Long-Base Laser Strainmeters: A Review

Duncan Carr Agnew

Frank K. Wyatt

Institute of Geophysics and Planetary Physics
Scripps Institution of Oceanography
University of California, San Diego

Scripps Institution of Oceanography Technical Report

6 January 2003

Abstract

This document reviews the history, design, and use of the long-base laser strainmeters developed at IGPP/SIO since 1970: the longest-running and most stable examples of such instruments. We describe the principles of operation, limitations on dynamic range the frequency response, and major sources of error, along with the design and construction methods used. We also provide results on reliability of operation and the methods used to produce final high-quality data from what is recorded in the field: both based on close to 100 instrument-years of experience. We describe the setting of the two long-running instruments in California, at Piñon Flat Observatory and at Durmid Hill; for both locations we summarize results on long-term strain, tidal strains, and various strain ‘‘events’’: the last almost all related to earthquakes, since the results of these high-quality measurements have shown few departures from steady strain accumulation with time.

Table of Contents

1. Introduction

This review aims to provide information about long-base strain measurements, for potential users of such data, giving them some insight into how the data are produced. We have gone into considerable detail because this class of measurements is pursued by few (in the United States, just ourselves), so that the basic principles are not widely known. Some of the material is revised from Agnew (1986), which can be consulted for a broader overview of different types of strain and tilt measurement; in presenting results from long-base strainmeters we have focused on more recent results.

1.1. The Role of Long-Base Deformation Measurements

We begin with an overview of how longbase strain measurements fit in with the more common methods for measuring deformation: GPS and borehole strain. Long-base strain (and tilt) measurements have capabilities not otherwise available, so that we should use these sensors if we are serious about studying deformations. Longbase strainmeters (or tiltmeters) should usually be a part of larger clusters of GPS and borehole strainmeters, in regions for which the aim is to measure transient signals over a wide frequency range.

There is nothing unusual in long-base instruments having both superior capabilities and higher costs: better performance usually comes at a higher price. Given a finite budget and many goals, we do not want to overspend on high-quality sensors when other ones would do; but equally we do not want to be too economical: if what we install cannot measure the phenomenon of interest, we would have been better off spending the money somewhere else.

Given that long-base measurements have lower noise than other sensors in the period range from months to a few years, are there phenomena present which can justify their use? An unequivocal answer is not yet possible, simply because so few transients have been measured. While it might be that all transients have time constants shorter and longer than this period range, it seems unlikely. And if aseismic transient deformations have the same size-frequency relation that seismic ones do, we will see many more of them if we use instruments with lower noise.

To appreciate the unique capabilities of long-base sensors, it is worth remembering that all deformation measurements are of differential displacement: long-base strainmeters measure over baselines of 0.1−1 km, borehole sensors over baselines 10^{-3} to 10^{-4} times as long, and GPS over baselines 10 to 10^3 times as long. For each system, the limits to performance come less from design details (though these are important) than from the environment in which the system operates: for borehole instruments a drilled hole at depth, for long-base sensors the surface of the Earth (but in a controlled setting); for GPS and InSAR also the surface and in the atmosphere.

What performance do these environmental differences translate into? For periods of days and less, the long-base environment is more stable than the GPS one, and the borehole one more stable yet. The two types of strainmeters thus have noise levels much lower than GPS. GPS can provide data at these frequencies, but these recordings are so

noisy as to be uninteresting. There is one exception, namely at the times of large coseismic offsets (and strong shaking). At such times the borehole environment may not be very stable: not uncommonly the large dynamic strains cause readjustments in pore pressure, something which the short baselength of borehole sensors makes them very susceptible to. Long-base systems, averaging over more of the volume, are less affected (Evans and Wyatt 1984), and GPS measurements even less.

At periods of years to decades the stability situation is reversed, partly from the environment, and partly from instrument design. The atmospheric changes that affect GPS matter less because the atmosphere is more stable averaged over long times: the mean air pressure, temperature, and water vapor vary less from year to year than from day to day. But at such periods a borehole is not necessarily a stable environment: hydrologic fluctuations can induce pore-pressure changes. Borehole sensors also depend on the long-term stability of materials: not just inside the instrument (which can be controlled during manufacture) but also the bonding material (less easily controlled), and the recently-drilled stressed rock nearby (not controlled). Long-base strainmeters, like GPS, measure electromagnetic radiation whose stability is tied to atomic standards: something far less susceptible to long-term drift. Like borehole strainmeters, they also have an interface to an uncontrollable material, namely the near-surface materials; but the need for stability is mitigated by the much longer baselength.

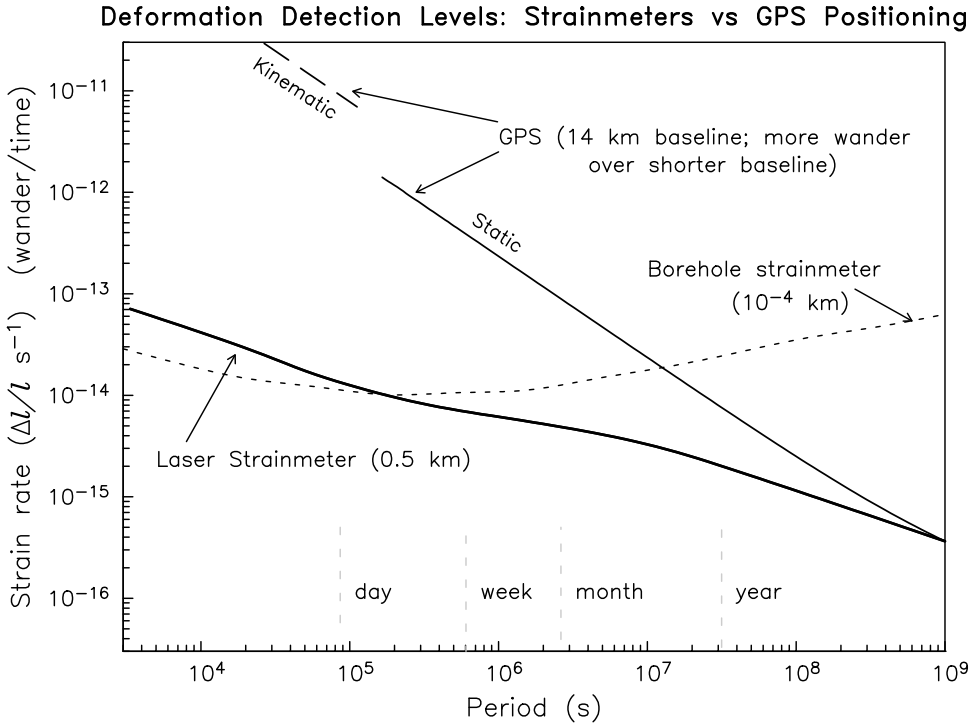

Figure 1

The tie to atomic standards makes GPS the best way to measure secular deformation, though this is also possible with long-base strainmeters: while the shorter baselengths of strainmeters decrease stability, this can be compensated for by the more-

controlled environment of the end points.

How this all works out for instrument performance is summarized in **Figure 1**, which shows estimated noise for GPS (over a typical scale for fault-induced motions), long-base strain, and borehole strain. The two curves for GPS and longbase strain were derived using the method of Agnew (1992); the borehole strain curve comes from the compilation included in the documentation for the Plate Boundary Observatory (PBO). That the two strainmeter curves agree at the higher frequencies implies that the different estimates are of the same quantity: ground noise. At periods of (say) a week, long-base strainmeters show fluctuations about 100 times smaller than a GPS baseline would. This figure shows that longbase strainmeter measurements are a better match to GPS in terms of covering a wide range of frequencies with two methods of comparable sensitivity.

1.2. History of Long-Base Strainmeter Development

In this section we outline the development history of long-base laser strainmeters, as we have pursued it. The first step was taken by Berger and Lovberg (1969, 1970), in demonstrating that it was indeed possible to make strain measurements over hundreds of meters with an optical interferometer. This was done in a prototype instrument installed at Camp Elliot, near UCSD (**Figure 2**); but it was realized even from the outset that this was not a particularly good location to measure tectonic strain changes. The next step, starting in 1971, was to construct three instruments at a place with flat topography, reasonable proximity to active faults, and no sediment cover: Piñon Flat Observatory (PFO), the setting of which we describe in more detail in section 6.1; **Figure 2** also shows the site location.

The initial measurements at PFO showed strain rates that by the standards of the time appeared to be reasonable; but with the accumulation of a few years of data it became clear that much of the apparent signal in the long-base strain data was caused by motions of the end-piers of the strainmeters (large monoliths set about 2 m deep in the ground). The first step was to measure the tilts of these end-piers, a correction that began in 1974. This produced some improvement for two of the instruments, but none on the NWSE strainmeter (**Figure 16**).

This result led in turn to the construction of the first ‘‘optical anchor’’ (Wyatt *et al*. 1982), using an auxiliary interferometer to record the lateral displacement of surface end-monuments relative to points at depth (about 25 m). Installation of one anchor at the NW end of the NWSE LSM (all that funding would allow at the time) provided some improvement, but of course to get the full improvement we had to anchor both ends. The other was subsequently done, and produced, starting in about 1984, a much more stable record.

We next investigated using optical fibers as the light path for the anchors, rather than evacuated pipes; this greatly simplifies the construction of such anchors (drilling especially), and makes them much cheaper to build and operate. We pursued this over the next several years, with initial results summarized in Zumberge *et al.* (1988), and later results in Zumberge and Wyatt (1998). Our conclusion is that fibers can be used for this

purpose—and using them is certainly better than not anchoring—but their temperature sensitivity limits them to being used in a differential mode, which is how they are employed for the lateral anchoring of horizontal strain measurement; as **Figure 16** shows, installation of a fiber anchor at the south (least stable) end of the NS strainmeter definitely improved its performance (there was, again, not enough funds to anchor both ends). Unfortunately, the aging of fibers also makes them inadequate for use in the most precise measurements; for our current (2002) installations we have reverted to an evacuated light-path for the anchors (Section 2.5).

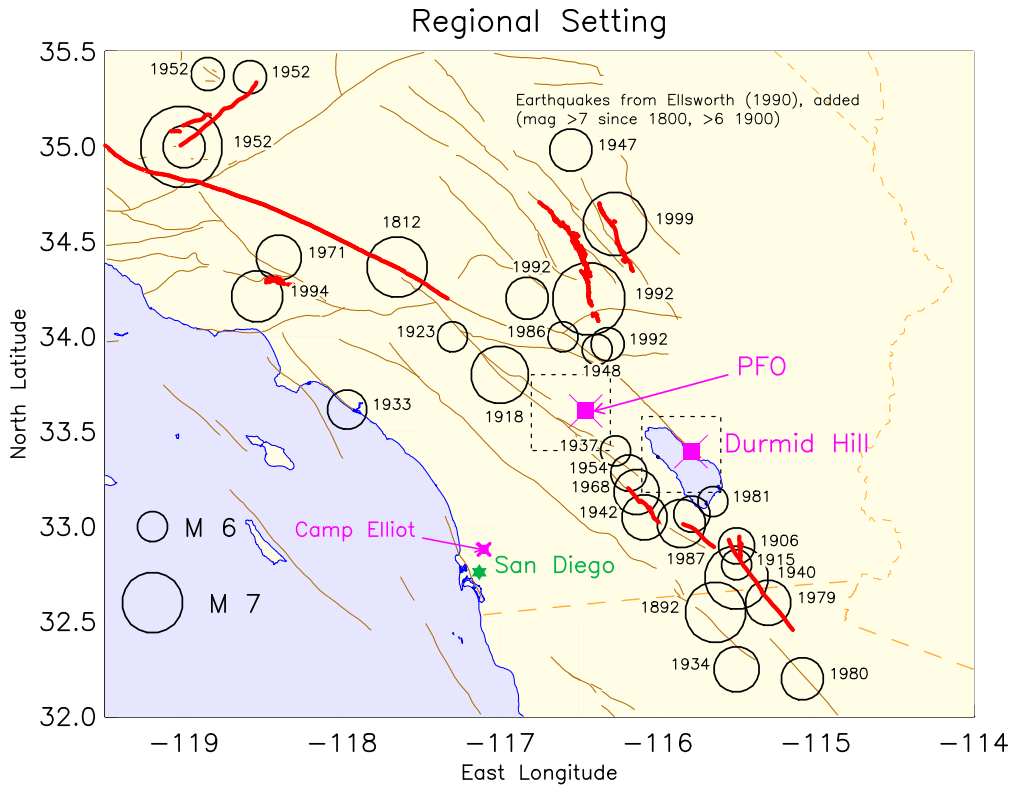

Figure 2

Once we had anchored the two end-monuments of the NWSE LSM, its performance improved to the level that the next most important noise source became apparent: the long-term frequency drift of the laser systems used in these instruments. As we describe in Section 2.4.2, lasers used at PFO were originally stabilized with reference to a physical length standard; these laboratory standards turned out to be less stable than the Earth. We modified the system so that we could monitor the behavior of these standards with reference to an atomically-stabilized reference laser; more recently we have added an atomically-stabilized laser to the NWSE strainmeter. With these combined improvements (in place by 1987) the stability of the NWSE laser strainmeter reached a level not seen, so far as we know, in any other continuous strain measurement. While these measurements were interrupted by a massive brushfire which swept through PFO in July 1994, this did not damage any of the critical facilities—though it destroyed much signal cabling, and the LSM vacuum pipes. The University of California provided funds

for rebuilding; reconstruction of the strainmeters was done piecemeal to ensure that there was always one instrument running, with priority going to the NWSE LSM.

The next stage in strainmeter development took place at a new location, Durmid Hill (DHL), described more fully in Section 6.2. The initial impetus towards work in this area, at the southern end of the San Andreas fault, came with a study in 1988 (Jones *et al.* 1991) of what monitoring was needed for this area, which has a relatively high probability of producing a large earthquake. We built a fiber-optic vertical strainmeter to study the stability of the local near-surface material, and found that it was feasible to use long-base instruments anchored to depth in this area (Zumberge and Wyatt 1998). In January 1991 we proposed the construction of a long-base strainmeter; construction began one year later, the final optics components were installed in February 1994, and the instrument has run since then, though the optical-fiber anchors were not installed and operating until December 1994. The main changes in the instrument were in the vacuum system and counting electronics, and automated beam-steering: these changes combined to produce a much lower-maintenance system than the instruments at PFO had been, and we have been retrofitting the improvements to the PFO systems.

We had originally planned that by 1997 we would be moving the DHL strainmeter to a site in the Los Angeles area, to provide data on possible aseismic strain in a quite different (compressional) tectonic environment from the strike-slip settings of DHL and PFO. In 1996, the Southern California Earthquake Center (SCEC) submitted proposals to NSF and to the Keck Foundation to fund a substantial densification (and expansion) of continuous GPS networks in southern California, focused on the Los Angeles area: the SCIGN project. The proposal to the Keck Foundation, which was funded, included full support for the construction and installation, and initial operation, of at least one long-base strainmeter in the Los Angeles area to complement the SCIGN array. Construction on this instrument, next to the Glendale Freeway in Verdugo Canyon, began in 2001, and was completed in September 2002.

In 1999 we were also funded (by the Department of Energy through the University of Nevada) to install a long-base strainmeter in the south adit of the Exploratory Test Facility of the proposed radioactive waste repository in Yucca Mountain, Nevada. This was started in 2000 and was also completed in mid-2002; in this case much of the construction has been done by contract personnel according to schedules set by DOE.

2. Principles of Operation

In this section we discuss, in some detail, the way in which the long-base laser strainmeter works. We begin with a general explanation of the basic measurement method, optical interferometry, since this is not something most geophysicists work with. We then describe the details of the implementation of this to the laser strainmeter; we provide a detailed description because an understanding of these details is needed to appreciate what the instrument performance can be, in terms of dynamic range and frequency response—which is what we address next. We then discuss three sources of error:

- Variations in the optical path length. These are straightforward to reduce, but they govern the need parameters for a vacuum system.
- Changes in the laser frequency. The stability we need remains a developmental issue, at least for a reasonable price. We give some general background on lasers, for a better appreciation of the problems.
- Motion of the end points. This has turned out to be the single biggest source of noise at long periods, and mitigating it (using ‘‘optical anchors’’) a major source of construction cost.

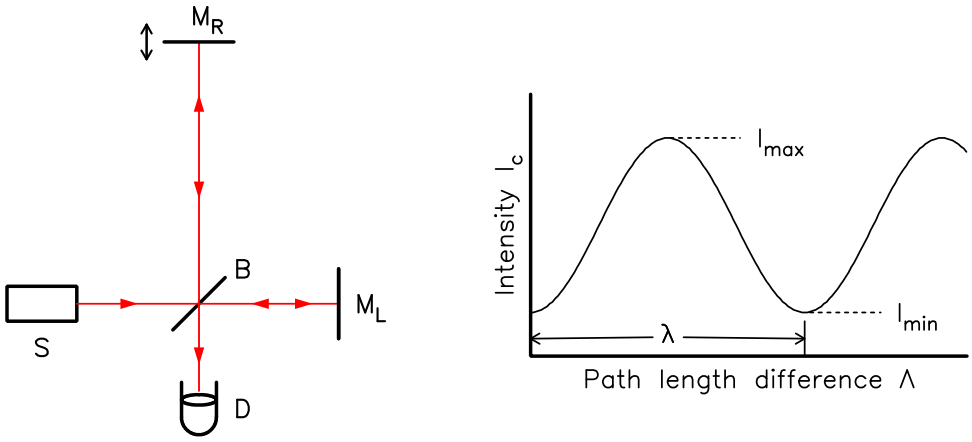

Figure 3

2.1. Basic Measurement Method: Interferometry

The optical system used in the long-base strainmeters is an interferometer, which depends on the wave nature of light. **Figure 3** shows a schematic version (Born and Wolf 1980) of a Michelson interferometer. A beam of light from a source S is sent to a beam-splitter B, where it is divided equally and goes to two reflectors M_L and M_R. The returned beams meet and interfere at B, the interfered energy going (in part) to a detector D. If the source is monochromatic with intensity I, the electric field of the beam incident at B from S will be the real part of $\Sigma = ae^{i\omega t}$ with $I = \frac{1}{2}\Sigma\Sigma^*$. If the optical path lengths BM_LB and BM_RB are Λ_L and Λ_R respectively, then the electric field of the interfering beams will be

$$\Sigma_c = \tfrac{1}{2}\, a\left(e^{i(\omega t + k\Lambda_L)} + e^{i(\omega t + k\Lambda_R + \Theta)}\right)$$

where Θ is a phase shift introduced by the multiple reflections and transmissions and $k = 2\pi/\lambda$ is the wavenumber of the light. The intensity seen by D is $\frac{1}{2}\Sigma_c\Sigma_c^*$ or

$$I_c = \tfrac{1}{2}\, I[1 + \cos(k(\Lambda_R - \Lambda_L) + \Theta)] \qquad (1)$$

which will vary with the path length difference $\Lambda = \Lambda_R - \Lambda_L$ as shown on the right-hand side of **Figure 3**. Equation (1) implies that the intensity should vary sinusoidally between zero and some maximum value: the well-known interference fringes. In fact

imperfections in the optics mean that some scattering occurs, so that the intensity varies from I_{max} to I_{min}: the ratio of this difference to what would be possible in an ideal system is called the depth of modulation. If this depth exceeds a minimum value, measuring the light to dark transitions (fringes) will allow us to detect movements of either arm of the interferometer. For the long-base strainmeter, the arm to M_L (local) is fixed on an optics table; the measurement is made over the longer remote arm, to M_R, whose movements are measured.

All this assumes perfectly monochromatic light; in practice the light will be a mix of frequencies. It is easy to see that for the fringes to be visible the maximum intensity for light at one frequency should not coincide with the minimum for another frequency; this means that the argument of the cosine in (1) should be nearly constant. If the range of wavenumbers in the light source is Δk, this implies $\Delta k \Lambda \ll \pi$. The *coherence length* L is $2\pi/\Delta k$; in terms of this quantity we get $\Lambda \ll L/2$. A more exact analysis of fringe visibility (Born and Wolf 1980) shows that for a source with a Gaussian line shape the fringe visibility will be 80% of its maximum value for $\Lambda < 0.25L$. In terms of bandwidth (the usual specification) this means that $\delta f < c/4\Lambda$, which for a one-kilometer instrument would mean a laser bandwidth of less than about 0.1 MHz. So, unless the interferometer arms are nearly equal in length, the source must have a narrow bandwidth (that is, a long coherence length). This means using a laser (of a rather special kind) as a light source. The main measurement of the laser strainmeter has Λ equal to the instrument length, which is hundreds of meters. For the optical anchor system (Section 3.3) the path lengths of the two arms are made more nearly equal, so the laser can have a shorter coherence length, and less stability.

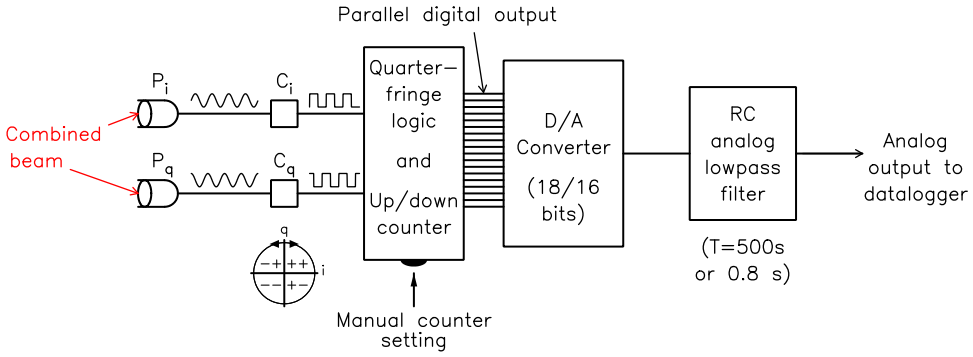

Figure 4

2.2. Recording of the Signal

The interferometer system, in its simplest form, generates changes of intensity in the combined output light beam as the optical path length changes. **Figure 4** shows, in block-diagram form, the strainmeter electronics, which converts these intensity changes into a recorded change in strain.

The first step is actually taken in the optical system: one half of the beam sent to the local reflector has its phase retarded by 90°, so that half of the combined beam is in

quadrature with the other half. Each half is then sent to a separate photodetector (P_i and P_q) so that the available input is two intensity signals, separated by 90°. These can be thought of as the in-phase and quadrature part of a single complex-valued intensity; as the simple sketch in **Figure 4** shows, the result of changes in length is to move this signal in a circle, with the direction of rotation determining if the path length to the remote mirror is increasing or decreasing. A complete rotation of the signal by 360° around the origin corresponds to a full-fringe change in optical path length; In practice, because of differing depths of modulations of the two beams and imperfect phase retardation, this Lissajous pattern is not an exact circle; but unless it is badly distorted its rotation can still be measured.

Because these in-phase and quadrature signals can vary with a high frequency (see the next section) they each are simply digitized with a single-bit system: a pair of comparators, C_i and C_q. The four possible outputs of these two systems then define, as shown in **Figure 4**, four quadrants of the complex intensity; at this level, the signal has been digitized to the nearest quarter-fringe.[1] These two one-bit outputs are then fed into some simple hardwired logic circuits which output a + or – signal to a digital counter. The output of this counter represents the strain change.

The fringe counting electronics is capable of operating at frequencies up to several MHz, but for digital recording at reasonable speeds some filtering of the signal is necessary to avoid aliasing. The output of the fringe counter therefore is used to drive a digital-to-analog converter, whose output is passed through a single-pole lowpass filter and then redigitized. The standard ‘‘tectonic’’ recording system uses a filter with a time constant of about 500 seconds; we also operate channels with shorter time constants, and different gains, for recording seismic waves: the gain is just a function of which bits from the counter are fed to the A-to-D, and the time constant just depends on the filter.

2.2.1. Dynamic Range and Frequency Response

We are now in a position (finally!) to discuss the capabilities and limitations of an optical interferometer for recording earth strain. As in seismic recording, the largest signals come from seismic waves, and it is useful to remember that the peak strain from a plane wave with displacement amplitude d and wavelength S is d/S; if we put this in terms of the wave frequency and phase velocity V we can find that the peak strain rate is $amax/V$, where a_{max} is the maximum acceleration associated with the wave (Gomberg and Agnew 1996).

Dynamic Range. This includes both the lower limit of what can be resolved, and peak values; the latter has to be discussed in terms of both strain and strain rate. First of all, a change in pathlength difference Λ (a full fringe at the detector) corresponds to a change in baselength of $\lambda/2$, since the light path makes a roundtrip to the mirrors. For an instrument of baselength B, this means that one quarter-fringe count corresponds to a

[1] This system in part reflects what was possible when the strainmeters were first built. Direct digitization, to more bits at the same speed, is now possible, and could provide resolution to a small fraction of a fringe—but the current resolution limit is more than adequate.

linear strain of $\lambda/8B$. For $\lambda = 6.328 \times 10^{-7}$m and $B = 500$ m, (the parameters we will use throughout this discussion) each count is thus 1.582×10^{-10} of strain change. As with any other digital system, what ‘‘resolution’’ this corresponds to depends on the sample rate; for 10 Hz sampling this corresponds to a power spectral density of −214 dB relative to $1\varepsilon^2\text{Hz}^{-1}$, which is well below ground noise except perhaps above 1 Hz.

In terms of the upper limit, for strain-rates low enough for the counter not to miscount, the dynamic range of the optics and electronics is infinite: though the counter may ‘‘wrap around’’ (from all bits on to all bits off) as the number of bits it can hold is exceeded, it will not lose track of the up/down counts. Though when the counter wraps around there will be an apparent offset in the analog voltage, such offsets will cancel out for a transient signal—so while a very large transient will be distorted, the strain change across it will be correctly measured. With a 16-bit D-to-A, the limit for no distortion is close to 10^{-5} in strain; of course, unless the analog signal is digitized at least at the same level as the A-to-D converter, the final recorded signal could be limited by the range of this digitizer. For a 1-Hz P-wave, this amount of strain would correspond to a 5 mm displacement. For sufficiently large strains, of course, the optics will no longer be in adjustment; we believe the current system would cease to work for strains much above 2×10^{-5}.

Rate limits and Frequency Response. The electronics of the photodetector are bandlimited at about 1 Mhz, which corresponds to a strain rate of about 10^{-3}s^{-1}: more rapid strains can cause the Lissajous pattern to become so small that the system does not count the fringes accurately. Using the rule of thumb above, this would correspond to an acceleration of 2–3 ms^{-2}, which is strong shaking. The actual limit from local shaking would be set by vibration of the optics, which will cause the aim of the beams to become incorrect, so that fringes will not be formed. In practice, the strainmeters have given accurate records for all teleseisms and regional earthquakes; Table 1 (updated from Wyatt 1988) shows that the cases in which the recording was *not* accurate were either moderate earthquakes very nearby, or large but still local events. As we note in Section 7.3.1, in such events GPS and InSAR record the coseismic effects quite well.

Table 1. Earthquakes Miscounted, Piñon Flat Observatory

Earthquake	Date	Mag	R, km	$a_{\max}$, m/s^2
Horse Canyon	1975:214	4.8	18	0.78
Buck Ridge	1980:056	5.1	19	1.40
Anza	1982:166	4.8	24	0.53
N. Palm Springs	1986:189	6.2	45	
Joshua Tree	1992:114	6.1	43	
Landers[1]	1992:180	7.3	66	0.48
Anza	2001:304	5.1	19	

[1]Power failure at site, still off at time of Big Bear event.

The actual frequency response of the signal as recorded is given by whatever analog filter is applied to the output of the D-to-A converter, and may be adjusted to be whatever is suitable for the recording system. However, there is also an intrinsic limit, since for high-frequency waves, the wavelength of the recorded energy becomes comparable to the length of the instrument, and hence spatially aliased.

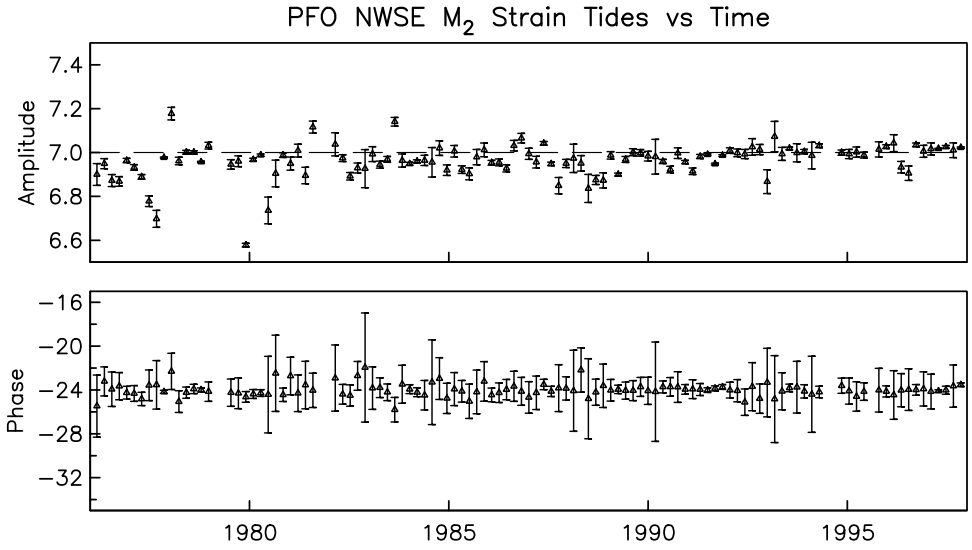

Figure 5

2.2.2. Linearity and Calibration

The strainmeters can be very precisely calibrated. The value of λ is known to within 10^{-6}, and the baselength B can easily be determined to within 10^{-5} using surveying methods: though it is probably not meaningful to give B more accurately than 10^{-3}, because the cross-section of the piers is about $10^{-3}B$. The number of datalogger counts per fringe can be measured to a few parts in 10^4 by manually setting the fringe counter to a value and comparing this with the value recorded on the datalogger; the only reason for a change with time in this would be a drift in the D-to-A converter. The time constant of the lowpass filter can be measured by fitting an exponential to the recorded response of the filter to a step, which can determine the time constant to within 1% (±5 seconds for the standard lowpass filter; this corresponds to an uncertainty in phase of $.004°$ at tidal frequencies). **Figure 5** shows the estimated M_2 tide at PFO over the last two decades; the lack of variation shows that the calibration, and the tides, have not varied significantly (about a 1% level) over this time. (The borehole strainmeter results reported by Hart *et al.* (1982) for the same location did show some significant secular changes in the tides).

Obviously, the interferometer and counting electronics, being digital, are intrinsically linear. Agnew (1982) examined the tidal signals from the NWSE laser strainmeter at PFO and found a nonlinear M_4 tide of amplitude not more than 5×10^{-12}, or .05% as large as the M_2 tide of 1.2×10^{-8} which would imply a squaring nonlinearity of 1 part in

1000, which could easily be present in the analog filter (the large capacitors used for the tidal filter can be slightly nonlinear).

2.3. Error Sources I: Path Length

In Section 2.1, Λ_L and Λ_R are the *optical* path lengths; if light travels through a medium of refractive index n, then $\Lambda_R = n\, l_R$, where l_1 is the physical length. Minimizing the effect of varying n plays a large role in the strainmeter design. At standard temperature and pressure, the temperature coefficient of air is not much worse than for Invar, but the sensitivity to pressure changes (which are harder to shield against) means that an open-air system will give large spurious signals unless both arms vary together. The simplest solution is to regulate the pressure by holding it near zero; this also reduces the temperature sensitivity and nearly eliminates bending and distortion of the light beam by temperature gradients. Mechanical vacuum pumps can easily reach 1 Pa.

It is thus the case that a longbase strainmeter interferometer has two arms: one long one in a pipe in which pressure is kept low but which is exposed to temperature, and the short air paths at the interferometer and retroreflector, which are at (fluctuating) atmospheric pressure, but temperature controlled. Because sliding couplings (or, more recently, bellows) are used to compensate for the expansion of the pipe, the lengths of all these paths may be taken to be constant. It should also be noted that the length of air path that matters in producing changes in the interference pattern is the difference in length between the interferometer arms; this is designed to be small.

For a path of length Λ through a medium with index of refraction n, the number of waves of a wavelength λ will be

$$N = \frac{2\Lambda}{\lambda}\, n$$

The index of refraction of air at the laser wavelength is (Owens 1967)

$$n = 1 + n_0\, \frac{T_0}{P_0}\, \frac{P}{T}$$

where $n_0 = 2.76 \times 10^{-4}$, P and T are the pressure and absolute temperature, and P_0 and T_0 are standard temperature and pressure (1.013×10^5 Pa and 298 K). If we define δ to be the ratio between the path length Λ and the strainmeter length B, we find that the dependence of strain on temperature and pressure is

$$\frac{d\varepsilon}{dT} = -\delta n_0\, \frac{T_0 P}{T^2 P_0} \qquad \frac{d\varepsilon}{dP} = \delta n_0\, \frac{T_0}{T P_0}$$

For the air path, the mean temperature is very nearly T_0, and the mean pressure about $0.86\ P_0$. The path length difference is about 1 cm, so $\delta \approx 1.4 \times 10^{-5}$. The pressure and temperature sensitivities are then

$$\frac{d\varepsilon}{dT} = -1.5 \times 10^{-11}\varepsilon/\mathrm{K} \qquad \frac{d\varepsilon}{dP} = 4.4 \times 10^{-14}\varepsilon/\mathrm{Pa}$$

The spectrum of atmospheric pressure shows that the pressure effects may be neglected;

even the semidiurnal atmospheric tides, with an amplitude of about 10^2 Pa, produce an apparent strain of less than 10^{-11}. The temperature fluctuations of the air path (in the temperature-controlled end buildings) have a diurnal variation of about .25 K, so that these also should not contribute significant apparent strain.

For the vacuum path δ is 1, P is about 1 Pa and T varies between 275 and 300 K. The temperature and pressure dependence of strain is then

$$\frac{d\varepsilon}{dT} = -10^{-11}\varepsilon/\mathrm{K} \qquad \frac{d\varepsilon}{dP} = 2.8\times10^{-9}\varepsilon/\mathrm{Pa}$$

The air temperature typically has a diurnal variation of 5 K, though the temperature of the vacuum pipe may vary by more than this because of insolation. However, increased pipe temperatures also lead to more rapid outgassing, which can result in increased pressures. For the system now in use (no pumping), the direct temperature effect probably is dominant; as the time since the last pumpdown increases, we see an increasing amount of daily cycling of the strain (though much less than the tides).

A small portion of the optical path is in the glass windows used at the end of the vacuum pipe. In addition, two of the strainmeters at PFO have a prism at the center to allow the vacuum pipe to follow the terrain. The temperature effects in this case arise both from a change in index of refraction of the glass and from an actual change in path length caused by thermal expansion. If the coefficient of thermal expansion is α, the apparent strain is

$$\delta\left(\alpha(n-1)+\frac{\partial n}{\partial T}\right)$$

For the glass used, $n = 1.515$, $\alpha = 7.1\times10^{-6}$ K^{-1} and $\frac{\partial n}{\partial T} = 1.3\times10^{-6}$ K^{-1}. The length of the prism (or of both of the windows) is about 5 cm, so the coefficient is

$$\frac{d\varepsilon}{dT} = 3.4\times10^{-10}\varepsilon/\mathrm{K}$$

The end window temperature fluctuates by about .25 K, producing an apparent strain of 10^{-10}, the prisms, though insulated, are not temperature controlled, and probably vary in temperature by several degrees, producing an apparent diurnal strain of 10^{-9}. This relatively large effect has encouraged us not to include prisms in our more recent instruments.

2.4. Error Sources II: Laser Frequency

Obviously, a core element of the optical system described above is the laser; and since these (especially stabilized lasers) are again not generally familiar to geophysicists, we again begin with some background to make clear what is needed for crustal deformation interferometry. We are then in a position to describe the different laser systems we have used, and are using.

2.4.1. Background on Stabilized Lasers

We need a laser which produces light that is as monochromatic as possible (Δk small, or L large) and whose frequency does not change with time. The second requirement comes from the appearance of the product $k\Lambda$ in equation (1); a change in k cannot be distinguished from a change in Λ. If we use a Michelson interferometer with one fixed arm, strain in the other arm and fractional changes in wavelength (or frequency) of the light used are numerically equivalent, so that wavelength stability must be 1×10^{-8} yr^{-1} or better. Stability is usually given in frequency, and for a helium-neon laser 1 MHz frequency change corresponds to a fractional frequency change (apparent change in strain) of 2.1×10^{-9}.

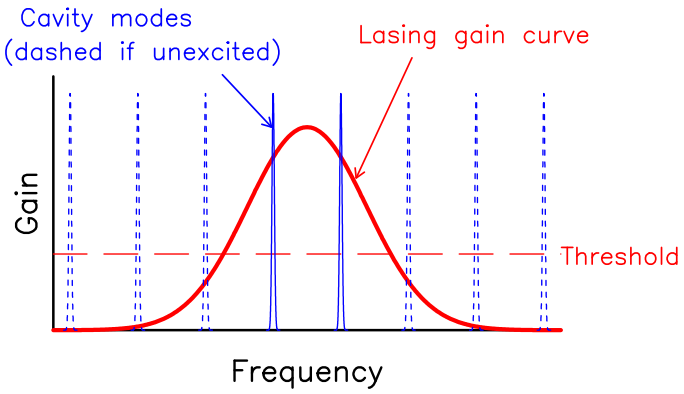

Figure 6

The laser universally used for the best frequency stability is a HeNe gas laser, which is a tube filled with an excited plasma and with highly reflective mirrors sealed on each end. These mirrors form an optical cavity resonator. The excitation causes the Ne atoms to be in an inverted population distribution, with more in a high-energy than in a lower-energy state. Stimulated emission can then take place: a light wave of frequency corresponding to the transition energy can cause an atom to lose energy by radiating exactly in phase with the incident wave. Thus, a light wave with the correct frequency will be coherently amplified as it passes through the tube. The mirrors reflect the light back and forth, with some amplification on each pass until a steady state is reached. The mirrors transmit a small fraction of the resonant light, and it is this light that emerges from the laser.

The wavelength range over which stimulated emission takes place is, for an individual atom, very narrow: $\Delta k/k$ can be 10^{-13} or less, which for the visible Ne line with $\lambda = 632.8$ nm would give a coherence length L of 6000 km—more than adequate. However, in practice the range of frequencies emitted can be much broader. The atoms in the tube are not at rest but have a range of velocities relative to it. A light wave with frequency slightly different from the transition frequency will interact with those atoms whose motion causes a Doppler shift that makes the light appear to have the transition frequency. Because the velocity distribution is Maxwell-Boltzmann, waves with increasingly different frequencies will see fewer atoms with the correct velocity and will be less

amplified as they pass through the tube. **Figure 6** shows the gain curve, or amount of amplification versus wavelength. Doppler broadening means that lasing can take place over a wavelength range of about 2 pm, or $\Delta k/k = 3 \times 10^{-6}$. Outside this range there are too few atoms (and hence too little amplification) to overcome losses: the laser is below the gain threshold and will not lase.

At a given moment the wavelength of light emitted by the laser is determined by the distance between the end mirrors. The wavelength must correspond to a resonant mode of the optical cavity formed by these mirrors. If it does not, the multiply reflected waves will not interfere constructively and there will not be sustained amplification. The approximate condition for resonance is $\lambda = 2\Lambda_c/N$, where Λ_c is the optical cavity length and N an integer, each value of N corresponding to a different longitudinal mode. While the cavity has losses, giving each of its resonances a finite bandwidth, this bandwidth is much less than the Doppler width, giving a gain curve for the cavity which is a series of unit spikes at the resonant wavelengths (also shown schematically in **Figure 6**). In many lasers, there are several modes whose wavelengths fall inside the gain curve; the output of these devices is a mix of frequencies, with a coherence length of 0.1 to 1 m. To get a long coherence length the laser must be made to be *single frequency*, with only one mode excited. (In the laser literature 'multimode' often refers to the transverse structure of the light beam; the modes considered here are longitudinal.) The spacing between modes is $\lambda^2/2\Lambda_c$ so that for Λ_c sufficiently small (less than about 15 cm) there will only be one mode excited; the frequencies of the others will fall below the threshold level of the gain curve.

A single-frequency laser is sufficiently monochromatic to be used in long-base interferometry, but the wavelength of the emitted light depends directly on Λ_c and thus will shift with changes in the physical separation between the mirrors or the refractive index of the path between them. These can both vary a lot (10^{-6} or more), the only restriction being provided by the width of the gain curve; once the mode wavelength moves too far from the center of the gain curve, the laser will stop lasing in that mode. The wavelength must therefore be stabilized in some way; how best to do this has been much studied since lasers were developed (Baird and Hanes 1974), and it remains an active field. Almost all techniques let λ depend on Λ_c, but add a feedback loop that varies the physical length of the cavity to keep the emitted wavelength 'locked' to some particular behavior. A wide range of methods set the lock point using the shape of the gain curve, for example by adjusting two modes on either side of the curve to have equal strength, with only one mode being used in the interferometer (Brown 1981). The lasers we currently use are of this type, the two modes being orthogonally polarized. The resulting stability, while much better than that of a 'free-running' laser, is nonetheless affected by shifts in the gain curve, which alters as the tube ages and as changes in the pressure and gas composition occur.

The most stable lasers use a passive saturated absorption cell placed within the optical cavity, in series with the gain cell which contains the plasma. The absorption cell is filled with a gas chosen to have an absorption line at a wavelength within the gain curve; this absorption line, like the laser gain, will be Doppler broadened but the

absorption will be less at the peak of the absorption curve, giving an enhanced laser output at this frequency. The gas used in the absorption cell, being purely passive, can be chosen for the insensitivity of its spectral peak to pressure or magnetic fields, rather than its lasing efficiency; and, because the cell does not contain a hot plasma, its properties will be more stable with time. What has become the standard system uses a HeNe laser radiating at 632.8 nm, with an iodine absorption cell (Layer 1980). Long-term comparisons between iodine lasers (Chartier 1983) show stabilities of 3×10^{-19} s^{-1} over 6 years, far better than is needed for crustal deformation measurements.

2.4.2. Lasers used in Laser Strainmeters

The lasers used at PFO were originally stabilized with reference to a physical length standard: a quartz Fabry-Perot cavity in a tightly-controlled environment. As noted above, this was still—surprisingly—less stable than the Earth in the long term. Working with Dr. Mark Zumberge, we installed an iodine-stabilized reference laser at PFO, and used optical fibers to send the light from each strainmeter laser to the location of the reference laser. We could then monitor the frequency of the lasers (effectively the length of the cavities) by observing the beat frequency of the combined light sources.

This system proved to be adequate but had two defects: it required considerable effort for the frequent comparisons (visits to the site), and the original cavity-stabilization scheme required a substantial amount of attention, both because of the amount of sophisticated electronics involved and its age. For the DHL strainmeter we needed a laser with adequate stability on its own, and so employed a polarization-stabilized laser which has proven to be stable enough and reasonably reliable; we are now also using them in the NWSE and NS strainmeters, and will also be installing them in our newer instruments. We continue to use an iodine system for occasional checks of the stability of these systems.

2.5. Error Sources III: End Point Motions

The largest source of noise in any good deformation-measuring instruments is how it is attached to the earth: a difficult problem, not least because it must be done in the field rather than in a laboratory. This location also means it is much messier, and much less amenable to quantitative discussion, than the errors discussed in the previous sections. We there make only a few general comments, and then describe the system—the ‘‘optical anchor’’—that is used for the long-base strainmeters.

While the surface of the earth is accessible at low cost, it is not, unfortunately, as stable as is needed. This is not to say that it is decoupled from deeper motions: seismic, tidal, and geodetic measurements made at the surface clearly reflect motion at depth. The difficulty with surface installations is not that signals from deeper motions are absent, but that these signals are masked by other motions of the near-surface layers. Very commonly, strain or tilt measurements made near the surface show a correlation with rainfall; such meteorological noise has been observed in a wide range of settings. One probable cause is deformations caused by varying groundwater pressure; another possibility is the

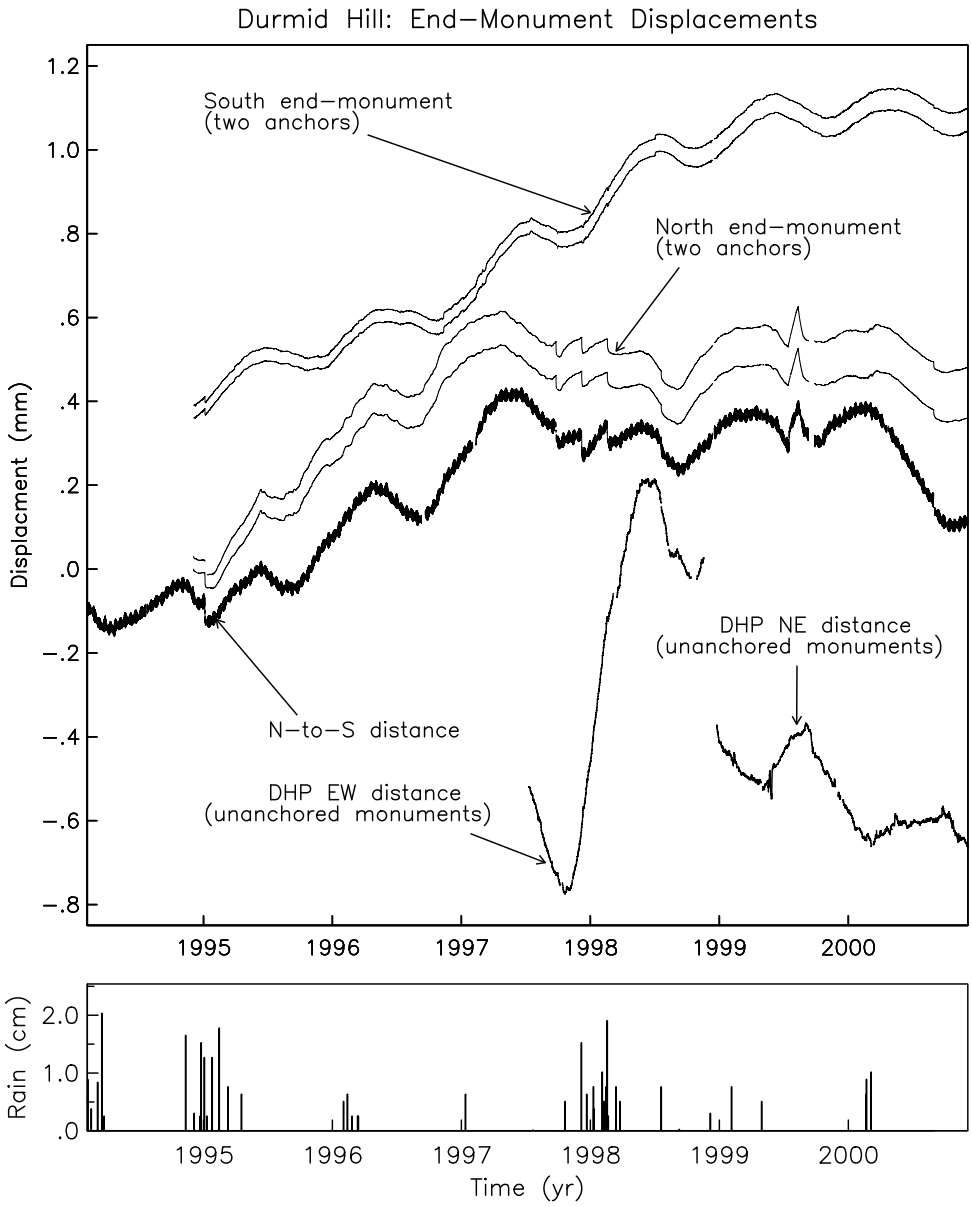

Figure 7

inflation of water-filled cracks (Evans and Wyatt 1984). Shallow installations are probably more affected by local motion of the regolith caused by wetting and drying.

Such motions can easily cause large signals in near-surface deformation measurements. For example, a relative motion of 0.1 mm yr^{-1} at one end of a 10-m instrument would give an apparent deformation of 3×10^{-13} s^{-1}; this is much larger than tectonic rates, but the soil creep causing it would be too slow to be measurable by geomorphological methods. Measurements at PFO (Wyatt 1982) showed that large monuments (1 m^3 or greater) at depths of 2-3 m moved .02 - 0.1 mm yr^{-1} horizontally; the horizontal displacements inferred from tiltmeters at a depth of 4.5 m were 0.002 - 0.01 mm yr^{-1}. The motion was episodic, being correlated with the infrequent rains. Measurements at PFO of

the vertical shifts of monuments 2 m deep show similar rates of motion (Wyatt 1989).

Figure 7 shows a similar results for a very different setting, namely Durmid Hill, which is described in Section 6.2. The near-surface materials are barely-indurated Pleistocene sediments, certainly very different from the batholith at PFO. **Figure 7** shows the displacements of two end-monuments, each monument being measured by a two pairs of fiber-optics anchors (see below). This figure also shows the differential displacements between the two ends, and the same for two pairs of unanchored monuments; of course both of these are just strain measurements, and so show tides. The two monuments associated with the NS system have total motions approaching 1 mm; these are physically anchored by the steel borehole casings, and in temperature-controlled underground vaults. The ‘‘unanchored ’’ monuments consist of inclined and vertical rods driven to 6 m depth, isolated from the material around them to 1.5 m, and inside an air-conditioned building. The relative displacement between these two marks has been much more rapid, and of the same magnitude.

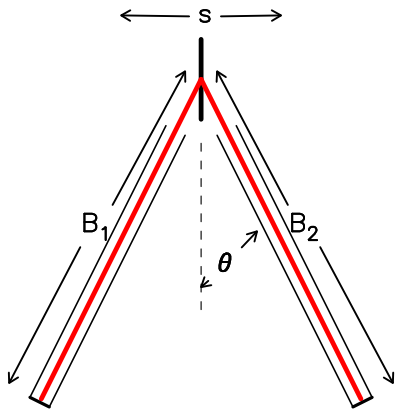

Figure 8

Our solution to the problem of endpoint motion is another optical system (another interferometer) to tie the surface measurements to depth: hence, ‘‘optical anchor’’ (Wyatt *et al.* 1982). This is, effectively, a shear strainmeter, measuring along two equally inclined boreholes, both lying in the same vertical plane as the anchored strainmeter and intersecting at a point on the surface. If, as shown in cartoon form in **Figure 8**, B_1 and B_2 are the distances from the point of intersection to the bottoms of the two holes, θ the angle of dip of the holes, and s the horizontal motion of the point of intersection in the plane of the holes, simple geometry gives $s = 2(B_1 - B_2)\cos\theta$. To measure $B_1 - B_2$, retroreflectors cemented at the bottom of each hole and a beamsplitter at the top form a Michelson interferometer.

Optical anchor response, and range. Because the interferometer arms in the optical anchor are of nearly equal length, an unstabilized multimode laser can be used as the source. The fringe-counting system is the same as for the main interferometer, and so has the same frequency response (flat to very high frequencies) and dynamic range (10^{-4} strain without wrapping, which would be a displacement of 0.5 mm). The limit on the

physical displacement is probably several mm.

Optical anchor stability. The long-term stability of the anchor depends on two factors:

1. Changes in the optical path length. The first anchors (at PFO) used an evacuated path; but this was so difficult to construct that we turned instead to optical fibers, which require neither a large nor a straight hole. In the terms used in Section 2.3, this means that the entire path is glass, so that the analysis for the center prism applies, with the length being the difference between the two arms. This can be made small, but avoiding temperature effects then requires that both arms experience the same temperature changes. Given the symmetrical design of the optical anchor, this is roughly true. A bigger problem with the fibers is that the index of refraction *n* varies with time as the fiber ages; and the rate of variation is not the same even for adjacent sections of the same fiber. Anchors using fibers thus drift somewhat with time; **Figure 9** shows the differences between the two setups shown for each end-monument in **Figure 7**, and shows rates of up to 2×10^{-5}m/yr, which for a 500-m strainmeter would give a strain rate of $4 \times 10^{-8}\,\mathrm{yr}^{-1}$: close enough to many tectonic rates that we have come to believe that, except in areas of high strain rate, an evacuated-pipe system is necessary.

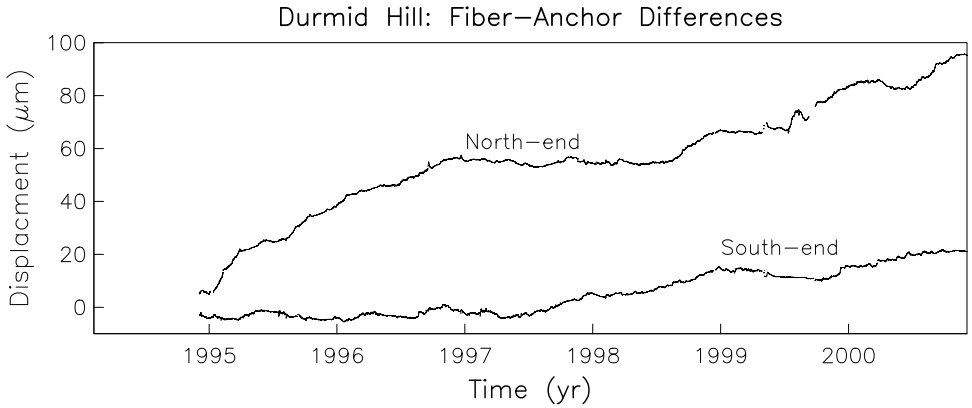

Figure 9

2. The coupling of the bottom of the anchor to the ground. We cement the remote relectors in with expansive grout, but in the end how stable this is can be shown only by the end result. As Section 7.1 shows, fully-anchored strainmeters show excellent long-term stability.

3. Construction Details

In this section we discuss some of the construction details needed to actually realize the systems described in the previous section, with a special emphasis on particular design choices which have to be made.

3.1. Physical, Vacuum System

The physical design of the strainmeter is dominated by the vacuum pipe connecting the two end buildings: these components dominate the weight (about 2×10^4kg). Except in the first (Camp Elliott) instrument, the vacuum pipe has been made from aluminum sections 24′ long, with flanges and O-ring gaskets between each section. For ease of adjustment the interferometer optics and remote retroreflector are not in the vacuum. Because of this, any motion of the end of the evacuated pipe changes the proportion of the path in air and in vacuum, and therefore the optical path length, with 1 mm of motion causing an apparent strain of 0.4 nε. Temperature changes cause the aluminum pipe to expand and contract up to 0.6 m. To eliminate this problem, the vacuum pipe is anchored in the middle, and joints are put at each end that keep the end of the evacuated path a constant distance from the optics. A prism at the midpoint of the pipe allows a break in the optical path, so that the pipe need not be in a single straight section. The optics at each end are inside heavily-insulated boxes in small air-conditioned buildings.

One byproduct of anchoring the strainmeter endpoints, unanticipated and at first not appreciated by us, is a much improved vacuum system. With unanchored endpoints, the ends of the vacuum pipe needed to be sealed with a fedback, telescopically-sliding (and, as it turns out, ‘‘leaky’’) joint, so as not to introduce stresses in the ground close to the ends. Deep-anchor referencing of the endpoints eliminates the need not to stress the ground, so we can instead use a passive bellows system cemented to the ground. The first full implementation of this, at DHL, showed that the vacuum in the main light path could be maintained at an acceptable level (< 10 Pa) for months without pumping; the three instruments at PFO had always required continuous pumping with large mechanical pumps. Not having to run the pumps makes the strainmeter a much lower-power system, reduces the electricity bill, and makes the data largely resistant to the inevitable power failures, since the entire instrument can be run on backup uninterruptible power. This is a substantial improvement. (Given this finding, we installed similar bellows systems into the NWSE LSM at PFO after the fire, with the desired beneficial results.) The vacuum in both strainmeters degrades at a rate that depends on temperature, suggesting that outgassing is the cause; the rates deduced at about 10^{-8} W/m^2, which is reasonable for an unbaked metal system.

3.2. Power and Electrical

We currently estimate the power consumption of the strainmeter at around 1200 W. Of this, the largest part (about 900 W total) is for the conventional air conditioners on the two end buildings, which provide the first stage of thermal control for the strainmeters; the actual optics table and end monument are heavily insulated, and will have a second-stage active temperature control (about 40 W). The three lasers for the various interferometers use about 90 W, and the strainmeter electronics about 50 W. The datalogger (a PC) and telemetry account for the remaining 100 W. In the older systems, the continuously-running vacuum pumps accounted for another 1.5 kW, effectively doubling the power requirements, but with the improved vacuum system these pumps are run only a few days each year, at times selected to be when there are no other problems.

We follow standard procedures for grounding our equipment (tying all grounds to a single point) and for isolation from power-line spikes. We have not had significant problems from lightning at PFO, despite regular thunderstorms in the summers. Occasionally some electronics is damaged by nearby strikes, though it is notable that a nearly direct hit in November 2001, which damaged many facilities, did not affect the strainmeters. Of course, the downhole components (and much of the uphole measurements) are optical systems, and so not affected by lightning; even a direct hit on the end buildings would not affect the downhole reflectors.

3.3. Optical Anchors

The physical layout of the optical anchor is, at each end of the instrument, two boreholes, inclined at 30° from the vertical, both in the vertical plane aligned with the azimuth of the strainmeter axis. These holes need to intersect at a point directly in-line with the main strainmeter beam, and as close as possible to the main beam-splitter (or retroreflector) this turns about to be about 0.3 m away. The casings for the boreholes extend into the cement columns which form the end-points for the strainmeter, with the intersection just above the optics table which is secured on these columns. Reflectors mounted near the bottom of the boreholes serve as the fiducial reference points for the instrument.

Proper orientation and alignment of the two boreholes are essential to the operation of the instrument, and a challenge to drill. For the first installations at PFO we used a water-well, top-head-drive drill rig, positioned above each strainmeter vault and locked at 30° from vertical. The holes were drilled to 23 m with an 16.5 cm air rotary percussion drill; to maintain straight holes the hammer bit was followed by a 6.4 m stabilized drill rod, made by welding three steel bars in a spiral along a standard drill section so that the overall diameter was 16.5 cm. Profiles of these holes showed nearly uniform curvature, with a deviation of less than 5 cm from the straight line connecting the end points. The final clear aperture, necessary for the laser to illuminate the center of both remote reflectors, was 3.5 cm, barely larger than the beam diameter (2 cm). More recently, in Los Angeles, we contracted with Malcolm Drilling, who used a track-mounted double-head rotary drilling rig (Klemm 860). This equipment was much more suited for the task, though it lacked a means to stabilize the drill bit, With this system we achieved adequately straight holes for anchoring to depths of 14 and 21 m, as planned for that installation.

Each of the boreholes is cased with PVC pipe, with a 1.2 m stainless steel anchoring assembly threaded onto the end of the casing. The assemblies are cemented into place using non-shrink grout. Rubber pads are attached to the bottom of the anchors to reduce axial loading caused by borehole rebound and deformation of the cement. The cementing grout is pumped through check valves in the bottom of the stainless steel assemblies and up the outside of the PVC casings until it reaches the surface.

Near the upper end of the stainless anchoring assembly, a tapered and threaded insert serves as the mating surface to guide and secure the retroreflector housing. This

housing is attached, by stainless steel bellows, to the end of a long stainless steel vacuum pipe. The bellows are required both to isolate the reflector housing from the vacuum pipe mechanically and to allow compensating pressure on both sides of the reflector. The last stage of assembly involves lowering the vacuum-pipe assembly into place and twisting it until the downhole reflector housing is secure. The vacuum pipe is left supported at the surface, assuring that the reflector housings, though still coupled through the bellows, are not subjected to excessive longitudinal forces from the weight of the tubing or from thermoelastic stresses.

We measure the monument displacement using (again) a Michelson interferometer; because the two optical paths for each anchor are very nearly equal, this system does not require a laser with particularly high frequency stability. The frequency stability of an inexpensive He-Ne laser ($\Delta k/k \approx 10^{-6}$) gives a maximum error of only 5×10^{-9} m for a representative path length different of 5 mm.

To reduce possible errors, the pressure in the anchor vacuum pipes is lowered by connecting each pair of pipes to the main vacuum system of the laser strainmeter. Because of the choking effect of this connection, the pressure in the optical anchor vacuum paths can be as great as 10 Pa—though this is more than adequate to eliminate temperature differentials in the two paths as a source of error. However, pressure differentials of only 1 Pa would correspond to a displacement of 8×10^{-8} m, which is roughly the limit of resolution of the instrument. Maintaining the absolute pressure below 20 Pa limits the magnitude of the differential pressures to an acceptable level.

4. Operations

4.1. Reliability

Because our data-editing procedures (Section 5) include creating files which describe all edits, it is not difficult to state the reliability of the complete system. Table 2 gives the statistics of 2495 days of data from DHL (a more modern system, run without many changes), from 1994 through 2000. The total down time is about 5.5%, but about half of this is in very long interruptions; the maximum gap is 17.56 days. The many shorter gaps require an efficient edits-handling system but do not pose any real problems; Section 5 describes the editing procedures we follow.

4.2. Required Maintenance

What do we have to do to get this reliability, what causes the gaps in the data, and what can we do to improve the situation? The laser strainmeter requires semi-regular visits as the optical alignment gradually drifts from proper adjustment. The principal cause for this, misalignment of the beam to the distant reflector, is something we can and do automatically correct to a large degree, but there are limitations to what can be done as some of the adjustments are coupled (non-orthogonal), making simple fedback controls inappropriate.

Table 2: Statistics of Edits for DHL

Gap length hr	# gaps	% of series	% series OK if not fixed
0-1	145	0.13	94.54
1-2	71	0.17	94.67
2-3	34	0.14	94.84
3-5	40	0.27	94.98
5-12	42	0.52	95.25
12-24	11	0.30	95.78
24-48	7	0.34	96.08
>48	18	3.58	96.42

Frequency of visits. For the new-generation instruments we anticipate the need to visit at least every other month (6 times a year), after the initial shakedown period. We currently attend to the instruments about every two to three weeks, with the bulk of these interventions being handled by on-site personnel. The biggest sources of data gaps, in roughly decreasing order of importance, are:

Power failures. Disruption of the power has proven to be the greatest source of trouble for the laser strainmeters, often requiring visits to reset various control systems. One of the more difficult issues actually stems from the automated laser-beam steering: this control system, while trying to improve the alignment, can actually make erroneous unrecoverable adjustments depending on how the power is lost and restored. Also, because the frequency-control circuitry for the laser has a limited control range, it often needs to stay off for about an hour to equilibrate before being reactivated.

For providing power during power outages we use an Uninterruptible Power Supply (UPS). Most outages are very brief (seconds to minutes) and are covered by the UPS power. (Note it used to be the case that the vacuum pumps were running continuously; with the motors for these being well beyond any battery-supplied backup, the vacuum was lost with every power outage. This is no longer the case.) Given the inescapable power demands of the LSM (starting with the HeNe lasers: 40 W each) and the capacity of affordable UPS's, we are generally limited to providing only a couple of hours of power backup. After an hour or so, the system will fail and often adjustments are be needed to recover from this (see Improvements below). Ironically, in the course of multiyear operations one not-uncommon reason for power outages has been the failure of the uninterruptible power supplies.

Datalogger. While not formally ''part of the strainmeter'' this is a contributor to the data gaps—actually a major one. This is not an uncommon experience: it seems to also be true for the global seismic network.

Laser. The laser we used is a limited availability state-of-the-art system, which we use near the limit of its capabilities. This requires working with the manufacturer to maintain the needed performance. Several advancements have been made in the last year and more are planned. Before we adopted the commercial equipment we

were maintaining our own systems which, in such limited numbers, were very difficult to support.

Air-conditioning. Traditionally, failure of air-conditioners has been a problem. We require they run continuously, and they usually develop problems after a few years of use. Temperature control is essential to (1) keep the optics well aligned and (2) reduce thermal contamination of the recordings. Failure of a unit, or worse yet, undetected, gradual degradation of thermal control has often meant losing days of recordings. Our new installations feature side-by-side units for ease in switch-over.

Electronics. Most of the circuitry is homemade. Some of the circuitry dates from the early 1970s and still works well. We have had some problems with aging, but our new installations will use all new circuitry and circuit boards.

Environment/physical plant. Extreme weather conditions can physically disturb the instrument; so would vandalism, though so far this has not been a problem. At times floods, lightning, fires, and snow have led to time-consuming problems. We hope we have learned from this; the design of the new systems incorporates many new features. Given the size of the instrument, it is inevitably exposed to the elements.

What would improve matters? Certainly, reliable telemetry, duplicate systems and spare parts on hand, and capable local assistance are key elements. Telemetry allows us to know about problems quickly, and having alternative components available for immediate use often helps identify and remedy instrument problems: in this respect, the main laser, the datalogger, and the air conditioners stand out.

4.3. Longevity and Preventive Maintenance

Since we have been operating the instruments at PFO for about 30 years, we have plenty of information about longevity of components. First of all, nothing except the anchor is inaccessible, and this consists of only vacuum pipe and a glass reflector, which will last a very long time. The main vacuum system has a lifetime of at least 25 years (how long the PFO systems ran before the fire), especially since we have replaced the sliding end-joints with bellows (though these last probably have a shorter lifetime than the rest of the system). The electronics has at least a 30-year life—we are still using most of the original systems. The shortest-lived component of the strainmeter proper is the lasers: these last about 3—5 years (typical for hard-sealed tubes). The other items which need regular replacement are the dataloggers (PC's), and air conditioners, every couple of years or so.

These lifetimes suggest that the primary need in maintenance would be replacing the laser at regular intervals (say, every two years) since the failure is gradual and difficult to detect at first. Failure of the air conditioner and datalogger is obvious, so on-site spares, used as needed, would be adequate in this case.

5. Data Processing

Geophysicists who are used to working with seismic and GPS data tend to be puzzled by the emphasis laid, in dealing with strainmeter records, on the importance of distinguishing between raw and processed data. Since we have worked with all three data types (one of us having for 18 years distributed data from the Project IDA seismic network) we may perhaps be well-placed to offer some explanations. Certainly, it never seemed necessary to process the IDA data before distribution; whereas it does seem very important to do this for continuous strainmeter and tiltmeter data, to keep users from being misled. Why is this?

One simple reason is that seismic data are, in between earthquakes (which is most of the time) not of much interest; the user can focus on only a small section, and any problems outside this will not affect the interpretation. This is something of a special case of a more general statement about seismic data: there is a well-defined expectation, based on 100 years of studying earthquake records (and interpreting them using elastic-wave theory) as to what the data ought to look like—which is enough to decide if the recorded data are in fact probably all right.[2] Somewhat the same thing is true for raw GPS data: the GPS system is so tightly engineered to produce a specific signal that it is straightforward to decide if recorded phase and pseudorange data are valid or not. And there is, besides, massive redundancy in going from raw GPS data to (say) a daily position, which also makes it relatively easy to judge the correctness of data values. Thus, GPS processing is now something that can be (and is) completely automated.

By contrast, for strain data our prior expectations are much weaker. especially for long-term behavior or aseismic events. If we are to get reliable results, it is important not to be misled by instrumental problems; but given the technical difficulty of the measurements, great care is often needed if we are to be sure that some fluctuation is in the Earth rather than the instrument. These same issues arise when looking for temporal fluctuations in GPS time series, for which close attention must be paid to possible system changes or reference-frame variation; in both cases part of the difficulty also is that the most interesting signals are often very close to the noise level. (Indeed, sometimes the most important result from strainmeter records is the absence of a signal—in which case the aim of processing is to make the resulting bound on a theoretical model as tight as possible.) Part of not having strong prior expectations about strain data is that we cannot be sure when there might be (or in retrospect turn out to have been) a significant fluctuation in a strain record—so in a quite literal sense, every point counts.

These reasons explain the importance of processing strain data as part of interpreting them. So it should be no surprise that we, like other groups concerned with producing long time series of strain data, have developed fairly elaborate systems for processing raw strain data. We describe ours briefly here, with reference to the flowchart shown in **Figure 10**. This chart may seem a bit complex; we would emphasize that the procedures

[2] It is notable that when such prior expectations are not available, as with the first lunar seismic records, the result was uncertainty over whether the unusual seismograms recorded might have reflected instrumental problems (L. Knopoff, pers. commun.).

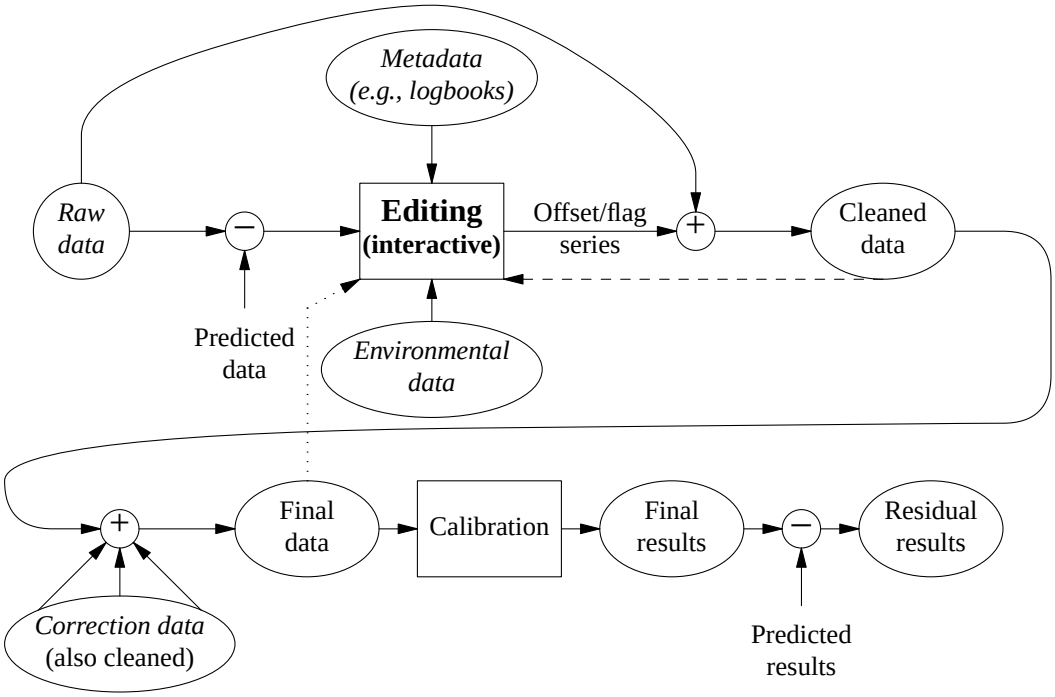

Figure 10

have been developed with the aim of utilizing the data for research, which has encouraged us to try to get the most information possible (including very small signals) from the recorded data—inevitably, not a simple task.

The upper part of **Figure 10** shows the procedure for an individual data series: we first remove as much ‘‘predictable’’ energy as possible, whether from a true prediction (as for the tides) or from an auxiliary series (as the air-pressure correction to water-level in a well). The resulting series is then edited using an interactive program which displays the data (raw or as edited) and allows it to be flagged as bad, and offset to remove jumps. This editing is of course a matter of judgement—but no automatic system can do as well. Obviously, this judgement is aided by experience; it is also helped by examining supplementary information, notably records of environmental data (rain, sun, temperature...) for the same time, and also what we have called ‘‘metadata’’ — the whole range of records of what was done to the instrument when, which we record on forms and in notebooks. This last category of information is of course the most difficult to transmit to a larger group of users.

The result of the editing process is information about which parts of the time series to discard and what offsets to add to it; this editing information could also (but does not for us) include time-varying scaling. When the editing information is combined with the raw data, we have cleaned data. The dashed line in **Figure 10** expresses a kind of feedback which plays a role in the process: not infrequently, only after the data have been edited can subtle problems be identified which call for further editing. The next steps, shown in the lower part of **Figure 10**, are to combine the data series as needed to produce a final estimate; for example, the final strain time series from the laser strainmeter

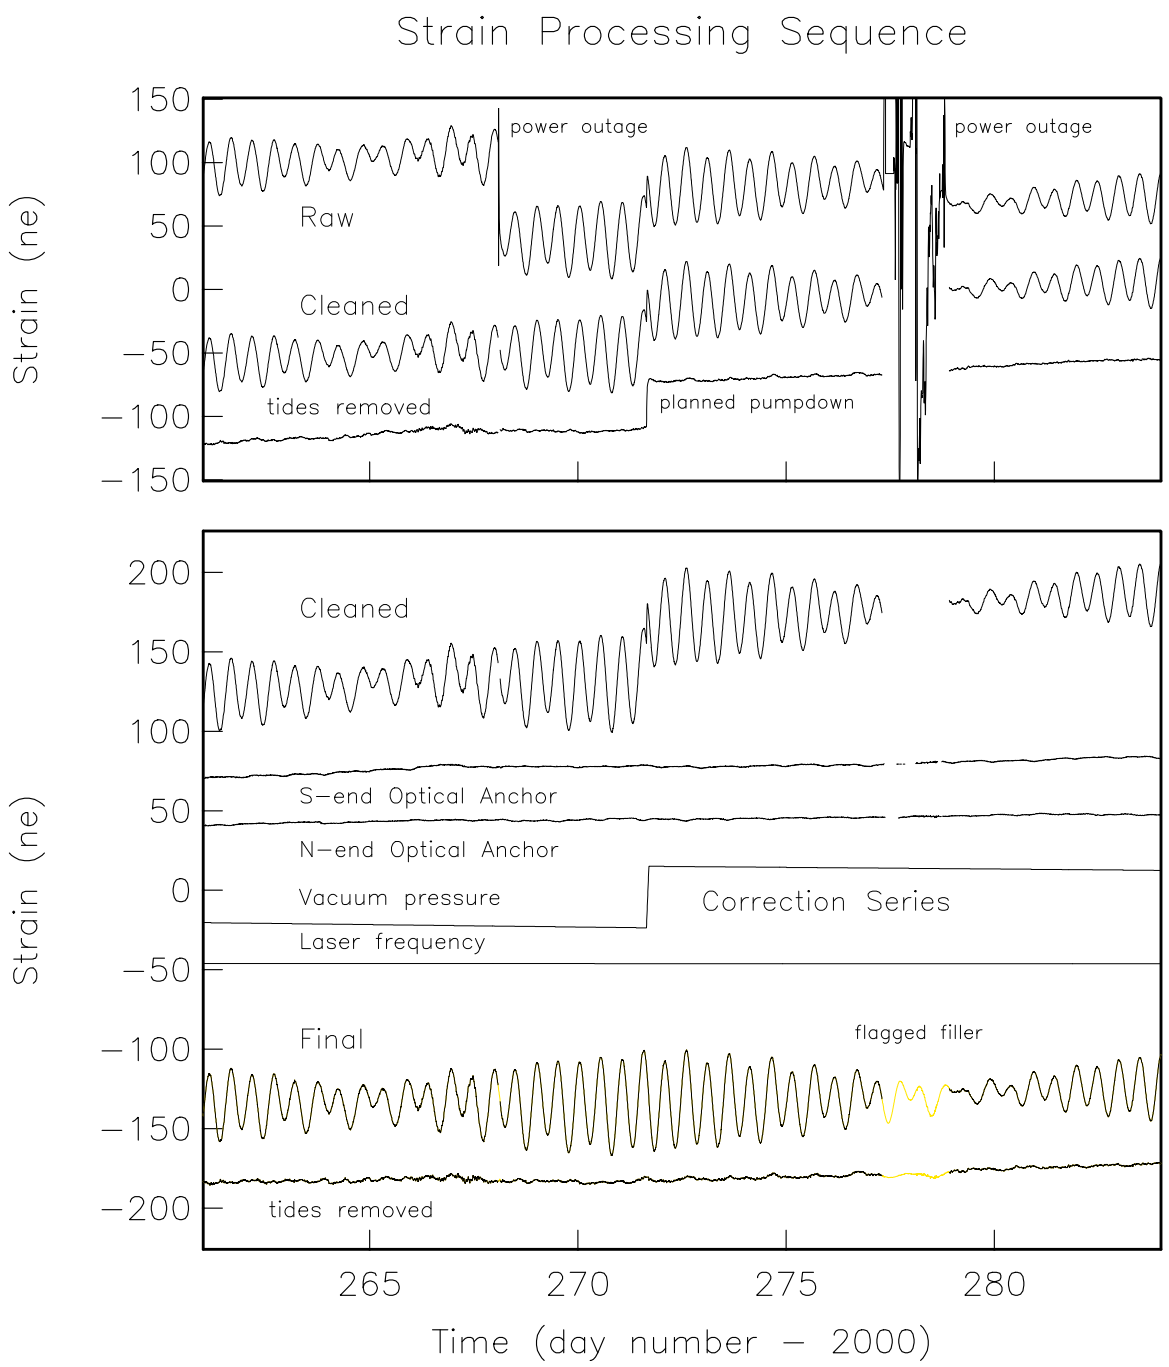

Figure 11

requires the cleaned strain to be combined with two optical anchor series and a correction for laser frequency. **Figure 11** shows this process for some recent data from DHL: in **Figure 11a** the raw and cleaned strain; in **Figure 11b** the strain, the correction series, the final combined series, and the ‘‘residual’’ series (the final series with predicted energy again removed).

It is sometimes argued that because laser strainmeters use interferometry, they cannot monitor long-term deformation, because if the fringe count is lost there could be an unknown amount of strain (analogous to a GPS cycle slip). This is not a problem in practice; if the fringes are lost from (say) a power failure there is little ambiguity about how to interpolate across a short gap—and as **Table 2** shows, most gaps are short. **Figure 1** shows the reason for this lack of ambiguity: for gaps up to 10^5 s, the expected wander of the strain is less than 10^{-9}, and so it is no difficult to estimate the amount of offset to this level; obviously, longer gaps are a challenge, though still (as **Figure 11** shows) not a major contribution to the uncertainty in the long-term series.

Because the distribution of data outages includes a wide range of time intervals and possible offsets, it does not appear possible to automate the data-editing procedure easily. There are certainly ways in which a rough edit could be done by assigning limits to the range of data (and its time derivative), especially by working with data from which the tide has been removed. However, without considerable programming effort it does not seem likely that an automated editing system can do as well as an interactive system; the person using the interactive system can look ahead by varying amounts and examine a range of information that would not be available to an automated package. Any simple system would thus not eliminate the effort required, but merely make a rough version available earlier. Since, to this point, the primary aim has been to produce the highest-quality results, we have not invested in developing a system for approximate editing.

6. Locations of Observations: Seismotectonic Background

In this section we describe the settings of the two longbase strainmeter installations. For each one we include a figure (e.g., **Figure 12**) showing the region around the strainmeters with velocities from a preliminary version of the SCEC Crustal Motion Map, Version 3 (arrows; we have omitted confidence limits, but the errors are typically 1 mm/yr); and also the locations of SCIGN sites (triangles with names). Because this preliminary version of the CMM included data only up to late 1999, some of the SCIGN sites do not yet have velocities associated with them, since there is not enough data. A constant velocity has been subtracted from all the results shown to make it easier to see any pattern.

6.1. Tectonic Setting of Piñon Flat Observatory

PFO is in the region of strike-slip faults that extends south from the Transverse Ranges into northern Baja California. A profile of these velocities, along a line perpendicular to the local strike-slip faults, shows that the deformation field around PFO is close to simple shear, with a velocity gradient of $2.4 \times 10^{-7}\ \mathrm{yr}^{-1}$. This is in accord with the expected deformation from the two closest active faults: the San Jacinto fault zone (14 km SW of PFO) and the San Andreas fault zone (25 km NE).

The San Jacinto fault near Anza (closest to PFO) has a slip rate of about 11-12 mm/yr; as the top panel of **Figure 1** shows, this section has not had a large earthquake in historic time. Paleoseismic studies show the last large earthquake on this part of the fault

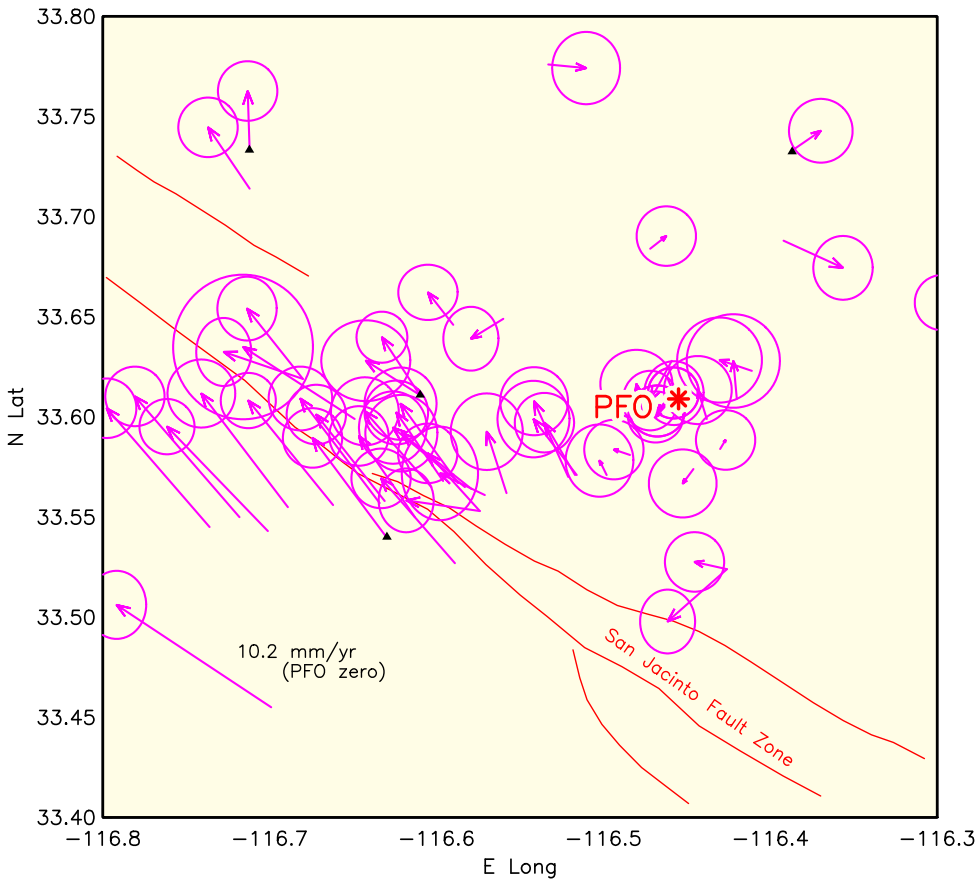

Figure 12

was about 250 years ago (Petersen and Wesnousky, 1994), giving a slip deficit of around 3 m. The most recent formal estimate of a 30-year large-earthquake probability for this part of the fault (WGCEP, 1995) is about 20%. The lack of a known large earthquake on this part of the fault, and the high activity the rest of it has shown in this century, are a large part of the reason why PFO was originally placed where it was.

The site for PFO was chosen partly as a location near this two ‘‘overdue’’ faults, and also because the size of the instruments planned required a large flat area not covered by alluvium. The flat area is Pinyon Flat: part of the southern California batholith (mid-Cretaceous granodiorite). The top meter of material at the site is nearly fully decomposed; below this level the weathered rock grades to highly competent grus at about 3 m, which has decomposed in situ. From 3–25 m the material grades from grus to grus with corestones and finally to jointed granodiorite; borehole logging shows that ‘‘basement’’ velocities (5.4 km/s for P and 3 km/s for S) are reached at about 70 m, implying that there is little deeper weathering. The degree of fracturing in this deeper material is low compared to batholith rocks closer to the San Jacinto fault (Fletcher *et al.*, 1990; Radzevicius and Pavlis, 1999).

Figure 13 is a sketch map of PFO, showing the main long-base instruments: three strainmeters and two tiltmeters; other labels refer to measurements discussed below. Our

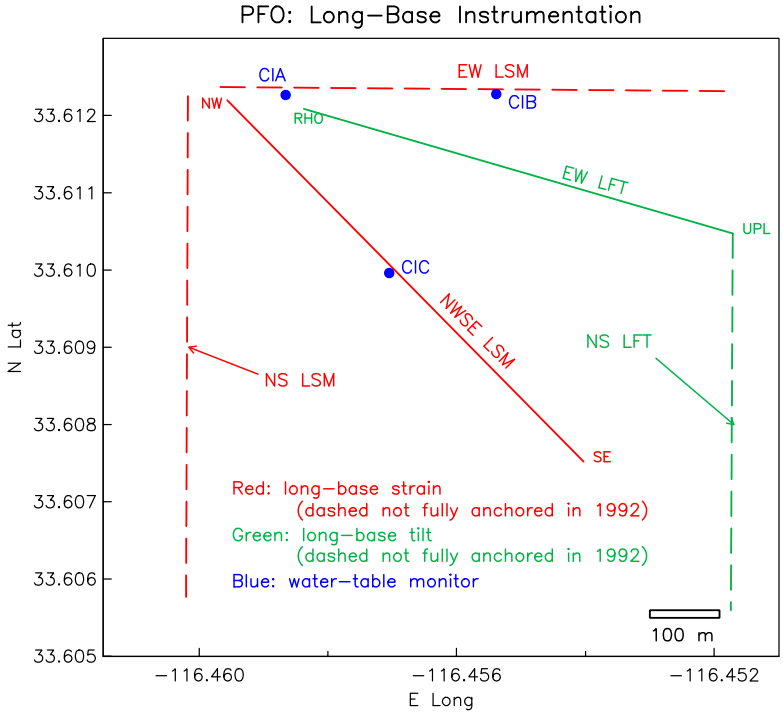

Figure 13

website, www-pfo.ucsd.edu/, includes a more detailed map and other information.

6.2. Tectonic setting of the Durmid Hill Laser Strainmeter

DHL is also in the region of strike-slip faults, and next to the most active of these faults, the San Andreas.[3] The southern section of the San Andreas Fault runs from San Gorgonio Pass to the Salton Sea near Bombay Beach, where it meets the northern extension of the Brawley Seismic Zone; DHL is located near this intersection. In many ways this fault segment is quite enigmatic. Its geomorphic expression is extremely clear so it is certainly active in some way; trenching near Indio has found evidence for four large slip events between 1000 and 1700 AD (Sieh 1986), giving an average recurrence time of some 220 years; it has been over 300 years since the last event. The seismicity is low but geodetic measurements show total motion across this segment of the fault of about 25 mm/yr (that is, 2−3 times the total rate of contraction across the Los Angeles Basin and a slip deficit currently greater than 7 m) and it is an area of active surface fault creep. The strains over the faults are predominately shear, though there are complications (as noted by Johnson *et al.* (1994)) around the Brawley Seismic Zone and to the east of it, including a region of pure dilatation.

The southern limit of geomorphic expression of the San Andreas is on the large, gentle topographic uplift known as Durmid Hill; The local geology (Babcock 1969; Ḧurgmann 1991) is interbedded claystones and siltstones, only weakly cemented together. Northeast of the fault the bedding is relatively undistorted and the topography gentle; southwest of the fault the strata are intensely folded and sheared. The fold axes have an

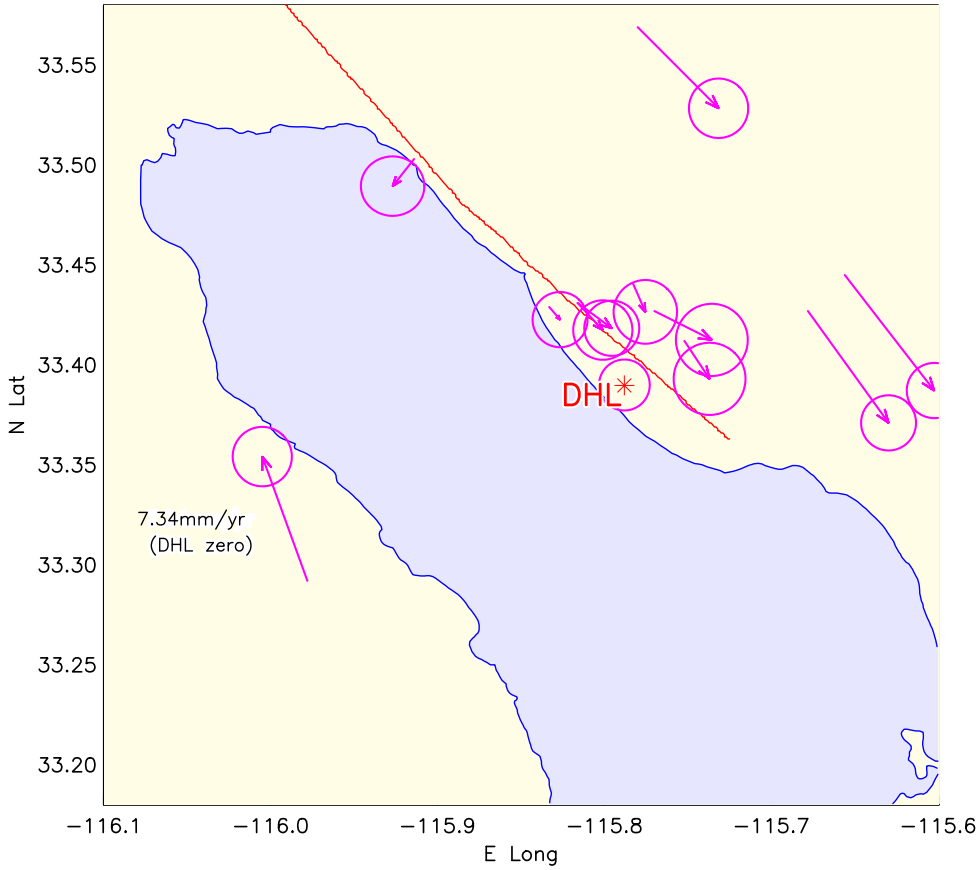

Figure 14

average trend about 20° away from the local strike of the San Andreas (N47.5°W). The surface trace of the San Andreas fault is clearly visible on Durmid Hill, suggesting that there is significant shear on the fault plane. Most recently, this part of the fault has shown triggered slip following several large earthquakes. There also is ongoing creep of 1−4 mm/yr, which Sieh and Williams (1990) show has been going on for at least the last 300 years (their nearest location to Durmid Hill, at Salt Creek, showed 2 mm/yr over the last 70 years).

Figure 15 shows some details of the region, with the strainmeter installation in an inset. The only fully-anchored system is the NS instrument, but we have also operated a ‘‘portable’’ long-base strainmeter at two different azimuths there to study possible distortions of the strain tide by the fault zone; Section 7.2 discusses the results.

7. Some Results

In this section we describe some of the results that have been obtained with the long-base strainmeters at PFO and DHL. We begin with the longest-term records, which bear on the stability of the instruments and the extent to which they are affected by hydrologic changes. We then discuss tidal results, primarily from DHL, but also from PFO; the PFO results indicate that the tides observed there are close to theoretical models. After

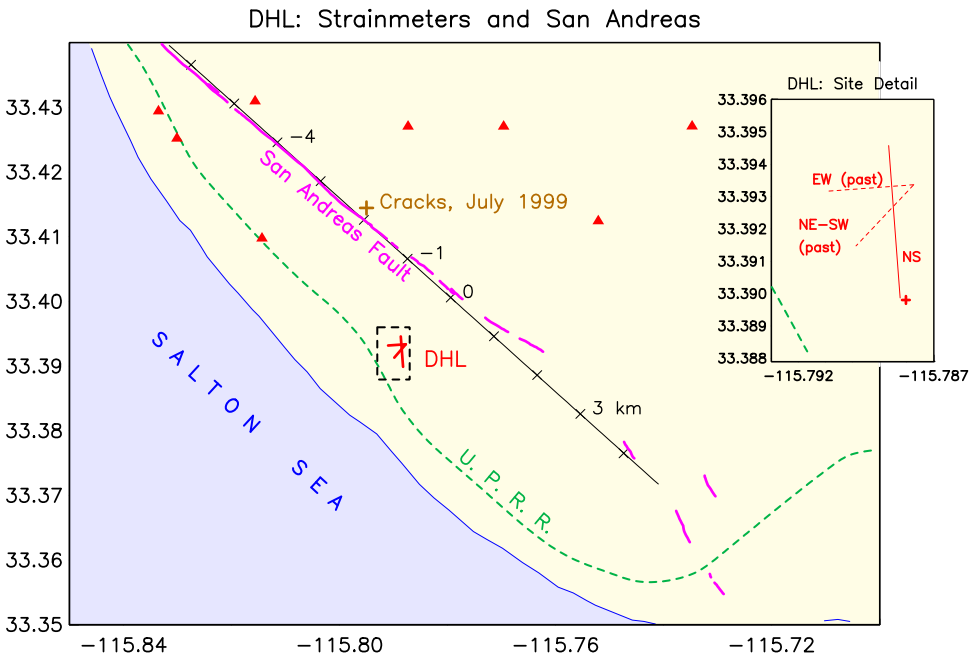

Figure 15

this we turn to transients, some associated with earthquakes, and others not.

7.1. Long-term stability

In this section we discuss the most severe test of any strain measurement: how stable it is over the long term (months to years). It is in this area, we feel, that long-base instruments have shown the best performance, providing complete coverage of the strain spectrum out to the longest-term secular strains.

7.1.1. PFO

What makes PFO unique in the measurement of crustal deformation are the five surface long-baseline sensors there: three strainmeters and two tiltmeters. We have learned much from operating these instruments: first, how to make good measurements of deformation, and second, how remarkably stable the rates of deformation generally are when they can be measured over baselines of hundreds of meters. **Figure 16** shows all the data that we have collected from the three strainmeters, from the beginning of corrected measurements up to about a year ago. This panel also shows, on the same strain scale, results from fixed GPS observations over an EW 14-km baseline with one end at PFO (G. Anderson, pers. comm., 1999). A clear secular trend is evident in the GPS data, but even the most recent of that data, using state-of-the art equipment and processing, has a scatter far above that shown by the strainmeters. The two lower panels in this figure show the much higher resolution that is possible; strainmeters can easily resolve the earth tides, (lower right) as well as seismic waves and small coseismic offsets (less than 10^{-9}).

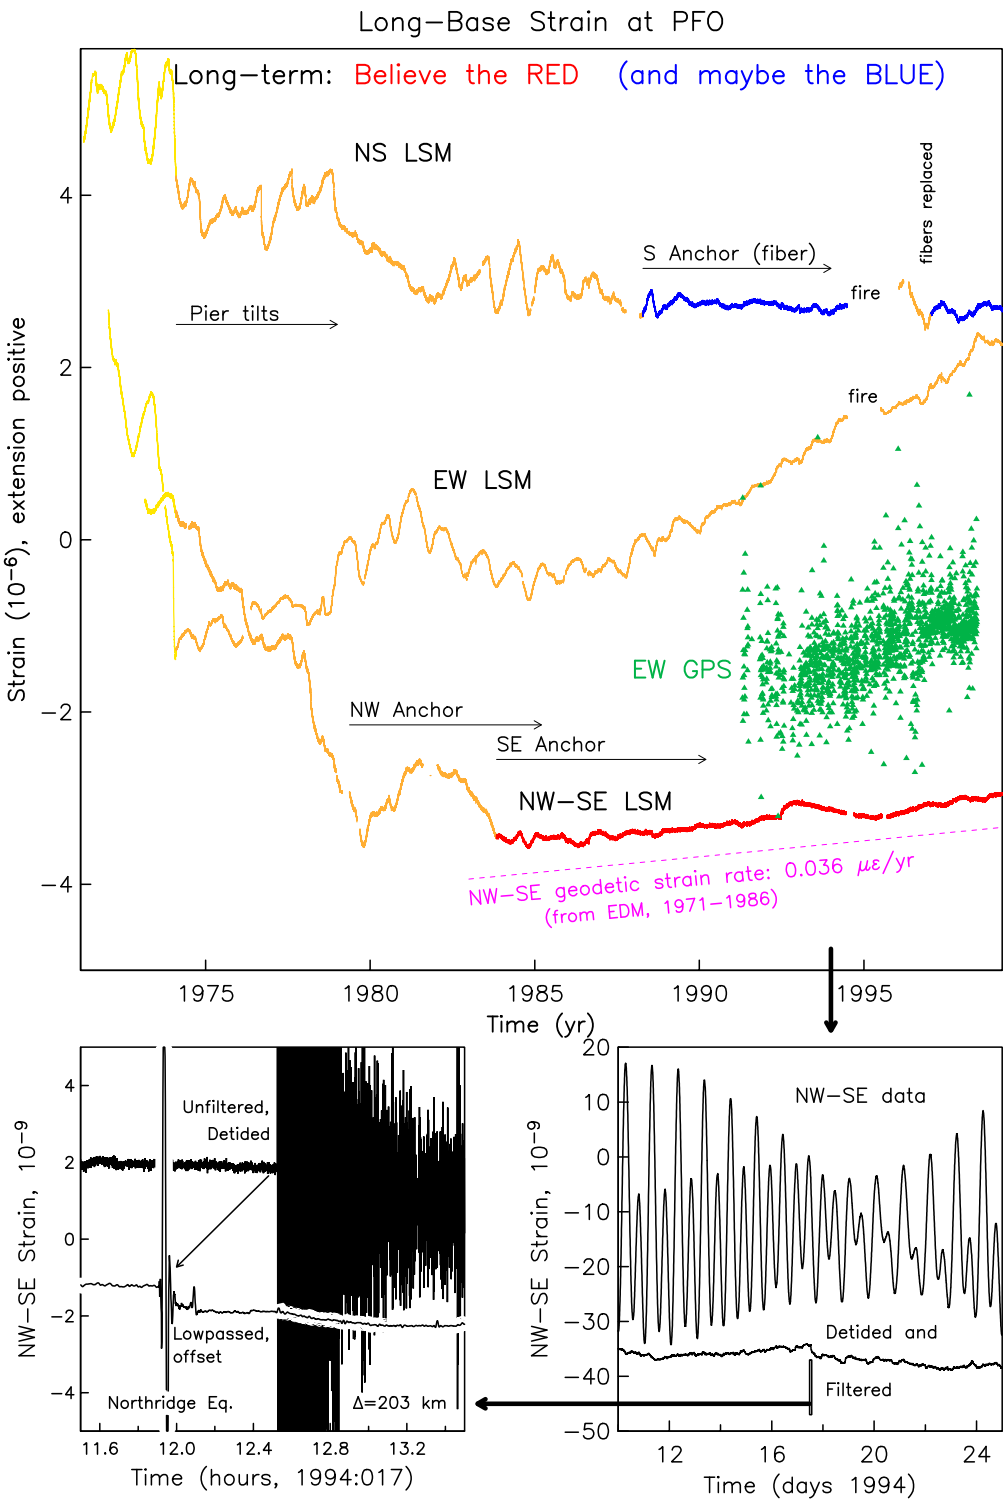

Figure 16

We have outlined the history of our developments in Section 1.2 above; the progression from yellow to orange, to blue and red, shows the improvements in strain measurement made by adding end-point anchors and stabilizing the lasers better. With these

combined improvements (in place by 1987) the stability of the NWSE laser strainmeter reached a level not seen, so far as we know, in any other continuous strain measurement. Most impressively, the pre-Landers strain rate from this instrument agrees with the 20-year average strain rate found using the Geodolite measurements of the US Geological Survey Crustal Strain group (J. C. Savage, pers. comm.).

The NWSE instrument is thus able to capture the actual secular strain accumulation over relatively short spans of time: while this is a severe test of any system to measures crustal strain continuously, it is the test we strive to satisfy, since an instrument which shows what the Earth is doing at periods of years is, as certainly as can be, providing an accurate record at all periods shorter than this.

Thanks to these improvements, we can detect interesting geophysical signals that would have been far below the original noise level. The largest such signal is just visible as a small ‘‘bump’’ on the NWSE strain record in **Figure 16** between 1992.5 and 1995. **Figure 17** shows this on an expanded scale, along with data from the EW long fluid tiltmeter (LFT) (at present the only other well-anchored instrument).

This ‘‘bump’’ is a long-term postseismic strain from the Landers earthquake. The coseismic offset from this earthquake has been removed from both records; the apparent offset at the time of this shock is actually rapid aseismic strain accumulation, which began immediately after the event. The deformation rate decreased with time: for the first six months postseismic strain accumulation was roughly proportional to the log of the elapsed time. However, in late 1992 the strain rate actually reversed sign. This reversal lasted until 1995, to about the point at which the immediate post-seismic strains had been completely recovered. In 1996, the rate returned to approximately its pre-earthquake value, though with time it has become clear that the post-Landers rate is noticeably higher. The data from the EW LFT is somewhat noisier than the strainmeter, but still an extremely stable record. In late 1990 one end-vault of this instrument flooded, destroying the optical anchor; we replaced this with an extension originally built for testing of end-monument anchoring using optical fibers. The fiber anchor is not as good as the original vacuum-path system. Keeping this in mind, the series shown in **Figure 17** suggests changes in tilt rate: from mid-1987 through mid-1992 the overall tilt is near zero (0.02 μrad/yr). After the Landers earthquake we see a very similar signature to what the strainmeter shows: an immediate postseismic response which slows, reverses, and then returns to close to the long-term rate, though because of the larger fluctuations in this series, all this is somewhat less clear than on the strainmeter.

Because the instruments have no elements in common, we can rule out internal instrumental problems as the source of these signals. Another possibility would be some effect ‘‘local’’ to Pinyon Flat, such as a poroelastic strain change of the type seen (at much larger amplitudes) close to the Landers rupture. To evaluate this, we examine hydrological records from PFO. The second panel of **Figure 17** includes water heights recorded in three boreholes at PFO (drilled in 1982-83; their locations are shown in **Figure 13**). All three show tidal responses; they also show long-term water-level changes, notably a relatively abrupt increase in 1993, in response to the high seasonal rainfall of

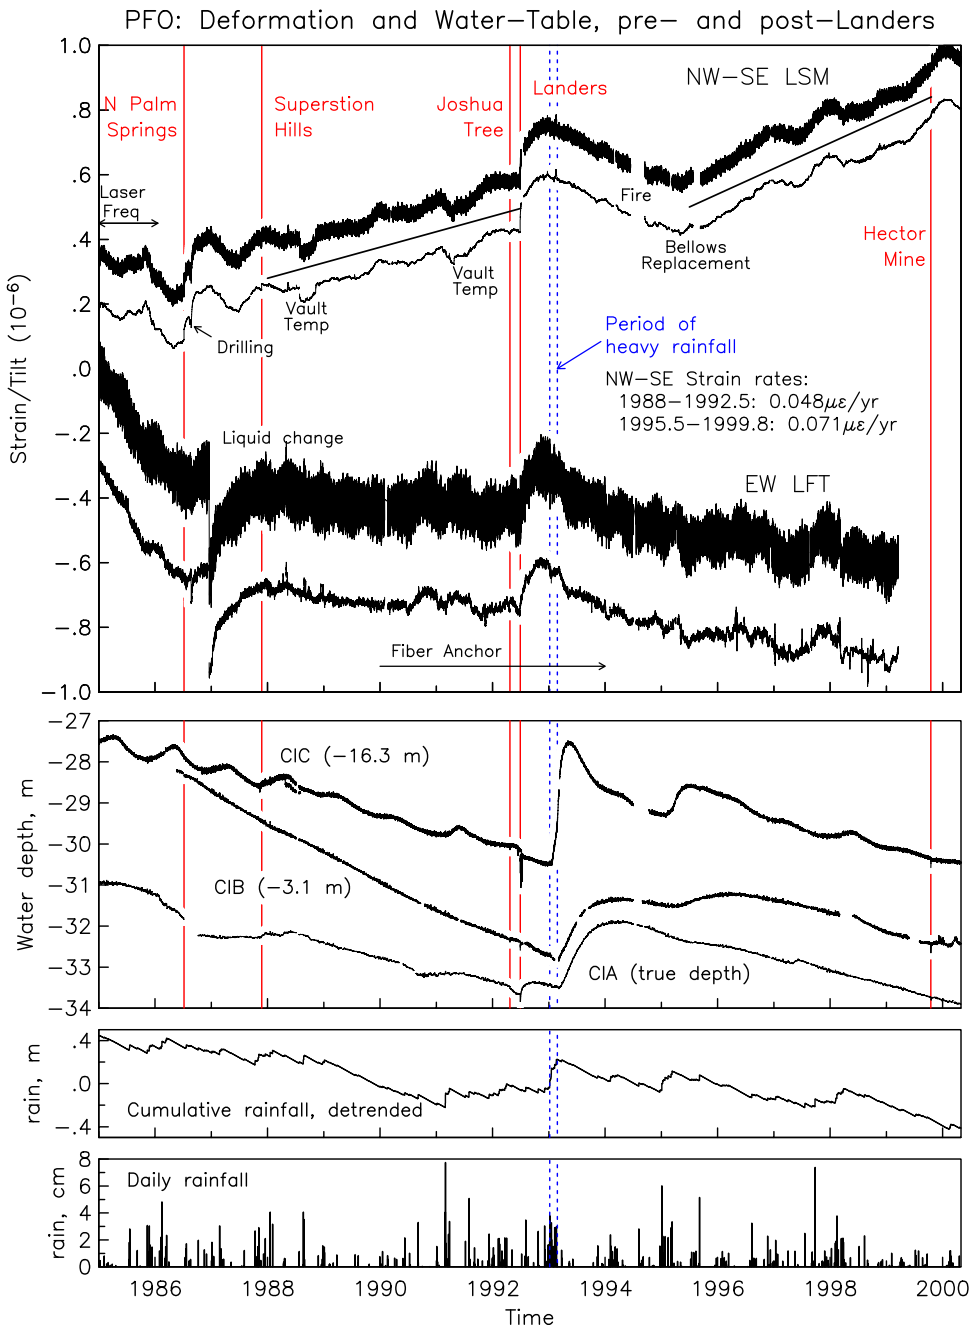

Figure 17

1992-93. This increase in the water table coincides roughly with the strain and tilt reversal, but a close look shows the strain reversal precedes the water-table change; certainly the rapid rise in water table does not create any parallel behavior in strain or tilt. There does not, therefore, seem to be any evidence for a link between pore-pressure (as reflected by the water table) and strain changes. Additionally, a local increase in pore pressure would create extension, not contraction. Evans and Wyatt (1984) discussed the effects of water-table changes in producing localized deformation, and concluded that the

effect of this on strain measurements would in many cases diminish with increasing baseline length; this appears to be borne out by the response of the long-baseline strainmeter to the large excursion in the water table.

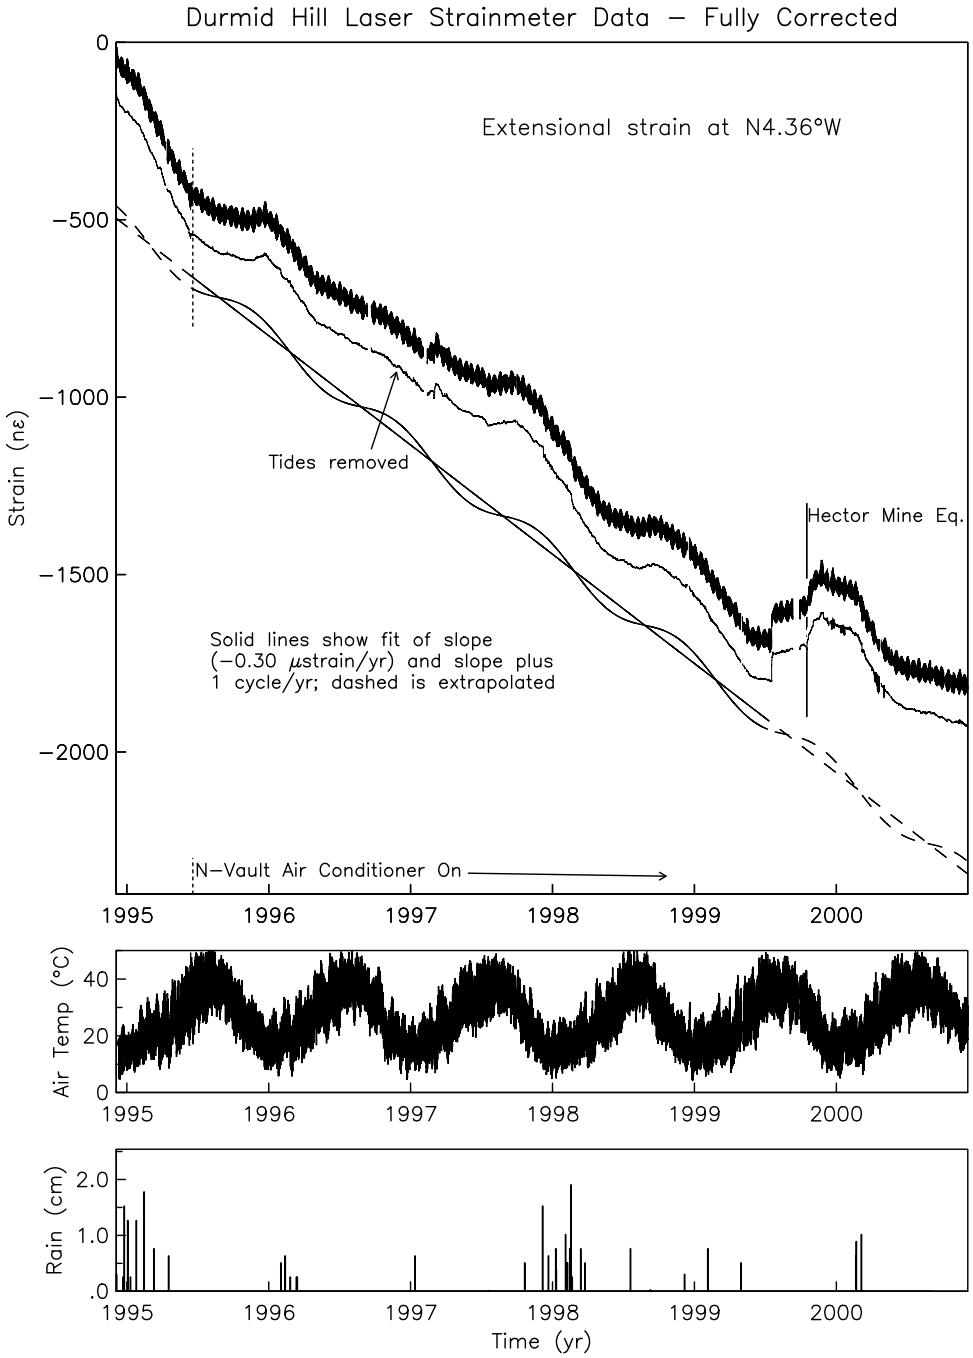

Figure 18

7.1.2. DHL

For DHL, the longest-term measurements are shown in **Figure 18**, namely the data from the NS (fully-anchored) strainmeter, from the time the anchors began running. Comparing this to the data without the correction provided by the anchors (**Figure 7**) indicates that the anchors do an excellent job of recording end-point motions: the corrected signal is free of spurious fluctuations, including some large rain-related events early in 1995. Indeed, the corrections change the sign of the secular rate from positive to negative.

The strain record is dominated by a secular trend. Fitting a slope over the span shown by the solid line gives a secular strain rate of −0.31 $\mu\varepsilon$/yr. We have two other estimates of secular rate to compare this to. The best estimate from fitting a uniform strain to the nearby geodetic stations (**Figure 14**) gives a rate of −0.41±0.09 $\mu\varepsilon$/yr; We have also computed the strain for a dislocation model of the San Andreas fault and Brawley Seismic Zone, using a deep slip rate at 25 mm/yr. Assuming a locking depth of 11 km for both gives a rate of −0.27 $\mu\varepsilon$/yr. We conclude the strainmeter is successfully recording the secular strain, even in the poorly consolidated material around Durmid Hill. To be able to measure a secular strain rate, and also to get the high resolution of a strainmeter at shorter periods, is something that, so far as we know, cannot be done using any other technique.

There is also an annual cycle visible in parts of the strain record; a fit for this gives an amplitude of 35 nε, with a phase of 37° relative to January 1. A similar fit to the air temperature gives an annual cycle of amplitude 10.7°C, and phase of −199.8°: 123° different from the phase of the strain. Whether this cycle comes from genuine thermoelastic deformation, or (quite possibly) incomplete correction of end-motion by the fiber anchors, we do not know. Compared to any other near-surface strain record (except the long-base instruments from PFO), it is small.

If we remove the secular rate and annual cycle we get the residual series of **Figure 19**. For large parts of this record, the strain fluctuations seen on this record are comparable to what we have observed at PFO, on hard rock well outside the fault zone. We conclude that, given an instrument capable of measuring the full range of strain changes, which we have, the strain fluctuations are, most of the time, not large. But in the last two years we have seen a number of exceptions to that ‘‘most of the time,’’ which we discuss in the section below on aseismic strain events.

7.2. Tidal and seismic measurements

We have looked for tidal anomalies at both PFO and DHL by comparing the measured tides with those expected theoretically: the theoretical tides can be computed with considerable accuracy thanks to the existence of good models of the global ocean tides.

The first question, which can be answered only by spectral methods, is how well we can measure the tides. We analyzed 74.5 months of data from the DHL NS strainmeter; for the portable instrument we have 15.9 months in the EW azimuth, and 23.1 months in the NE-SW azimuth. The data from the portable instrument shows a very large

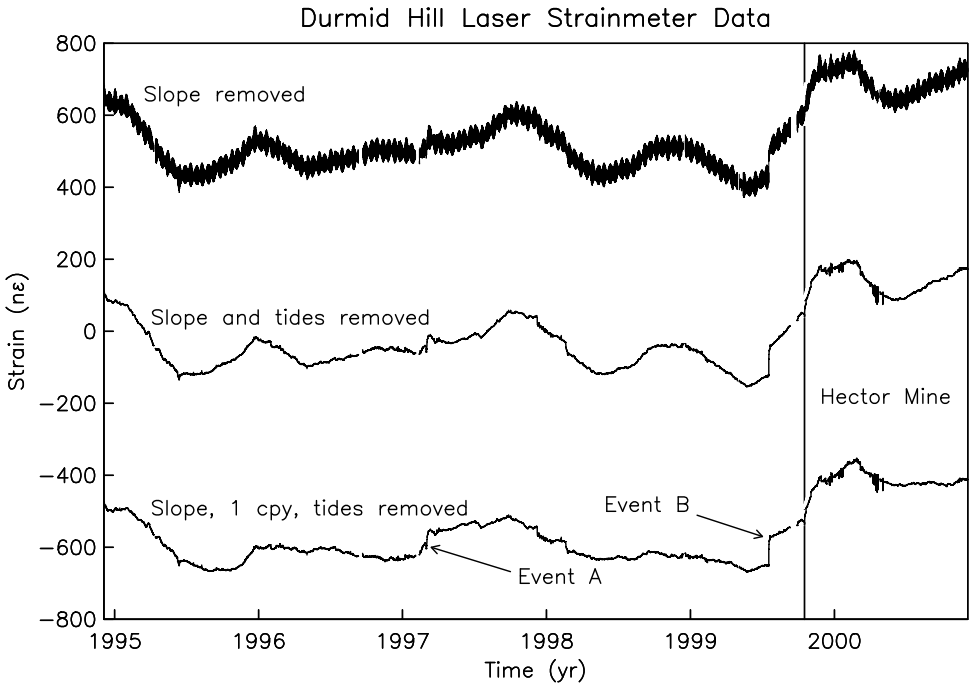

Figure 19

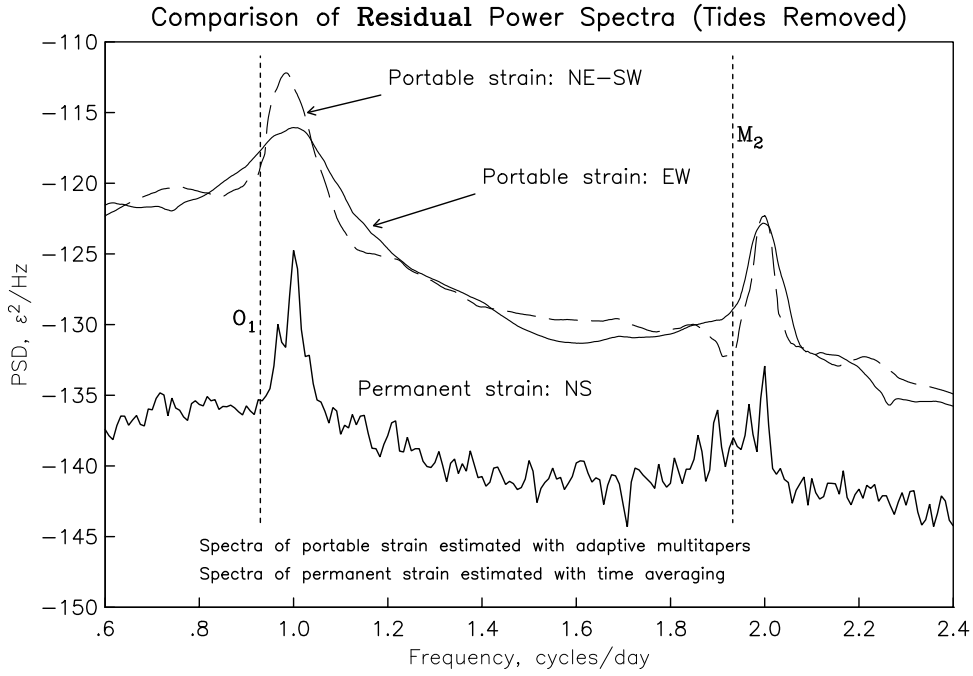

Figure 20

daily cycle, caused (we think) by temperature coefficients of the system, and thermoelastic deformations of the ground around the vault. We have been able to remove most of this effect with a slowly-modulated sinusoid with a frequency of 1 cycle/day. A spectral

analysis of the data from the three DHL instruments (**Figure 20**—which plots the residual after the tidal analysis) shows that in the tidal bands the portable instrument has about 3 times the noise level of the permanent system. But for both types of instrument the signal-to-noise level of the largest tides (M_2 and O_1) is adequate for our purposes. The permanent DHL instrument (the NS) is entirely buried (at a few meters depth for the end vaults, and about 0.5 m for the vacuum pipe) has a very low noise level: low enough, in fact, that the spectral level in the semidiurnal tidal bands is limited by the presence of small unmodelled tides (degree-3 harmonics).

Table 3 gives the results for the largest usable tides (those near one cycle/day are contaminated by thermal effects, and their modelling is complicated by the core resonance). The PFO tidal observations are from the cross-spectral analysis in Agnew (1979).

Table 3: Observed and Theoretical Tides

Place	Az.	Tide	Observed				Theory	
			Amp	err	Phase	err	Amp	Phase
DHL	−4.3°	M_2	16.4708	0.0085	−4.96	0.03	13.7155	−3.1236
DHL	−4.3°	O_1	4.4212	0.0121	11.19	0.16	4.3469	4.9305
DHL	−94.3°	M_2	5.0219	0.0489	31.70	0.56	6.6415	22.3274
DHL	−94.3°	O_1	5.0505	0.1948	−8.73	2.21	6.3839	−3.3321
DHL	−139.3°	M_2	6.7977	0.0405	1.90	0.34	6.8216	11.7112
DHL	−139.3°	O_1	4.2731	0.1282	18.07	1.72	5.7399	14.1774
PFO	0°	M_2	12.2429	0.0245	−0.90	0.11	12.3181	-0.0447
PFO	0°	O_1	3.5873	0.0359	8.58	0.57	4.3971	8.2523
PFO	90°	M_2	5.3376	0.0107	20.03	0.11	7.6044	17.0608
PFO	90°	O_1	5.0800	0.0254	−8.93	0.29	6.4301	-6.3205
PFO	135°	M_2	12.6341	0.0253	0.06	0.11	13.1770	-1.2025
PFO	135°	O_1	4.7298	0.0236	−11.40	0.29	4.9176	-10.2501

Amplitudes are 10^{-9} strain (extension positive), phases in degrees relative to the tidal potential, lags negative. Errors are estimated from the spectrum of noise. Theoretical tides include ocean loading from the CSR3.0 ocean model, (and for M_2 the model of Stock for the Gulf of California), using the continental-structure Green function of Farrell.

It is clear from this table that the theoretical and observed tides are in closer agreement at PFO than at DHL; but because of the different azimuths of measurement the results are not easy to compare. To get a better interpretation, we use the approach of Berger and Beaumont (1976) and Hart *et al.* (1996) to compare the observed tidal strain tensor $\mathbf{e}^{\mathrm{O}}$ with the ''theoretical'' tidal tensor $\mathbf{e}^{\mathrm{T}}$, the latter being the tides predicted for a spherical elastic Earth, but including the ''load tides'' induced in the Earth by the ocean tides.

These two tensors are related through the fourth-order strain-strain coupling tensor $\mathbf{C}_e$ (King *et al.*, 1976), defined by $\mathbf{e}^{\mathrm{O}} = \mathbf{C}_e\mathbf{e}^{\mathrm{T}}$. Given tidal strains along three different azimuths we can find the complete matrix of components for this tensor. A least-squares fit of the largest tidal components gives the matrix components—though we should note that the misfit significantly exceeds the errors in the observed tides. We have tried a range of ocean-load models, but the differences between these appear to be too small to

be a possible solution to this misfit. Increasing the errors to give a reasonable value of χ^2 for the misfit indicates that the errors in the components of $\mathbf{C}_e$ are about 1% at PFO, and 1−5% at DHL.

The actual values of the matrix components for $\mathbf{C}_e$ depend on how we parameterize the strains $\mathbf{e}^{\mathrm{O}}$ and $\mathbf{e}^{\mathrm{T}}$ (unless these strains are the same and $\mathbf{C}_e$ is the identity matrix). The most useful parameterization, for interpreting the results in terms of fault distortion, is to parameterized both strains as the tensor components for the 1-axis parallel to the fault, since then the components are extension perpendicular to the fault, extension parallel to the fault, and fault-parallel antiplane shear. For this parameterization, the coupling matrices are

$$\mathbf{C}_e(\mathrm{PFO}) = \begin{pmatrix} 0.896 & -0.038 & -0.168 \\ 0.040 & 0.757 & 0.084 \\ -0.039 & -0.065 & 1.070 \end{pmatrix} \qquad \mathbf{C}_e(\mathrm{DHL}) = \begin{pmatrix} 0.841 & 0.357 & -0.077 \\ 0.198 & 0.473 & -0.294 \\ 0.098 & -0.285 & 1.436 \end{pmatrix} \quad (2)$$

Noting that for ‘‘no distortion’’ $\mathbf{C}_e$ would be the identity matrix, we see that the strains within the fault zone (at DHL) are significantly more distorted than those outside it (at PFO); note that at PFO some distortion is expected from topography, while the topography at DHL is so subdued that we expect very little distortion from this cause. The distortions in the fault zone can be up to a factor of two, and often more than 25%: certainly large enough that this effect would need to be taken into account in interpreting other precise geodetic data.

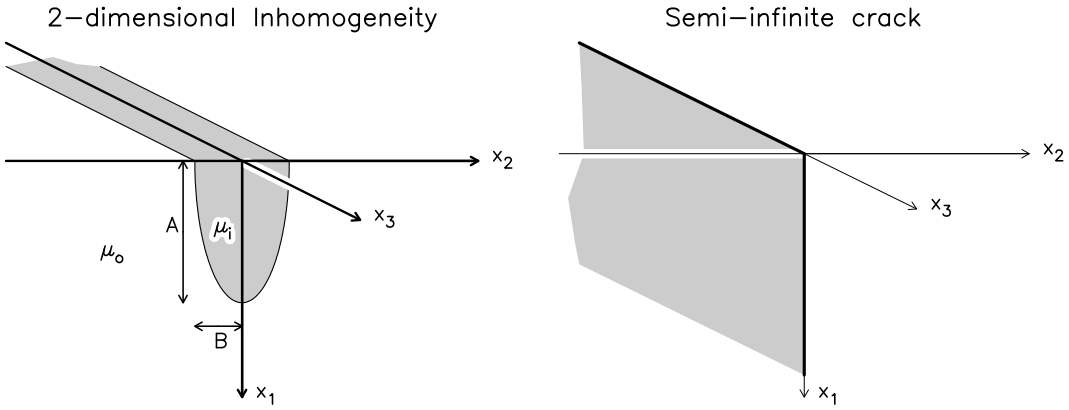

Figure 21

The result at DHL also illustrates the difficulty of modelling the strain-strain coupling with a simple model. **Figure 21** shows two simple models: one, a two-dimensional model of an inhomogeneity with a different shear modulus; and the other, a semi-infinite crack in a halfspace: the latter might be thought to be more reasonable given the location of DHL near the end of the mapped San Andreas (though of course the fault zone does continue). Both models would give enhanced fault-parallel shear, as observed (the (3,3) component of $\mathbf{C}_e$). The problem with both models is that an applied fault-parallel extensional stress would, assuming isotropic elasticity, not itself be modified nor couple into the other two components of strain: so the second row of $\mathbf{C}_e$ would be (0, 1, 0)—which is far from being the case. That the inhomogeneity is not two-dimensional also argues that

a more complex model will be needed, either of the fault zone, or of the elastic inhomogeneity of the Salton Trough (or, probably, both).

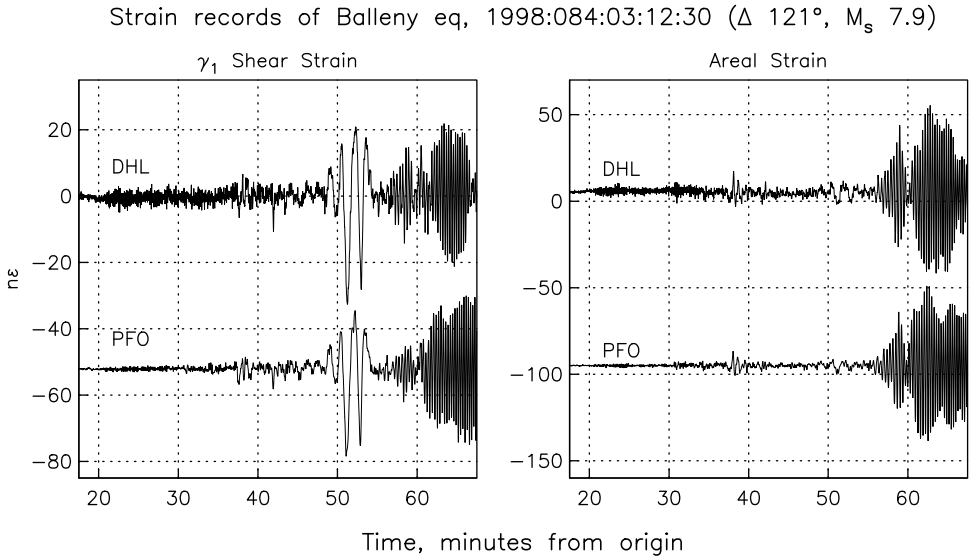

Figure 22

Elastic inhomogeneity should affect seismic waves as well as tides, though the comparison is more difficult because we do not have as good a model for the ‘‘theoretical’’ signal. Also, we need to look at waves with a long-enough wavelength that we would expect the strain field to be (nominally) uniform over the distance between these locations; we should see the kind of amplification implied by the tidal results. **Figure 22** shows some strain seismograms recorded while we were operating orthogonal components at DHL; we show these, not in the original components, but as shear strain γ_1 and areal strain e_a. The shear component shows an amplification of both the *G* phase and the subsequent surface waves, as well as what we think are basin reverberations at DHL associated with the body waves. The areal strain shows cross-coupling from shear, as it has a small amount of the *G* wave visible, which should in theory be zero.

7.3. Other strain events

In this section we review some observations of various other types of ‘‘strain events’’. Some are related to earthquakes, either the well-understood coseismic offsets or the more mysterious postseismic ones; and some are not—which makes them truly enigmatic, aseismic, events. We have not yet observed any clear preseismic strain changes.

7.3.1. Earthquake-related: coseismic

The recording of coseismic offsets using strainmeters has a long, and somewhat checkered, history: a great many measured offsets were probably, in retrospect, nonlinear effects on the instrument. Both longbase and borehole sensors have, since 1970, shown that properly installed systems give coseismic offsets in agreement with elastic dislocation theory; see Wyatt (1988) for the results from PFO through 1982. Later results for, e.g. the 1987 Whittier Narrows (Linde and Johnston 1989) and 1994 Northridge (**Figure 16**) earthquakes have produced similar results; as noted in Table 1, most of the more

recent and nearby events have been so large that the strainmeters did not give an accurate measurement.

One recent exception was the second large earthquake close to PFO in a decade: the Hector Mine event, an M_w 7.1 shock 110 km away, on 1999:289:09:47 (day 289.4076). **Figure 1** shows the location and surface rupture, to the east of the 1992 Landers event.

Table 4: Hector Mine Coseismic Offsets

What	NS Strain	EW Strain	NWSE Strain	SWNE Strain
PFO Obs	289	167	161	
PFO Theory	412	−72	113	
DHL Obs	~1200			~1200
DHL Theory	35			−43

Theory for a source at 34.59°N 116.27°W, 13 km deep, pure right-lateral slip with strike N29°W, dip 77°, moment 5×10^{19} N-m. Units are 10^{-9} strain.

While the excellent geodetic coverage of the Hector earthquake means that the coseismic strain and tilt offsets recorded at PFO cannot be expected to have much importance for determining the actual earthquake mechanism, these offsets do provide a (severe) test of the ability of the sensors to record the complete strain history during large dynamic strains.

As noted in Section 2.1.1 above, so long as the alignment of the laser beam is not interrupted and the fringe-counter rate is not exceeded, the system can maintain lock during rapid strain changes. **Table 4** shows that two of the three LSM's at PFO have offsets in reasonable agreement with the predictions of a source model based on regional seismograms. The EW LSM does not: a consequence of poor alignment of the beam at the time of the earthquake, a problem to be remedied by installing automated beam steering (Section 9). This earthquake was thus within the range of what can be recorded reliably—which in turn means that the instruments at DHL—at 140 km, more distant from the epicenter—should have given reliable measures of the coseismic offset. In fact the offsets recorded at DHL are very much larger than anything seen at PFO; we explain this as caused by triggered slip on the nearby San Andreas fault. Field examination of the San Andreas fault immediately after the Hector shock showed surface slip of 3-6 mm from Salt Creek SE to about the ''−2 km'' point in **Figure 15** (D. Yule, pers. commun., 1999); InSAR interferograms (D. Sandwell, pers. commun.) also showed evidence of widespread triggered slip, again not extending quite as far south as DHL. There is thus good independent evidence for the large anomalous coseismic strains at DHL having been caused by local fault slip. The Hector earthquake also triggered seismicity at the northern end of the Brawley seismic zone, most especially an earthquake swarm from 1 to 12 hours after the Hector event, with 14 shocks of magnitude 2 and above. Hough and Kanamori (pers. commun.) have shown that these earthquakes were probably aftershocks of a magnitude 4.6 earthquake which happened at the time of arrival of the seismic energy from the Hector shock. A dislocation model of this earthquake, assuming a strike-slip fault plane coincident with the aftershocks, and a moment corresponding to magnitude 4.6, would produce a coseismic strain change of 1 nε at DHL—not enough to

explain the large offset shown. As shown in the next section, the postseismic strains at DHL were also large for this event; again, most likely a reflection of slip on the nearby San Andreas.

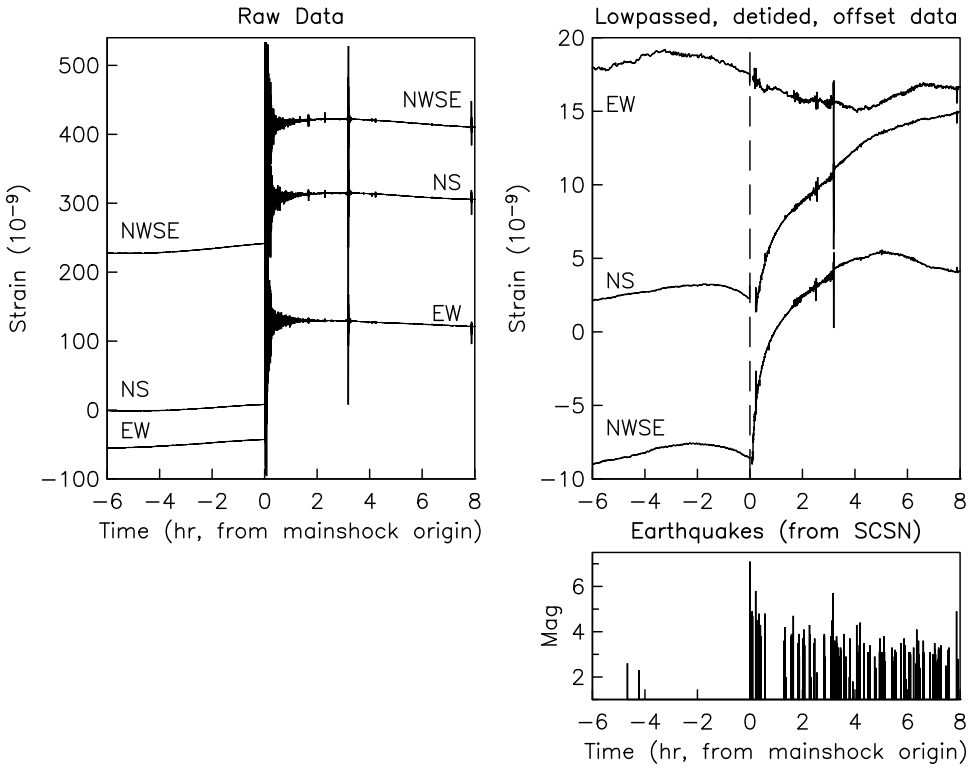

Figure 23

We think it is very important to note that strainmeter measurements of coseismic offsets are much less important now than they were even 15 years ago, since the advance of GPS and InSAR methods means that there is now a wealth of geodetic measurements in the nearfield of any earthquake, which will have much more weight in any estimate of coseismic fault slip than more distant strainmeter measurements. This will certainly be true for any future earthquakes also, with the exception of smaller events: for these, the lower noise of strainmeters (of whichever type) will allow detection of coseismic offsets which will not affect GPS.

7.3.2. Earthquake-related: Postseismic

In the previous section we have described the long-term postseismic motion from the Landers earthquake; here we discuss postseismic effects on shorter time scales, for which we have observations from a larger number of earthquakes.

One case of postseismic motion (or perhaps the lack of it) was for the Elmore Ranch and Superstition Hills events, which occurred 12 hours apart in late 1987 (Agnew and Wyatt 1989). The PFO data for this pair of events showed no obvious postseismic deformation (at the level of 10% of the coseismic offset) from the first event: a constraint

on possible mechanisms for triggering of the second, larger earthquake. For the Superstition Hills mainshock, the data from the NWSE strainmeter show about $1.5\ n\varepsilon$ strain change—about 10% of the coseismic offset— in the first 5000 seconds after the earthquake, with smaller changes over the next 6 hours, including a small step (~ 0.14 nε) just after 328:13.5, at the time of an M_L 4.8 aftershock. The slip at the surface increased substantially over the next few days (Kahle *et al.*, 1988); this can be reconciled with the PFO results by supposing that it reflects propagation of deeper slip through unconsolidated material to the surface; the low modulus of the overlying material means that slip in it does not produce significant deformation in the far field.

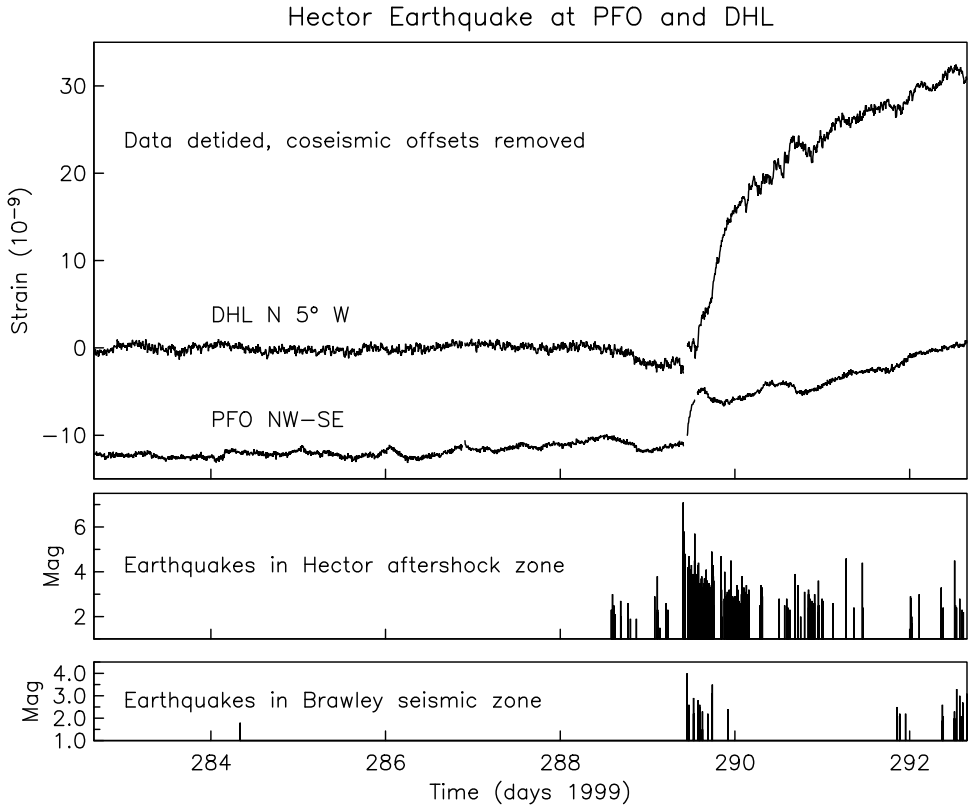

Figure 24

The next event to cause postseismic motion was the 1992 Landers earthquake; Wyatt *et al.* (1994) describe the results. Here we note only two points. The first is the extent to which the strainmeter data were important for evaluating short-term hazard. The day after the Landers earthquake began with an early-morning telephone call from Dr. Lucile Jones of the USGS to Dr. Agnew (at home), asking about the strain data from PFO. In the previous 24 hours aftershocks of the Big Bear event had extended as far south as the San Andreas fault, leading to concern that rupture on that fault might be imminent. An inspection of the telemetered records showed that a strain change of unprecedented rapidity was in fact in progress, with the amount of change in the previous day being about what we normally see in a year. Later that day we were asked to send PFO strain data to a meeting of the California Earthquake Prediction Evaluation Council,

meeting in emergency session, since the PFO strain records were the best clue as to whether or not some instability might be developing. Some hectic data editing culminated in an evaluation that the rate of postseismic strain, while still rapid, was decreasing. This reduced the immediate concern, though not the interest in how these strains were evolving. As shown in **Figure 25** below, they did indeed slow down with time: for the six months after the earthquake the accumulated strain varied roughly as the logarithm of the elapsed time. Our other note would be that the borehole strain data, both at PFO (as shown in Wyatt *et al.* 1994) and at PUBS, in the western Mojave desert (Johnston *et al.* 1994), showed much higher rates of postseismic deformation in the short term, perhaps reflecting hydrological readjustments following the large dynamic strains.

Finally, we describe the postseismic response for the Hector Mine earthquake observed on the strainmeters. Simple plots of the raw data (**Figure 23**, left) are not very enlightening, since this is dominated by the dynamic strains and the offsets, and (once these are removed) the tides. The right panel of this figure shows the strains with these signals removed: the seismic energy by lowpass filtering with a corner period of 60 seconds, and the tides somewhat imperfectly because of problems in modeling thermal effects with such a short span of data. For a useful sense of scale, note the coseismic offset (~0.5 nε) from the M_L 5.7 aftershock about 3¼ hours after the mainshock.

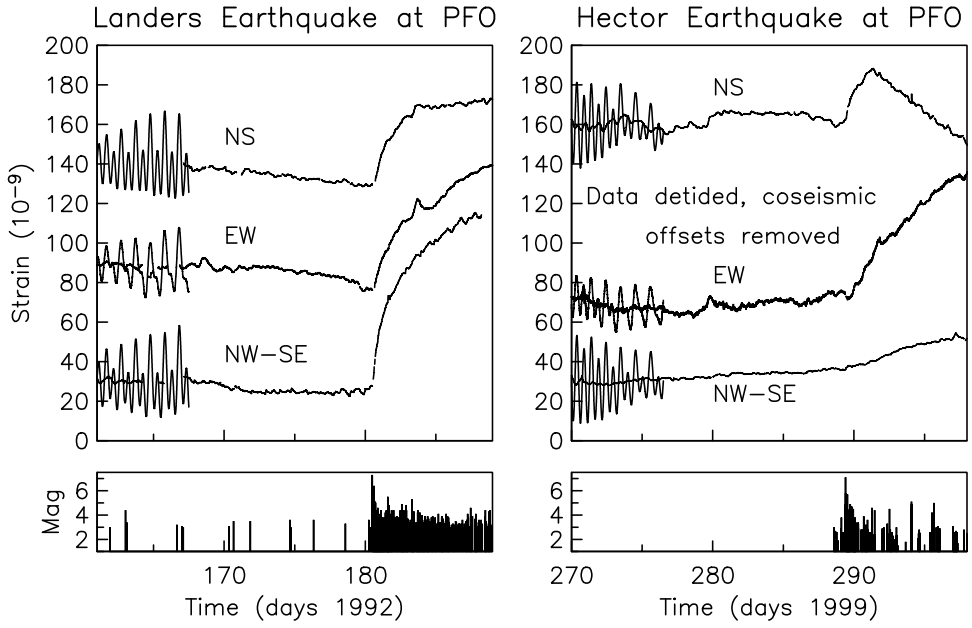

Figure 25

Clearly the rapid postseismic deformation shown, while small relative to the coseismic offset, represents as much or more moment release than the aftershocks. What is perhaps most notable about this postseismic strain is the absence (on this short time scale; **Figure 25** suggests a delayed response) of any response from the EW strainmeter—also true for the aftershock. An examination of the response at PFO to slip along different parts of the mapped rupture shows that a null response for EW strain would be seen for

slip on the southernmost part of the rupture plane. Afterslip confined to this area cannot be the whole story, however, since this would give a larger signal on the NS than the NWSE.

Figure 24 and **Figure 25** show the deformation from the Hector earthquake at PFO over somewhat longer times, in both cases with another record for comparison: in **Figure 24**, the same earthquake at DHL, and in **Figure 25** a different earthquake (Landers) at PFO. Note that in comparing these figures with **Figure 23**, the offsets in the hours after the earthquake tend to be reduced in the longer time series because of the way in which the less-frequent sampling of these data (for our long time series) combines with the editing of the series; also, the detiding is different (and more effective for the longer series).

What is clear from these figures is that the postseismic deformation pattern is quite complex, and cannot be explained by simple decaying afterslip on one part of the fault. For example, the EW instrument, which shows no immediate response, does begin to show a definite signal within a day—though with a sign opposite to what would be expected for afterslip on the rupture plane. Similar complexity is also implied by the record from the NS instrument, which reverses sign after 2 days—again, the rate after this is not consistent with afterslip on the rupture. None of this resembles the uniformly monotonic decay seen in the Landers postseismic signals—and these differences provide reassurance that this signal is not just some kind of local response of Pinyon Flat (the site, not the instruments) to strain changes or strong shaking.

Finally, **Figure 16** shows the strain over the longest time, for the one strainmeter we believe is reliable over these long times. Clearly, there is no large response; indeed, unlike the Landers shock, the Hector earthquake does not seem to have had any significant effect on the strain. Only additional data will show if this remains true over even longer times. At this point it seems clear that either rapid bulk relaxation or some kind of triggered slip on other faults will be necessary to explain these complex signals.

7.3.3. Aseismic

The least understood strain signals are the rapid (but aseismic) strain changes seen at DHL; we have never seen such events at PFO. The first aseismic events were on days 42 and 60 of 1997, when the NS instrument recorded two +27 nε, multiday exponentials (A1 and A2 in Table 5). Examination of the various auxiliary records collected at DHL showed no other disturbance; in particular, there were no meteorological events that could have caused apparent strains. Given the absence of any corresponding change at PFO, we could only deduce that the source was closer to DHL. We were naturally inclined to believe that this strain was caused by local fault slip; unfortunately, creepmeter measurements are no longer being made on this part of the San Andreas.

The next, and very clear event of this type came on days 198 and 199 of 1999, with three abrupt strain changes (B.1, B.2, and B.3) the left-hand panel of **Figure 26** shows the first one. These events were observed on both strainmeters, which share only line power and a datalogger: given this independence, and the lack of any disturbance on the many other instrument-related channels recorded, these strain changes cannot be instrumental

Table 5: Strain Events at Durmid Hill

Event	Time	Offset (nε) NS	NE-SW	Time constant	Secular change
A.1	1997:042:11	22	†	10 days	(rate change)
A.2	1997:060:10:49	17	†	1.5 days	no
B.1	1999:198:18:47	43	21	2-4 min	rate change
B.2	1999:198:21:24	10	1	~1 min	
B.3	1999:199:12:31	4	1	~1 min	(overall ~80 nε)
C.1	1999:289:09:46	~1200	~1200	abrupt	rate change
	(Hector Mine Eq.)	20	60	12 hr	rate change
D.1	1999:293:10:10	7	−58	~50 sec	rate change
E.1	1999:305:13:22	10	9	~20 sec	rate change
F.1	1999:321:18:01	4	0	~40 sec	rate change
G.1	1999:331:18:00	−8	−61	~1 min	rate change

†Portable strainmeter (now NE-SW) not operating until 1997:184.

artifacts. These events had durations of up to 10 minutes, and amplitudes of up to 40 nanostrain, with predominately extensional strain change. Unfortunately there are no creep measurements on the fault here; a field check for cracking along the fault trace, showed no clear evidence of surface fault slip. We did find a zone of tension cracking just off the fault (**Figure 15**), but believe this was more likely related to desiccation during the summer. InSAR data for this segment of the fault (D. Sandwell, pers. commun.) suggests ongoing creep, but also indicates that this creep stops somewhat north of the DHL site. Buried slip of the amount we infer would not produce a measurable InSAR signal.

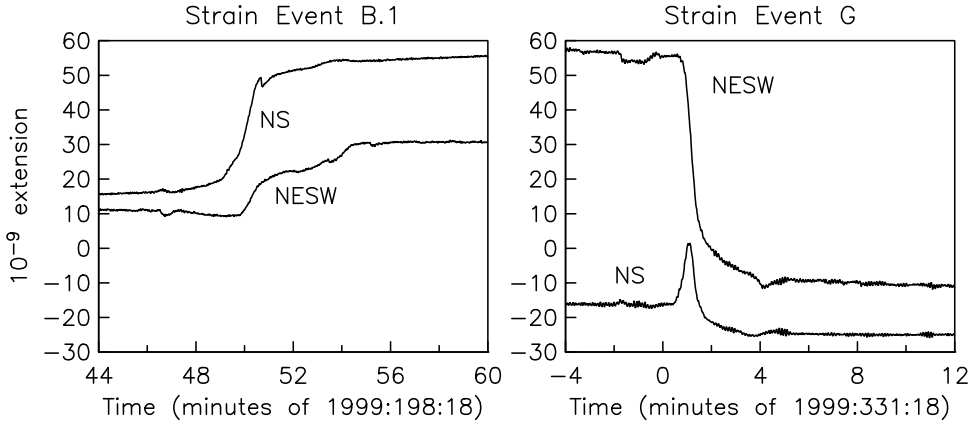

Figure 26

The ratio of strain on the two strainmeters has not been the same for the different events; this ratio powerfully constrains possible locations for slip—provided that we make the reasonable (though restrictive) assumption that what we are seeing is strain caused by aseismic slip on the adjacent fault, and that this slip is, like the geological slip, horizontal. The slip cannot all be located far from the instruments, since a source far

away would give the same waveform on both sensors—not what is observed. We computed the response of the strainmeters to slip at different points along the fault trace (simplified as a single plane, shown in **Figure 15**) assuming a small dislocation patch in a halfspace, and allowing for the fact that the strainmeters actually measure differential displacement. Given these responses, we can then ask where along the fault plane a source of slip would give the observed ratio of strains, and also positive (extensional) strain along both azimuths. We found that the different events must occur in limited, and fairly shallow, regions of the fault plane (at around the 1-km point in **Figure 15**.) Assuming this location the observed strain gives the moment release: event B.1 had a moment of 3×10^{13} to 3×10^{14} N-m; the equivalent moment magnitude is **M** 2.9−3.6. Given the somewhat different time histories of strain on the two components, the event must have involved slip propagating along the fault. This also would imply that different events would have different time signatures as different segments of the fault slipped; as **Figure 26** shows, such variation is in fact the case. As shown in **Figure 18**, all of these aseismic events bracket the time of the Hector Mine earthquake (1999:289:09:46, M_W 7.1, 140 km NNE of DHL); though we would not want to claim any association, we have seen none since the end of 1999.

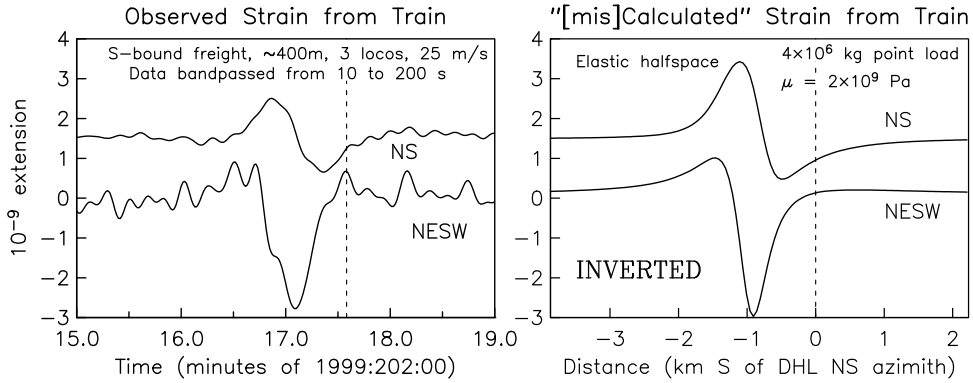

Figure 27

7.3.4. Local loads: "The Change in Strain comes Partly from the Train"

Given that our aim is to understand the local elastic structure using the response of the strainmeters to known signals, we would also be remiss if we failed to mention one case in which such understanding has proved elusive, namely the response of the strainmeters to large vertical loads applied nearby. The loads in question are those created by trains going past the site on the Union Pacific Railroad, whose track is shown in **Figure 15**. This is a major freight route between Los Angeles and the southern US, so there is extensive traffic. Each passing train produces a small but easily detectable signal on the strainmeters; the lefthand panel of **Figure 27** shows an example.

The simplest model for this is Boussinesq's problem, with a vertical point load applied to a halfspace. The righthand panel of **Figure 27** shows the response of the strainmeters to a point load as function of distance along the tracks; this calculation is

done for the differential displacement along each azimuth. The x-axis (distance) has been scaled to match the time axis from the measured speed of the train, and the signal scaled to match in amplitude (the resulting train weight is quite reasonable). The problem is that the modelled signal has the correct form, but the wrong sign. The Boussinesq solution gives extensional strain around a point load—but we observe compression. Of course the trains are not a point load, but the expected signal is just the convolution of the point-load response with the mass distribution along the train; since this is everywhere positive such a convolution will not be able to change the sign. And, any polarity error in the instrument or recorder is ruled out by the agreement with the theoretical tides. The implication of these data is that a vertical load causes displacement away from the load, rather than towards it as in the elastic halfspace. We believe that this can be explained by elastic layering in which the shear modulus increases rapidly with depth, which is reasonable for this area.

8. Future Improvements

The discussion above has mentioned a number of improvements we have made in the newer installations, and which we are now retrofitting to the older ones. The most notable is certainly the anchoring, and the (unanticipated) improvement in the vacuum system which it makes possible. This was a major change; others are more incremental; for example, the two most recent instruments have improved thermal control, partly from better insulation of the end enclosures, partly from adding of a second (inner) stage of control, to make the temperature variations around the end optics even less.

Perhaps the most important future improvement will be to introduce more automatic, and remote, controls, to minimize the number of visits required to the instrument. We have already implemented an automated system for steering the laser beam on the main interferometer, though as with all such fedback systems there is a limit to how much the system can be perturbed before the controllers fail to function properly. To deal with such situations we need to introduce remote control. The design for our newest datalogger includes a system for adjusting various parts of the system remotely; such remote operation is now very much the norm for many small astronomical observatories, and will definitely make the instruments easier to operate and increase the data return. Of course, visits would still need to be allowed for those times when extreme environmental effects cause massive failures: floods or lightning strikes do happen.

There are other, more radical changes which might improve the system, but which would require an extensive development effort. Most of the cost of the laser strainmeter comes in the physical plant (pipes, anchors, and buildings) not in the actual measuring system—so any improvements in the latter could be retrofitted at a relatively small cost. We have considered two, either of which would greatly simplify the editing described in Section 5. The first would be to double the number of interferometric measurements made between the endpoints by adding a second interferometer to the remote end, making the system symmetrical and providing two votes on how the distance between the end points has changed. As with the redundancy provided by multiple GPS satellites, this would make it much easier to deal with possible cycle slips.

A more difficult task would be to make an absolute measurement of the distance between the ends. Even the best conventional EDM's, such as the Mekometer, have errors of 0.1 mm (2×10^{-7} over 500 m), too large to be really useful. It would be possible to install a purpose-built system of somewhat higher accuracy, but a potentially much more powerful approach would be to use a variable-frequency laser (such as a tunable diode) as for a second (but coaxial) interferometer: counting fringes while scanning between two known frequencies can give an absolute measure of length, potentially to the 10^{-8} level. Occasional measurements of this type, combined with the routine high accuracy of the regular measurement, would create a system of unparalleled stability and precision for measuring crustal deformation.

References

D. C. Agnew, “Strain tides at Piñon Flat: Analysis and interpretation,” Ph. D. Thesis, University of California, San Diego, La Jolla (1979).

D. C. Agnew and F. K. Wyatt, “The 1987 Superstition Hills earthquake sequence: strains and tilts at Piñon Flat Observatory,” *Bull. Seismol. Soc. Am.,* 79, pp. 480-492 (1989).

E. A. Babcock, “Structural geology and geophysics of the Durmid area, Imperial Valley, California,” Ph.D Thesis, University of California Riverside (1969).

K. M. Baird and G. R. Hanes, “Stabilization of wavelengths from gas lasers,” *Rep. Prog. Phys.,* 37, p. 927−950 (1974).

J. Berger and C. Beaumont, “An analysis of tidal strains from the United States of America, II, The inhomogeneous tide,” *Bull. Seismol. Soc. Am.,* 66, p. 1821−1846 (1976).

J. Berger and R. Lovberg, “Earth strain measurements with a laser interferometer,” *Science,* 170, p. 296−303 (1970).

J. Berger and R. H. Lovberg, “A laser earth strain meter,” *Rev. Sci. Instr.,* 40, pp. 1569-1575 (1969).

M. Born and E. Wolf, *Principles of Optics,* 6th ed., p. 808, Pergamon, New York (1980).

N. Brown, “Frequency stabilized lasers: Optical feedback effects,” *Appl. Opt.,* 20, p. 3711−3714 (1981).

R. Bürgmann, “Transpression along the southern San Andreas Fault, Durmid Hill,” *Tectonics,* 10, pp. 1152-1163 (1991).

J.-M. Chartier, “Results of international comparisons using methane-stabilized He-Ne lasers at 3.39 μm and iodine-stabilized He-Ne lasers at 633 nm,” *IEEE Trans. Instrum. Meas.,* 32, p. 81−83 (1983).

K. Evans and F. Wyatt, “Water table effects on the measurement of earth strain,” *Tectonophysics,* 108, p. 323−337 (1984).

J. Fletcher, T. Fumal, H.-P. Liu, and J. C. Carroll, “Near-surface velocities and attenuation at two boreholes near Anza, California from logging data,” *Bull. Seismol. Soc. Am.,* 80, pp. 807-831 (1990).

J. Gomberg and D. C. Agnew, “The accuracy of seismic estimates of dynamic strains: an evaluation using strainmeter and seismometer data from Piñon Flat Observatory, California,” *Bull. Seismol. Soc. Am.,* 86, pp. 212-220 (1996).

R. H. G. Hart, M. T. Gladwin, R. L. Gwyther, D. C. Agnew, and F. K. Wyatt, “Tidal calibration of borehole strainmeters: removing the effects of local inhomogeneity,” *J. Geophys. Res.,* 101, pp. 25553-25571 (1996).

M. J. S. Johnston, A. T. Linde, and D. C. Agnew, “Continuous borehole strain in the San Andreas fault zone before during, and after the 28 June 1992, M_w 7.3 Landers, California earthquake,” *Bull. Seismol. Soc. Amer,* 84, pp. 799-805 (1994).

L. M. Jones, K. E. Sieh, D. Agnew, C. Allen, R. Bilham,, M. Ghilarducci, B. Hager, E. Hauksson, K. Hudnut, D. Jackson, and A. Sylvester, “Short-Term Earthquake Hazard Assessment for the Southern San Andreas Fault, Southern California,,” USGS Open-File Report 91-32. (1991).

J. E. Kahle, C. J. Wills, E. W. Hart, J. A. Treiman, R. B. Greenwood, and R. S. Kaumeyer, “Surface rupture of Superstition Hills earthquakes of November 23 and 24, 1987,” *Calif. Geol.,* 41, pp. 75-84 (1988).

H. P. Layer, “A portable iodine-stabilized helium-neon laser,” *IEEE Trans. Instrum. Meas.,* IM-29, p. 358−361 (1980).

J. C. Owens, “Optical refractive index of air: Dependence on pressure, temperature, and composition,” *Appl. Opt.,* 6, p. 51−58 (1967).

G. Peltzer, P. Rosen, F. Rogez, and K. Hudnut, “Poroelastic rebound along the Landers 1992 earthquake surface rupture,” *J. Geophys. Res.,* 103, pp. 30131-30145 (1998).

M. D. Petersen and S. G. Wesnousky, “Fault slip rates and earthquake histories for active faults in southern California,” *Bull. Seismol. Soc. Amer.,* 84, pp. 1608-1649 (1994).

S. J. Radzevicius and G. L. Pavlis, “High-frequency reflections in granite? Delineation of the weathering front in granodiorite at Piñon Flat, California,” *Geophysics,* 64, pp. 1828-1835 (1999).

I. S. Sacks, J. A. Snoke, R. Evans, G. King, and J. Beavan, “Single-site phase velocity measurements,” *Geophys. J. R. Astron. Soc.*, 46, p. 253−258 (1976).

K. E. Sieh and P. L. Williams, “Behavior of the southernmost San Andreas fault during the past 300 years,” *J. Geophys. Res.,* 95, p. 6629−6645 (1990).

Group on California Earthquake Probabilities Working, “Seismic hazards in Southern California: probable earthquakes, 1994 to 2024,” *Bull. Seismol. Soc. Am.,* 85, pp. 379-439 (1995).

F. Wyatt, “Displacement of surface monuments: horizontal motion,” *J. Geophys. Res.,* 87, p. 979−989 (1982).

F. Wyatt, “Measurements of coseismic deformation in southern California: 1972-1982,” *J. Geophys. Res.*, 93, pp. 7923-7942 (1988).

F. Wyatt, “Displacements of surface monuments: vertical motions,” *J. Geophys. Res.,* 94, pp. 1655-1664 (1989).

F. Wyatt, K. Beckstrom, and J. Berger, “The optical anchor − a geophysical strainmeter,” *Bull. Seismol. Soc. Am.,* 72, p. 1707−1715 (1982).

F. K. Wyatt, D. C. Agnew, and M. Gladwin, “Continuous measurements of crustal deformation for the 1992 Landers earthquake sequence,” *Bull. Seismol. Soc. Am.,* 84, pp. 768-779 (1994).

M. A. Zumberge, F. Wyatt, D.-X. Yu, and H. Hanada, “Optical fibers for measurement of earth strain,” *Appl. Opt.,* 27, pp. 4131-4138 (1988).

M. A. Zumberge and F. K. Wyatt, “Optical fiber interferometers for referencing surface benchmarks to depth,” *Pure Appl. Geophys.,* 152, pp. 221-246 (1998).