First Results on Milagro Water Attenuation Using Measurements from the Upgraded TUBE

D. Coyne & M. Schneider, August 2002

History: There has been a continuing need for precise measurements of the transmission characteristics of Milagro pond water. Prior to 2000, we used measurements by MS from the UV spectrophotometer at SLAC (the "PIPE"), anticipating that the water would have an attenuation length L of about 5 m. The PIPE had a sample cell about 0.2 m in length, so even this instrument was marginal. Then we realized that we could obtain much higher quality water, with L of order 10 to 30 m, and that imprecision in this number could affect the results of our Monte Carlo calculations of detector efficiency. At this point a new instrument was proposed and developed at SLAC by MS to obtain a one-wavelength measurement of L with better precision. Changes to an automated mode suggested by DC promised to make the device super-precise and these were incorporated at UCSC by MS. We chose to develop this instrument rather than buy commercial equivalents with short cells on the basis of robustness of the technique and our better control over systematics of the device. As matters turned out, this was crucial, though the realization of these advantages was far from trivial. The device was set up at Fenton Hill in summer 2000 by DC and was used to measure water from both pond and circulator for many months. During this time, though most of the results seemed to make sense, we gradually became suspicious that the instrument was plagued by systematic errors which could contribute to the variations in water quality we were seeing. By mid-2001 this became a certainty as the systematics were tracked down and a major rebuilding of the device was done. This report, in August 2002, gives the causes of the known problems and describes how they were eliminated. It also gives the first measurements of water quality that we now consider reliable, thus setting the baseline for the future. Future tests will be done to see if this optimism is justified.

The TUBE Device: The original idea was to use a sample cell of length $X_o \sim 1$m in order to get a precise measurement of T_o, the transmission of the light through an appreciable distance X_o of water. Suitably corrected for effects of reflections in the endcap windows and for variations of the light source, this translates to a measurement of L by the relation $L = - X_o / \ln T_o$. (Even more information is lodged in the breakdown between the part of L due to scattering and that due to absorption in the medium, but the TUBE was designed to first measure only the combined attenuation). We've reported previously on the basic technique illustrated in Fig. 1, below. Passage of light is measured with the cell both **empty** and **full** of water, as well as with the cell **out** of the light path entirely. The ratio of **empty/out** shows the condition of the windows and lineup, while **full/empty** contains the information needed to extract T_o. The laser wavelength was chosen to be 325 nm, because that wavelength is very near the peak of the distribution formed from the Cherenkov spectrum of light in the pond convoluted with the photomultiplier sensitivity. The laser light is split before entering the cell and part sent to a reference detector, which automatically normalizes each of the three quantities and cancels out temporal fluctuations of the source. The ratio of the photodiode currents is read out once per second and a set of ten such measurements logged into a computer file. An operator manually prepares each of the three configurations, with realignments of the beam between each step. The data are then semi-automatically analyzed and a result displayed.

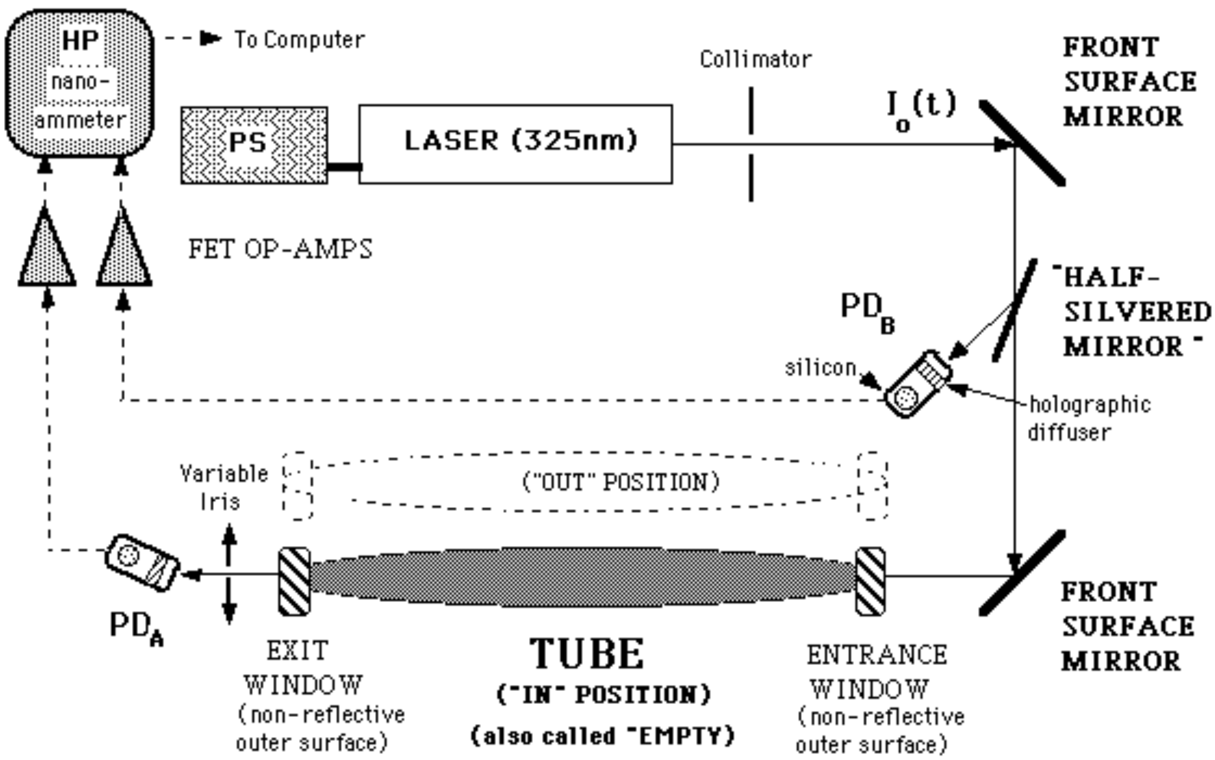

Figure 1: The TUBE Schematic.

Requirements: We set out to measure water attenuation lengths to about 10% precision for 30m water. This was just a guess. Recently, we have analyzed what we need more precisely. We use the measured attenuation length L not only to monitor impurities in the pond, but to evaluate the parameter needed for the Monte Carlo simulations of the rates and sensitivities. What we really want is to calculate transmission as a function of path length in the water, T(x), and make the relative error in this quantity, $P=\Delta T/T$, negligible with respect to other inaccuracies of the Monte Carlo. Given some desired P(x), this will constrain the ΔT_o required in the measurement (and alternatively the requirement on $\Delta L/L$). Appendix A gives the pertinent error formulas and a table of these relationships for various path lengths x and water qualities L, for typical values desired for P. The result is that even if we agree to incur ~5% inaccuracies in T(x) in the Monte Carlo, then for light reaching the muon layer we need a precision measurement with $\Delta T_o = \pm 0.0045$ from our TUBE device (water with L about 10m). This is equivalent to $\Delta L/L = \pm 5\%$, so our original choice was close, but probably not tight enough. Fortunately, things are not so bad for the important air shower layer, where the corresponding numbers are like $\Delta T_o = \pm 0.02$ or $\Delta L/L = \pm 25\%$. Still, it is clear that the instrument must suppress systematic and statistical errors at a level of order 1% or better. Our initial calculations for the "super-precise" TUBE suggested that statistical errors in T_o of order 0.002 could be achieved, two-to-ten times better than needed. The game was going to be all in suppressing systematics.

Systematics: The main problem with devices used to measure transmission is getting a good-geometry setup in which only a well-collimated beam passes through the sample and enters the detector, wherein it is wholly and consistently measured. To this end the TUBE uses a monochromatic 1mm diameter laser beam threaded through a pyrex tube with quartz end windows of diameter 50.8 mm, and incident on a photodiode of diameter 11.3 mm. The position of the beam on the photodiode can be repeatedly reset to within about 0.2 mm using a scan across an 1mm iris which is later opened to about 8 mm for the actual measurements. Every reflection from every surface is tracked and localized in space to make sure that only known quantities enter the photodiodes used. In order to eliminate the huge number of zeroth, first and second-order reflections which could contribute to the quantitative corrections, we devised geometries (tilted photodiodes and windows) wherein the strongest reflections would be intercepted and absorbed outside the tube.

We quickly confirmed that the statistical noise in the measurements of T_o was indeed 0.2%, with hourly drifts of 0.5% to 1%. But we also found that whenever the apparatus was adjusted, as it must be to return the beam to the same spot on the photodiode after windows and/or water have been inserted into the beam, huge steps would occur in the identical measurements. These steps of 5 to 10% in the baselines would often render the results nonsensical and useless. Since reflections and positioning of the beam seemed under good control, something else was changing.

After several frustrating journeys from UCSC to Fenton Hill, a chance observation revealed the problem. Each endcap window was acting individually as a Fabry-Perot interferometer between its two faces, and each photodiode was internally reflecting as well among its three surfaces (quartz window and shiny silicon detector), producing another complex interference pattern at the silicon pickup. These patterns had peak-to-peak variations of 15% in the integrated signal recorded, as a function of a change in beam angle of only about 0.4 milliradian. This meant that our needed stability was ruined within about 4 microradian, or in about 4 microns of beam movement on the photodiode! There was no way to hold such tolerances in the repeated angular adjustments required in a water measurement. Fig. 2 shows the accidental (almost!) discovery of the interference when data were taken while hand-pushing the tube stage across the beam; slight torques on the glass tube changed the beam direction systematically by microradians and produced the pattern.

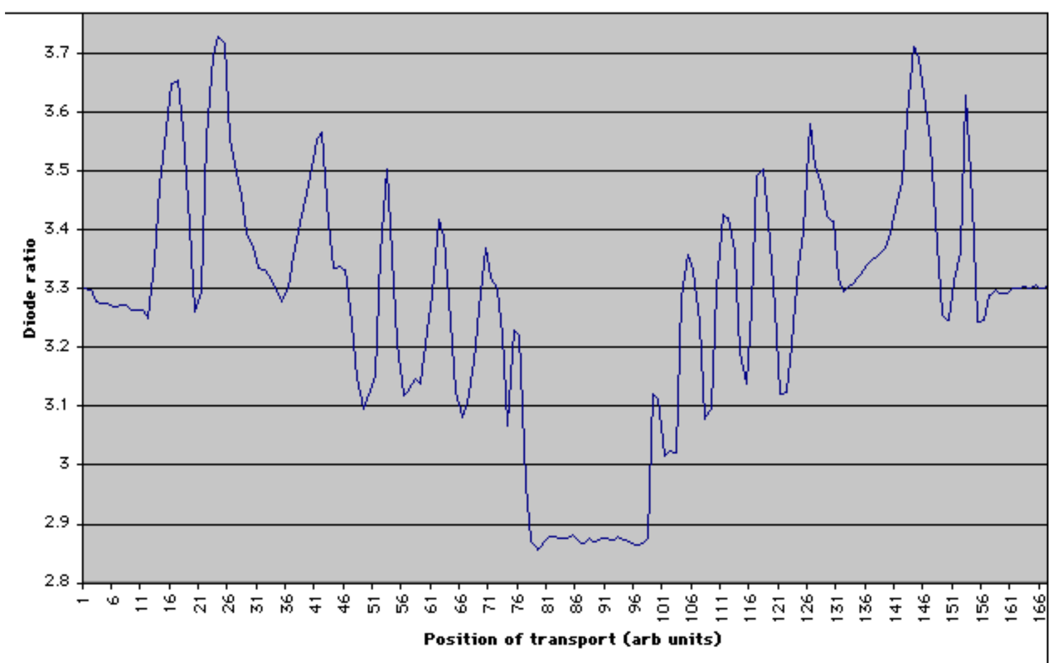

Figure 2: Scan of beam across the entrance window (and back). The butterfly pattern results from the approximate repetition of the pattern as the scan direction is reversed.

Once seen, it was easy to calculate that this effect had to be present, and in this magnitude. Fig. 3a shows a simplified calculated effect from one photodiode alone. The higher-frequency components of Fig. 2 show that in actuality other interference patterns are superposed.

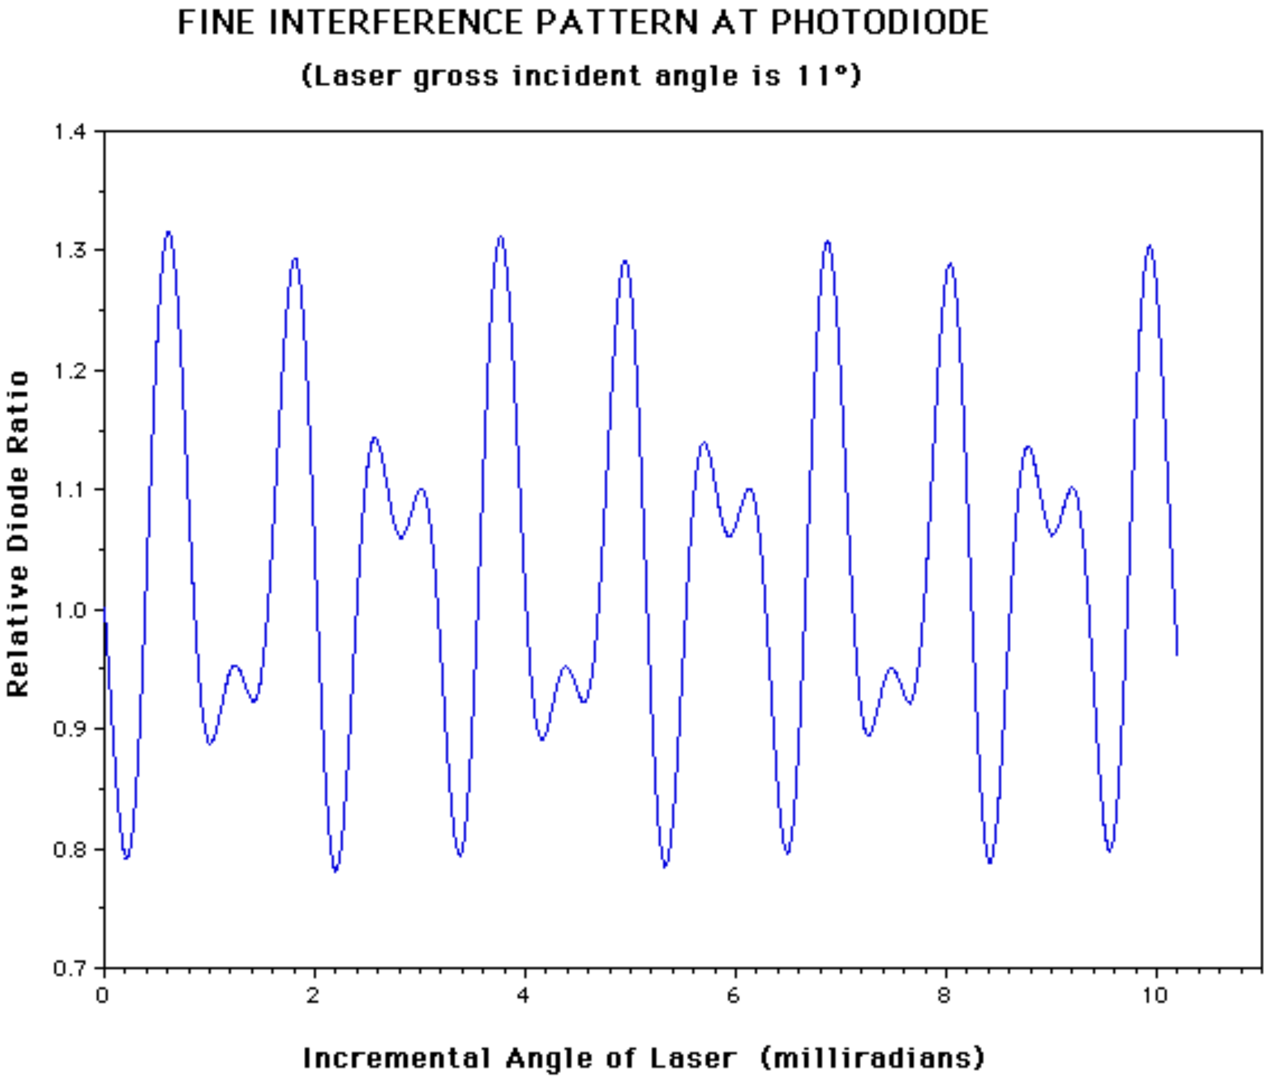

<u>Figure 3a: Calculation of interference of primary beam and the two leading reflections within a single photodiode, as a function of small variations around an incident angle of 11°. The diode was tilted to avoid strong reflections from the exit window.</u>

Note that the reason the interference pattern was so easy to see (and the reading difficult to repeat) is that one gets a much more rapidly varying oscillation with the photodiode tilted than with the "natural" position of the photodiode looking directly down the beamline (0°). The tilt had been devised to get rid of many secondary reflections which were very hard to control and quantify. Previous incarnations of the TUBE had used a 0° angle and had better repeatability. Fig. 3b (overleaf) shows such a situation was really a "fool's paradise" because although the variation of pattern is gradual, one has no absolute knowledge of where one <u>is</u> on it, and thus cannot achieve absolute answers for transmission. Thus, like so often in experimental work, the removal of one problem (the reflections) revealed another even more serious underlying problem (the interferences).

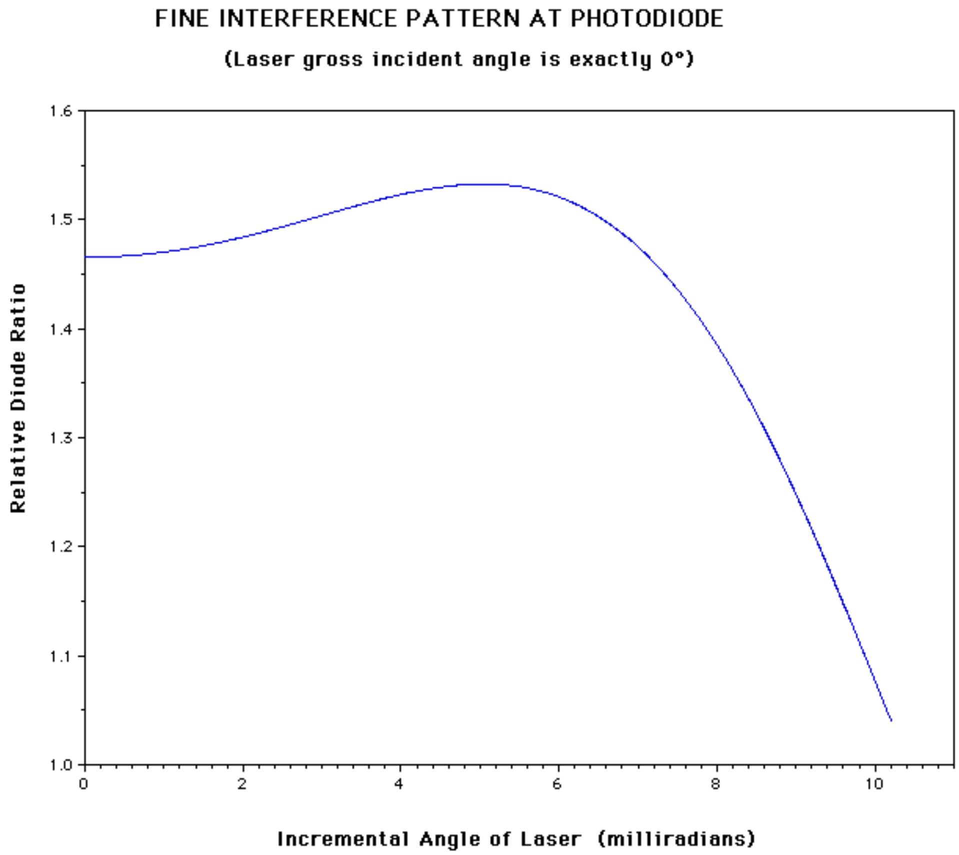

Figure 3b: Same calculation as in 3a, except for an incident angle of 0°.

Control of Systematics: One might think that the solution to all these interference effects would be to destroy the coherence of the laser beam, so that during the time it takes for the data acquisition to sample the photodiodes, all effects would time-average out. But the reason we used a laser was to get intense UV light in a finely-collimated spot, so these goals are in conflict. In fact, a poor laser (such as a laser-pointer, with coherence length a fraction of a mm), doesn't show much interference in windows but is too red, too weak and spreads the beam too much (not to mention that we already paid 7k$ for a fine UV laser optimized for our situation). Besides, many interference effects persist with incoherent light.

At the photodiode, where we no longer care about keeping a collimated beam over a long distance, we can insert a diffuser which scatters the beam over angles large compared to microradians, but still small enough that in the remaining pathlength of ~6mm will not diverge the beam beyond the limits of the diode. Typical diffusers such as opal glass or even your everyday Mystic Scotch tape help but have a lot of absorption (which hurts the measurement statistically) and spread the beam with significant long tails, defeating the good-geometry criterion. To solve this problem we used a relatively new product called a "holographic diffuser", a product of Physical Optics Corporation. This clever device is just a hologram of the beam output of a complicated laboratory device, which in its original lab gave a Gaussian beam of sigma 5°. When illuminated by our laser, the hologram simply produces a picture of the output of its parent device, which of course is the divergent beam we desire! Thus in a mm or so of space we can pack in all the information processing some large device originally produced. (Gerard t'Hooft thinks that the entire universe acts this way, but let's not get too far beyond our immediate concerns...). Even better, this diffuser transmits an amazingly large fraction of the incident intensity—we measure 84%.

The holographic diffuser works beautifully to kill the interference effects within the photodiode: we see a suppression factor of over 100. Fig. 4 shows what happens if we shoot the beam right at the photodiode without the tube inserted (**out**) and then slowly rotate the photodiode by ±1°. The ordinate is the normalized photodiode ratio and the abscissa is a sequence number for the measurement roughly proportional to time. Numbers 0-12 and 27-35 have no rotation: they show the inherent peak-peak noise of 0.2% (note greatly suppressed zero) and an rms of about 0.06%. Numbers 13-16, 17-27 and 35-44 show rotation scans: the first two are almost indistinguishable from the no-rotation sequences and the third shows a peak-peak of about 0.6% and rms 0.2%. (Glitches near the starting and stopping points of some sequences are artifacts of the different manual pressure exerted by the operator to start and stop the rotation).

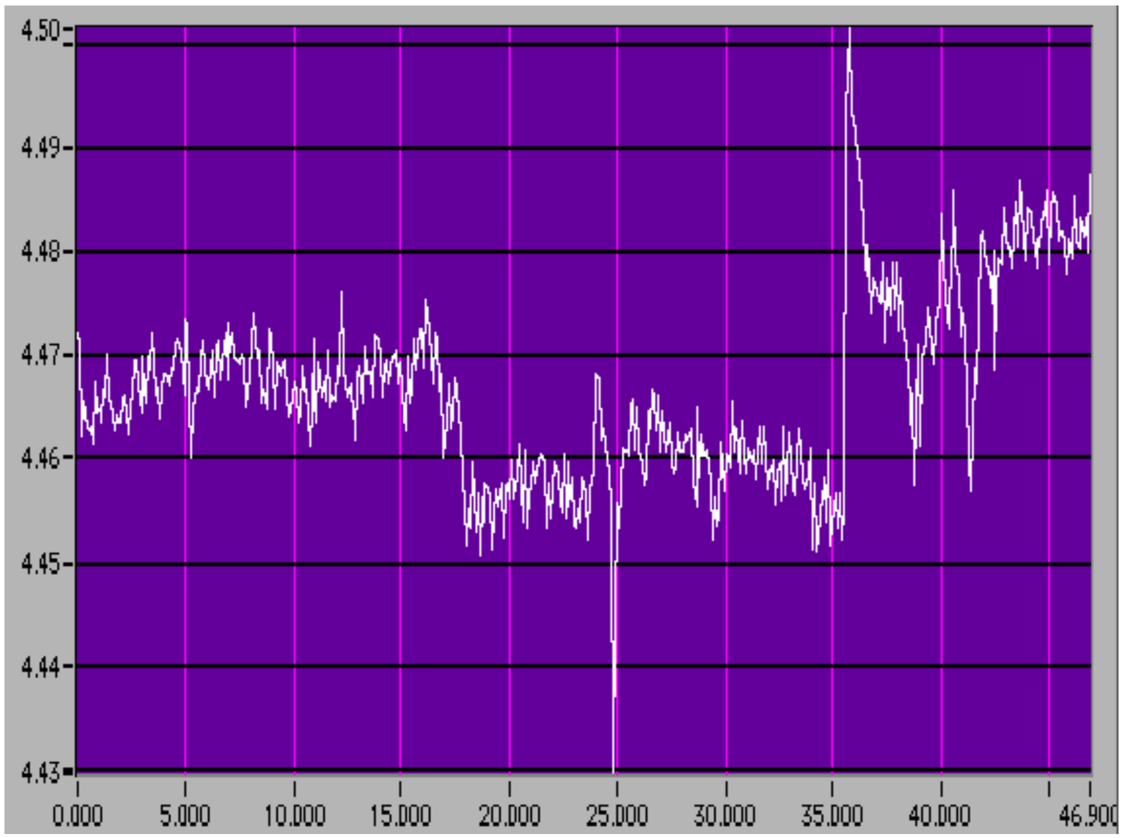

Figure 4: The stability of the measurements using a holographic diffuser.

This was such a good result we wondered if the interference effects had already mysteriously vanished since the last trip. Fig. 5 (overleaf) shows what happens when we removed the diffusers and did the same test. Numbers 0-20 are the no-rotation baseline and 20-26 the angular scan. Not only is the 15% peak-peak interference back (note scale change), but the system is so unstable that even with "hands-off" we see a classic relaxation oscillator behavior, as some tiny fed-back effect (beam heating?) affects the beam lineup at the microradian level.

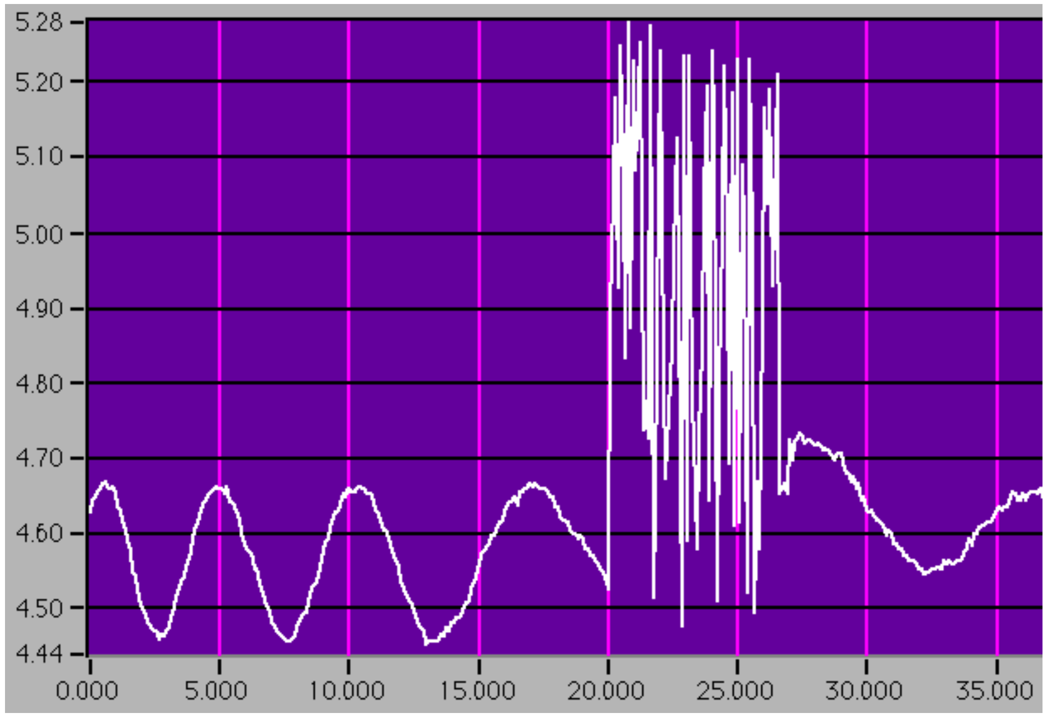

Figure 5: The effect of removing the holographic diffuser.

So holographic diffusers were installed on both prime and reference photodiodes. But recall that this was only half the problem—the endcap windows are also interferometers, and we can't use the diffuser solution on them. We had contemplated using non-reflective (NR) coatings on the windows, but had rejected this because as soon as the TUBE is filled with water, the quartz-water interface behaves differently than does the quartz-air interface and the surface becomes reflective again. A little further thought about time-reversal invariance gave rise to the following idea: suppose one coats only the outer surfaces of the endcap windows? Then, at the entrance window, reflections from the inner surface go back to the coated surface and should transmit through the coating perfectly, thereby getting lost from further interactions. A tedious calculation for a NR coating on a window, using the boundary conditions on the E and H fields, confirmed that this is true. A similar suppression of interference occurs at the exit window. But what then of the two uncoated surfaces, ~1 m apart, acting as a Fabry-Perot interferometer in themselves? The condition for interference (of the type which actually changes the direction of energy flow) is that the various reflections overlap where you are applying the boundary conditions. The laser is highly collimated, and just a 1-2° relative tilt of the windows makes the beam reflection non-overlapping at the critical surfaces. Thus singly-coated windows and a bit of careful alignment for them should do the job. In practice we tilt the windows enough to send the first reflection completely out of the system, and absorb it.

We found that Oriel Corp. would coat new, 11mm thick fused silica windows with a multi-layer non-reflective coating optimized for 325 nm, and purchased two such, plus two spares. We rebuilt the endcaps to conserve space for the much thicker windows. After the tests of the diffuser, the entrance window was installed and aligned to control the position of the reflections. We then did a transverse stage scan following the procedure of Fig. 2; the results are shown in Fig. 6. Readings from 0 to 4.0 correspond to the stage just sitting untouched; again, the inherent noise is about 0.2%. At a value of 4.0 the manual scan is commenced. If the coating is uniform and NR at 325 nm, there should be no variation in the signal. Instead, between 4.0 and 8.2 there are various dips and peaks amounting to losses between 1 to 4 %. In particular there are two strong dips. From values 8.2 to 9.8 the stage is at rest in the center of the window--no variation occurs. From 9.8 to 16.0 the stage is moved very slowly back to its original position. The pattern repeats, roughly, expanded because the motion is slower.

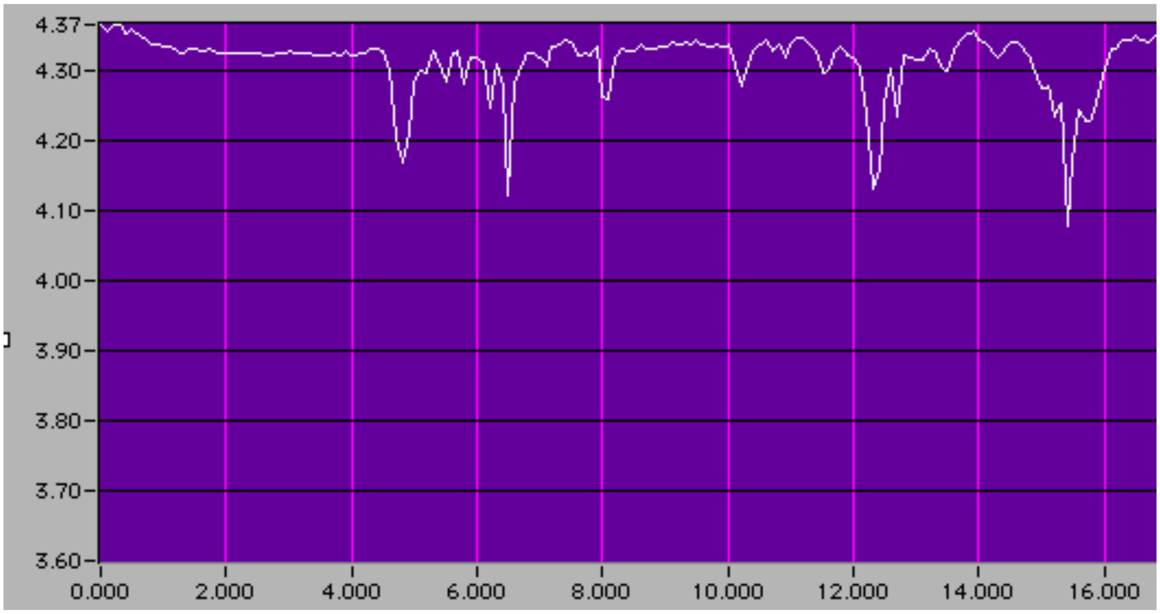

Figure 6: A scan as in Fig. 2, with an entrance window only. Note the scale and amplitude changes from Fig.2, a net factor of about five.

The NR coating has clearly helped, but not eliminated the prior 15% variation uniformly over the surface. We observed the window during the scan in a dark room. When the two very bad regions were encountered, one could visibly see reflections at the coated surface and the uncoated surface change in intensity and color, indicating the presence of pinholes. The laser beam is larger than the pinhole, so the pattern does not revert back completely to a 15% interference term, but does get to 4%. However, by sheer luck, the portion of the windows we use in the normal position of the stage has a variation of < 0.5%, which is tolerable. (So the window installation must be tuned).

Further measurements with the glass either in or out of the beam (no water, **empty/out** plotted in Fig.7 overleaf) reveals that two windows in series (thus two NR surfaces and two reflective) have a net transmission of 0.9083±0.0015. The prediction is 0.9267, correcting for the singly-reflecting surfaces and absorption in the 22mm of fused silica. At > 10-sigma, there is 2.0% of missing light. Two possibilities leap to mind: could there still be a small reflection from the film, leading to an interference cross term which gives us a dip of this magnitude, and would thus change and be an uncontrolled systematic when water was added? Or could it be that the coating, while perfectly nonreflective in the region of use, could also be absorptive, leading to less than 100% transmission at the two coated surfaces? If it is the latter, then in the **full/empty** measurement leading to T_o it should cancel out, being present identically in both modes.

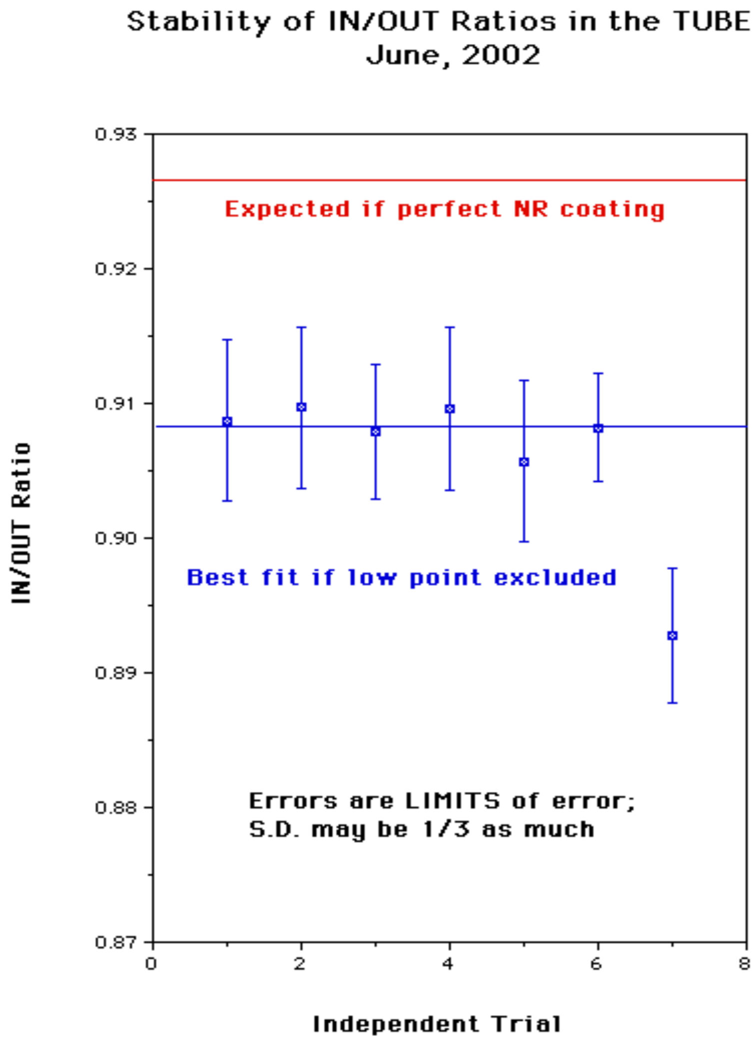

Figure 7: Repeated measurements of the **empty/out** ratio.

It is rather easy to distinguish between these two alternatives. We have seen that a particular interference-induced amplitude can't be reproduced with independent lineups of the beam, yet the data in Fig. 7 reproduced very well over 3 days involving 6 different sets of alignments, instrument disassembly for cleaning, varying temperatures, etc. Only on the 7th and last measurement, done hurriedly before leaving, did we see one variation of the order we would associate with an interference effect. (It is likely that this was just a dirty window or a spot of water left on the window. The statistics above were done with this point thrown out.) The constancy of the measurement in Fig. 7 argues that this missing 2% is not an interference effect. Furthermore, when multiple measurements of T_o were done, using different water samples from the same source, the results were in agreement within errors (see next section). This would not have happened with an interference effect because there are two independent and non-reproducible correction factors per measurement (two phases introduced by two different lineups during the measurement). They would also introduce gross changes between water samples, unlike the constancy predicted by a single absorption parameter. Finally, intense grilling of the Oriel people got them to admit that 1% absorption per surface of their NR coating was reasonable, and that pinholes could exist at the level at which we saw them. Our conclusion is that everything is consistent with, and our T_o measurements independent of, an absorption factor in the NR film.

Water Measurements: Thus finally we arrive at the point where we should expect consistent and correct absolute values of the water attenuation length. After all the fuss above, the actual measurements are rather anticlimactic, except for one curiosity, explored below. The measurements were made of recirculator water, gathered from the clean tap of the last filter (#4), and of pond water, gathered from the pickoff pipe from the main return line of the pond, before any filter is encountered. A sample of each was collected on 6/28/02 (7 pm), and then again on 7/1/02 (noon). The first samples were tested on 6/30/02 (9am to noon) and the second set on 7/2/02 (9am to noon), which allowed the water to degas and settle for a period before being decanted into the TUBE. Measurements proceeded rather routinely, and though hundreds of values of transmission were collected, so that statistical errors are completely negligible, the remnant systematic errors still dominate the precision. They are most likely from the thermal drifts of the electronics or "aging" of the detector surfaces. The results can be summarized succinctly by the following ratios: (using throughout $X_o = 0.9858$m)

Recirculator Water, first sample:

empty/out = 4.23 ± 0.02/4.65 ± 0.02 = 0.9097 ± 0.006
full/empty = 4.20 ± 0.01/4.23 ± 0.02 = 0.9929 ± 0.005

Best value for transmission through two NR coatings = 98.2% ± 0.6%

Best value for T_o (all corrections included) = 0.9243 ± 0.0048

Best value for L (all corrections included) = **12.52 +0.9, -0.8 m**

Recirculator Water, second sample:

empty/out = 4.70 ± 0.01/5.175 ± 0.02 = 0.9082 ± 0.004
full/empty = 4.70 ± 0.03/4.70 ± 0.01 = 1.0000 ± 0.006

Best value for transmission through two NR coatings = 98.0% ± 0.4%

Best value for T_o (all corrections included) = 0.9309 ± 0.0056

Best value for L (all corrections included) = **13.76 +1.3, -1.1 m**

Pond Water, first sample:

empty/out = 4.275 ± 0.01/4.72 ± 0.02 = 0.9057 ± 0.004
full/empty = 4.29 ± 0.01/4.275 ± 0.01 = 1.0035 ± 0.0035

Best value for transmission through two NR coatings = 97.7% ± 0.4%

Best value for T_o (all corrections included) = 0.9341 ± 0.0033

Best value for L (all corrections included) = **14.50 +0.8, -0.7 m**

Pond Water, second sample:

empty/out = 4.66 ± 0.02/5.22 ± 0.02 = 0.8927 ± 0.005
full/empty = 4.64 ± 0.02/4.66 ± 0.02 = 0.9957 ± 0.006

Best value for transmission through two NR coatings = 96.3% ± 0.6%

Best value for T_o (all corrections included) = 0.9269 ± 0.0056

Best value for L (all corrections included) = **12.98 +1.1, -1.0 m**

SUMMARY OF RESULTS:

Best combined value for recirculator: **L = 13.0 ± 0.7 meters**
(CL for consistency of the two results =48%)

Best combined value for pond: **L = 13.8 ± 0.6 meters**
(CL for consistency of the two results =28%)

Consistency of best values for pond and recirculator: CL = 40%

Discussion of results: The results for the heretofore unknown transmission through the NR coating largely agree with the larger multiplicity of such measurements mentioned earlier, except for the second pond sample. If this anomaly is caused by a dirty window in the last test, then the value of L should not be affected much, and it was not.

The consistency of the results from different water samples taken under different conditions on different dates is gratifying. Note that the absolute numbers obtained for **out**, **empty** and **full** can be quite different, while the ratios repeat well. The change in the absolute numbers reflects the long-term drifts in behavior of the laser, electronics and the transducers (the temperature range in the room is from 65°F to 95°F!).

There are only two remaining puzzles about these results. One is the curiosity that the Pond came out with slightly better water than the Recirculator, or, statistically speaking, they have equivalent water. One might expect the Pond to be worse, as past (but unreliable!) results have indicated. One can conjecture and come up with various possibilities: that the recirculator or that particular filter has become contaminated; that the pond water, because of suspended circulation, has settled its particulates, etc. This worries us enough that a test will be done with the refurbished TUBE using laboratory reagent-grade water. Unfortunately, the absolute value of such water is also unknown, but at least should have no particulates and might be expected to have L > 15-20m.

The other puzzle has not been discussed above, but has been noticed repeatedly. It is that once the full measurement has been done, we sometimes move the water-laden tube out of the beam and take an **out** measurement. We often, but not always, see a systematic difference between **out**(no water) and **out**(with water). In principle this is not expected, unless some strange reflection, weight-dependent distortion, or capacitative pickup in the nearby electronics is playing games with us. Some tests have seemed to eliminate all of these hypotheses. Since **ou**t(with water) is not used in the calculation, who cares? We do, because who is to say which is the affected measurement, if they differ? We will also pursue this slight glitch, which has the potential to displace results by about ±1-2 m.

In final conclusion, it appears that many of the frustrating troubles are behind us, thanks to better understanding and nice modern technology. We will know for sure only if the future measurements make sense and continue to be robust. The proof of the Pond Pudding will be in the Testing!

Appendix A: Formulae and Table of Values Used with this Report

One can use either L or T_o (transmission/m) interchangeably. For a sample cell of length X_o, the basic interrelations are:

$$T_o = e^{-(X_o/L)}, \quad T(x) = e^{-(x/L)} \quad \text{or} \quad T(x) = T_o^{(x/X_o)}$$

We slightly prefer to use L because T_o is constrained to be less than one, and imprecision in its measurement can push it into an unphysical region. However, one cannot avoid this and it leads to distorted probability functions for L as well. The likelihood function $\mathcal{L}$ for L as a function of the measurement of T_o (including the non-physical region) is:

$$(\mathcal{L}\,; T_o^*, \Delta T_o^*, X_o) = \{ [e^{-(X_o/L)}]/[L^2] \} e^{-[T_o^* - e^{-(X_o/L)}]^2/2(\Delta T_o^*)^2}$$

The formulae showing needed precision follow from the definition of P, the desired relative error = $\Delta T/T$:

$$\Delta T_o/T_o = P\,(X_o/x), \quad \Delta T_o = P\,(X_o/x)\, e^{-(X_o/L)}, \quad \Delta L/L = (LP)/x$$

These formulae are useful for interpolations and extrapolations of the Table of Values:

x(m)	L(m)	T_o (1 m)	T(x)	ΔT_o (required for P=5%)	ΔT_o (required for P=1%)	$\Delta L/L$ (required for P=5%)	$\Delta L/L$ (required for P=1%)
2	5	.8187	.6703	±0.0200	±0.0041	±12.5%	± 2.5%
2	10	.9048	.8187	±0.0220	±0.0045	±25.0%	± 5.0%
2	30	.9672	.9355	±0.0260	±0.0048	±75.0%	±15.0%
10	5	.8187	.1353	±0.0040	±0.0008	± 2.5%	± 0.5%
10	10	.9048	.3679	±0.0045	±0.0009	± 5.0%	± 1.0%
10	30	.9672	.7165	±0.0050	±0.0010	±15.0%	± 3.0%