately 1 min which is sufficient to study the evolution of solar fine-structures. The typical lifetime of photospheric fine-structures is in the range of 5 to 10 min, i.e., near real-time in this paper refers to a time-scale which is short compared to the evolution time-scale of solar features. The other important time-scale for speckle imaging is given by the fact that the object should not change while we take a sequence of short-exposure images. An upper limit for photospheric proper motions is about 2 km s^{-1}. Since a typical pixel is about 0.08 arc sec or 60 km, the data has to be acquired within about 30 s.

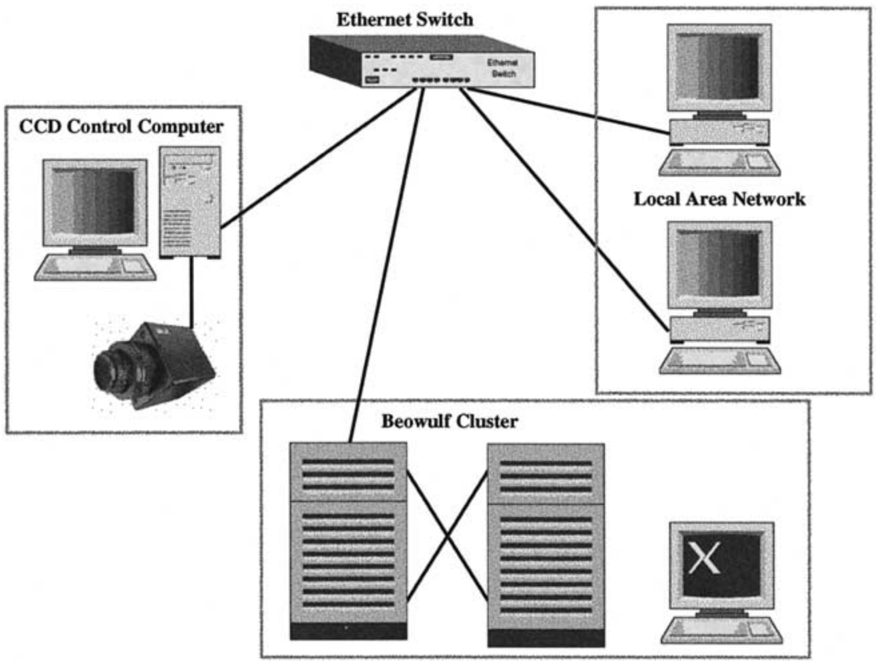

Figure 2. Schematic sketch of the hardware configuration.

3. Hardware Configuration

Beowulf-class computers (e.g., `http://www.beowulf.org/`) are cost-effective since they utilize off-the-shelf, mass-market technologies to provide high computational performance for scientific calculations and applications. The program development is currently performed on a two-node Beowulf cluster manufactured by Paralogic, Inc. Each node has two 500 MHz Pentium III processors, 512 MByte of memory, and 11 GBytes of disk space. The nodes are connected by a 100 Mbit Ethernet. The operating system is Linux with a 2.2.13 kernel for Symmetric Multi-Processing (SMP). The cluster is a hybrid of a distributed and shared memory parallel computer. The parallel computer is connected via a separate Ethernet link to a switch which serves the Local Area Network (LAN), camera control computers, and telescope computers. Figure 2 shows a schematic sketch of the hardware configuration. In the final set-up, the camera control computer will be part of the Beowulf cluster to enhance I/O performance and avoid network conflicts.

In the following sections, we use the aforementioned 1024×1024 pixel data to study various implementation aspects of near real-time image reconstruction. However, the first system that we employed at BBSO used a 256×256 pixel, 12-bit, 100 fps Dalsa CA-D1-256T CCD camera.

4. Software Implementation

The image-reconstruction program was originally developed in the Interactive Data Language (IDL) by Research Systems, Inc. Due to the interpreted nature of IDL, the program does not run very efficiently, e.g., on an Intel Pentium III 500 MHz computer, it takes about 30 min to reconstruct a 256 × 256 pixel image out of 50 short-exposure frames.

The image-reconstruction algorithm relies heavily on Fast Fourier Transforms (FFTs). The Fastest Fourier Transform in the West (FFTW, Frigo and Johnson, 1998, `http://www.fftw.org/`) is a freely available C code which very efficiently calculates FFTs. FFTW adapts the computation of the FFT automatically to any particular computer hardware. This self-optimizing approach yields not only an excellent performance but also allows for easy scalability. A first serial implementation of the image reconstruction algorithm for the same 256 × 256 pixel image, as mentioned above, takes only about 3 min which is still not sufficient for near real-time image reconstruction. However, speckle-masking imaging is naturally suited for parallel processing. The data have three dimensions (one temporal and two spatial dimensions). This data cube can be sliced in two different ways, either along the temporal axis or in stacks of small images with sizes according to the isoplanatic patch. In the following paragraphs, we discuss our implementation of a parallel version of the algorithm.

The parallelism of the code is achieved by explicit message passing. There are two common libraries for distributed computing: Parallel Virtual Machine (PVM, Geist *et al.*, 1994) and Message Passing Interface (MPI). PVM is the existing de facto standard for distributed computing, while MPI is an industrial standard for writing parallel programs. Compared with PVM, MPI provide more features, such as more point-to-point and collective communication, parallel I/O, etc. However, since PVM is relatively easy to implement, we chose PVM 3.4.3 as our message passing library. MPI may be tested in the future for a comparison of its (dis-)advantages in our specific application.

Multi-threaded programming allows very fast shared memory communication and synchronization between concurrent sections of a program on SMP systems. Therefore, we decided to 'break-up' the algorithm into several tasks and run multiple threads on the dual CPU motherboards. One thread, for example, handles message passing between different nodes. Another requirement of multi-threading is the isolation of local from global data, so that inefficient data usage is eliminated. However, multiple node systems still require a significant amount of data transfer. Each node of our system duplicates the calibrated, normalized, and aligned 3D data cube and subsequently computes the speckle masking bispectrum for the different isoplanatic patches. This implementation is not fully optimized, but does not currently pose a problem for the 100 Mbit Ethernet connection. However, the larger the image size and the more nodes involved, data transfer, load balancing, and message passing become more intertwined. Contrary to intuition, increasing the

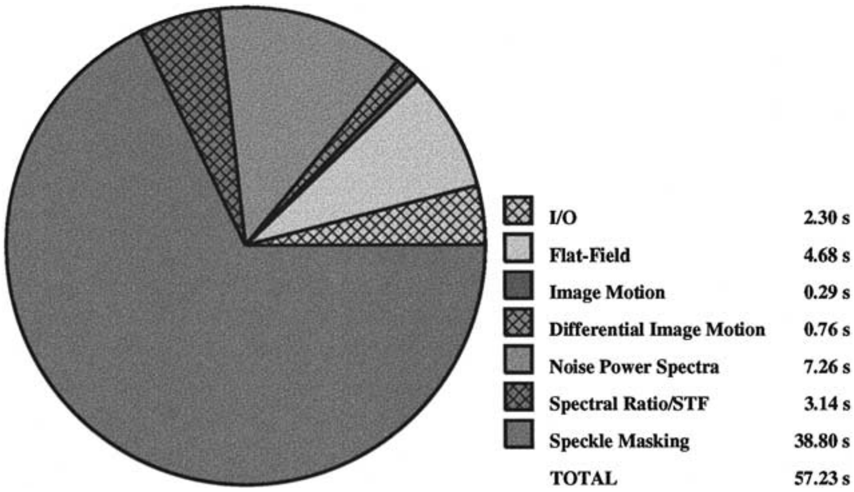

Figure 3. Pie chart of the computing time for the individual parts of the image reconstruction algorithm.

number of nodes and CPU speed, does not necessarily improve the performance of the program. An I/O-dominant program reaching the network capacity might illustrate this fact. On the other hand, Gbit networks are already available even though they are still expensive.

This still 'crude' parallel program reconstructs a 256×256 pixel image in about 57 s. The pie chart in Figure 3 illustrates the amount of computing time spent on different sections of the image reconstruction algorithm. If the camera control is handled by the host node, the time for data I/O would essentially be zero. When used with a correlation tracker, image motion no longer needs to be compensated by the algorithm. However, the corresponding time savings are only minor. At the moment, the noise power spectra are calculated for each speckle reconstruction individually. This is not necessary for the given detector and wavelength setting since in this case we use pre-calculated noise power spectra. Two-thirds of the computing time is currently spent in calculating the speckle-masking bispectrum. The computing time as well as the memory requirements are a function of the truncation of the bispectrum as shown in Figure 4. The technical aspects of the speckle-masking phase reconstruction algorithm and the role of the truncation parameter are discussed in Pehlemann and von der Lühe (1989). We use a truncation parameter of $t = 5$ which still provides reliable phase information. We are studying different ways to utilize the symmetries of the four-dimensional speckle-masking bispectrum to further improve the efficiency of this part of the program.

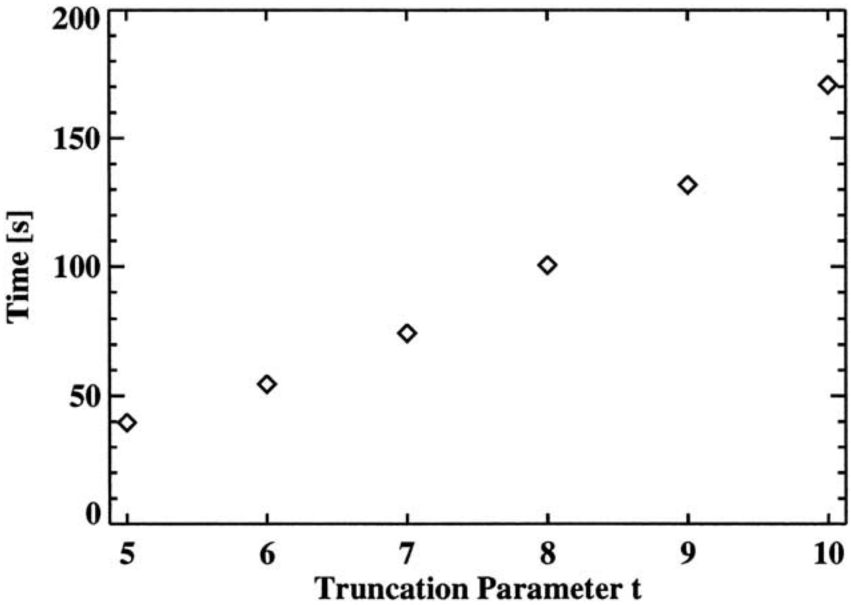

Figure 4. Computing time as a function of the truncation parameter of the speckle-masking bispectrum.

5. Discussion

In principal, there are no limitations to the field of view. At the moment the limiting factor is the detector size of fast CCD cameras which is currently about 2k $\times$ 2k pixels. With the current two-node system, a 1024 $\times$ 1024 (1k $\times$ 1k) pixel reconstruction takes about 1220 s, a 768 $\times$ 768 pixel reconstruction takes about 660 s, and a 512 $\times$ 512 pixel reconstruction takes about 270 s, which shows that the program scales efficiently on this small cluster. The next step is a detailed study on how to adapt the parallel program to clusters with more compute nodes. We expect that a Beowulf cluster, of 32 CPUs at 1 GHz each, is capable of handling 1024 $\times$ 1024 pixel reconstructions at a 1-min cadence. According to Moore's law, processor speeds have been doubling every 18 months. It is often argued that the increase in processor speed should continue for another 5–10 years until we reach fundamental limits of today's silicon technology, i.e., we can expect an increase of at least a factor of eight. 2k $\times$ 2k pixel reconstructions should be within reach over the course of the next few years.

At the moment, direct speckle imaging is limited to relatively broad (a few tenth of a nanometer) wavelength regions. However, filter systems with high transmission, e.g., Fabry–Pérot filters, and detectors with high quantum efficiency will allow narrow-band (less than 0.025 nm) applications such as speckle reconstruction of Hα center-line filtergrams or speckle spectro-polarimetry of solar magnetic structures (e.g., Koschinsky, Kneer, and Hirzberger, 2001). Parallel computing in combination with frame selection and image reconstruction will give us immediate access to the physics of active regions on the Sun. In addition, we avoid saving huge amounts of raw data (100 MByte min^{-1}), and we can check immediately the data quality. Visualization of active region evolution is extremely important in predicting solar flares, filament eruptions and disappearances, as well as coronal mass ejections.

The work described in this paper will be beneficial in exploiting the capabilities of the proposed Advanced Technology Solar Telescope (ATST). In addition to speckle-masking imaging, phase diversity image reconstruction and multidimensional spectro-polarimetry are some other obvious applications. Scalable numerical algorithms and computer architectures ensure that we are in step with the demands for higher spatial, spectral, and temporal resolution, and larger fields of view.

Acknowledgements

This work was supported by NSF under grant ATM-00-86999, by ONR under grant N00014-97-1-1037, and by NASA under grant NAG-5-9682.

References

de Boer, C. R.: 1993, 'Speckle-Interferometrie und ihre Anwendung auf die Sonnenbeobachtung', Ph.D. Thesis, Georg-August Universität Göttingen.

de Boer, C. R.: 1996, *Astron. Astrophys. Suppl. Ser.* **120**, 195.

Denker, C.: 1998, *Solar Phys.* **180**, 81.

Denker, C. and Wang, H.: 1998, *Astrophys. J.* **502**, 493.

Denker, C., Spirock, T. J., Goode, P. R., and Wang, H.: 1999, in T. Rimmele *et al.* (eds.), *High Resolution Solar Physics: Theory, Observations, and Techniques, ASP Conf. Ser.* **183**, 124.

Fried, D. L.: 1966, *J. Opt. Soc. Am.* **56**, 1372.

Frigo, M. and Johnson, S. G.: 1998, *Proc. ICASSP* **3**, 1381.

Geist, A., Beguelin, A., Dongarra, J., Jiang, W., Manchek, R., and Sunderam, V.: 1994, *PVM: Parallel Virtual Machine: A Users' Guide and Tutorial for Networked Parallel Computing.* MIT Press, New York.

Goode, P. R., Wang, H., Marquette, W. H., and Denker, C.: 2000, *Solar Phys.* **195**, 421.

Korff, D.: 1973, *J. Opt. Soc. Am.* **63**, 971.

Koschinsky, M., Kneer, F., and Hirzberger, J.: 2001, *Astron. Astrophys.* **365**, 588.

Labeyrie, A.: 1970, *Astron. Astrophys.* **6**, 85.

Pehlemann, E. and von der Lühe, O.: 1989, *Astron. Astrophys.* **216**, 333.

von der Lühe, O.: 1984, *J. Opt. Soc. Am.* **A1**, 510.

Weigelt, G.: 1977, *Opt. Comm.* **21**, 55.