Chapter 5

Summary

The LIGO community has been studying thermal noise in interferometers since the beginning of the project, and will likely continue to do so in the near future. Every year, new, previously unexpected types of thermal noise are being discovered, and experiments are only starting to catch up with theory. Structural damping noise was given the generic name "thermal noise" when it was the only expected noise source in the mirrors. Since then, the catalog has grown to include those described Chapter 2, as well as many other noise sources that aren't intrinsic to the mirrors.

The first success at measuring thermal noise was in 1994, when the old LIGO 40 m prototype interferometer measured a noise floor consistent with structural damping at the 10^{-18}m/rHz level [55]. Recently, Harry *et al.* [60] have put a handle on coating structural damping, and Kenji Numata [94] has measured thermoelastic noise and structural damping at the 10^{-18}m/rHz level at high frequencies ($\geq$ 1 kHz) in high-noise substrates. But we have very little data on the interactions between the mirror coatings and the substrate, which has been predicted by some to be the greatest limit to improving LIGO's sensitivity [26, 108].

Using the photothermal data from Chapter 3, we can provide a hint that something interesting happens between the mirror coating and the substrate. The data from sapphire are consistent with the model for the photothermal effect derived in Chapter 2, which assumes that the dielectric coating on the mirror is a low-conductivity glass. Relevant to LIGO I mirrors, the photothermal data for BK7 and fused silica both suggest that the coating makes a significant contribution to a mirror's thermal expansion, and that it increases at higher frequencies. The implication that the coating is not necessarily a slave to the substrate's temperature suggests that models for coating thermoelastic noise need to account for fluctuations originating in the coating as well as in the substrate.

There are two aspects to photothermal noise: thermal expansion and index of refraction. Are test mass surface fluctuations due to these correlated? If so, then coatings could be designed in which a negative dn/dT cancels thermal expansion, thus completely removing coating thermoelastic noise as a problem for LIGO. Pursuing this will require more experiments to understand coating properties, as well as direct measurements of coating thermoelastic noise.

Further development on the photothermal experiment should focus on evaluating more coatings, particularly Ta_2O_5/SiO_2 dielectric stacks like those used in LIGO, and on extending the measurements to higher frequencies to observe the high frequency behavior of the coating. It would be especially interesting to observe the photothermal response above 10 kHz, where the thermal diffusion length is less than the coating thickness.