

LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY  
- LIGO -  
CALIFORNIA INSTITUTE OF TECHNOLOGY  
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

<b>Document Type</b> <b>LIGO-T980092-00 - E</b> 10/8/1998
<h1>Microway Alpha Workstation Speed Testing</h1>
James Kent Blackburn

*Distribution of this document:*

LIGO LDAS Group

This is an internal working document  
of the LIGO Project.

**California Institute of Technology**  
**LIGO Project - MS 51-33**  
**Pasadena CA 91125**  
Phone (818) 395-2129  
Fax (818) 304-9834  
E-mail: info@ligo.caltech.edu

**Massachusetts Institute of Technology**  
**LIGO Project - MS 20B-145**  
**Cambridge, MA 01239**  
Phone (617) 253-4824  
Fax (617) 253-7014  
E-mail: info@ligo.mit.edu

WWW: <http://www.ligo.caltech.edu/>

# Microway Alpha Workstation Speed Testing

*James Kent Blackburn*

California Institute of Technology  
LIGO Data Analysis Group  
June 30, 1998

The LIGO interferometers will begin continuous collection of data sometime after the start of the new millennium. Analysis of this data will be carried out in such a manner as to allow both near real time on-line filtering of the instruments' response, and off-line searches requiring more resources than are planned for the sites to detect gravitational waves from a range of astrophysical sources. Many sources can be characterized in terms of our present understanding of the physics producing the gravitational waves LIGO intends to detect. These include such catastrophic events as the supernova end point in the life of a massive star, exploding in a largely undetermined burst of gravitational radiation to the highly deterministic inspiral as two massive neutron stars or black holes coalesce, or the ringing of a large spinning black hole that has been excited by interactions with infalling matter much like a bell that has been struck. LIGO will be opening a new window on the universe, and the possibility of seeing serendipitous sources of gravitational radiation in the data from the two interferometers is highly likely and will require less specialized, and more robust means of detection.

Carrying out searches for known sources of gravitational waves will rely heavily on the use of optimal filtering methods. These methods require that the data which is acquired as a time ordered collection of signals, be recast into the frequency domain, preserving amplitude and phase information and then correlated with expected frequency domain waveforms in order to achieve the maximum signal to noise ratio for those particular astrophysical sources where we do have insight into the physics and evolution of the systems. Chief among these are the binary inspiral of two neutron stars or black holes and the ringdown of excited massive black holes. In these cases it is possible to estimate the number of independent waveforms each segment of data must be compared with in order to have a reasonable chance of seeing a gravitational wave signal of sufficient strength above the instrumental noise. For example, on a per interferometer basis it will require nearly 35 thousand filters to look for the binary inspiral of two neutron stars or black holes where each companion has a mass greater than or equal to 1.2 solar mass units. It will require nearly 1000 ringdown filters to look for massive spinning black holes with quality factors less than 30. It is estimated that 80 percent of the computation associated with these filters is in carrying out the Fast Fourier Transforms (*FFTs*). If a computer could deliver 200 million floating point operations per second (*MFLOPS*) while carrying out the FFT algorithm, then LIGO would need roughly 35 such computer nodes working in parallel to perform these two classes of filtering per interferometer data stream (*the binary inspiral search makes up about 90% of this*).


Demonstration of such performance is often difficult on the basis of published benchmarks. Performance of a particular piece of computer hardware depends on many aspects of the overall system design such as processor performance, memory cache, operating system and compilers, just to name a few. In order to understand if the target of 35, 200 MFLOPS computers is in scope for the LIGO on-line data analysis needs, an Alpha workstation (see Figure 1 below) from Microway, Inc. ( *their URL is <http://www.microway.com>* ) has been tested by LIGO.

**FIGURE 1. Microway Alpha based “Screamer” Workstation**




The workstation was purchased in Spring of 1998 and consisted of an Alpha LX motherboard (see Figure 2 below) with a 600 MHz Alpha CPU, a 4MB 9ns SRAM cache (*the largest available on the market at that time and capable of going up to 8MB*), 128 MB of dynamic RAM and a 4 GB ultra wide SCSI hard drive. The unit was shipped with a modified version of the Redhat Linux 5.0 operating system. The motherboard also includes 32 and 64 bit PCI bus interface. The workstation includes a standard DEC keyboard and a Iiyama 17 inch monitor and a three button Logitech mouse.

**FIGURE 2. 600 MHz Alpha LX Motherboard**



A series of FFT benchmarks were carried out on this workstation using the ANSI C FFTW 1.3 library available from MIT (*their URL is <http://theory.lsc.mit.edu/~fftw>*) using the GNU C compiler that was shipped by Microway, Inc. This was compared with a Sun Ultra 30 with a 300 MHz ultrasparc processor using both GNU's C compiler on the Sun and Sun's Workshop 4.2 C compiler. The FFTW package includes its own benchmark algorithms which was used as the "standard" for comparing both machines. Since the Sun compiler was expected to demonstrate superior optimization over the GNU C compiler on the Ultra 30 workstation, it was decided that an additional benchmark using the DEC C compiler to build the FFTW benchmark routines would be a good idea if it was determined to be possible to run DEC compiled binaries on the LINUX Alpha operating system. It was learned that the LINUX kernel on the Alpha workstation could be modified to allow statically linked binaries from DEC Unix built using the DEC C compiler to run on the LINUX system. A second vender, Alta Technology offered to provide LIGO with static versions of the FFTW benchmark program compiled with the DEC C compiler under DEC Unix. These benchmarks were made with both floating point precision and double precision complex FFTs. The results are shown in Figure 3 below.

**FIGURE 3. Performance comparison for complex FFTW benchmark code**



Looking at the results we see that peak performance for the FFTW occurs for very small data size (*order 32 complex data points*) for both the Sun and the Alpha workstations. The

most likely reason for this is the number of data registers in the CPUs. There is a sharp fall off beginning at above 1024 data points which then tails off at about 16384 sample points where the Alpha is delivering about 250 MFLOPS of FFT performance. This is short of the earlier design estimates that suggested that working with million point FFTs would be optimal if the performance was flat in the region between 128K and 2048K samples. The FFTs could be shortened to roughly 256K samples without modifying the underlying convolution algorithm and its performance for the optimal filtering techniques. This is driven by the time that the binary inspiral waveforms are in the interferometers' most sensitive frequency band (*a 1.2, 1.2 solar mass binary will be in initial LIGO's band for 70 seconds*) and by the significance of higher frequency structure in the waveforms to the detection process. In this region of interest to LIGO data analysis the Alpha is delivering the best performance but only at the level of about 150 MFLOPS for FFTs (*down by a factor of 5 from peak performance*). The added cache size on the Alpha workstation helps to maintain the large sample size performance to twice that of the Sun with its 2MB cache. Clearly the larger cache is needed for mega-point FFTs.

The Alpha workstation was also tested by the end-to-end group for its performance in that application. It gave the best performance of any computer tested by a factor of between 2 and 3 over others. These shows the tremendous numerical performance provided by the Alpha CPU to a broad scope of problems (*the end-to-end model is not dominated by FFTs*). However, there was a clear downside to using the Alpha workstation, at least with the Redhat LINUX 5.0 operating system. The Alpha port of LINUX was both incomplete (*many system header files were missing*) and unstable in its behavior and ability to compile system level code. Redhat LINUX 5.1 is available and will soon be installed on the Alpha workstation. It is hoped that this new version will fix many of the software problems that have been encountered when using the Alpha workstation. Greater optimization of the GNU C compiler for this platform would also be extremely beneficial.

Microway has already dropped this particular unit's price by 25%. A 633 MHz unit is now available and a 700 MHz system delivering over twice the integer and floating point performance of the 633 MHz unit has been announced for delivery the third quarter of 1998. The other Alpha vendor that LIGO has communicated with, Alta Technology, is commercially producing mini-Beowulf boxes containing 8 533 MHz Alpha motherboards loaded with 128 MB of memory, 4 GB disks, intelligent power supply and an internal switch for under \$2000 per node. If this trend holds and the Alpha Linux stabilizes, then distributed compute engines like the Beowulf on Alphas represents a respectable performance per cost advantage in today's computer market.