

Eyeglass Large Aperture, Lightweight Space Optics FY2000 - FY2002 LDRD Strategic Initiative

R. Hyde

February 10, 2003

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U. S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

This report has been reproduced directly from the best available copy.

Available electronically at <http://www.doe.gov/bridge>

Available for a processing fee to U.S. Department of Energy
and its contractors in paper from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-mail: reports@adonis.osti.gov

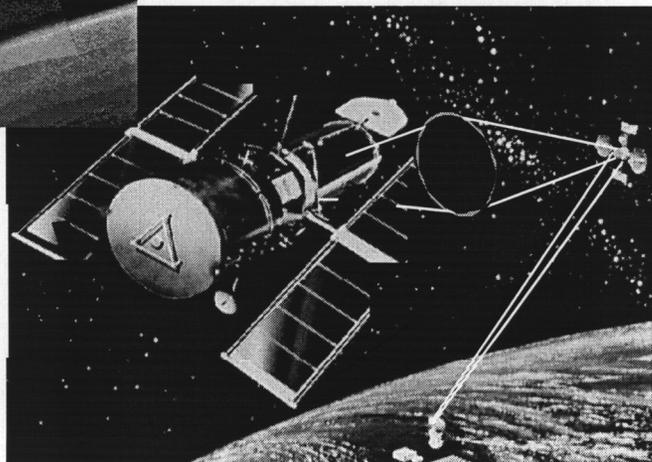
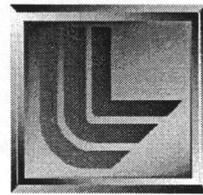
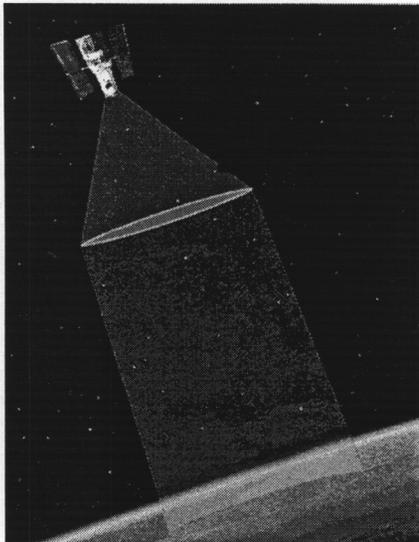
Available for the sale to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/ordering.htm>

OR

Lawrence Livermore National Laboratory
Technical Information Department's Digital Library
<http://www.llnl.gov/tid/Library.html>

Eyeglass

Large Aperture, Lightweight Space Optics



Final Report

LDRD Strategic Initiative Project : 00-~~777~~-003

SI

Large Aperture, Lightweight Space Optics FY2000-FY2002 LDRD Strategic Initiative

Directorates: NIF, P&AT, C&MS

PI: Rod Hyde (NIF, P&AT)

Co- Inv.'s: Charles Bennett, Jerald Britten, Sham Dixit, Jim Early,
Mark Eckart, Mark Hermann, Mike Rushford, Leslie Summers,
John Toeppen, Ian Thomas, Andrew Weisberg

Introduction

A series of studies by the Air Force, the National Reconnaissance Office and NASA have identified the critical role played by large optics in fulfilling many of the space related missions of these agencies. Whether it is the Next Generation Space Telescope for NASA, high resolution imaging systems for NRO, or beam weaponry for the Air Force, the diameter of the primary optic is central to achieving high resolution (imaging) or a small spot size on target (lethality).

While the detailed requirements differ for each application (high resolution imaging over the visible and near-infrared for earth observation, high damage threshold but single-wavelength operation for directed energy), the challenges of a large, lightweight primary optic which is space compatible and operates with high efficiency are the same.

The advantage of such large optics to national surveillance applications is that it permits these observations to be carried-out with much greater effectiveness than with smaller optics.

For laser weapons, the advantage is that it permits more tightly focused beams which can be leveraged into either greater effective range, reduced laser power, and/or smaller on-target spot-sizes; weapon systems can be made either much more effective or much less expensive. This application requires only single-wavelength capability, but places an emphasis upon robust, rapidly targetable optics.

The advantages of large aperture optics to astronomy are that it increases the sensitivity and resolution with which we can view the universe. This can be utilized either for general purpose astronomy, allowing us to examine greater numbers of objects in more detail and at greater range, or it can enable the direct detection and detailed examination of extra-solar planets. This application requires large apertures (for both light-gathering and resolution reasons), with broad-band spectral capability, but does not emphasize either large fields-of-view or pointing agility.

Despite differences in their requirements and implementations, the fundamental difficulty in utilizing large aperture optics is the same for all of these applications: It is extremely difficult to design large aperture space optics which are both optically precise and can meet the practical requirements for launch and deployment in space.

At LLNL we have developed a new concept (Eyeglass) which uses large diffractive optics to solve both of these difficulties; greatly reducing both the mass and the tolerance requirements for large aperture optics. During previous LDRD-supported research, we developed this concept, built and tested broadband diffractive telescopes, and built 50 cm aperture diffraction-limited diffractive lenses (the largest in the world). This work is fully described in UCRL-ID-136262, Eyeglass: A Large Aperture Space Telescope.

However, there is a large gap between optical proof-of-principle with sub-meter apertures, and actual 50 meter space telescopes. This gap is far too large (both in financial resources and in spacecraft expertise) to be filled internally at LLNL; implementation of large aperture diffractive space telescopes must be done externally using non-LLNL resources and expertise. While LLNL will never become the primary contractor and integrator for large space optical systems, our natural role is to enable these devices by developing the capability of producing very large diffractive optics.

Accordingly, the purpose of the Large Aperture, Lightweight Space Optics Strategic Initiative was to develop the technology to fabricate large, lightweight diffractive lenses. The additional purpose of this Strategic Initiative was, of course, to demonstrate this lens-fabrication capability in a fashion compellingly enough to attract the external support necessary to continue along the path to full-scale space-based telescopes.

During this 3 year effort (FY2000-FY2002) we have developed the capability of optically smoothing and diffractively-patterning thin meter-sized sheets of glass into lens panels. We have also developed alignment and seaming techniques which allow individual lens panels to be assembled together, forming a much larger, segmented, diffractive lens. The capabilities provided by this LDRD-supported developmental effort were then demonstrated by the fabrication and testing of a lightweight, 5 meter aperture, diffractive lens.

Organization of this Report

Introduction	2
The Mission	5
Why are large aperture space optics vital?	5
Why haven't they been fielded yet?	9
A New Approach: Diffractive Optics	11
The concept	11
The challenges	13
What has been achieved so far?	21
The Proposed Strategic Initiative	25
Our goal: To make a 5 meter lens	25
Basic approach to making the 5 meter lens	26
Choice of materials and patterning method	26
Assembling panels into a full-size lens	39
Testing the lens	46
What Was Accomplished	49
Demonstration of a Foldable Segmented Lens	50
Thin, Uniform Sheets for Diffractive Lens Panels	59
Design and Patterning of Diffractive Lens Panels	78
Precision Assembly of Large Diffractive Lenses	83
Fielding and Testing of the 5 meter Lens	118
Summary	134

The Mission

Why are large aperture space optics vital?

Space optics are already crucially important, both for Earth observation and for astronomy.

The national security is vitally dependent upon the information provided by earth observation satellites. Referred to as "national technical means," our satellite reconnaissance capability has been central to determining the intent and capabilities of others and to insuring treaty verification. This is widely considered to be the greatest benefit of the national space program, and, by itself, to have been worth the entire investment the nation has expended therein.

Space-based Earth observation is also critical for many other applications (weather prediction, resource planning, science, etc) and is beginning to emerge as a new commercial field-of-endeavor. For astronomers, the virtue of space-based telescopes is that they don't have to look through the Earth's atmosphere. This provides such benefits (both in opening up new spectral windows and in eliminating atmospheric distortions) that the Hubble Space Telescope is considered one of the "crown jewels" of astronomy.

As useful as existing space optics are, their value could be increased significantly.

Our existing space telescopes are small aperture (few meter) instruments. This limitation severely limits their utility in existing (Earth observation and astronomy) applications, and also effectively precludes desirable new ones (such as space-based laser weapons, or direct extra-solar planet detection, to name just a few).

Larger space telescopes, with apertures tens-of-meters across, offer not just evolutionary quantitative improvements, but also open up revolutionary new capabilities and missions.

National Security

A drastic increase in optical apertures will clearly improve our ability to examine sites on the Earth. Larger apertures permit one to either improve resolution, look from a greater distance, or extend given-quality observations to longer wavelengths. Large gains in aperture can be translated into significant improvements along one or more of these avenues, with the potential of greatly enhancing our observational capabilities.

To date, the primary national security role of space assets has been the strategic one of intelligence acquisition. This has been essential and will continue to be so. There are, however, other valuable things that could be done with space-based optics; missions that are infeasible now, but will become possible with the introduction of large (tens of meters) aperture optics. We'll illustrate such opportunities by discussing two new roles

for space optics; one for the support of tactical operations and the other for directed energy weapons.

The importance of local, timely, surveillance in tactical operations has been apparent from antiquity. Desert Storm and subsequent events have reinforced this lesson. At present, there is a major thrust to learn how to use our existing strategic space assets to gather and rapidly disseminate information to local commanders for use in tactical operations. Having a good, continual, view of the battlefield can be more valuable than occasional (though higher quality) glimpses. The ability to “hover” over a battlefield, looking at it all the time, makes GEO-based, large aperture, telescopes truly enabling for tactical surveillance.

This utility is not confined to only passive observation. Tactical commanders want to use surveillance information not just to identify opponents, but also to attack them. High resolution space platforms can not only see the presence of opponents, but can also pinpoint their locations; this information can then be used to attack and destroy targets. Such targeting can be either passive, identifying the enemy’s precise geolocation, or active, illuminating the enemy with a laser beacon. Both approaches are attractive and can be used either independently or in concert. So, large space optics enable tactical commanders not only to see the enemy, but to selectively destroy him as well.

Space optics are inherently reversible; they can not only gather light, but can also deliver it. This capability can be used indirectly (as mentioned above) to illuminate targets, enabling them to be identified and thence killed by separate weapons. At high enough intensities, however, the delivery of photons can have more direct effects, actively damaging and destroying targets. The virtue of space-based (laser) weaponry is both its remote application and its immediacy; one doesn’t have to get close to the target to kill it, and it can’t move (much) between the time it’s seen and when it’s attacked. As discussed in the Space Based Laser Concept of Operations, such weapons can potentially be used in many different ways, either strategically (missile defense, etc.) or tactically (enforcing air superiority, suppressing enemy troops or other assets, etc.) depending upon the available power and circumstances.

One specific application that is the subject of considerable interest and research is the use of space-based lasers (and/or optical relays) to defend against ballistic missiles. This is an application in which both remote kill and immediacy of effect are crucial, space lasers are one of the few options available for such a (supremely important) mission.

Technically, the means of killing an enemy missile is the delivery of a given amount of fluence to it. The amount of fluence required to kill a missile depends on the specifics of the target and the laser, but the process of remote transmission leads to fundamental relations between the strength of the laser, the size of our space-optics and the range-to-target: The delivered fluence is proportional to the strength of the laser, to the square of our optical aperture, and to the inverse square of range.

Large aperture space optics can be crucial in two respects; enabling new, more attractive, kill-mechanisms, and making the overall system affordable. The currently

planned chemical battlestations rely upon thermal kill, requiring the beam to dwell on target for a sufficient length of time. In combination with new, short-pulse, laser systems, large aperture optics might permit single-pulse kill-mechanisms; these allow qualitatively superior fire-and-forget attacks, in which the beam does not have to remain focused on its target.

As attractive as this different kill-mechanism would be, large aperture space optics offer an enormous practical benefit to any and all space-based laser systems; they make them much more affordable. Destroying ballistic missiles with laser weapons is a difficult, challenging, and expensive task; requiring fleets of powerful lasers. Here, the effects of power and range of individual lasers can be combined into a general, global figure-of-merit, namely the total deployed power needed for the fleet. Regardless of the specifics of target interaction and/or laser system, the total deployed power in the fleet scales inversely with the square of the optical apertures. The adoption of large aperture optics can be used either to enormously expand the capabilities of a system and/or to cut its cost. A mission cost which is expected (for the currently investigated chemical battle stations) to be dauntingly high could, for instance, be cut ten-fold by the replacement of 8 meter optics with 25 meter ones.

While ballistic missile defense via space based lasers may not technically require large aperture space optics, their use will most likely be truly enabling in a practical, financial, sense.

Astronomy

The history of observational astronomy is dominated by the continual quest for larger, higher quality, telescopes. This compulsion is active still (the newest ground-based telescopes have 4 times the aperture of Hubble), and will continue into the foreseeable future in space as well (the Hubble follow-on, NGST, is planned to have an 8 meter aperture, more than 3 times that of Hubble).

There are two benefits to having larger aperture astronomical telescopes; increased collecting area allows us to gather more light, and increased resolution allows us to more clearly distinguish different objects. Both of these improvements are, by themselves, of major importance: Increased collecting area shortens observation times, thereby allowing the telescope to be used to examine many more objects. Increased resolution permits us to view finer features on nearby objects and to sort-out distant targets (for instance, extrasolar planets from their stars, and extragalactic supernova from their hosts).

The combination of these two advances, increased area and resolution, is of fundamental importance in allowing us to look at ever-further (and hence dimmer) objects. Given the expansion of the universe, distant objects are ancient ones. Looking further away allows us to look deeper into our past, into the earlier, more primitive, history of the universe. Seeing faint, distant, objects is not simply a matter of collecting a finite number of photons, but also of separating the desired signal photons from unwelcome noise ones. The limiting noise for high-quality space telescopes comes from distributed sources such as zodiacal light, so finer resolution leads to less noise. Larger

telescopes provide a double advantage in seeing faint, distant, objects, not only does their area increase signal, but also their resolution decreases noise.

One of the most compelling (philosophically as well as scientifically) new missions in astronomy is the search for extraterrestrial life. This mission is central to NASA's Origins program and success would galvanize both the scientific community and the overall public as few other conceivable achievements. While the search for life in our own solar system continues, statistics argue that life is more likely to be found elsewhere. A key component in the search for extrasolar life is the detection, and subsequent observational inspection of extrasolar planets.

In recent years, a number of Jovian-class planets have been detected orbiting nearby stars. These planets have not been directly seen, but rather have been found indirectly (detecting the wobble of their parent stars). While of some scientific interest, such detections have little relevance (either scientifically or in the garnering of political support) to the extrasolar life issue. The desired goal here is the direct detection (and subsequent examination) of light from Earthlike planets. In the near term, direct detection would provide an independent, and more compellingly obvious, proof-of-detection; in the future such measurements could provide much more planetary data through analysis of the light's spectroscopic and/or spatial content. This detection and subsequent examination can be performed with a variety of wavelengths, but the most interesting alternatives seem to be either in the visible (seeing reflected light from the planet) or thermal-IR (looking at thermally re-emitted energy).

Direct detection of Earthlike extra-solar planets is, however, quite difficult. Not only are the incoming photon fluxes from such planets very weak, but they are angularly located close to the very much brighter signals from their parental stars and so are easily masked by their glare. The problem is not so much that planetary signals are intrinsically weak, but rather that they must be distinguished from nearby (and very much stronger) noise sources. There are two main sources of noise photons, zodiacal light from dust, and direct emission from the planet's parental star. Zodiacal light forms an extended source, it can be "burned through" via a high-angular-resolution instrument (either a large aperture telescope or a long baseline interferometer). More severe is glare from parental starlight; for visible wavelengths this is $\sim 10^{10}$ times brighter than an Earthlike planet, while for thermal-IR imaging the glare ratio falls to a still daunting value $\sim 10^7$. The defeat of parental glare is, as with zodiacal light, dependent upon very tight angular resolution instruments. For systems employing visible-light, photons are sufficiently bullet-like that a large (properly apodized) aperture telescope can isolate, even at 10^{-10} levels, the light from the planet and its star. With thermal-IR, diffraction effects are too strong for this option, but the glare is weaker, nulling interferometers, employing a few (relatively small aperture) telescopes can be used to detect planets. However, while this thermal-IR interferometric approach, (currently embodied in NASA's baseline Terrestrial Planet Finder) can detect planets, the actual spectroscopic inspection of them (key to detecting atmospheric life-signatures) demands the light gathering power provided by large aperture telescopes.

Large apertures are also essential when, after finding exo-solar planets, we decide to actually image them. Here, since we need sub-planet resolution at stellar ranges, we require very wide optics; for 100-1000 km resolution at a range of 50 lyrs, this requirement is 1-10 Mm. Clearly, this size can not be provided by a monolithic telescope (certainly not the ones proposed here). However, it is quite feasible to entertain long base line optical aperture synthesis. It might initially be thought that the requirements for aperture synthesis of apertures as large as 10 Mm would be impossibly demanding, but this is not the case. The requirement that must be levied on the wavefront quality produced by the individual sub-apertures in order to produce relatively efficient interferometric combination is that the imaging performance be dominated by diffraction. Thus if the individual sub-aperture telescope systems can produce nearly "diffraction-limited" performance, they can, in principle, be combined interferometrically. The absolute pointing requirements of the sub-apertures are also relatively benign. It is merely necessary to have the field of view of corresponding pixels in the image planes of the individual telescopes point to within a reasonable fraction of the single aperture diffractive point spread function. These relatively undemanding requirements parallel the analogous case of very long baseline radio astronomy; it is only necessary to point the individual radio telescopes in the array to within a reasonable fraction of the single dish antenna pattern.

Accordingly, arrays of separate, individually large-aperture, telescopes can actually image exo-solar planets. This is an exciting astronomical prospect, and has been identified as a goal of NASA's Origins-of-Life mission.

So, large aperture space telescopes are essential enabling tools for the detection of life on extra-solar planets. They provide both the angular resolution and light gathering power necessary to detect planets and light with visible-light, while their light gathering power is essential for thermal-IR based life-detection.

Why haven't they been fielded yet?

The benefits of having large aperture space optics are well known and there are extensive efforts aimed at learning how to field such systems. The reason they don't yet exist is simply that fielding large optics in space is exceedingly difficult.

The two fundamental difficulties preventing the fielding of large aperture space optics are meeting the tight optical tolerances necessary to achieve high resolution images (furthermore, doing so across large apertures), and simultaneously dealing with the weight, packaging and deployability limitations of space implementation. Either of these challenges is, by itself, quite severe; in concert, they have proven to be unsolvable.

Given these two extremely difficult tasks, existing research efforts typically try to start from a system which solves one of them, and then attempt, as best as possible, to deal with the other.

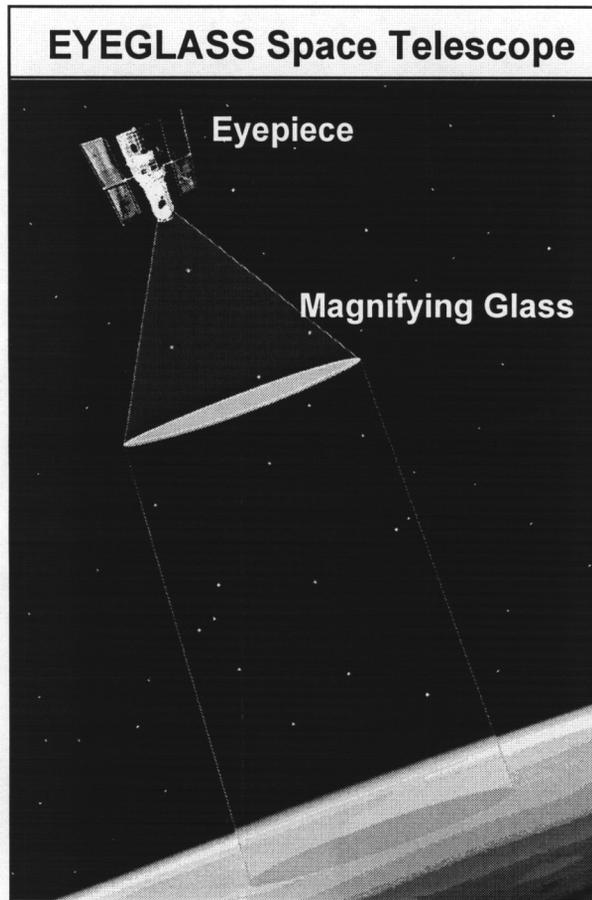
One approach (taken by most of the community) is to first concentrate on meeting the optical tolerances, and then try to deal with weight, packaging, and deployment issues.

Optical tolerances can be held over several meter apertures simply by use of well-figured, rigid materials; this has been done in ground and space telescopes for years. It is not feasible, even on the ground (where weight and deployability issues don't arise), to directly extend this approach up to 25+ meter apertures. However, it should be possible to design segmented telescopes which can hold tight tolerances. Here, the optical surface is split into segments (typically a few meters across), each of which is small and rigid. The overall aperture consists of a suite of such segments, mounted together by a stiff, lightweight (generally composite-based) truss. In order to approach practical launch weights, the segment masses must be substantially reduced from those of traditional optics. There has been substantial progress in doing so, both with purely passive segments, as well as the more promising method of using locally-stiff, actively-controlled, shells. Despite the research and progress in this approach, it remains challenging. The shells, support truss, and actuators remain relatively heavy and the packaging and reliable-deployment of such optics present enormous practical difficulties. Space deployable optics based on this technology are currently at the 25-30 kg/m² range, far too heavy for systems with 25+ meter apertures.

The alternative (and less mainstream) approach to implementing large aperture space optics is to first concentrate on a system which is light and fieldable, and then try to achieve the optical tolerances. The starting point here is to use thin membranes to reflect light, noting that a thin metal film (even when implaced on a polymer substrate) is both extraordinarily lightweight, and (given its thinness) flexible enough to be easily packaged. The difficulty, of course, is that such membranes have essentially no local stiffness. In order to be useful, they must be first deployed, and then held, in an optically precise shape; both tasks are quite challenging. In order to pull-out packaging wrinkles, and to achieve an optically smooth operational surface, the membranes must be held in tension. Furthermore, to function optically, the taut membrane must be precisely curved. This combination, a curved taut surface, can not be achieved from the edge supports, but also requires the distributed application of a pressure-like force across the entire surface of the membrane. Applying such a force, (typically attempted either by actual gas pressure or via electrostatic or electromagnetic forces) greatly complicates this approach. Of course, even once a basic shape-forming mechanism is in-place and successful, the membrane must be kept in this shape to optical ($\lambda/20$) levels of precision. Researchers attempt to employ adaptive optics to achieve this precision. However, the absence of any local stiffness in the membrane makes this quite challenging; the adaptive optics system must function with both high spatial frequency (because of the absence of mechanical stiffness) and high temporal frequency (to combat surface vibrations). So far, the difficulties of meeting optical tolerances have plagued this approach, preventing its successful adoption, despite its inherent fielding advantages.

A New Approach: Diffractive Optics

Recently, we have introduced a new approach, based upon the use of diffractive optics, to the fielding of large aperture space optics. This optical approach (labeled Eyeglass) is designed to simultaneously solve both of the fundamental difficulties facing large aperture space optics.

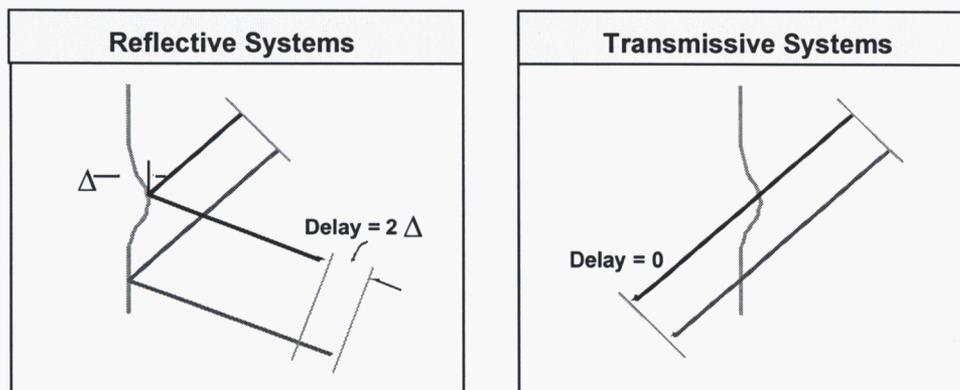


The concept

An Eyeglass space telescope consists of two distinct spacecraft, separated in space by approximately 1 kilometer. The telescope's large aperture is provided by a transmissive diffractive lens, acting as a Magnifying Glass. This large lens gathers and gently focuses light, directing it to a distant focal surface, where it is collected by a separate, much smaller, space telescope serving as the Eyeglass's mobile Eyepiece.

The use of a transmissive diffractive lens for the Eyeglass's large aperture is the key to solving the optical tolerance and fielding problems generally associated with large space telescopes.

Our new approach to large aperture space optics stems from one crucial fact; reflective systems are extraordinarily sensitive to surface bumps or ripples, while transmissive systems are not.



The reason for this fundamental difference is shown above: Mirrors reflect light, thereby doubling the effect of surface errors; in contrast, pathlength errors for a thin transmissive film are cancelled (not reinforced) as light arrives-at and then departs-from its surface.

The Eyeglass seizes upon, and vigorously exploits, this natural advantage. Its Magnifying Glass is not only a lens, but optically slow as well; this combination allows enormous, 10,000 to 100,000-fold relaxation in surface-figure tolerances. This advantage arises from the fact that in a reflector, surface ripples result in a phase distortion which is twice the magnitude of the ripple. In a transmissive optic, light passing through a ripple experiences the same optical path as light passing next to the ripple. Quantitatively, a ripple on a reflective membrane designed to focus light through an angle θ changes the light path by a $(1+\cos \theta)$. If the same membrane is used in transmission, then the effect of the ripple is to multiply the light path by $(1-\cos \theta)$. Lenses therefore have a $(1+\cos \theta)/(1-\cos \theta)$ tolerance advantage over mirrors. By making the lens weak, i.e., by keeping θ small, this tolerance gain becomes large, scaling as $4/\theta^2$, i.e., as $16(f/\#)^2$. A slow thin, $f/50$ lens can tolerate 40,000 times greater shape errors than a similar mirror. This relaxes surface figure tolerance requirements from hundreds of angstroms up to millimeter levels.

The reason for using a diffractive lens as the Eyeglass's large Magnifying Glass is simply that a conventional refractive lens would be thick; it is far too heavy and bulky to be fielded. Diffractive optics (i.e., coherent Fresnel lenses), however, do provide a means to field truly thin transmissive lenses, since they chop the applied phase correction into 2π -thick sections. This sized phase correction can be performed (in the visible) with a series of $1 \mu\text{m}$ deep quadratically-blazed surface cuts; even with an underlying substrate, the overall film is thin, packagable, and extraordinarily lightweight; enabling optics with an areal density less than 1 kg/m^2 .

The second challenge to large space telescopes, the difficulty of fielding objects in-space, is naturally solved because the Magnifying Glass is a thin flat membrane. The membrane's thinness ($< 200 \mu\text{m}$) makes it both lightweight and, because it is flexible, easy to package. Unlike reflective membranes, the lens's focusing power comes from a diffractive surface profile, not a curved shape; this membrane lens operates while flat, and is therefore easy to deploy and hold taut.

One unfortunate feature of an Eyeglass telescope is that it is inherently long; the product of its large aperture and high $f/\#$ (responsible for its tolerance advantages) dictates primary focal lengths of approximately a kilometer. Such lengths are impractical for a single spacecraft, so the Eyeglass telescope is split into two separate vehicles, a Magnifying Glass and an Eyepiece.

The Magnifying Glass vehicle is responsible for supporting the large aperture diffractive lens, holding it in the proper shape and for swiveling it in order to point towards desired targets. Incident light is gathered by the Magnifying Glass's diffractive lens and gently focused towards a distant focal point. There it is collected and imaged by a free-flying Eyepiece vehicle. This compact spacecraft has a modest (meter-scale) aperture; it is quite similar to existing (conventional) space telescopes. This compact spacecraft has a modest (meter-scale) aperture, and is, in most respects, similar to existing space telescopes; but does face two special requirements due to its cooperation with the Magnifying Glass. It is responsible for the stationkeeping necessary to keep the overall Eyeglass telescope in focus and properly pointed, and also performs the chromatic correction necessary to turn this diffractive telescope from a monochromatic to a broadband instrument.

The challenges

In principle, the Magnifying Glass offers a superb way to provide large optical apertures in space. Because it is an optically-slow lens, the traditional, very tight, surface tolerances faced by conventional apertures are virtually eliminated. Because it is a thin flat membrane, it is inherently quite fieldable (lightweight, packagable, and deployable).

An Eyeglass telescope also offers clear programmatic virtues; it decouples the cost, mass, and developmental risk of large space optics. The cost and mass are concentrated in a conventional vehicle that can be built (with very little risk) by existing contractors and procedures. The novel component, the large diffractive Magnifying Glass, dominates the Eyeglass's developmental risk (and possibly developmental cost), but not its operational mass or cost.

Actually implementing an Eyeglass telescope does, however, require us to face and solve significant challenges. These can be summed up by three questions:

Can we build it?

Can we field it in space?

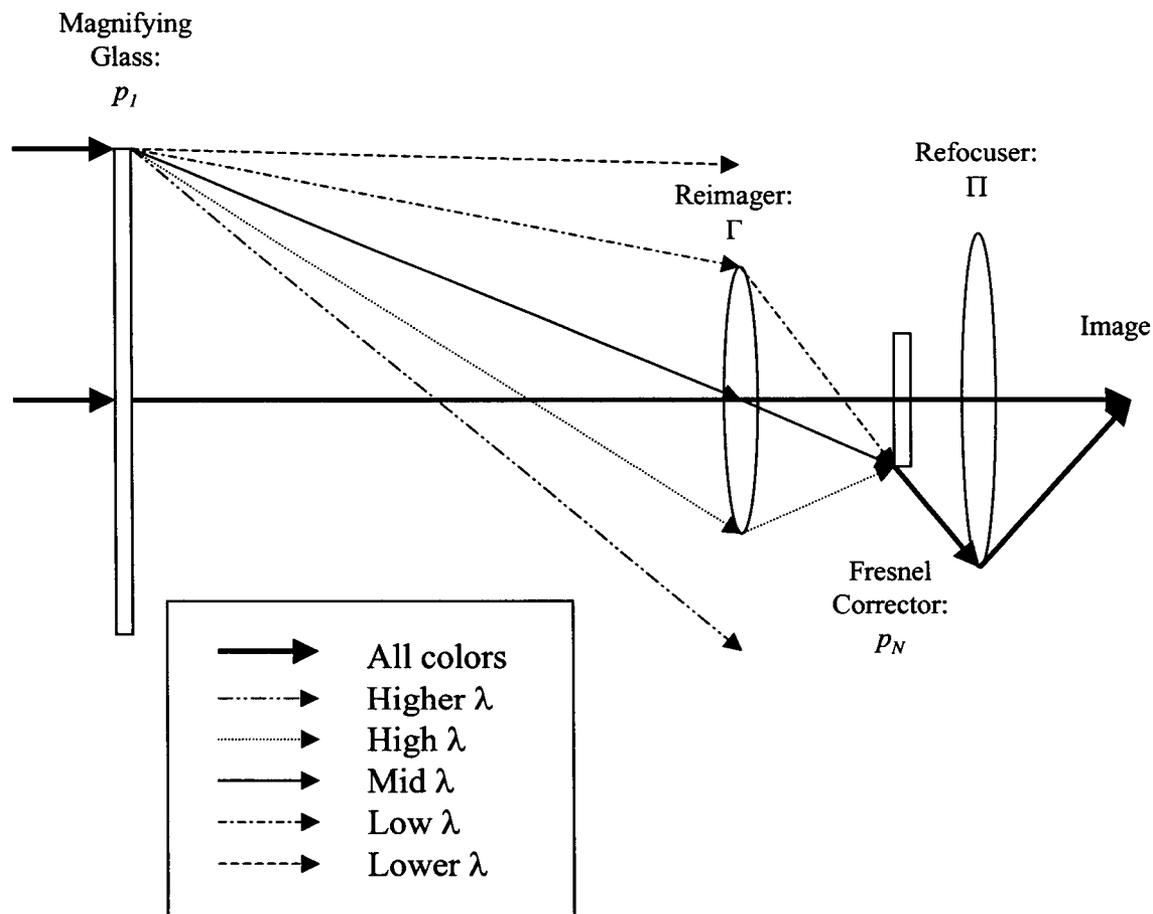
Will it work?

Will it work?

Basing a telescope upon a diffractive lens does pose a fundamental optical challenge:

Diffractive optics are strongly chromatic, with a focal-length inversely proportional to wavelength. Because of this, they are generally used only for monochromatic (i.e., laser) applications. Can we use diffractive optics in broad-band telescopes (for surveillance or astronomy), or must this approach be limited to laser applications, such as the space-based laser weapon?

To prevent the overall Eyeglass telescope from sharing the micro-bandwidth of its Magnifying Glass primary, we must employ some form of chromatic correction. We can not make the Magnifying Glass itself achromatic, so instead will cancel its chromatic aberrations with correcting optics in the separate Eyepiece vehicle. A way of doing this was invented 100 years ago by Schupmann for refractive systems. Fortunately, his technique is general, and applies just as well for diffractive optics. The general layout for a Schupmann-style diffractive telescope is shown below:



Physically, the reason Schupmann correction works for a diffractive telescope such as Eyeglass is clear; light leaves each point of the Magnifying Glass's diffractive lens in an angular spray, each color being sent into a different direction. As the light from this site travels towards the Eyepiece it spreads apart, diverging both spectrally and physically; both effects must be corrected. The physical reassembly is achieved first, by making the light pass through a reimaging telescope as it enters the Eyepiece. This internal telescope focuses the surface of the Magnifying Glass onto that of the Fresnel Corrector, thereby physically recombining rays which left each site on the first diffractive lens to a matching site on the second one. Now each site on the Fresnel Corrector sees an incoming angular/color spray corresponding to that from the departure site on the Magnifying Glass; by employing an inverse (defocusing) diffractive profile, it can remove this angular/color spray. As a result, each ray bundle is now both physically and spectrally recombined; the set of bundles can then be brought to a common achromatic focus. This is the same idea employed in the pulse stretchers found in chirped-pulse amplified lasers (one dimensional correction of spectral divergence).

In principle, this chromatic correction is perfect. In practice, with careful design and effort, it can be made nearly so, allowing diffraction-limited performance over wide spectral bandwidths.

The first challenge involved in achieving ideal performance is that the reimaging from Magnifying Glass to Fresnel Correction must be near-perfect; the relay optics themselves be achromatic and aberration-free. This prevents large operational telescopes from using a simple refractive lens relay as shown above. Instead, we have designed sophisticated, off-axis, dual-element, reflective relay systems which will allow diffraction-limited imaging with large Eyeglass telescopes.

The second challenge arises from the fact that the Eyepiece collector can only correct the light that it sees. Light begins spectrally spreading apart as soon as it leaves the Magnifying Glass, and that which is sufficiently far off-color will simply miss the Eyepiece. So, an Eyeglass telescope is naturally a single spectral-window device, having a natural bandwidth determined by the aperture ratio between the Eyepiece and the Magnifying Glass.

Multi-spectral diffractive telescopes are made possible by replacing the standard, mod- 2π , quadratically-blazed phase-profiles used in the Magnifying Glass and Fresnel Corrector of a single-band Eyeglass with higher-order, mod- $2\pi N$, profiles. From a fabrication standpoint this is a fairly benign adjustment; the surface features become N -fold deeper and wider, while the tolerance requirements remain unchanged, set by the shortest wavelength to-be-viewed.

Optically, this approach draws upon the fact that a quadratically-blazed diffractive lens designed to focus at some wavelength λ_0 also does so (with full efficiency) at its harmonics, i.e., at $\lambda_0/2$, $\lambda_0/3$, $\lambda_0/4$, etc. Because the Eyeglass's chromatic correction depends only upon the geometrically-fixed, achromatic mapping between the Magnifying Glass and Fresnel Corrector, this crucial process works simultaneously for all of the wavelengths.

The introduction of multi-band lenses therefore enables the full Eyeglass telescope to (simultaneously) deliver diffraction-limited performance at multiple spectral windows, centered about each of the harmonic wavelengths. Each of these spectral windows has the same fractional bandwidth, set by the Eyepiece/Magnifying Glass aperture ratio. The fact that the window's relative widths stay constant, while their central wavelengths are (for increasing harmonics) getting closer together means that eventually the windows overlap; for all higher harmonics the windows merge together, resulting in fully-continuous spectral coverage for all wavelengths below some value. In this fashion, we have designed Eyeglass telescopes (for astronomical applications) which are diffraction-limited from 0.4 – 5.0 μms .

The answer to the crucial optical question facing the Eyeglass concept is clear: Despite the strongly chromatic behavior of its individual diffractive elements, an Eyeglass telescope will be achromatic, allowing diffraction-limited imaging either for a single spectral window, for multiple separated windows, or over continuous spectral regions.

Can it be built?

Clearly, building a large, thin-film, diffractive lens will be a challenging task, and until actually demonstrated, will be a major source of concern. Nonetheless, it does look feasible.

In building a large, thin-film, diffractive lens, we can take advantage of two things; the required feature sizes are large, and the lens can (if necessary) be seamed together from smaller sections.

The smallest zone in a diffractive lens is given by $2\lambda f\#$. Since the large diffractive lens is optically slow, $f\# \approx 50$, the smallest zones will be $\approx 50 \mu\text{m}$, easily achievable with current technology. The challenge is presented by two other tolerances; the need to control the phase thickness of the lens material, and to place the Fresnel zones in their correct radial positions. In order to achieve $\lambda/10$ imaging, the phase profile (pattern plus substrate) of the film must be accurate (for typical materials) to $\lambda/5$; for visible light this is about $0.1 \mu\text{m}$. Note that this is purely a manufacturing challenge, fielding issues will primarily change the shape of the film, not its thickness. Typical commercially purchased thin sheets, of glass or polymers, are not globally accurate to this level, although can be over meter-scale regions. The required phase-thickness precision can, however, be reached in several different ways: The preferred approach is simply to improve the processing, giving thickness uniformity a much higher priority than it is typically assigned. Another, independent procedure, is to incorporate total phase-depth control into the diffractive patterning step. A third, and also independent step, is to achieve local accuracy (over ~ 1 meter scales) by the first two methods, but to correct larger-scale nonuniformities in-the-field, using an adaptive optics system in the Eyepiece.

At present no capabilities exist, anywhere in the world, for producing the large aperture diffractive lenses required for these missions. However, as a result of developing the gratings for the Petawatt laser and the phaseplates for Nova, LLNL has substantial infrastructure and expertise for fabricating diffractive optics at meter-size apertures.

Our current meter-sized diffractive optics are made by a lithographic process similar to that used by the semiconductor industry. We start with an optical flat, which is used as the initial surface. LLNL has developed a meniscus coating technique which can lay down a very precise thickness liquid coating on a surface. This meniscus coater can deposit flat ($\lambda/10$) photoresist coatings on apertures up to $1\text{m} \times 2\text{m}$ in size. Once this is done, an optical pattern is produced in the photoresist via illumination through a mask, and this pattern is then etched into the optical substrate. For complex, multilayer patterns this process must be repeated several times. The mask production, alignment, and exposure capabilities can currently handle about 1 meter substrates.

We have used these lithographic-based patterning capabilities to create diffractive lenses on two types of thin-film materials; a $40 \mu\text{m}$ polymer and a $380 \mu\text{m}$ glass. The polymer lens was produced by first creating a silica lens and then replicating it into a UV-

curable polymer. The glass lens was created with our standard etching processing, slightly modified to deal with glass instead of silica.

One should, therefore, be able to use the meter-class diffractive optics fabrication capability developed at LLNL for fabrication of multi-meter class diffractive lenses. The strategy is to produce ~1 meter x 1 meter size sections of the Fresnel lens using the existing LLNL capabilities. One can then sequentially seam together these panels to form a large aperture diffractive lens. The alignment accuracies necessary do not appear more severe than we have already developed for the multi-mask exposures used in some of our current diffractive optics assemblies.

Can it be fielded in-space?

After being successfully built, the thin-film diffractive lens must be delivered into space and successfully fielded there.

While diffractive lenses do greatly improve the launchability and out-of-plane optical tolerances of large space telescopes, they do face relatively tight in-plane tolerances. In order to focus coherently, the radial zones of a diffractive lens must be correctly located; zone borders must be globally, not just locally, precise. In principle, this requirement holds for conventional lenses and mirrors as well, not only for diffractive lenses; the location of a given phase change matters as much as its size. In practice, however, the effects of radial tolerances for mirrors are masked by the (much tighter) out-of-plane ones; this issue becomes apparent for diffractive lenses simply because out-of-plane tolerances have been relaxed so much. The size of the radial tolerances is easily found, for $\lambda/10$ wavefront errors the location of the Fresnel zones must be accurate to 1/10th of their width. For a typical f/50 visible-wavelength lens, the outermost zone widths are 50 μm , so the in-plane tolerance is 5 μm .

The in-plane tolerances of our diffractive lenses are, of course, very much looser than the normally encountered out-of-plane requirements for reflectors, 5 μm compared to 0.025 μm . However, they are much tighter than the corresponding 1 mm value allowed for the diffractive lens's out-of-plane tolerance. So the tolerances required for a diffractive lens are tight in-plane and loose out-of-plane. From a structural standpoint, this situation is well matched to the natural characteristics of thin-film membranes, which are inherently stiff in-plane and floppy out-of-plane; in contrast, reflective membranes are required to achieve extremely tight tolerances in their naturally floppy direction.

While the need to satisfy in-plane tolerances for a diffractive lens is a much easier task than trying to meet out-of-plane tolerances for a reflector, it is nonetheless a major challenge. Meeting in-plane tolerances will be one of the two major challenges to fielding large aperture diffractive optics (the other being the practical need to avoid "setting" creases in the film while handling and packaging it).

After being deployed in space, a thin-film diffractive lens will be kept flat by being held in tension. This flattens out elastic packaging wrinkles and raises the membrane's

out-of-plane vibrational frequencies, making it less floppy. The dominant challenge in meeting the in-plane radial tolerances is not due to this stretching, (which is predictable and hence can be allowed for when writing the Fresnel pattern), nor is it from the presence of manufacturing seams; the largest distortions instead come from thermal strains. It is essential for the diffractive lens system (the Magnifying Glass) to be thermally “clean”, that is, to avoid gradients in heating or in material properties that lead to nonuniform thermal stretching. Most of the standard mechanically-intensive tensioning schemes (such as extendable trusses or inflatable rims) are not thermally clean, and so will be hard-pressed to deliver the necessary in-plane precision.

Axial rotation is a particularly simple tensioning mechanism; slow, few rpm, rotational rates are sufficient to tension the membrane. In addition to its virtue of simplicity, a rotational system is both self-deploying and thermally clean; it requires only the optical membrane, with no extrinsic support structure. Uniform temperature changes (due to the emplacement in space and/or by changes in the intercepted solar irradiation) are therefore not a problem, since they lead only to areally uniform stretching. For a diffractive lens this just changes the focal length and hence is accommodated by a slight motion of the Eyepiece vehicle. With rotational tensioning, the largest source of radial distortion results from thermal stretching due to nonuniformities in the properties of the lens material. These effects can most readily be kept small (without requiring extreme material uniformity) by employing thermally rigid materials for the lens, either glass or specialized low-CTE polymers. An alternative method of dealing with in-plane distortions (if the required material uniformities cannot be met) is to take advantage of the fact that they have long spatial and temporal periods, so their effects can be removed by adaptive optics in the Eyepiece.

The potential drawback to axial rotation, of course, is gyro-stability; spinning makes it harder to swivel the Magnifying Glass, both to switch between targets and (for GEO-based Earth observation) simply to maintain a basic Earth-looking orientation during the telescope’s orbital travel. Our Magnifying Glass designs deal with this problem by placing a counter-rotating gyrowheel inside a small core-body at the center of the membrane. This central gyrowheel cancels the overall angular momentum of the Magnifying Glass, allowing it to be agilely reoriented (without external torques) simply by tilting the gyrowheel within its housing. The potential problem with this approach, that the transmission of reorientation-torque outward throughout the membrane deforms it away from being a flat disk, is not, in fact, significant; the distortions are too small to be optically troublesome.

The act of starting and stopping swiveling-maneuvers will generate out-of-plane membrane vibrations; if these grow to centimeter-scale sizes they will interfere with the lens’s optical performance. For astronomical applications, as well as most GEO-based ones, such vibrations are easily avoided; maneuvers can be performed slowly enough compared to the 2-4 second periods of the most troubling vibrations, that they are not strongly excited. Vibrations are a greater issue in SBL applications where large, rapid, reorientations are often necessary. Controlling vibrations in this case is more challenging and will require techniques such as damping at the housing-membrane interface, and optimizing the temporal profiles of the starting/stopping torques. Another practical

concern, namely jitter fed into the membrane by the central gyrowheel, has been investigated and found to contribute in-plane distortions well below 1 micron.

The Eyepiece of an imaging system is basically similar to an existing (conventional type and sized) space telescope, except that instead of looking at virginal light, it sees light that has already been focused by the diffractive Magnifying Glass. While this causes changes (discussed previously) in the vehicle's internal optics, it has little effect on the spacecraft aspects of this vehicle. The Eyepiece part of a large aperture diffractive telescope can be built by existing contractors with minimal developmental effort or risk.

Most stationkeeping tasks for Eyeglass telescopes, both moving between image-sites and preventing orbital drift, are handled by the Eyepiece vehicle. While the acceleration needed to prevent orbital drift is small (for instance, 4 μ gee to maintain a 2.5 km separation at GEO), it accumulates to large ΔV values over the course of a several-year mission. This combination of low acceleration, yet high fuel efficiency, is naturally met by electric propulsion. Tethers would also be well suited to this task, and eliminate any fuel cost. The concern with tethers is that they provide a channel which could feed vibrations into the diffractive lens from the Eyepiece vehicle; electric propulsion is a "cleaner" approach.

The second aspect of stationkeeping is, of course, precision; the two separate vehicles of an Eyeglass telescope must be positioned accurately enough so that the overall telescope functions properly. Fortunately the precision values which the vehicles must meet are modest, tens of centimeters in position and tens of arc-secs in tilt. As long as the Eyepiece is positioned and oriented to these levels of accuracy, then internal optical shifts and tilts within it can be used to maintain the telescope's imaging quality.

At present, it appears that if we are able to fabricate a thin-film Fresnel lens, then we will also be able to deploy and use it in a large aperture space telescope. The caveat is that this conclusion is based solely upon thought and analysis at LLNL. While this work has been reviewed by a number of outside agencies, it has not been subjected to the detailed, full-bore, examination of, say, an aerospace design team. Eventually, such studies must and will be done. After this, will come the even more crucial step of actual flight experiments. Only following such in-space demonstrations can success in implementation of real, full-sized, space telescopes be compellingly promised.

What has been achieved so far?

This approach of using diffractive lenses for large aperture space optics was created at LLNL in early 1996. The Strategic Initiative was begun in FY00. Beforehand, this project (under the label Eyeglass) was studied for three years, using both internal, LDRD, funds as well as external, government provided, support.

LDRD ER Research:

Eyeglass received (separate) ER funding for the three previous years, FY97, FY98 and FY99.

The purpose of the initial FY97 Physics ER project was to study the feasibility of this concept and to perform rudimentary systems designs. The basic optical feasibility of diffractive telescopes was established (on paper and computer) with single-band, on-axis, diffraction-limited designs. We also developed the current space-fielding concept, i.e., a flat, rotationally tensioned Magnifying Glass controlled by a central, counter-rotating gyrowheel. During this study we did substantial analysis of the effects of seams, swiveling torques, transverse vibrations, and nonuniform material properties on the Magnifying Glass.

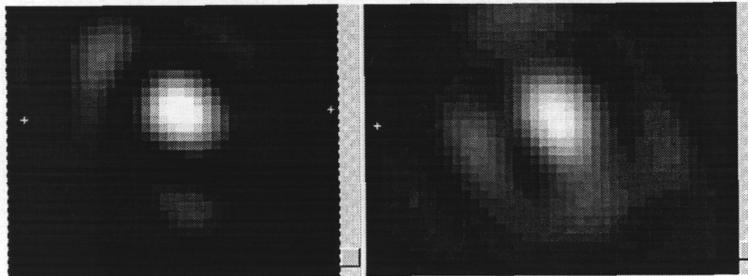
The emphasis in a FY98 joint Physics and Lasers ER project was to explore fabrication options for the diffractive primary. Most of our effort was focused towards replication; as the lithographic tooling and experience of the Diffractive Optics Group was directly available for making optical masters. We first used these facilities to build high-precision, 4 inch aperture, $f/100$ Fresnel lenses (both linear and circular) in fused silica. These Fresnel lenses were then used as masters for replication (via a UV embossing process) onto a thin photopolymer film. We successfully replicated a number of 4 inch Fresnel lenses from the fused-silica masters onto thin polymer films. The replicated copies accurately matched both the surface profile and optical performance of the original silica masters. Both the master and the replicated lenses were optically tested, and produced well-focused image spots.

Following the FY98 fabrication effort, Eyeglass was the focus of a FY99 Physics ER project. The main purpose of this effort was to experimentally demonstrate chromatic correction, which is basic to the optical feasibility of diffractive telescopes.

We built a diffractive telescope based upon a 20 cm, $f/100$ lens. The lens was fabricated by a lithographic process, generating a 4-level, digitally blazed, phase profile. In order to chromatically correct this diffractive lens, we combined it with a 4 cm relay lens (a refractive achromat) and a 2.2 cm, $f/11$, inverse diffractive lens.

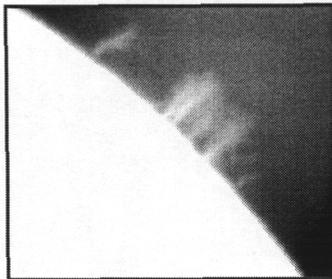
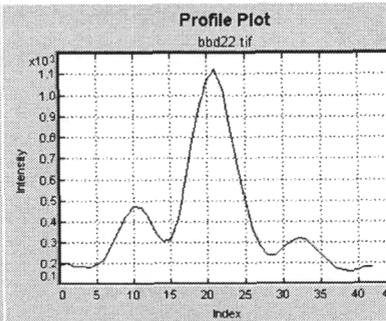
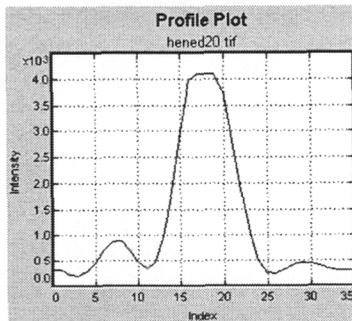
This experiment successfully demonstrated our ability both to chromatically correct diffractive telescopes, and to do so with diffraction-limited performance. Below we show the telescope's PSF for 633 nm laser light and for broad-band (470-700 nm) white-light.

The broad-band light is clearly focused into a tight spot, which has the predicted 95 μm Airy width. We also used this telescope (piping in light via turning mirrors) to look at real astronomical objects, taking pictures of the Moon, Saturn, Jupiter, and the Sun.

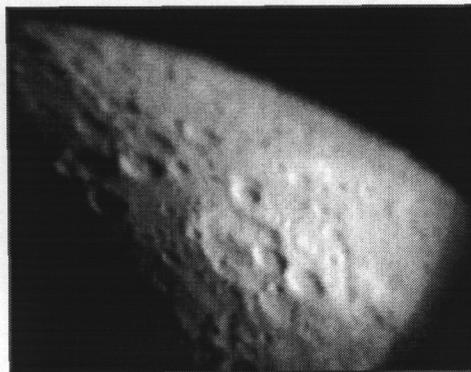


633 nm; Bldg AC off; doors open.

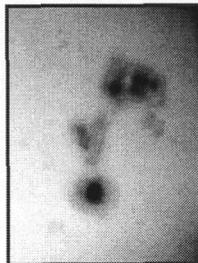
470-700 nm; Bldg AC off; doors open.



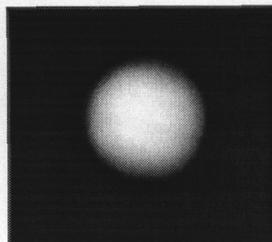
Solar Flares



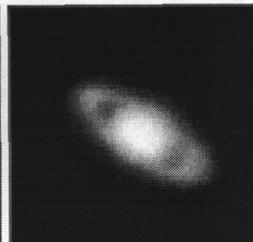
Lunar Surface



Sun-Spots



Jupiter



Saturn

The main result of the FY99 LDRD research was to convincingly demonstrate that diffractive lenses can be color corrected, and hence be used for broad-band imaging. In addition, however, we did investigate several other matters.

A full-sized space-based Eyeglass must, of course, use a thin-film diffractive lens, not a conventional, thick silica one. To demonstrate thin-film lenses, we built a 40 μm thick, 20 cm wide, polymer lens and successfully tested it in the color correcting telescope; it also gave color-corrected, diffraction-limited, optical performance.

A detailed final report discussing these results (and other optical-design and space-implementation matters) from our three year ER program has been published:

Eyeglass, A Large Aperture Space Telescope, Rod Hyde, Ian Barton, Sham Dixit, Frank Patterson, Mike Rushford, UCRL-ID-136262

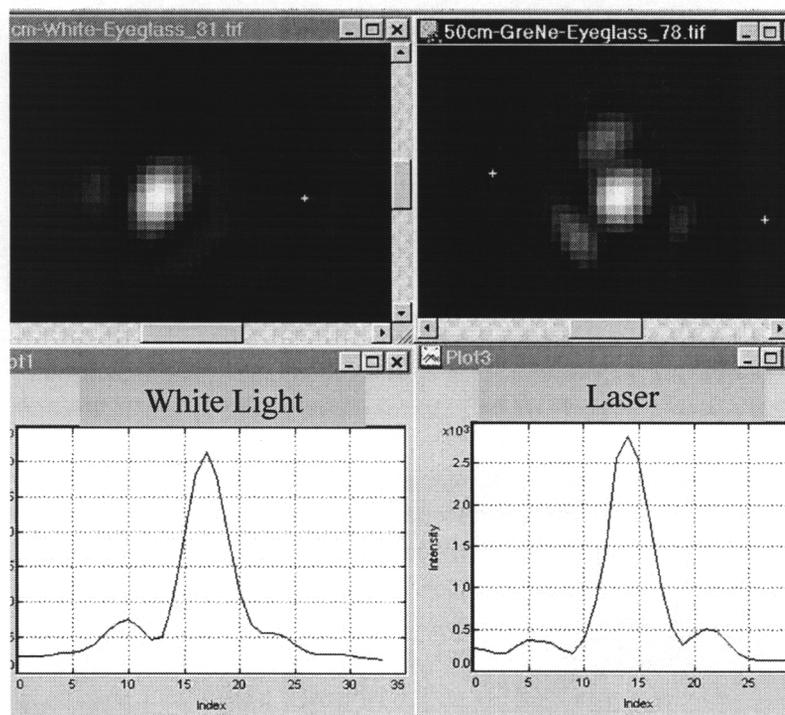
Externally funded research:

In FY99-00, Eyeglass first received external funds. This sponsorship was provided to scale-up fabrication and optical demonstrations to 50 cm sizes. The first task was to create a diffractive lens in silica, the second was to demonstrate color correction with this lens in a larger size telescope than our original LDRD one, and the third activity was to scale-up our thin-film polymer replication capabilities to this larger size.

We used our existing photolithographic tools to make a 50 cm diameter, $f/100$, 4-level diffractive lens. This required the writing of two separate binary masks, one forming an $f/100$ zone-plate pattern, and the other creating an $f/50$ pattern. The binary pattern from each mask was then converted into a corresponding pattern on the surface of a common, high-quality, silica substrate; the first profile was chemically etched to a phase depth of $\lambda/2$, while the second one superimposed a $\lambda/4$ pattern. Together, these masked etchings created a 4-level depth profile, which provided the digitally-blazed, phase-profile for our silica diffractive lens. To our knowledge, this is the largest, optically precise, phase-profile diffractive lens in existence.

We used this silica lens as the primary in a 50 cm aperture, $f/100$, color-corrected, diffractive telescope. The large diameter and slow f /number of this lens, leads to a 200 foot long design for the telescope. We used the NOVA laser bay provide the large, vibrationally and environmentally controlled, location needed for fielding this diffractive telescope. The optical design was basically a scaled-up version of the one used for the 20 cm LDRD telescope. The system had 3 parts, a plane-wave generator, the color-corrected diffractive optics, and a final focuser. The plane-wave generator simply consists a point-source (laser or white-light) which sits at the focus of a 36 inch $f/5$ parabolic mirror; the resultant plane wave then directly enters the diffractive telescope. The primary for our diffractive telescope is the $f/100$ silica lens described above. This is focused (at the 600 nm central-band wavelength) onto a 10 cm, custom-built, refractive achromat, which images the large diffractive primary onto a much smaller Fresnel Corrector. This corrector is the same one used in our 20 cm LDRD telescope; it is an e-beam written, continuously-blazed, diverging lens. After passing through the three components of the diffractive telescope, the light has been chromatically corrected, but is diverging from a virtual focus. A final optical system, a dual mirror Schwarzschild setup, converges light to a real, physical focus at the surface of a CCD camera.

As with the original 20 cm telescope, this larger one clearly demonstrates color corrected focusing. The figure below shows line-outs of the focused spots achieved with incident plane waves, one using a 543 nm laser, and the other composed of broad-band white light (470 to 700 nm). Both plane waves are tightly focused to spots of the diffraction-limited size, 62 microns, i.e., 9 pixels for our CCD.



The final task for the external funds was to scale-up our photopolymer replication process to 50 cm diameters. We retained both the same basic procedure, and the same polymer material developed during earlier, LDRD sponsored, work. The major development required in this scale-up was the acquisition and use of a large, 50 cm, stainless steel flat to serve as the bottom layer of our casting mold. The size and weight of this large plate required the design and construction of special equipment to mount and handle the parts, as well as to thermally release the film from the mold after it was cured. We successfully used our replication techniques to create 50 cm wide, 60 micron thick, films. The uniformity of these films have been interferometrically measured to be 1 wave (600 nm) over the full area of the part, and 0.2 waves in the central 25 cm region; these errors are dominated by those of our stainless steel flat.

The Proposed Strategic Initiative

Our goal : To make a 5 meter lens

The ultimate goal of our work is to create a new class of very large aperture, lightweight, space optics, based upon the use of diffractive lenses.

This is an ambitious undertaking, involving much more than simply solving the technical challenge of fabricating a huge, 25-100 meter diffractive lens. Not only does the lens have to survive and function in the space environment, it must first be compactly packaged, delivered there, and then deployed and held-in-place to optical tolerances. More complications arise from the fact that the large lens is just part, albeit the enabling portion, of a full-bore, space-based, optical platform. This will likely require tradeoffs complicating implementation of the lens itself; it also brings onboard all the infrastructure involved in designing, testing, and fielding large, complex, spacecraft. Clearly, the creation of this new class of large space optics requires both capabilities and resources far beyond those available to us here at LLNL.

Unfortunately, in both the national security and the astronomical communities, large aperture space optics (let alone diffractive ones) are viewed as a highly-desirable, but far-off goal. As a result, the resources needed to develop large diffractive optics are steered in other directions. To correct this situation, it is essential to convince the appropriate agencies that large diffractive optics are not just feasible, but in fact can be realized within decent time scales. The best, most effective way for LLNL to help make this case, is to actually build and demonstrate a large diffractive lens.

This is the reasoning which led us to make the fabrication of a 5 meter diffractive lens the primary goal of this Strategic Initiative.

Basic approach to making the 5 meter lens

At the start of our Strategic Initiative, we had settled upon the goal of making a 5 meter diffractive lens. This was intended both as a technological milestone along the way toward full-scale (25 – 100 meter) lenses and as a practical demonstration that large diffractive lenses could be made.

However, at the time, the state-of-the-art (set by LLNL's Diffractive Optics Group) was for sub-meter lenses constructed from thick fused-silica substrates. The only high-precision lightweight optics were 50 cm lenses (developed for Eyeglass) made from 60 micron thick polymer films. Both these optics were far smaller than the 5 meter lens desired for this Strategic Initiative. So our first need was to choose a scaling path which could produce a 5 meter sized lens.

We adopted a segmented approach to producing the 5 meter lens. Rather than building a full-sized, monolithic lens, we chose instead to create the overall lens by forming, and then joining together, a number of smaller lens segments.

The virtue of this approach is that it splits an extremely difficult task into two, much more tractable, parts. We use optical engineering to create lens segments, and then call upon mechanical engineering to accurately join them together. With this segmented approach, we can utilize meter-class optical engineering for producing lens panels, without having to extend these capabilities up to 5 meter scales. The price for this is, of course, that we need to accurately place, align, and join the individual lens panels together into an overall lens; at few- μm precision over a 5 meter throw. Fortunately, these tasks can be addressed via modern high-precision mechanical engineering. This segmented lens fabrication approach required us to stretch the state-of-the-art in both optical and mechanical engineering, but did not demand that we shatter present capabilities in either field.

Furthermore, to help serve as a convincing proof-of-principle for Eyeglass space optics, our 5 meter lens needed to address other properties necessary for use in space. We wanted to fabricate a thin lightweight lens from a space-suitable material, and do so in a form that could compactly packed for launch into space.

Choice of materials and patterning method

In order to produce the individual panels of a large diffractive lens, we must first identify an appropriate material, and then develop a technological route to form it into thin, optically patterned, sheets. The two issues are generally intertwined, with different fabrication approaches available or precluded depending upon the material choice. An acceptable material must satisfy several different conditions; it must be optically suitable, must survive and function in space, and enable patterning of the Fresnel zones.

The basic optical requirements for a material stem from the fact that it's being used in a transmissive lens. So, clearly, the material must be transparent at the telescope's operating wavelengths. We must concentrate upon dielectric materials instead of metals.

The other optical requirements stem from the fact that in order to correctly focus light, the lens must apply the proper, precise, phase correction to the light passing through it. For the type of lens we're developing (a surface-relief diffractive phase plate) this means that we need to cut a series of precisely made (depth and location) grooves in the surface of a uniform-phase substrate. So, in order to fabricate the lens, we need uniform sheets (to $\sim \lambda/5$ precision) of an accurately-patternable material.

The lens will be used in space, so it must function during exposure to vacuum, meteoroids, sunlight, and radiation; it does not, however, face severe loading conditions such as gravity or wind. The most fundamental requirement is certainly that the material survive the space environment; there are many materials that simply fall apart when exposed to either vacuum or radiation. A second requirement, of course, is that the material must not lose its transparency during space exposure; there are unfortunately a large number of materials which darken and become opaque in space.

Beyond raw physical and optical survival, however, is another hurdle; the material must survive without significant deformation. The optical performance of the lens depends upon it continually applying the precise phase profile built-in to it during fabrication. The dominant optical challenge is posed by the accumulation of small in-plane distortions over the global extent of the lens. The lens's Fresnel rings continue to exist, but don't stay in the right position, and so the lens no longer delivers a tight, coherent focal spot. In order to survive optically as well as physically, the lens must be built from a stable, low deformation material.

There are many effects which will cause the lens material to laterally stretch or contract during its operation in space. Exposure to vacuum can cause outgassing, and hence shrinkage. Radiation and UV exposure can break or rearrange chemical bonds, leading to dimensional changes. There are other, less permanent but still troubling, sources of distortion; ones caused by thermal or structural effects. If the lens's temperature changes, then its material will swell or contract by an amount depending upon the thermal expansion coefficient (CTE). Since the lens temperature is governed by solar heating, then (short of hiding behind a sun shield) thermal changes will occur whenever the lens is repointed, so thermal distortions must be dealt with even for environmentally robust materials. While a thin-film lens is not subjected to severe loads (such as gravity or wind) it must be kept under some tension (to pull it taut and to provide vibrational stiffness); this loading will be accompanied by stretching, by an amount depending upon the material's structural parameters ν and E .

All materials will deform under the operational conditions faced by the lens; the keys to functionality are the size of the deformation and its uniformity. As long as lateral distortions are small enough, they have little optical effect; adopting an optical error-limit of, say, $\lambda/10$, leads to a distortional threshold of 1/10th the local Fresnel zone size. This

depends upon the wavelength and f /number of the lens; for a visible-band lens operating at $f/100$, the distortional threshold is about 10 microns. Since we are interested in large lenses, with apertures > 20 meters, this distortional limit requires ppm or better precision. Given the environment and thermal cycles faced in space, it will be very hard for any material to hold this level of precision.

Fortunately, this extremely tight requirement is only imposed upon relative, not absolute, precision. Uniform expansion of a Fresnel lens only changes its focal length, not its focal quality. A diffractive telescope can easily accommodate reasonable changes in the primary lens's focal length by adjusting the location of the secondary optics (for instance by varying the separation between separate Magnifying Glass and Eyepiece spacecraft).

Our ability to tolerate dimensional changes in the lens material depends upon their nature. Virtually any-sized distortion can be handled, as long as it is known and predictable; one simply takes the anticipated changes into account when building-in the lens's diffractive profile. This allows us to deal with static (and spatially uniform) effects due to outgassing and differences between the lens's fabrication and mean-operational temperatures. It also can be used to handle static (but spatially nonuniform) distortions such as structural swelling in response to lens tensioning. This allow-for-it-when-building approach can not be used for distortions which are either not accurately known, or are time-varying. Distortions caused by thermal cycling, by accumulation of radiation dosage, or uncertain material properties can only be handled if they are small enough. Here is where the difference between spatially uniform effects (which can be tolerated to ~ 0.001 level) and nonuniform ones (which must be held below ~ 0.000001) becomes crucial. As long as the lens is designed and used properly, then thermal and radiation effects will be uniform and so are relatively tolerable; the tight ppm-level precisions apply only to nonuniform distortions. These are driven by the fractional variations in material properties across the lens and by the absolute level of the material distortion. Naturally, meeting the ppm-level nonuniformity tolerances will be easiest for a material with low uniform distortions, since a high distortion material will demand greater precision in material uniformity across the large lens aperture.

So, in principle, any uniform material which physically survives and remains transparent in space can be used for a diffractive lens; in theory any dimensional changes can be handled. In practice, however, there is a strong preference for dimensionally stable materials; we want stiff, low CTE, materials which don't swell or contract in the space environment.

There is a final, strongly practical, requirement on our lens material; it must be available and affordable. We want large, lightweight lenses, so their material must be available in the form of large, thin sheets. Materials which are too expensive, or require too much technical development, to be available in this form are not suitable for a project with the budget and schedule constraints of this Strategic Initiative.

In our search for a lens material, we have focused upon two different classes of materials, inorganics and organics. The inorganics are typified by silica or glass, while the organics by polymers such as Kapton and its derivatives. The table below lists and contrasts some of their basic properties relevant to our application.

	Transmission	CTE	Space suitability			Flexibility
			UV	Electron	Dimensions	
Silica	✓	✓	✓	✓	✓	X
Glass	✓	✓	✓	✓	✓	X
Kapton	X	X	✓	✓	?	✓
Adv. polyimide	✓	✓	?	?	?	✓

Inorganic materials

As is obvious from a glance at the preceding table, inorganic materials such as silica or glass have nearly ideal properties for our lens. They are optically transparent, mechanically stiff, and thermally rigid. Furthermore, these excellent properties are retained in the space environment. There are only two drawbacks to inorganic materials, their brittleness and their availability.

Brittleness:

Both glass and silica are brittle. This is no intrinsic problem for their operational use in one of our lenses, since the stiffening loads which will be employed are extremely weak. The brittleness concerns are primarily ones of handling while the lens is being built down on the ground, damage while it is being launched into space, and the effect of meteor impacts once it is deployed in space.

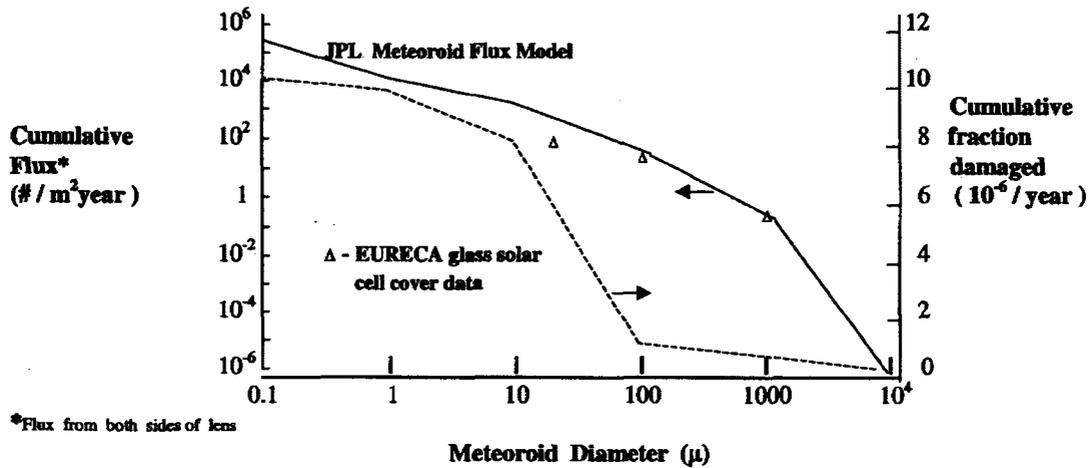
Insuring that the lens survives ground handling is basically just a matter of being careful enough with it. We have already accumulated some experience with 15 inch sized, 75 μm thick, glass micro-sheets. It is quite possible to either handle them successfully, or to break them. The keys to successful handling are the common-sense steps of not bending them too much, and of applying distributed loads, not concentrated ones.

Booster launches are a notoriously violent and vibration-rich environment; there are obvious concerns whether a lens built from thin glass sheets can survive launch. Until a lens is actually packaged and launched (or at least shaken on a launch simulator) we can't guarantee survival. The approach which appears to offer the best route to launch survival is simply to separate glass panels with soft, disposable, packing material so that lens panels don't touch one another, and to then pack the assemblage tightly, so that nothing can move. As will be discussed later in this report, there are lens folding techniques which do result in this desirable form of dense, layered, packaging. Using such techniques, we expect that glass lenses can be successfully launched; but launch simulator demonstrations will certainly be required.

In the event that handling or launch problems are more severe than anticipated, there is another approach which should greatly improve the robustness of the lens panels. This is to apply sacrificial plastic coatings to both sides of each glass panel. These plastic layers greatly protect the brittle glass panel in-between them, both by stiffening against flexure, and by preventing the glass from being scratched. The material used in the plastic layers is chosen to disintegrate in space, so that soon after the lens is deployed, the plastic overcoat disappears, leaving behind only the desired lens material. In effect, these plastic coatings are simply high quality packaging material, thrown away after delivery.

Micrometeoroids:

Thin glass structures will certainly be subject to damage from micrometeoroids. The magnitude of this meteoroid flux has been well characterized, and the nature of the resulting damage is understood. Data from numerous satellites, from atmospheric meteorite tracks and from light scattered to form the zodiacal light have been combined to form a consistent data set. While the details of velocity and angular distributions are not as well documented, the total meteoroid fluxes are believed to be well characterized. The graph below gives the JPL-NASA model of the cumulative flux of meteoroids larger than a given size. The European REtrievable CARrier (EURECA) satellite was returned from space after 11 months in LEO. The EURECA data on meteoroid cumulative flux is included in our graph and is in general agreement with the JPL model.



Thin (150μm) glass cover slides on the EURECA solar array provide a good database for the damage done to thin glass sheet by meteoroids. For small impacts the damage observed on the glass surfaces showed a central crater close to the meteoroid diameter and a damaged area with about eight times this diameter. When the meteoroid sizes reach about one-third the glass thickness, the meteoroid penetrates through the glass. The diameter ratio of the damaged area to that of the meteoroid decreases as the meteoroid size approaches the size necessary to penetrate the glass. The EURECA damage observations are generally consistent with laboratory damage test data.

Combining the flux model with the damage sizes observed on EURECA, we can generate the cumulative fraction damaged shown in the graph. After a ten-year exposure only 10^{-4} of the Eyeglass lens surface would be damaged. There would be hair-sized holes in the lens with spacings of about one inch.

The impact sites (either craters or holes) would have a damaged area which includes many cracks. These cracks are formed during the impact as the shear waves locally exceed the yield stress in the glass. Fortunately, these cracks will not continue to grow after the impact. The low tensioning loads in our lens do not generate significant stress in the glass. Also, the lack of water vapor in space makes glass much more resistant to crack propagation. In addition to permanent damage via cracks and holes, meteoroid impacts can also setup long-lasting vibrations in a thin-film optic. While these vibrations might severely trouble a membrane reflector (given its tiny, ~ 20 nm, tolerances), they are no concern at all for our diffractive lenses (which can tolerate millimeter-sized ripples). Thus, micrometeoroid damage (holes and/or vibrations) will not be a significant issue for an Eyeglass lens.

Assuming that brittleness concerns can be dealt with, then the remaining challenge to use of inorganic lens materials is the matter of availability. Can we acquire (and afford) large, thin, optically patterned, glass or silica sheets for use in our lens?

Acquiring thin sheets:

While glass and silica have broadly similar properties, and we could comfortably use either in our lens, silica is (if available) the ideal. Glass properties tend to mimic, with some penalties, those of silica. So, for instance, glass does not offer quite the broad-spectrum transparency of silica, has several-fold higher thermal expansion, and presents some (probably minor) radiation concerns. However, because it can be processed at much lower temperatures, glass is generally cheaper and easier to acquire than silica. So, while silica would be our ideal material choice, we might find glass a more attainable choice.

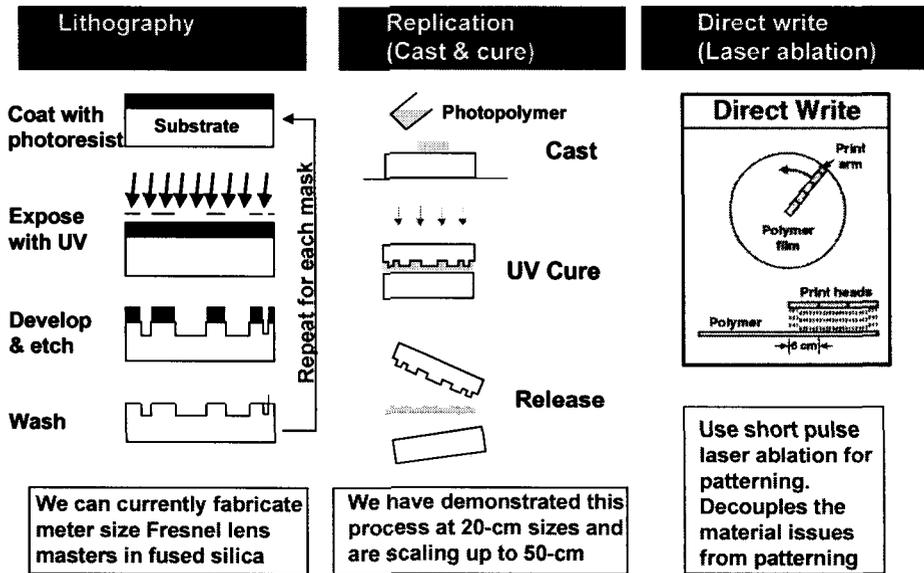
There are three aspects to fabricating inorganic lens panels. The first is acquiring large, thin sheets, the second is insuring that they are optically uniform, and the third is applying a diffractive phase profile to their surface. These tasks can be performed either sequentially or simultaneously.

Patterning:

Of these three fabrication tasks, the actual lens patterning appears the most straightforward. We at LLNL have already developed the capability of writing (by photolithography) meter-sized diffractive patterns, and have been successfully doing so for several years.

One complication for the thin lens patterning, in relation to our standard process, is the decreasing thickness of the base substrate material. For the patterning procedures under consideration, this is simply a handling issue, not a fundamental one. We have already successfully patterned both glass and polymer substrates at $<100\mu\text{m}$ thickness.

There are three main fabrication methods for the patterning process that we have could employ: photolithography, replication and direct-write. These different processes are illustrated below. Photo-lithography is most directly applicable to inorganic materials, and replication to organic ones; direct write would be applicable to all types of materials.



Photolithography is a mature technology and has been used by us to pattern diffractive structures in silica substrates up to $\sim 1\text{m}$ in size. It involves transferring the desired pattern into a light-sensitive photoresist film coating on the substrate that acts as a masking layer for subsequent etching steps. This process is suitable for virtually any material, as long as an appropriate uniform etching process can be found.

The photolithographic process that we routinely use for producing large diffractive optics in glass is detailed above. In this process the substrate is coated with a thin positive photoresist film, using a meniscus coating process. The desired pattern is printed in the resist using proximity exposure through a binary chrome mask. The photoresist layer now itself acts as a mask, with the exposed substrate (usually fused silica) being etched using a buffered HF acid solution. Depending on the pattern, a chrome layer can also be applied to the substrate underneath the photoresist layer that acts as a more resilient masking layer in the etching step. Multilevel structures are made by the repetition of this process, with an alignment step required to overlay subsequent mask patterns over the first with micron-level precision.

Replication, the second process illustrated above, is most directly applicable to the patterning of polymers. We developed this process (for up to 50 cm sizes) during the course of previous Eyeglass work, and have already discussed it earlier in this report.

While it will not be needed for the patterning of inorganic sheets, a variant of replication will be discussed later as an approach to making such sheets optically uniform.

Direct-write techniques involve the pattern being directly tooled into the substrate by either mechanical or ablation methods. The advantage of this type of process is that it could be configured for virtually any substrate material (inorganic or organic) and can be directly scaled to larger sizes. However, at the moment, the largest diffractive structures that have been produced using this type of process are ~1cm in size, so considerable development would be required to scale up.

Because of our vast experience with meter class diffractive optics fabrication and the compatibility of this type of patterning with the envisioned seamed Fresnel lens, we believe photolithography is the most well advanced (and perfectly adequate) procedure for patterning thin inorganic sheets.

Forming thin sheets:

The forming of large area, thin sheets (particularly with adequate uniformity) of glass or silica is expected to be a larger challenge than patterning them.

There are two basic approaches to producing such sheets, sculpture and growth. The first route is based upon acquiring a sheet which is either too thick, or insufficiently uniform, and then properly removing material to attain a thin, uniform, sheet. The second approach is based upon the fact that we actually need only a thin layer of material. Here we directly grow the sheet to the desired thickness, hopefully with the proper uniformity.

Historically, virtually all large silica or glass based optics have been produced by sculpture. The simplest (and oldest) such procedure is to start with a large boule of material and then grind and polish it until the desired shape is attained. Although this is wasteful of material, and generally quite expensive, it does work.

Since this process works as well for silica as for glass, there is little reason to apply it to glass sheets; we might as well produce the superior, silica, material. Large boules of silica are typically grown using a flame hydrolysis process; current capabilities at Corning consist of boules approximately 1.5 m in diameter, yielding 1-1.3 m size optical substrates. The challenges for using polishing to produce large, thin, sheets of silica are that conventional polishing techniques get increasingly difficult as the substrate thickness is reduced. Recently J. Burge at University of Arizona has developed a clever way for polishing thin sheets down to sub-mm thicknesses, using a thick substrate as a backing surface during the polishing process. Burge has demonstrated high quality ($< 0.1 \lambda$ transmission wavefront), thin sheets (~1 mm thick) of fused silica and BK7 glass at up to 1m size. He feels confident that the thickness can be reduced even further. This process however, is very labor intensive and hence expensive. Therefore it is our opinion that this approach ought to be used only as a last option if other approaches discussed below prove to be infeasible.

Instead of polishing down thick substrates, the thin sheets required for Eyeglass can be produced by direct growth. It is possible to directly (and cheaply) purchase large thin sheets of the mixed component borosilicate glasses. We have identified several candidates that appear promising for Eyeglass applications. Schott produces thin float glass, Borofloat, in meter sizes and thicknesses down to 700 μm . Corning produces meter-sized sheets of 1737 and 7059 glass by a fusion process in thicknesses as low as 500 μm . Down draw processes are used by Corning (0211 & 0214 glass) and by Schott (D263T & AF45) to produce much thinner (although smaller) glass micro-sheets. Both companies produce sheets about 35-40 cm wide; Corning goes down to 75 μm thickness and Schott to 30 μm .

These sheets are all borosilicate glasses. As such they deliver silica-like transmission over the visible regime, suffering extra absorption primarily in the UV. Their greatest drawbacks compared to silica are UV-induced browning under long-term exposures, and substantially higher thermal expansion ($4\text{-}7 \times 10^{-6}/^{\circ}\text{C}$ compared to $5 \times 10^{-7}/^{\circ}\text{C}$). The former problem can be eliminated by addition of Cerium (as is done with the Corning 0214 mix), while the higher CTE values are probably still acceptable.

Smoothing:

In principle, it is certainly possible to live with as-built, wavefront non-uniformities in our lens, and handle them by adaptive optics techniques. Here the wavefront aberrations of the primary optic would be measured and then be compensated for in the secondary section of the telescope by using an adaptive optics corrector system. The corrector can be either a metallic membrane or the more advanced liquid crystal modulator type correctors. The technology in this area, driven by their need for astronomy and high power laser systems, keeps improving continuously. Our application benefits from the fact that as-built errors are static ones, and do not require high temporal frequency correction. However, if wavefronts are nonuniform over several-cm size scales, then high spatial frequency correction will be needed, requiring megapixel correction systems. These are not out of the question for future adaptive optics systems, but will certainly be challenging in their own right. This approach also significantly dilutes our qualitative (and marketing) advantage over other, mirror-based, large optics proposals. We feel this adaptive optics approach should be adopted only as a last resort after other options for wavefront reduction have been exhausted.

To avoid having to retroactively smooth our lens panels, we should certainly try to improve the initial sheet fabrication quality. If this can be affordably done (in terms of both time and money) it is preferable than having to post-process the sheets. But, assuming that sufficient improvements are not feasible, there are a number of smoothing techniques which appear promising.

One such approach is a variant of the polymer replication process developed by us in the past, and discussed earlier in this report. This technique allows us to start with a coarse sheet (glass or polymer), and then use controlled etchback to planarize it. The

starting sheet can be produced either here or off-site by deposition on a precision flat substrate. We coat the raw sheet with the previously selected UV-curable polymer, and cover the sandwich with a transparent, precisely flat, top-mandrel. The diffractive surface profile can be introduced by pre-patterning either the top or bottom mandrels of this mold. The overlying photo-polymer is UV-cured in place, leaving us with a precisely uniform and patterned polymer film; the only problem being that it is a sandwich of the desired lens material and a space-unsuitable photo-polymer. At this point we can then remove the upper photo-polymer by ion etching, using a controlled etch having the same rates in the two different materials.

This type of approach is, in principle, applicable to a wide variety of sheet materials, including most inorganics as well as polymers such as the polyimides. There are two limitations, one technological and the other physical. The first concern is that the size of the patternable sheets is limited either by that of the flat mandrels (top or bottom) or the ion-etcher; currently we would be constrained to sizes slightly less than a meter (fine for a 5 meter lens). The second concern is that we need to match the etch rates in the top photo-polymer and the desired underlying lens material. In principle, this matching need not be precise, since the process can be iterated multiple times.

The other smoothing methods we're considering are more active; they require us to first measure the sheet's topology and then provide means to selectively remove material where necessary.

One such approach is clearly laser ablation. This was discussed earlier for lens patterning, and can also be adopted to smooth the lens panels. The challenges remain the same; developing a large laser ablation machine will be a substantial undertaking, a Strategic Initiative scale effort by itself.

Other selective smoothing approaches appear more tractable. The ones we will concentrate upon are based upon spatially selective etching. One method uses sequential processing, in which we apply etchant (at any given time) only to a small region of the part; dwelling for more-or-less time at a given site depending upon how much material must be removed there. A second, potentially faster, approach uses simultaneous etching. Here the entire part is etched at once, but the etch-rate is locally controlled to match removal needs. A number of process-knobs (temperature, wettability, acid concentration, surface access, etc.) can potentially be utilized to provide the necessary spatial control over the etching rate.

Organic materials

Organic materials such as polymers provide the backup option for our lens material. They have both advantages and disadvantages compared to the inorganic materials discussed above. The primary drawback is that organic materials generally degrade significantly in the space environment. Most polymers will not even survive in space, let alone with enough transparency and dimensional stability to be useful for a large diffractive lens. The reasons for maintaining an interest in polymeric materials are that

some have been successfully and extensively used in space, that they are readily available in large, thin sheets, and that they form flexible, easy-to-handle, films, without the brittleness of typical inorganics.

Material selection:

We are not, of course, interested in the vast majority of polymers that do fall apart in space; our focus is upon those that actually survive. Over the 4 decades that polymer films have been used in space, most have been one of three materials, either Mylar, Teflon, or Kapton. Unfortunately, none of these are directly suitable as a lens material. Mylar and Teflon are transparent but degrade fairly rapidly upon exposure to space radiation. Kapton is quite radiation tolerant, but not very transparent. In addition, all three of these materials have, as do almost all polymers, a high thermal expansion coefficient and a low stiffness, so will deform more than typical inorganic materials.

So, in searching for an organic lens material, we cannot simply use one of the existing space polymers; we must search for more advanced materials. We'll demand two properties, radiation resistance and transparency, and strongly desire two others, low CTE and high stiffness.

There are two types of polymers which appear most promising in satisfying the two essential requirements, transparency and radiation resistance. One candidate is polystyrene; it is both transparent and extremely radiation tolerant. Unfortunately, it has almost no actual space experience, and is not promising in terms of the desired virtues of low CTE and high stiffness. Because of this, we have focused on another class of polymers, the polyimides. Kapton, which we've already rejected, is of course a polyimide; while it is not suitable, other more advanced polyimides may well be.

The virtue of the polyimides is that they constitute a broad family of polymers, with many members. We can search for polyimides which retain Kapton's virtues of radiation resistance, space suitability, and good (for a polymer) stiffness, while improving upon its transparency and CTE.

Polyimides are a polymer chain whose monomer is formed by combining two different chemicals, a diamine and a dianhydride. The breadth of the polyimide family is due to the fact that many diamines and dianhydrides can be utilized; the number of possible pairings (and hence the number of available polyimides) is huge. Fortunately, this expanse is not uncharted territory; many advanced polyimides have already been developed with properties superior to Kapton. In particular, there are polyimides which are transparent, and also ones with very low, glass-level, CTE values. Our task would be to combine these two properties into an advanced polyimide that can be used in space for a diffractive lens.

Fabricating lens panels:

Assuming that a polymer is found which offers the desired material properties, we must next be able to pattern it into large, thin sheets with accurate diffractive phase profiles on the surface. The basic task of creating thin, large sheets is more advanced for polymers than for inorganics; we can readily purchase continuous polyimide sheets up to about 6 meters in size, and at ½ or 1 mil thicknesses. So a single thin sheet large enough for our 5 meter lens is available. The problem is that, as with inorganics, the thickness uniformity of commercially available sheets is inadequate; current precisions are about 1-2 microns over multi-meter sizes. While not bad, these values are far above the 0.1 μm level which we need for a lens. Nor, of course, do purchased films come with embedded diffractive profiles. So, production of actual polyimide lens sheets will be challenging, and require active development.

As with inorganic sheets, we have a choice between two basic techniques, either applying accurate thickness profiles to existing sheets, or producing them from scratch.

One technique which could be used directly on a commercially purchased, multi-meter, sheet would be direct write via laser ablation (as illustrated earlier). Given a large-throw beam delivery system, coupled with interferometric thickness measurements, we could both mill the sheet to a uniform base thickness and write-in the proper diffractive surface profile. Since the overall mass removal is only microns, the actual laser energy required to write a full-sized lens is quite reasonable. The challenge to this approach is the development of the integrated measurement, feedback, and beam delivery system; a key advantage is that such a machine would provide a largely material-independent patterning tool, which could be used on whatever material (inorganic or organic) we eventually select.

Other patterning options are less developmentally aggressive, although more limiting in both the size and suite of materials for which they can be used. We have, for instance, already developed one method to successfully produce and pattern thin, polymer, diffractive lenses. This involves casting a polymer between a patterned master and a transparent top-mandrel; the polymer is then cured in-place by shining UV light through the top surface of the mold. Unfortunately, this procedure is based upon use of a special low-shrinkage, UV-curable polymer; it is not applicable to the polyimides needed for in-space use. A variant of this approach, the replication-etchback procedure discussed earlier, does however, appear feasible for all materials, including the polyimides.

There is one fabrication technique which does appear capable of directly growing accurately uniform polyimide sheets. This involves physical vapor deposition (PVD) in which the two polyimide monomers are evaporated separately and combine on a substrate to form the desired polyimide sheet. This has been used on-site (by Tony Bernhart's group) to produce films with 1% uniformity, so would be adequate for our use. There are two drawbacks to this approach. The first is simply that a large scale-up would be necessary to go from their current 10 cm sizes up to the 60-100 cm sizes we'd like. Fortunately large PVD systems have been made in the past, and this scale-up should be straightforward. A larger concern is that PVD is conformal, so can not be used to apply the diffractive pattern to the polyimide surface. A subsequent patterning process would

be needed, and if either the etchback or laser ablation techniques had to be developed for this task, they might as well be used for the entire process.

Dealing with in-space property changes:

Because of the fact that even radiation resistant polymers such as Kapton suffer much greater property changes than do robust inorganics such as silica, it will be crucial to not only determine the changes in our selected polyimide, but also to have a means of dealing with those property changes which do occur. The particular polymers which we are investigating were invented for microelectronic applications, and so have never been space-qualified. However, a number of other polyimides (including transparent varieties) have been radiation tested for space applications. These species do share the basic radiation and space resistance of Kapton, and support the expectation that advanced transparent, low-CTE polyimides will be basically suitable for use in space. The type of degradations seen in the most recently published tests, corresponding to 5 years in the GEO environment, showed less than 0.05 increase in solar absorption. They also show modest, $\sim 10\%$, changes in structural stiffness. The changes in transparency are clearly acceptable, but the stiffness variations are more disturbing. The materials survive handily, and have no difficulty handling loads, but their stretching will change by $\sim 10\%$ amounts during the course of a mission. Since the resulting dimensional changes can easily exceed the in-plane optical tolerances necessary to insure coherent focusing of our diffractive lens, this problem must be actively dealt with.

In practice, fortunately, this issue appears quite tractable. The lens is weakly loaded, carrying only slight, $\sim 10^{-4}$, tensioning strains. So an unplanned 10% variation in these strains causes only 10^{-5} -level errors; these cause about 1λ of phase error, about 10 times our desired tolerances. The errors are serious, but not devastatingly so. They can be most readily dealt with by simply adjusting the tensioning loads by the proper 10% values to counteract the stiffness changes and thereby keep actual tensioning strains at their designed-for values.

Summary

At present, it appears quite possible that a polymer material, for instance an advanced polyimide, could be used to build a large diffractive lens. Such materials are likely to be lighter-weight and easier to handle than will be inorganics such as silica or glass; however, actual panel sizes may be limited to the same near-meter values due to limitations in patterning techniques.

Despite the potential advantages to organic thin-films, we decided to concentrate the efforts of this Strategic Initiative upon organic materials. There are three basic reasons for this. The first two are tactical; we don't have the resources to maintain two parallel panel fabrication efforts, and we have much more experience here at LLNL upon forming glass/silica optics than polymer ones. The final reason is more strategic; we would like success in this Strategic Initiative's task of fabricating a 5 meter lens to lead, as smoothly as possible, into an externally-funded, space-flight demo. At that time, however, an

organic lens would still face major uncertainties with regard to material properties in space; for a glass/silica lens, the dangling uncertainty is that of surviving launch. Since the latter issues are easier to address (via careful packaging and vibration-testing) than are basic material-property uncertainties, we decided that the step to a space flight-demonstration would be smoother if our SI built a glass lens than a polymer one.

Assembling panels into a full-size lens

Fabrication of large Fresnel lenses (for instance, the 5 meter size we're building for this project) as single elements is not feasible using current technology and certainly does not scale up to the larger sizes, 25-100 meters, desired for space applications. Also, since any lens must be packaged in a spacecraft for launch, it must be able to be folded to a much smaller size. We solve these two problems by designing the lens to be composed from individual panels of approximately 1m size, that are seamed or hinged together to form a larger optic. The lens would then be packaged by folding at the joints between individual panels.

We must investigate different types of seams and hinges to find an approach that will satisfy the optical tolerances and space conditions, as well as design an instrument which can seam individual panels together with high precision.

Packaging for launch

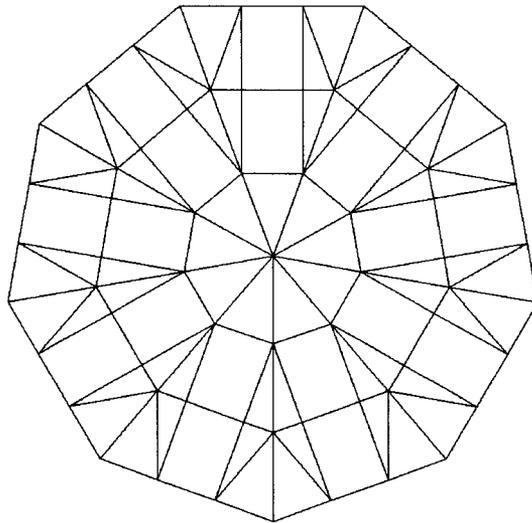
Packaging-for-launch is one of the fundamental challenges for any large aperture space telescope. In principle, the use of thin-film optics, such as our diffractive lens, should alleviate this problem. The actual volume (and mass) of these optics is small, and, since thin films should be flexible, packaging ought to be simple. In practice, this is not the case, since our lens can not simply be wadded-up. For glass films this is obvious; they are brittle and would immediately shatter. But even flexible polymer films must be handled carefully to avoid sharp creases; these lead to permanent distortions that ruin the in-plane precision which the lens must have in order to work correctly. So, whether we use glass or polymer, the fundamental packaging challenge is the same; the large aperture film must be compacted to fit inside a launch shroud, yet if the material is folded too sharply it will fail.

One potential solution is to roll the film instead of fold it; all of the thin-film materials of interest (even glass) are flexible enough to be rolled as long as the roll radius is not too tight. Launch shrouds are typically several meters wide, so a simple 1-D rug-roll, which is then stacked vertically in the launcher, seems attractive. The problem, of course, is length; a rolled-up lens would need a launch shroud as long as the lens's final aperture. This approach would be fine for our current 5 meter size, but does not scale to the 25-100 meter apertures of interest for actual space systems; launch shrouds are not this long. Fundamentally, the problem is that we need 2-D packaging; since rolling is only a 1-D process it cannot solve our problem. The lens must be folded.

There are a vast number of ways that we might fold and package a large aperture film for launch. We have been considering three approaches (fold-then-roll, umbrella-folds, and flat-folds); other methods may well exist.

Fold-then-roll combines folding and rolling. The film is first z-folded in one dimension to form a flat rectangular package, the width has been reduced to 5-10 meters, but the package still has the full length of the lens's aperture. This long package is then rolled-up against the wall of the launch shroud. The attraction of fold-then-roll is its simplicity and the small number of folds required. The concerns with this approach are; how to insure deployment, that the fold-lines don't have radial-like symmetry, that they must be compliant enough to handle the rolling, and that there is no convenient way to accommodate a central gyro-containing core.

Umbrella-type packaging is attractive since it does use radial-like folds. A simple, radial-only, umbrella pattern only shortens packing length two-fold, but the introduction of circumferential folds as well allows arbitrary length reductions. After a film has been umbrella-folded, it can be packaged in two different ways; one resulting in a flat package, and the other in a cylindrical one. The attractions of umbrella-folding are the symmetry of its folds, and its natural deployability; the drawback is that some folds are nested, complicating the design of these joints. A typical, cylindrically packed, umbrella fold pattern is shown.



There are also folding strategies in which a flat circular film is repeatedly folded on top of itself, occupying an ever-smaller area. This results in a cylindrical stack having the cross-sectional area of the launch shroud. The attraction of this approach is that it can be done with only single, un-nested, folds; the drawback is that these stacks are tightly interleaved and probably hard to deploy.

While the specific folding pattern which will be chosen to package large lenses is not yet clear, the essential point is that a lens must be folded. Since a fold is just a 180°

cylindrical roll; any material (even glass) can be folded as long as the radius-of-curvature of the folds is big enough. In practice, of course, we want folds to be as tight as possible. Unfortunately, even very thin glass films are too brittle to fold tightly; fortunately, they don't have to be folded at all. We can avoid the need to fold lens materials (such as glass) by breaking the lens up into separate panels, joined together along the fold lines. The joint has to fold, but the lens material does not.

Multi-panel lenses

So, the need to fold the lens for launch stowage is one factor driving us to a seam-together-separate-panels strategy. Another, equally compelling, reason is that we can't form and pattern large sheets. At some point, currently ~ 1 meter, it is much easier to build and join together separate sheets than it is to learn to build larger sized pieces.

It seems clear then that large aperture lenses must be built by joining together separate panels. Some of these joints will have to be foldable (for packaging reasons), while others (probably the majority) need not be, being present simply because of our sheet-size limitations.

Many types of joints can be considered, and certainly the folding and non-folding joints need not be the same. Whatever joint we do adopt, however, must insure that a joined-together lens has, and holds, an optically acceptable figure.

Placing joints or seams in a lens certainly has the potential to destroy its optical quality. Our concern is not simply the direct loss of optical throughput due to the physical area occupied by the joints; this area small and causes little performance loss. The real concern is that joints can destroy the optical precision of the rest of the lens. They change a simple, uniform, thin-film into a structure with sharp discontinuities in mass, thermal properties and load-paths. While we can't expect such a system to act precisely the same as a virginal film, it must do so accurately enough so that the lens still works.

Obviously, the one key advantage we have, is that our aperture is a flat lens, rather than a curved reflector. Out-of-plane errors, which would be fatal for a reflector, are not a concern for us. Our concern is with in-plane errors; does a seamed-together lens retain the (roughly 10 μm) lateral precision needed to insure high quality focussing?

There are two aspects to this task, first achieving precision in the as-built lens, and then maintaining it in-space. Clearly the individual lens panels, which have been carefully patterned to act as parts of an overall diffractive lens, can only do so if they are put in the right place. To insure this, our assembly process must first accurately place and align the individual panels, and then join them together in a way that maintains their precise location. Once the lens is fabricated, the second, probably harder, task is to keep the panels precise once the lens is deployed and functioning in space. We have two concerns; the first is that the seams should hold the panels in the right places, the second, more subtle requirement is that the joints surrounding each panel must not distort it.

In the first instance, we are concerned with deformation of the joint itself; as the joints move, so will the panels. This problem should not be crippling; not only is the absolute motion of a joint small (because the joint itself is narrow), but as long as this motion is predictable, it is ok. For instance, suppose a joint thermally expands more (or less) than the panel material. The overall effect of this swelling is moderated by the fact that we don't have much joint material (cumulative joint-width is $\sim 1\%$ of the overall lens aperture); furthermore, we can anticipate such joint motion, and build-in allowance for it when placing our panels.

The second concern is likely more worrisome; the presence of joints surrounding a panel can distort it, either structurally or thermally. To prevent structural distortion, we want joints that grip, but not stress, their panels. To prevent thermal distortion, we want joints to be at the same temperature as panels. When either of these conditions are not met, then the joints disturb the perimeter regions of each panel; the overall distortion of the panel depends both on the size of the border strains and how far into the panel this strained region extends.

Panel alignment

The actual alignment of the panels with respect to each other is quite straightforward and similar to the mask alignment problem. We can currently align a mask to a patterned substrate with micron-level accuracy and expect the same accuracy in positioning the panels for seaming. The key to achieving such precision is simply to write specialized alignment marks into the panels and then use them when placing the panels together.

The other issue here, is whether alignment and assembly should be done locally or globally. Can panels be assembled one-at-a-time, or do they have to be fit together all-at-once? In principle, the latter technique should give the best-possible results, but it will also be far harder to actually implement than would panel-by-panel assembly. In practice, we don't need best-possible precision, just good-enough accuracies (few μm levels). If carefully controlled, panel-by-panel assembly should be accurate enough for our needs. Since we consider all-at-once assembly to be impracticably complex given our current knowledge and resources, we will focus upon panel-by-panel assembly.

Given this, there is still the question of whether alignment metrology should be global or local. Again, full global measurements are better in principle, but local measurements are easier to do. To implement global alignment, we need to do precision metrology over the full extent of the assembly table, 5 meters for our current project, and 25-100 meters for future lenses. Building and maintaining such an assembly facility will be a daunting task. Accepting local-only measurements allows a sewing machine approach to assembly; here each panel is placed accurately with respect to its neighbor when being joined, but afterwards the seamed-together portions of the lens are not tightly watched or controlled. This approach obviously permits a much simpler and more affordable assembly facility, and is also well matched to our current mask-alignment experience and capabilities. The question is whether a sewing machine technique will insure good-enough global precision; if it does, it is clearly the way to go. In principle, as long as the panels are

made and joined precisely, then local alignment is sufficient and will insure global precision.

Joints

We considered two different types of joints between the lens panels; one a flexible seam and the other a rigid hinge.

To join two lens panels with a seam, one accurately positions them (close together, but not touching or overlapping), places a narrow tape of some flexible material over the gap, and then bonds it to each of the panels. The seam's precision is determined by the width, stiffness, and thermal rigidity of the narrow tape. The tape's thinness permits the seam to be folded, but this must be done without breaking or creasing the material. What makes this more feasible for a tape than the bulk lens material is simply that tape is a specialized component. It doesn't have to function optically, nor to be available in large, precisely uniform sheets. The main feature which make a good, precisely bendable, tape is for it to remain elastic for as large a strain as possible; this allows tighter folding for a given tape thickness. Several other desirable characteristics are high stiffness, low thermal expansion, and (to better match the panel temperatures) as low-as-possible an α/ϵ ratio. Obviously, the tape must be made from a space-suitable material.

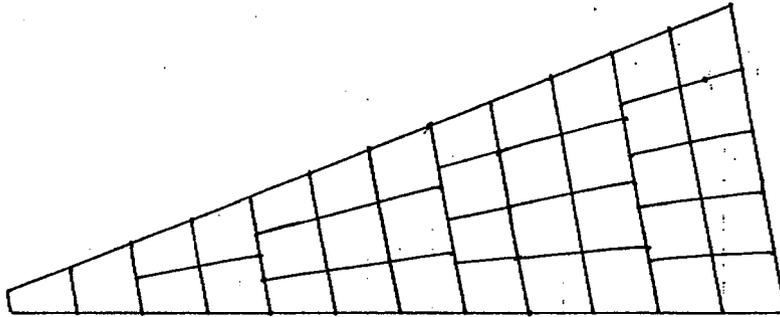
Hinges provide another type of joint option. These are obviously more complex than a simple taped seam, but do offer attractions, permitting tighter folds and eliminating creasing concerns. Our task is aided by the fact that, after deployment, the hinge does not have to bend again; we don't need a high duty-cycle hinge. While one can create a joint using as little as a single hinge, we are most interested in continuous, piano-style, hinges. These spread the load better (both in the hinge and in the panels) and should also produce a tighter, more precise, joint. The key issue here, is whether we can make miniature, yet very precise, hinges. Other issues, such as space-suitability, thermal expansion, and operational temperature, arise and are similar to those for taped-seams. We are considering several hinge options, with materials ranging from glass to metal. The precision of the hinge is crucial, and is set by the pin-to-hole tolerance; we can enhance this by keeping the dimensions small and by taking steps which insure tight-fitting pins.

Regardless of the seaming or hinging mechanism, bonding the seam to the panels will require a space-suitable glue, or some sort of welding technique. We investigated this, and identified possible adhesives. Fortunately, there is a lot of previous experience of using glues in space, so finding a suitable bonding method is relatively straightforward, particularly if (as in our case) we can cover the glue with a UV shield.

In-space precision of a joined-together lens

While there are compelling reasons for our Strategic Initiative's 5 meter lens fabrication to use the multiple-lens-panel-approach, we should also be concerned whether this is appropriate for an actual, space-based lens. Does the presence of seams and panelization preclude the global precision needed for a large aperture lens?

In practice, the panel/joint geometries will likely be complex, governed both by the layouts needed to fold the lens and by the restrictions on producable panel sizes. A typical layout for a 20° section of an umbrella-folded, 25 meter lens, implemented with 1 meter panels, is shown below.



Most of the joints and panels are dictated by our panel-size limitations, not the folding requirements. Of particular concern is the complexity of the nodes, where multiple panels and joints come together. The main factor in preventing severe thermal and structural complications at these, and all other, nodes is simply to physically eliminate them. The joints between panel edges should not be continued all the way up to multiple-intersection nodes; hence, neither panels nor joints must deal with such complex interactions.

Given their long, skinny nature, and the fact that they are not extended all the way to nodes, the behavior of joints (particularly of taped-seams) can be modeled fairly accurately by 1-D methods. We can predict the structural deformations incurred as the tape transfers load between panels, as well as the more important thermal distortions. These arise because the joint and panels reach different temperatures under the effects of solar heating, thermal reradiation, and lateral conduction. Knowledge of these temperatures allows us to predict the thermal expansion both of the joint itself, as well as that of the conduction-heated border region in the panel. More accurate calculations can, of course, be done by finite element modeling. The main physical driver in both the joint and panel distortions is the difference in temperatures between the joint material and that of the panels. We obviously want as low a mismatch as possible. Since the panels are optically transparent, and since polyimides and glass both have high emissivities, the panels will be quite cold (typically under 200°K). Making a similarly cold joint is challenging. Our calculations of glass panels separated by 8 mm and joined together by a white-paint coated, steel tape predict a 50°K temperature difference.

Under these conditions, detailed finite element calculations of seamed panels (by Tom Kuklo) show that significant stresses and distortions develop in the glass. These are caused both by the temperature mismatch as well as by metal-to-glass CTE differences. While stress and distortions are acceptable over the bulk of the panels, they are too high near the corners and edges of the panels. Too great a distortion simply imposes an optical

penalty; since the affected area is small, this can be tolerated. But too high a stress can cause cracks that then shatter the entire panel. Fortunately, our finite element calculations show that several strategies can be employed to reduce these stress and distortion peaks to safe levels. These steps include increasing the bond thickness, segmenting the seams, and properly rounding-off the panel corners. An independent approach is, of course, to adjust the tape properties so as to greatly reduce the thermal and CTE mismatches that cause the problems in the first place. Fortunately, better matched joints should be attainable, using either lower temperature tapes (silica coated metal tapes, transparent polyimide tapes, or glass hinges) or lower expansion ones (Invar or polyimide tapes, or glass hinges).

One final concern is what the cumulative effect of joint discontinuities is. We've discussed the fact that joints don't expand the same as panel material, and that the presence of the joints will also change the distortion within the panels themselves. So on a local, panel-scale, a joined-together panel does not act the same as a virgin panel section would. These distortions can be calculated, and by the methods discussed above, can be minimized; we believe that they can be held below the optical tolerances required for an effective lens. The question then becomes; What are the cumulative effects of the joints?

Our expectation is that as each panel swells (or contracts), the overall effect is simply a similar motion for the full lens, magnified by the number of panels. In this case, the overall displacements in the lens will most likely far exceed our optical tolerances. This, however, need not be fatal, since we already know that uniform expansion (even if very large) can be optically handled very well, simply by shifting the focal position of the Eyepiece image-collecting telescope. So the issue is not, whether or not the lens distorts, but rather to what extent the cumulative effects of multiple panels act as a uniform radial expansion. Here, we expect the use of umbrella-folding with radial-like fold layouts to help, as will keeping the panel sizes as uniform as possible. The basic effects for a geometry such as shown above can be calculated analytically. Joint swelling can be modeled as a position-dependent thermal expansion. The lens therefore has different radial and circumferential expansion coefficients; because the panel sizes change throughout the lens, so will the expansion. Nevertheless, this behavior can be analyzed analytically, and the result is as anticipated. The radial displacements of the lens are described primarily by uniform expansion, superimposed with smaller, panel-sized, local distortions. This simple analytical analysis has been backed up by detailed finite element calculations by Tom Kuklo. So, we believe that the optically troubling distortions caused by joints, are those which occur within the panels. Cumulative effects act simply as an overall uniform expansion of the lens and are dealt with by the Eyepiece vehicle.

Summary

In summary, we decided that it would be possible to assemble a large diffractive lens by joining together individual panels. The techniques required to assemble and align the panels are relatively straightforward. The largest challenges come in the joint mechanisms themselves; not only do the joints have to be precise, but their properties

must be well matched to those of the panels in order to limit interior distortions as well. On a global scale, a joined-together lens acts much the same as a homogenous one; the differences occur on the local, panel-sized, scale and are reduced by well matched joints.

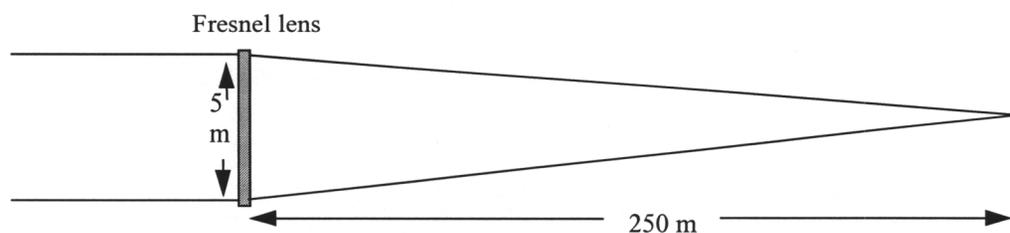
Testing the lens

The most complete and impressive demonstration of our 5 meter lens would involve the optical verification of its high precision focussing capability.

The most challenging programmatic aspect of such a test is the large beam-line required for the test. The combination of a large aperture size and long focal length (250 m for the $f/50$ lens) pose unique challenges for any optical demonstration.

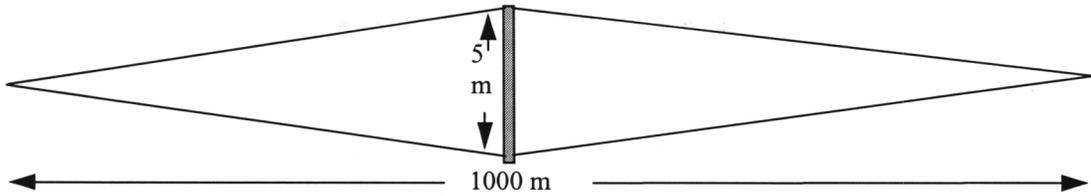
A simple (and relatively short) scenario is to use the Fresnel lens to focus a plane wave; this test requires at least 250 m of beam path. Unfortunately, of course, we did not have such a large aperture plane wave at our disposal.

A. Collimated 5-m source, single Fresnel lens

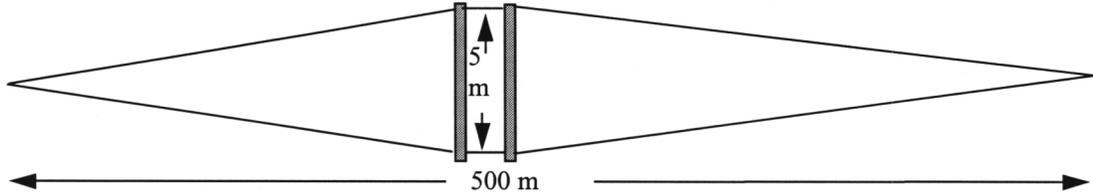


Since we didn't have access to a 5 meter plane wave, then we made plans to use a point source. While point sources are much more available than large plane waves, using one entails a considerably longer beam path. If our test is to image a point source through a single 5-m Fresnel lens, then we require a 1 km long beam path. If two 5-m Fresnel lenses are used, back-to-back, the minimum required beam path reduces to 500 m. This approach requires us to tradeoff the cost of the extra beam path against that of building an additional Fresnel lens.

B. Point source, single Fresnel lens



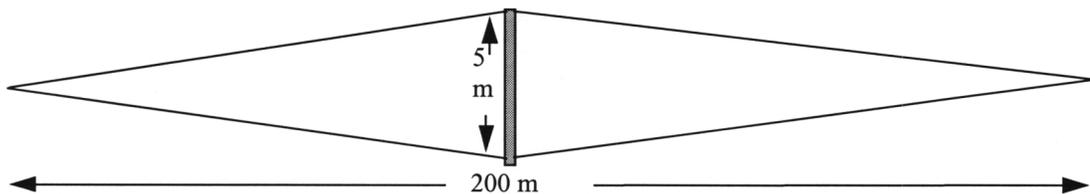
C. Point source, two Fresnel lenses



There are alternative testing routes, exploiting special characteristics of Fresnel lenses, which enable us to greatly shorten the beam path requirements.

One approach is to test the lens at an IR wavelength, rather than its visible-wavelength design value. Since the focal length of a Fresnel lens scales inversely with wavelength, we can achieve large reductions in the beam path. For instance, a $2.5\ \mu\text{m}$ test (instead of $0.5\ \mu\text{m}$) would reduce the single lens, point source, beam path to 200 m. There are a couple of drawbacks, however, to this technique. The first problem is that our lens is a phase-plate, rather than zone-plate, design. Because of this, the lens applies only 0.2 waves of phase shift to the $2.5\ \mu\text{m}$ light. The lens does still deliver a tight, coherent, focus, but at fairly low efficiency; most of the light passes through the lens without being focused at all. In practice, this efficiency loss is not crippling, since a laser point-source will provide plenty of signal strength for testing purposes. A more significant issue is that we are building a visible-wavelength lens; optically testing it only in the IR will dilute the impact of our lens demonstration.

D. Point source, single Fresnel lens, 5th harmonic



Another approach allows similar focal length reductions, while not deserting the visible-band design wavelength. Here we exploit the fact that Fresnel lenses feature a multiplicity of higher order foci. The first order focal length of our lens is 250 m, but it also has a suite of foci at harmonics of this value, i.e., at 125 m, 83 m, 62 m, 50 m, etc. As was the case above, the main cost to using these alternative foci is that they are less

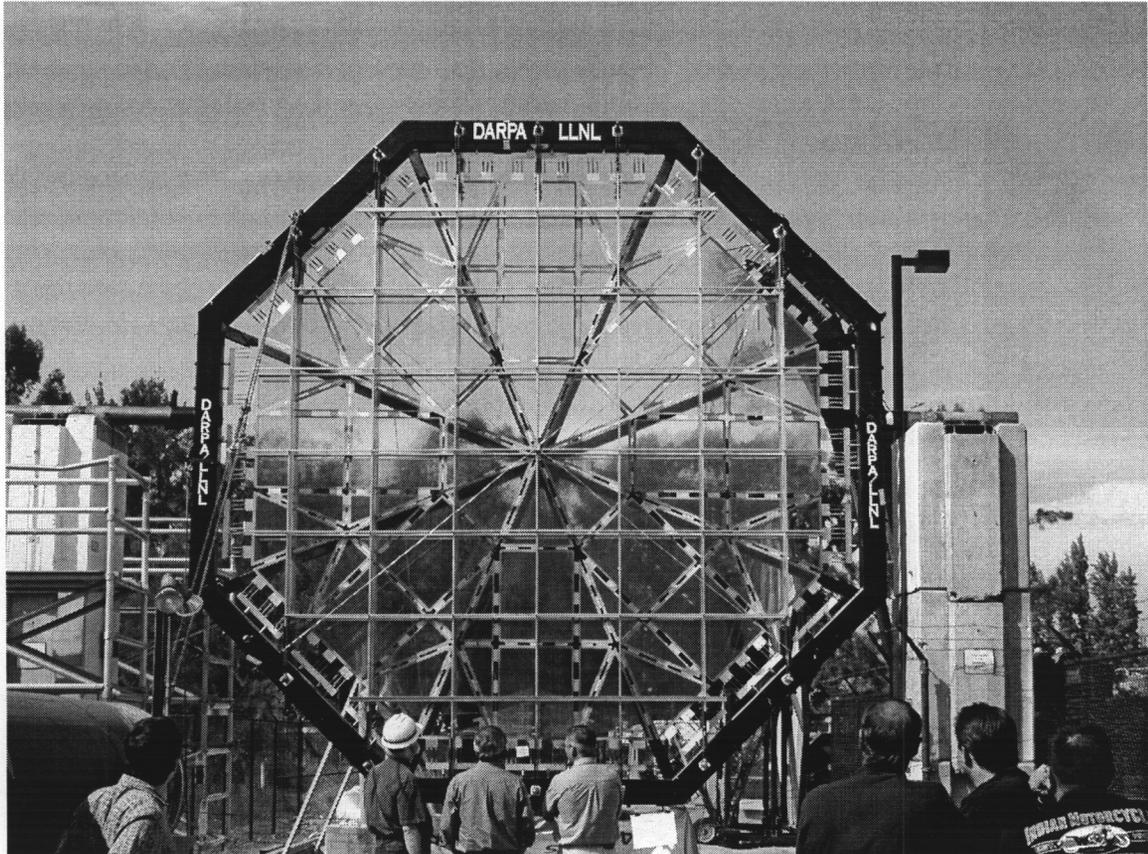
efficiently populated. However, since we won't be signal starved, this is not a serious drawback. As an example of this approach, use of the 5th order focus enables us to shorten the beam path of a single lens, point source experiment to 200 m.

We could also shorten the beam path by use of flat turning mirrors. Unfortunately, significant folding of the beam requires large aperture flat mirrors. For example, in case B, midway foldings at 250 m in front and behind the lens would require a pair of 2.5 m aperture folding mirrors. While this approach would shorten the beam path to 500 m, two such mirrors will be very expensive. On the other hand, use of affordable 1-m class folding mirrors would only shorten our beam path to 800 m.

It is clear that testing our 5 meter Fresnel lens is a significant undertaking. The effort will be dominated by both the need to mechanically support the large lens, and the need to use a very large beam-path, featuring both a large, 5 meter, aperture, as well as a long, up to 1000 meter, length. In contrast to this beam-line, the other aspects of an optical test (particularly if only the lens itself is involved) are quite straightforward.

What Was Accomplished

The primary goal of this Strategic Initiative was to build a 5 meter diffractive lens. This was successfully achieved.



Above we show the 5 meter diffractive lens fielded in September 2002 at the conclusion of this Strategic Initiative.

The funding necessary to assemble, field and optically test this lens outdoors was provided in FY2002 by DARPA. However, the ability to do so, was the culmination of the developmental research and optical patterning carried out during the course of this Strategic Initiative.

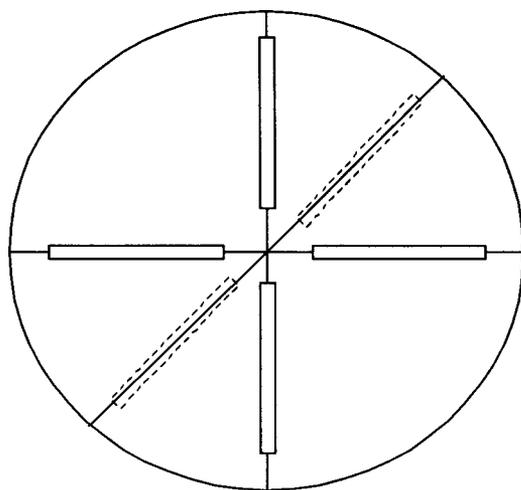
The research and development conducted during this program falls into 3 basic categories: development of the basic, segmented-lens approach, development of diffractive lens panels, and development of precision lens-assembly techniques.

Demonstration of a Foldable Segmented Lens

We decided to adopt a segmented panel approach to building the 5 meter diffractive lens. Since this was a new, untried method of building diffractive optics, we felt the need to build a sub-scale lens using this approach early in the program. We intended this as an early precursor to the 5 m lens, serving both as a learning tool and as a validation of the segmented lens approach.

Our demonstration lens was designed to have a 75 cm optical aperture (80 cm physical diameter) and a focal length, at 532 nm, of 39 meters. This represents an f-number of 52, and requires an outermost Fresnel zone-width of 55 μm . From an optical standpoint, this lens is more challenging than the ones we've fielded in the past, the best of which had a larger feature size (120 μm) and a smaller aperture (50 cm). The greatest challenge to this lens of course, was not the more difficult optical patterning, but instead was the use of a segmented, rather than monolithic, design.

The purpose of the 75 cm lens was to demonstrate a segmented diffractive lens, showing that we could successfully assemble, align, and join-together, a collection of individual lens segments into an overall, larger-aperture, lens. We selected a foldable, 6 segment layout as shown below.



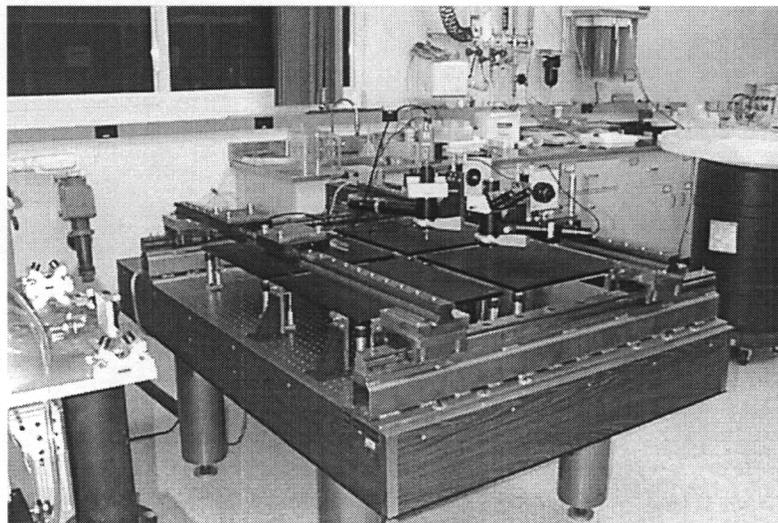
The reason for this design, rather than a simpler 4 quadrant layout, was so that only single-layer, un-nested, folds would be necessary when folding the lens. This configuration results in 6 foldable seams; 4 of which have one polarity, and 2 the opposite polarity.

The purpose for this demonstration lens was to demonstrate that we could successfully field a segmented lens; our emphasis was upon the assembly of lens panels into an overall lens, not the production of the panels themselves. Accordingly, we used an in-hand, conservative, method of producing the panels.

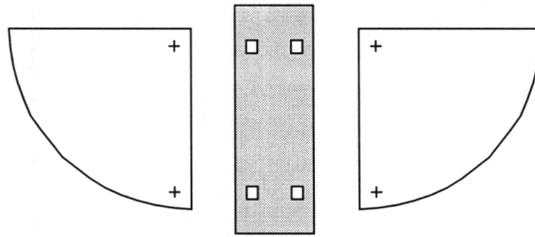
The raw panels for this lens were produced for us commercially from silica using a traditional grind and polish technique. In order to serve as lens segments, each panel must be optically flat, with ripples and/or/wedges held to $\sim \lambda/5$ (100 nm). Furthermore, the absolute thickness of each panel must be the same, to the same level of accuracy. To achieve these tolerances, the lens panels were produced together, by polishing an 80 cm circular disk, and then separating it into 4 quadrant-segments. We considered this approach much better suited to assuring global uniformity than the alternative of separately polishing each segment.

After receiving the lens segments, we used a two-mask lithography process to apply a 4-level diffractive lens profile to their surfaces. Two of the quadrants were then cut in half, leaving us with six finished lens panels.

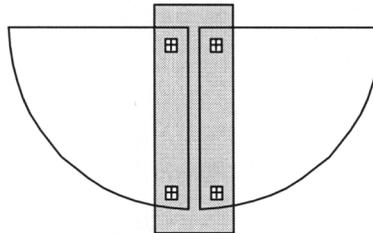
In order to assemble our segmented lens, we had to position, accurately align, and then seam-together the six lens panels. This work was performed on a specially designed and built assembly table, shown below.



The assembly surface was split into 4 support quadrants, 2 of which were fixed and 2 allowed precise, controllable, 2-D motion. The lens panels were placed on the table, aligned, and joined together, one seam at a time. We used a fiducial-to-fiducial alignment technique illustrated below.



Each panel contains alignment fiducials. A mask contains corresponding sets for each panel edge



Each panel is aligned to a set of fiducials on the central mask, which does not move. The panels are now separated by a precise distance.

To align two panels together at a precise separation, we align fiducials on each with respect to those on a central alignment mask. By patterning fiducials on the alignment mask at known separations, once each panel is aligned with respect to the alignment mask, they are therefore registered properly to each other. To implement this alignment technique on the assembly station, we place the two lens panels to be joined together, on the 2 movable support quadrants. The alignment mask is then placed in the gap between these two quadrants, at a level very slightly below that of the lens panels. One-at-a-time, we then shift each lens panel (by moving the support quadrant it's mounted on) over the edge of the alignment mask, so that its border fiducials precisely line-up with those on the alignment mask. After both panels have been aligned to the alignment mask, they are precisely registered to each other and are ready to be seamed together.

The seaming procedure was performed in a two stage operation. The initial step was to rigidly join the panels together with wooden support tabs. These provided a temporary (and removable) scaffolding, rigidly holding the panels together during both the attachment of the permanent, flexible, thin-metal seams, and during the handling steps needed to reposition the lens during its seam-by-seam assembly and its final mounting in a support ring. Once the temporary wooden tabs were securely glued in place across the seam, we attached the permanent thin-metal films. These were 1 inch wide, 1.5 mil Ti-alloy, films; $\frac{1}{4}$ inch of each was glued to the proper lens panel, leaving a $\frac{1}{2}$ inch gap between panels for folding.

After each panel-to-panel joint was aligned and seamed, the partially assembled lens was rotated on the assembly table to position the next joint-to-be-made into position straddling the two movable support quadrants. The wooden scaffolding tabs were crucial

here in preserving alignment of already-completed seams during this handling, as the foldable thin-metal films had very little compressive rigidity.

After all six lens panels were joined together, the lens assembly was complete. The lens was then mounted in a rigid metal ring and secured by 16 pressure clamps. At this point the lens was ready for optical testing.

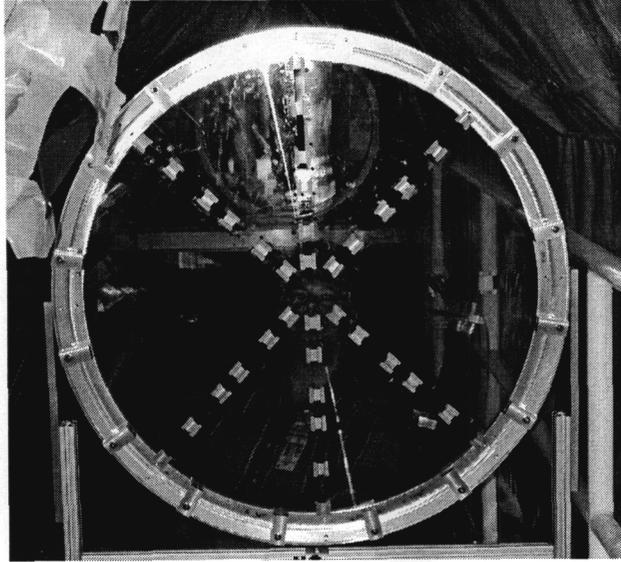
Optical testing of the lens was performed in our NOVA optical-beam-line by Mike Rushford. Our segmented diffractive lens was designed for infinite-conjugate focusing of a plane-wave. We used a 75 cm aspherical lens (previously used by the Diffractive Optics Group for holographic-exposure patterning of large diffraction gratings) to provide this large, 532 nm plane-wave. Given this plane-wave and beam-path, the segmented diffractive should, if successful, deliver a tight (~ 60 mm) spot at a 39 meter focal distance.

We used three types of measurements to determine the quality of the lens. The most basic was to image the focused spot with a CCD camera, determining its PSF. This provides an all-up, operationally-relevant measure of the lens's ability to focus light. We also used the lens produce images of standard Air Force test patterns; showing resolution capability. Finally, we used point diffraction interferometry to deliver more detailed information about the quality of the lens. This tool generated a 2-D optical fringe pattern covering the full aperture of the lens, basically measuring the wavefront-profile produced by the lens. Point diffraction interferometry of a perfect lens will generate a series of straight fringe lines. Deviations from this within a lens panel indicate errors in the lens panels, while jogs in the fringes as they cross the inter-panel joints denote that the panels are misaligned.

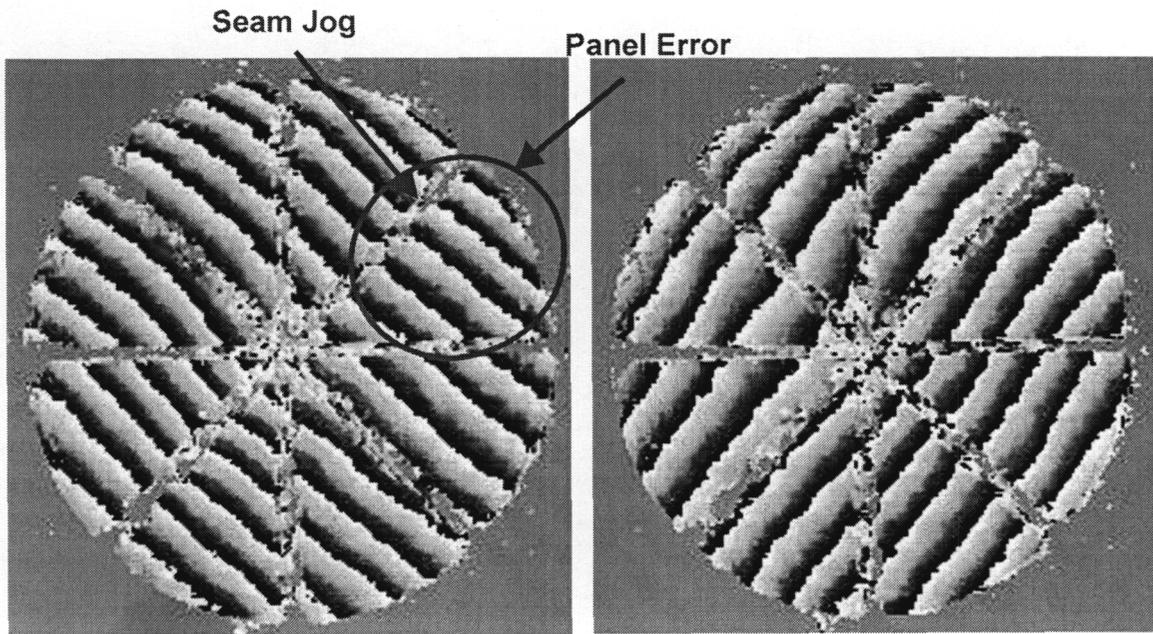
We measured the optical performance of the lens several times, corresponding to different lens configurations.

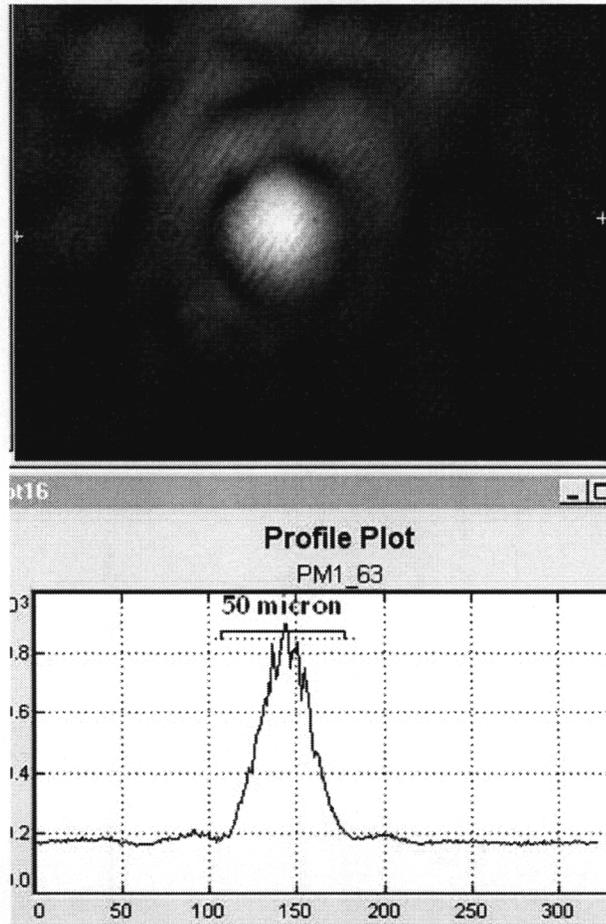
The first set of measurements was taken to characterize the quality of the as-built lens, before it was folded. In this set, we examined the lens twice, first while the wooden scaffolding-tabs were still in place, and then later after they were removed. The measurements with the wood-tabs characterize the lens as-it-was-built, showing the quality of the diffractive lens panels themselves, and our panel-to-panel alignment capability. After measuring this as-built configuration, we removed the wooden-tabs and retested the lens. Here, the panels are supported only by the flexible thin-metal films across each seam, so differences in lens performance would indicate that these joints did not maintain precision when supporting the lens panels.

Below we show a photograph of the lens in its mount during this configuration; the wooden scaffolding-tabs have been removed, and panels are joined together only by the flexible thin-metal films.



The lens performed the same in both of these measurements. There was no indication of changes in the inter-panel alignment when we removed the wooden-tans and gave full joint responsibility to the thin-metal films. Below, we show the optical results, both the PSF focal-spot and the interferometric fringe patterns.



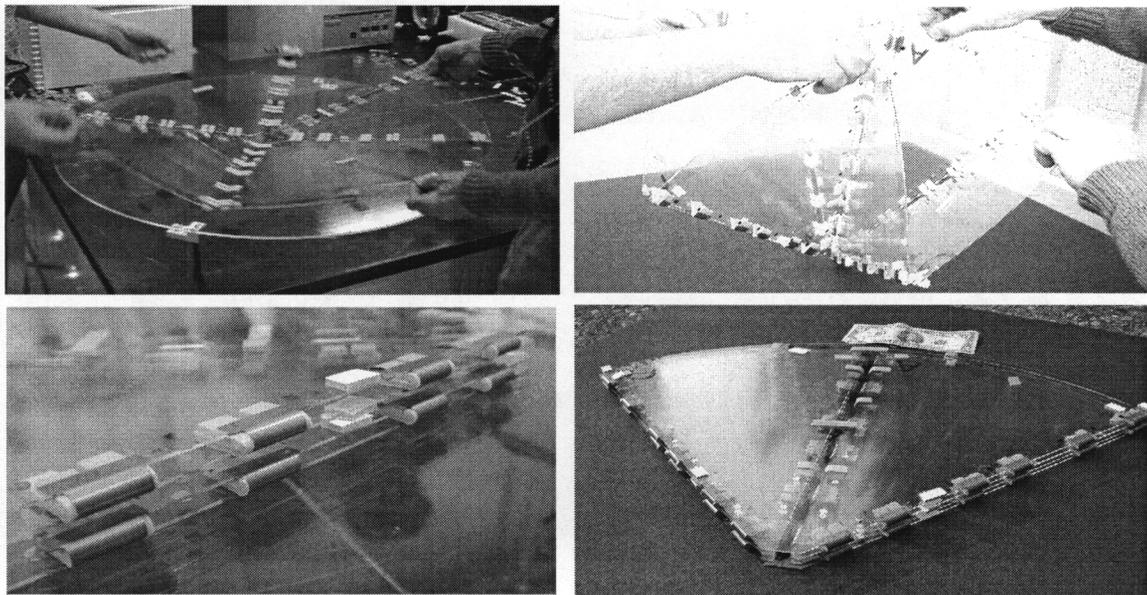


The PSF measurement shows that this segmented lens does optically perform as a full-aperture lens, delivering a tight, diffraction-limited, focal spot. The lens patterning, panel alignment, and inter-panel seaming are accurate enough to deliver a high-quality optical lens.

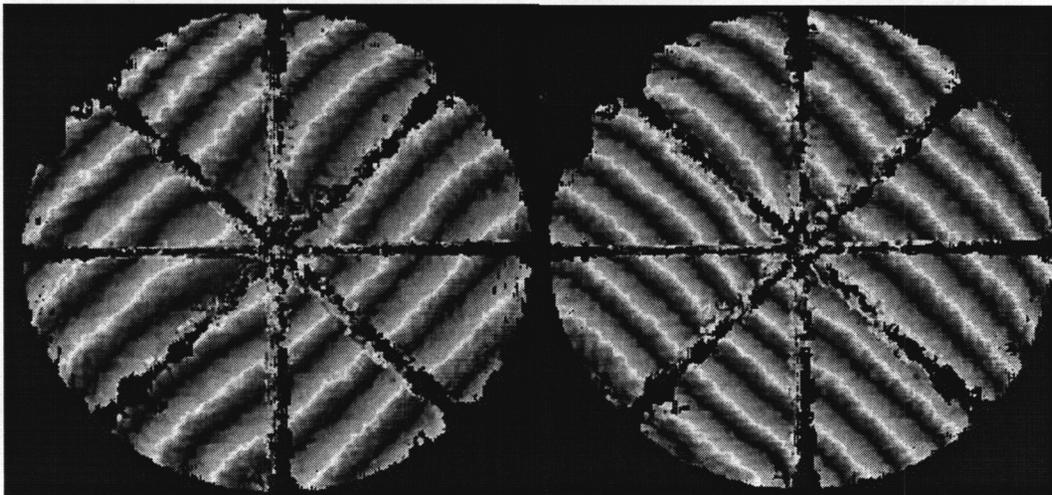
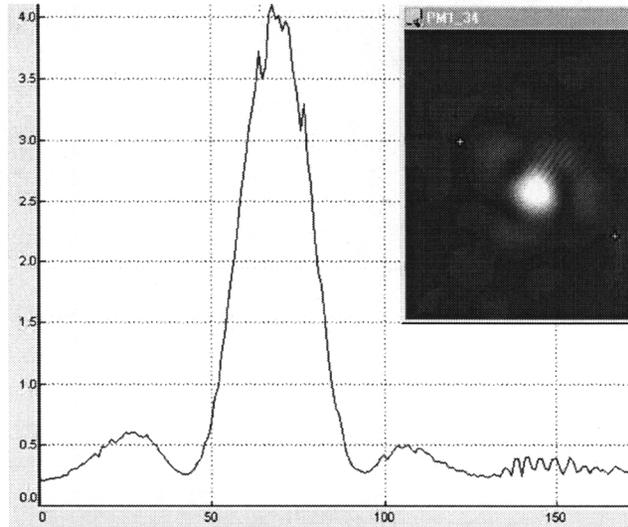
The point diffraction interferometer provides us with more detailed information on both the lens panels and the joints between them. Examination of the fringe patterns reveals that most of the panel area and inter-panel seams have excellent quality. The data does identify, however, two problems, pointed-out on the fringe picture; a portion of one lens panel is distorted, as is the seam bordering it. Our belief is that both of these errors (which from the fringe-jog across the seam, peak at $\sim 0.4\lambda$) are caused by a processing error during the fabrication of the right-most lens quadrant. This error appears to have been introduced by a thermal gradient across the panel during the mask-exposure step during its lithographic patterning. Subsequent tests have shown that this thermal gradient was probably caused by differential cool-down during application of a fluid-spacer used during the exposure step. As a result, the lens pattern was nonuniformly applied to this one lens panel, resulting in both the intra-panel and inter-panel wavefront errors shown in the fringe plots. While these ($\sim 0.4\lambda$) fabrication errors exist in the lens, they effect such a limited area of the aperture that the lens's overall focussing, demonstrated by the PSF, is excellent.

After measuring the as-built performance of the lens, and demonstrating that this segmented lens performs similarly to a monolithic lens, we then concentrated on folding the lens. Demonstrating that a segmented lens can be successfully folded and unfolded is essential for the eventual very-large aperture space-based systems. Since the folding/unfolding process also requires that the lens be removed from and then reinserted into its lens-mount (without the support of the wooden scaffolding-tabs) we performed two lens tests. In the first phase we simply removed and remounted the lens without folding it. This was done so that any errors after the entire fold/unfold process could be isolated to either mounting or folding. The optical tests of the lens after remounting showed that this process could be performed successfully, the lens still worked well after this procedure.

The actual folding/unfolding operation did involve risk. The concern was not so much breaking either the glass panels or the thin-metal seams, but more that we might distort the metal films, hence ruining the lens's optical alignment. The metal films and joint-spacing were designed so that the folded-over seams would remain elastic and not undergo any plastic deformation; so we expected the pure act of folding a single seam to be safe. Our worry was the non-ideal loading implicit in the folding of a 6-seam, 2-D, object. Specifically, the problem was that the thin-metal foils buckle when loaded in compression. Such buckling can impose small-radii bends in the films, thereby plastically deforming it. Practice operations with a surrogate lens did experience these problems. We dealt with this by using foam strips as temporary seam supports during folding and unfolding. These foam strips were applied to the inside-fold side of each thin-metal foil before the lens was folded. During the folding/unfolding process they served two functions; they resisted any inadvertent compressive loading (preventing buckling of the films) and they also provided gradual-resistance to the folding, insuring that the folded seams did settle into the desired max-bend-radius configuration. The photographs below detail the folding/unfolding process (performed by Mike Rushford and Ian Thomas).



After the lens was folded-up and then unfolded again, we reinserted it into the optical mount and removed the foam supporting strips. The lens was then remeasured in the NOVA test facility. Below we show the lens's post-folding PSF and fringe patterns.



These test results clearly show that the lens performs very well, the same after folding/unfolding as it did beforehand. The act of folding did no damage to the lens or the precision of its seams. The distortion evident in the post-folding fringe patterns is that seen before, in the pre-folding measurements. These errors were caused during panel fabrication, not introduced by lens assembly or by folding or unfolding of the lens.

We consider the 75 cm segmented lens demonstration to be a very successful validation of the multi-panel approach we planned to use to produce the 5 meter lens. This effort has shown that multi-segment diffractive lenses can be built, assembled, and folded; delivering an optical performance near that of an ideal, diffraction-limited, monolithic lens. In addition to demonstrating a successful lens, this sub-scale fabrication

exercise provided us with valuable experience; the lessons learned were later drawn upon in the fabrication of our 5 meter lens.

Thin, Uniform Sheets for Diffractive Lens Panels

Our earliest and most far-reaching decision here, was to decide upon using glass rather than plastic for our lens panels. The primary reason for this involved our exit strategy; after this Strategic Initiative shows that we can build a 5 meter lens, we would like to take part in a space-based, flight demonstration. In order to first win approval and then carry-out such a mission, it is essential to minimize programmatic risk; since we believe that risks due to uncertainty in polymer space-properties are greater than those due to glass brittleness, we preferred a glass lens. In addition to this strategic calculation, we faced the practical consideration that we at LLNL have much more experience in glass-based optical fabrication than we do with polymer films. Accordingly, we chose to build a glass lens, not a polymer one.

Organic materials:

While we have chosen not to use polymer films for our Strategic Initiative's 5 meter lens, they remain an attractive option for the full-size (25-100 m) space-based lenses. For a full-size telescope, their superior thinness (i.e., low mass) and launchability will likely be more relevant than our immediate concerns regarding lack of processing and material-property knowledge. In fact, our concern over space-property values is less a technical one that the properties are bad or intolerable, than it is a programmatic one that they aren't currently known.

Accordingly, before we decided to concentrate fully on glass panels, we did carry out developmental research on polymer films. Our emphasis was upon finding a polymer which would not only survive the space environment, but would be both transparent and thermally rigid. After extensive literature review, we focused our research among polyimides, because of their general thermal and radiation stability, their extensive history of use in space, and due to the great design flexibility offered by changing the specific polyimide monomers.

Kapton was the original polyimide to be developed, and is fairly opaque, with an orange-brown color. The absorption in polyimides is spectrally dependent, featuring a characteristic threshold wavelength; on the UV side of this wavelength the material is strongly absorbing, while on the IR side it is quite transparent. In Kapton, the transition occurs at about 550 nm, so blue light is absorbed, while red light passes. So, even Kapton could actually be used for our lens, if we were willing to give up the blue half of the visible spectrum. In the mid-1980s, researchers at NASA discovered that the introduction of CF₃ groups allowed the threshold wavelength to be shifted downward towards the UV, allowing blue light to also transmit through films, hence forming the first colorless polyimides. These polyimides have since been licensed to industry as CP1 and CP2; they have been space tested and are available in 6 meter sizes. In the last decade many other colorless polyimides have been created, with some having threshold wavelengths as low as 280 nm. So, at present, there are a number of transparent polyimides available for consideration as lens materials.

The development of low-CTE polyimides began about a decade ago in Japan, driven by the desire to expansion-match polymer layers to semiconductors. Researchers found that by using diamines and dianhydrides with rigid backbones they could create polyimides with very low, few- $10^{-6}/^{\circ}\text{C}$, values of CTE. Fortunately, such monomers are quite common. For instance, Kapton's high CTE stems from the flexibility of its diamine, ODA; its dianhydride, PMDA, is quite rigid. By replacing Kapton's flexible diamine with a more rigid one, DPTP, one can form a polyimide with CTE of only $2 \times 10^{-6}/^{\circ}\text{C}$ instead of Kapton's value of 2×10^{-5} . In this manner, Japanese researchers have developed (and Hitachi sells) a number of low CTE polyimides.

It would appear then, that the key to simultaneously achieving high transparency and low CTE in a polyimide is to use rigid monomers containing CF_3 groups. Some such polyimides have been developed and studied in Japan. The primary one combines a CF_3 bearing diamine, TFDB, and the common dianhydride, PMDA. While TFDB-PMDA is not completely transparent, it is much more so than Kapton, having shifted the transmission threshold down to about 430 nm. The remarkable feature of TFDB-PMDA is that its CTE is reported to be small and negative, -5×10^{-6} . This raises the interesting possibility of combining TFDB-PMDA with a positive CTE polyimide to form a very low, near-zero, expansion material. Another avenue of interest would be to find a similarly rigid, but lower absorption, dianhydride than PMDA. This could be combined with TFDB either to form a complete alternative to TFDB-PMDA or to blend with it to form a zero-CTE hybrid.

Our activity was aimed at obtaining the basic TFDB-PMDA polyimide, and characterizing its properties. The diamine TFDB is not commercially available, but Ted Baumann in C&MS successfully synthesized it for us. Ian Thomas then combined the TFDB with purified PMDA to form the TFDB-PMDA polyimide. He has spun-cast thin films of this material for characterization purposes. We sent several samples of this polyimide to an outside contractor, PMIC, to have high precision, non-contact, dimensional measurements done to determine the dimensional changes due to demoisturization, outgassing, and temperature (from 100°K to 350°K). These measurements showed that TFDB-PMDA does indeed have a negative, low-magnitude, CTE of about -3×10^{-6} per $^{\circ}\text{C}$. They also showed very little ($\sim 3.5 \times 10^{-5}$ dimensional change due to out-gas/vacuum emplacement. Stress-strain measurements by Steve Letts's group at LLNL showed that TFDB-PMDA is quite stiff, with a Young's Modulus of ~ 840 ksi. These films are not completely transparent, having a slight yellowish tinge; but the 50% transparency point (for 60 μm films) is at about 420 nm; this is a distinct improvement over basic Kapton, and adequate for most in-space optical applications.

This preliminary work has confirmed that TFDB-PMDA is a good starting point for Eyeglass polymer development. Future research will produce even better polyimides, having greater blue-to-UV transparency, and zero CTE values. The most crucial research needed however, is to measure the effect of space radiation (primarily UV and electrons) on these polyimides and their material properties.

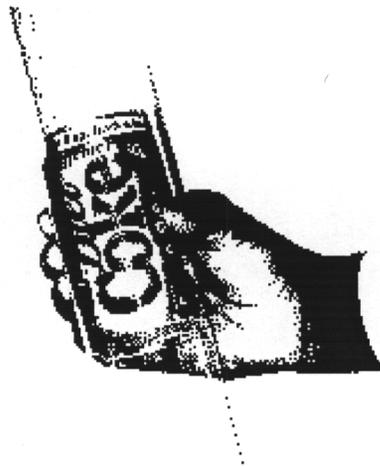
Inorganic materials:

Our main metric in choosing an inorganic material was not so much material properties (either silica or glass should be fine) but instead the ability to acquire large, thin, uniform sheets of it.

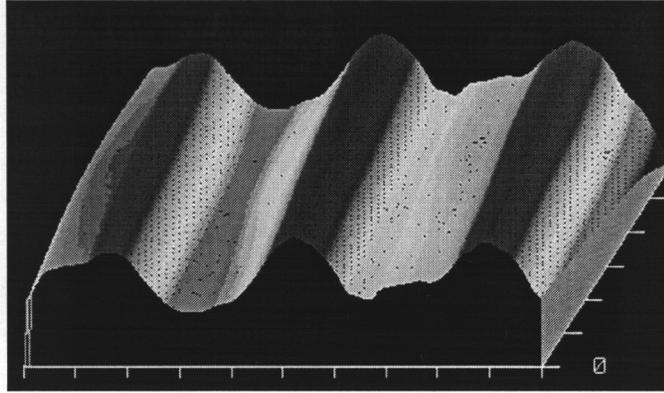
Traditional, grind-and-polish techniques certainly can be used to produce smooth, meter-class, sheets of glass or silica. However, these methods are expensive, and become increasingly more risky for thinner and thinner sheets. This method is useful for small systems (such as the 75 cm test-bed lens discussed later), but is not attractive for larger lenses such as our 5 meter one, or for eventual very large space-based lenses. As a result, we have pursued alternative methods to fabricate thin glass or silica sheets.

Early in FY00 we examined a number of approaches; hot-pressing of thin glass, ion-etchback of glass and silica, and CVD deposition of silica films. However, because of resource limitations, we soon chose to focus on a single thin-film technology; the commercially available glass micro-sheets.

We bought a large number of these micro-sheets (in 3 different thicknesses) from Corning, and used them for several purposes. The first, was simply to gain experience at handling large thin glass sheets. We have acquired a practical understanding of how we can and can't handle lens panels. Below, we show a picture illustrating the flexibility of a 75 μm glass sheet.



Our second use of these glass micro-sheets was to determine their properties. We've measured both their transmission and uniformity. Examination of thickness profiles (for a large number of sheets) has shown us 3 things; that most variations are across the pull-direction, not along it, that only the inner portion (about 50% of the width) of each sheet has reasonable uniformity, and finally that the variations that do exist in the inner part of the sheet are basically sinusoidal with few- λ magnitude.



Unfortunately, the existing glass micro-sheets are both too small and too non-uniform to simply be used as-is; they must be improved in both respects. Wider sheets will be essential, particularly given the fact that the edge regions are unusable. The good news about the thickness ripples in the inner portion of the sheets is that they are basically just that, 1-D sinusoidal ripples. We believe that most of this nonuniformity can be removed, either by us or by the manufacturers.

The best solution here, of course, would be to simply have the manufacturer (be it Corning, Schott, or Pilkington) sell us improved sheets. Unfortunately, their current sheets are about the best they have; improvements will take development and money. The quandary is that they have existing customers (which is the reason the sheets are available, and at such a small price) who are perfectly satisfied with the sheets as they are; we don't represent a large enough market for them to invest much effort in improving the sheets.

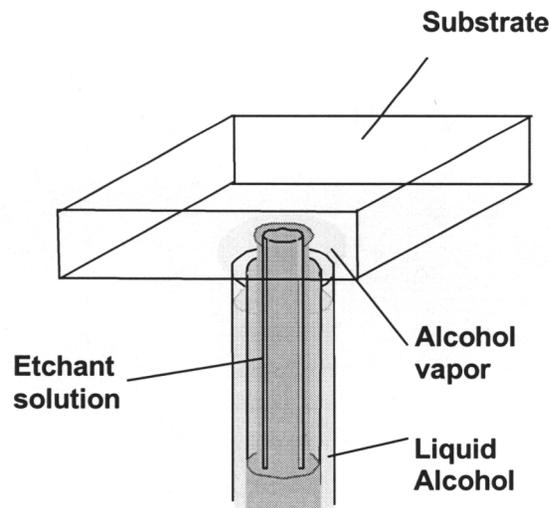
Accordingly, we investigated a number of techniques for smoothing the as-bought glass micro-sheets. The most promising approach was spatially controlled etching. This has two major advantages; low force and high-rate. Unlike physical polishing techniques, wet etching does not apply force to the glass sheets; hence the thinness and relative fragility of the sheets is not a concern. The other virtue is that wet-etching is potentially a very fast process, as $\mu\text{m}/\text{min}$ etch rates can be achieved.

The challenge, of course, is the need to spatially control the etch process. Wet-etching is traditionally used as a bulk process, to uniformly remove material. To be of use here, in smoothing glass sheets, it has to be controllably nonuniform, removing more material in some places than others, delivering a uniform thickness sheet at the end of the etching process.

There are several possible approaches to achieving spatially controlled etching. These can be split into direct and indirect methods. The direct approach is to simply apply more etchant to some places than to others. The alternative, indirect approach, is to apply the etchant uniformly, but to spatially control the rate at which it acts. During the course of this Strategic Initiative, we investigated both approaches.

Our first investigations involved the direct approach to spatially controlled etching, in which etchant is confined to a limited region of the surface. This etching region is then moved over the surface of the glass sheet, spending more dwell time in areas where we want more thickness removed, and less time in areas where less etching is required. Key to success with this technique is that when we leave an area of glass, we need to remove the etchant from it; we do not want to leave behind a thin layer of etchant which would continue to etch the surface.

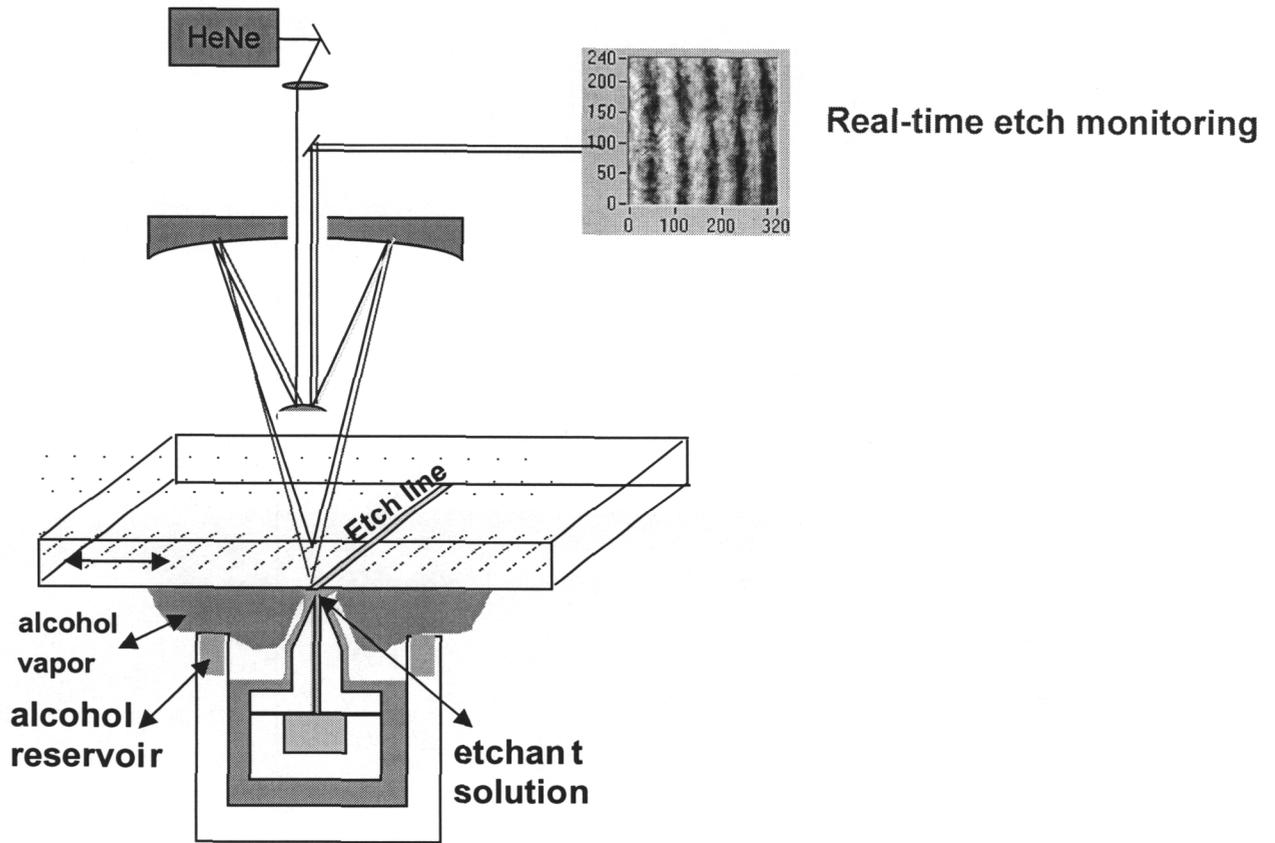
Jerry Britten and Mike Rushford developed a etchant confinement method to solve this dewetting challenge. Their method uses the Marangoni effect to actively remove etchant from the surface. The Marangoni effect uses forces produced by gradients in the surface tension to “pull” liquid off the surface. These gradients are induced by evaporating alcohol into the outer perimeter of the liquid etchant, and using the fact that the addition of alcohol greatly reduces the surface tension of the (water rich) etchant. As a result, we can introduce a small region of etchant via a small-tool etcher, and then actively remove it from the perimeter of the spot, resulting in a tightly confined region of etchant.



. Britten and Rushford’s machine makes use of numerically controlled, chemically-assisted, small-tool etching. Once the panel’s thickness profile has been mapped, we know where and how much material must be removed. A small-tool etchant deliverer is mounted on a controllable translation stage, and directed to remove material in the proper places. As long as the lateral size-scale of the roughness is reasonable, then this tool does not require extreme positioning accuracy. Since the depths of material which must be removed are relatively small, then such a micro-etcher can achieve reasonable smoothing times.

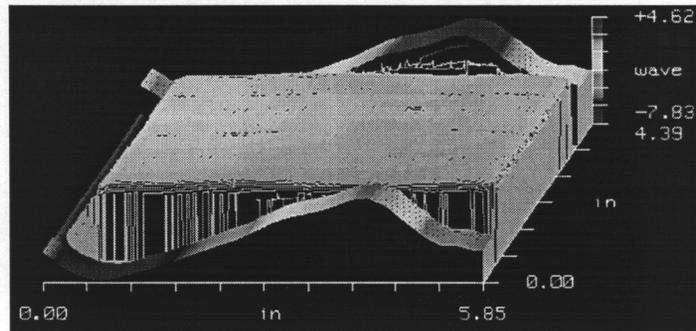
As noted above, the ripples in our glass micro-sheets are predominantly 1-D. This spatial variation was exploited by the development of a 1-D micro-etcher. This tool featured a 1-D etcher in which dilute HF acid was delivered through a slot, formed an etchant-line on the glass sheet, and was then actively removed (dewetted) via the

Marangoni effect. The location of this etch-head over the glass was controllable with a 1-D translation-stage, and was moved about so that local dwell-time could be used to selectively etch more-or-less glass at given locations. This etch process was actively controlled using feedback from a real-time thickness-measuring interferometer developed by Mike Rushford.



This approach has proved to be quite effective for removing thickness ripples from our as-bought glass micro-sheets. Below, we show a flat plateau etched into a rippled glass micro-sheet. The original (largely 1-D) ripples were smoothed down from their initial, 12 μm , depths to a flat plateau having ~ 100 nm residual variation in about 10 hours of processing.

Etched 4" x 5" Plateau



After developing this line etcher at small scale, we scaled it up to a larger machine, capable of controllably etching the 80 cm wide glass sheets used for our 5 meter lens.

In addition to the physical scaleup of the etching machine and interferometer, we also had to switch to a different glass composition (initial work used Corning 211) in order to purchase wide enough thin-sheets of glass. The new glass was Corning 1737, which is a borosilicate glass available in 840 mm x 1140 mm sheets of 700 μm thickness. The only complication to using this particular glass is that it contains sizable concentrations of barium and calcium oxides, which cause precipitates when etched with simple HF acid solutions. These precipitates deposit on the glass and prevent the achievement of smooth etched surfaces. In order to smoothly etch this glass, Ian Thomas developed a dual-component etchant, containing hydrochloric acid as well as hydrofluoric acid.

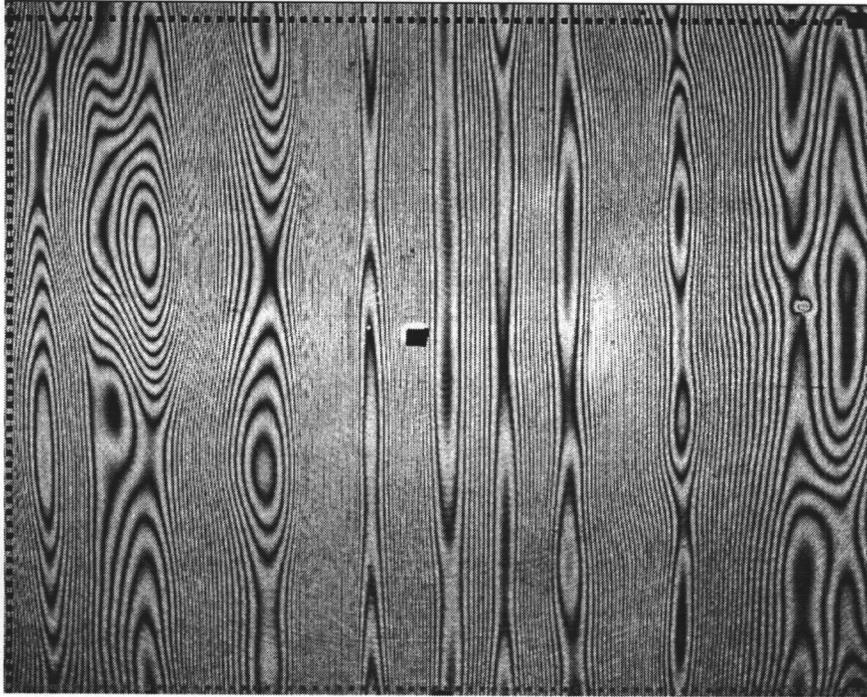
Below we show a photograph of the scaled-up 1-D etcher.



This machine operates by moving the 80 cm wide glass-sheet back and forth over a fixed 1-D etch-head. The rate of motion is controlled in concert with an interferometer so that dwell-time is matched to the required thickness-removal rate. The glass is gripped and moved from its edge, and supported over its area by a low-friction air cushion. The

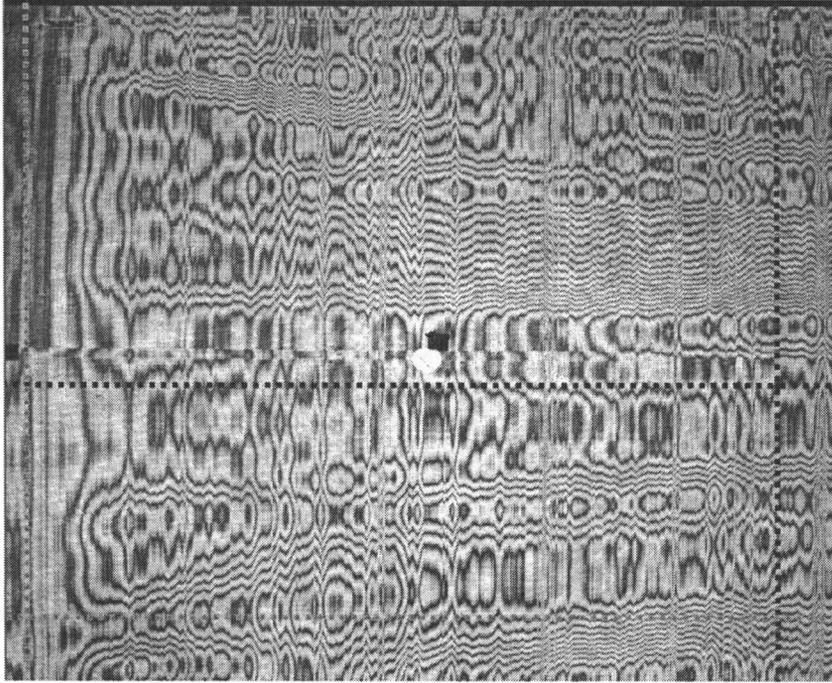
fixed line-etcher delivers and removes etchant from the underside of the glass, with the interferometer mounted above the table. The machine operates with multiple passes, in which only a small amount of material is removed each pass; continual real-time thickness measurements are used to determine the next pass's removal profile, and hence control the process.

The following two photographs show before-and-after 2-D interferograms of an 80 cm sheet of glass processed by this machine.



The above picture shows the sheet as-purchased, before etching. As usual for thin sheets of 1737 glass, the bulk of the thickness ripples are 1-D, aligned with the draw-dimension of the glass. Superimposed upon the primary ($\sim 10 \mu\text{m}$) 1-D ripples, however, are smaller ($2\text{-}3 \mu\text{m}$), 2-D, thickness variations.

Below, we show the same sheet after processing with the 1-D etcher. This view, which is rotated 90° from the before-etching one, shows both the strength and weakness of the 1-D etcher. The machine is designed for, and only capable of, removing 1-D ripples. The etching is controlled to smooth the glass along a reference line; the dotted vertical red line in the photograph below. Along this line, the 1-D etcher does an excellent job of smoothing the glass sheet; ripple was reduced in 2 days from an as-purchased value of $\sim 10 \mu\text{m}$, down to 30 nm . The concern is, of course, that this line etcher can only remove the truly 1-D portion of the ripple pattern. The post-etching view below shows that away from the 1-D reference line there are residual $2\text{-}3 \mu\text{m}$ 2-D ripples; these ripples can not be removed by our line etcher.



Accordingly, in order to fully smooth the glass sheets (to the 100 nm level needed to permit diffraction-limited performance) we need a 2-D smoothing capability.

During the course of this Strategic Initiative we considered 5 different approaches to doing 2-D spatially-controlled etching, and actively investigated three of them.

Two approaches, which were not pursued in this project, involved making the above Marangoni-based direct etcher into a 2-D machine. The most straightforward way of doing this, would be to use a point-etcher tool-head instead of the line-etcher. This was in fact done, at smaller size-scale, by Britten and Rushford for another project. Their machine, combining a point-etcher with a 2-D interferometer, does allow the etching of designated 2-D profiles into glass or silica. It can be used either to smooth glass to a uniform thickness (as needed for us) or to impose specific thickness profiles (for other optical missions). The drawback to this approach is that it is slow. This direct etching method is a serial process, working on only a portion of the surface at any given time. Obviously, the process-time increases as the area ratio between etchant-spot and overall sheet decreases; the 1-2 day times for the 1-D line-etcher get prohibitively slow for a 2-D point-etcher. Hence this, technically quite feasible, approach was pursued for other projects, but not for Eyeglass. Another variant, also not pursued, was to modify the 1-D line-etcher into a tool that could do 2-D smoothing. The approach being considered was to make use of the fact that the ripple profiles which had to be smoothed were mostly 1-D, having only $\sim 20\%$ worth of 2-D ripples. Accordingly, we considered ways in which the etch-rate along the linear etch-head could be slightly varied (to the tune of 20%). If this were done, then a suitably-modified 1-D machine could perform 2-D smoothing while retaining the process-time of the 1-D machine. While a number of methods were considered for varying the etch rate along the line (thermal, concentration, injection-rate, removal-rate, etc.), we ultimately did not pursue this approach. The reasons for not doing

so were partly due to schedule and resource concerns (such an effort could occur only after the 1-D etcher was successfully developed) and partly to the fact that it was still a serial-processing machine, limited to the etching rate of the 1-D machine.

We chose to focus our 2-D etcher development on parallel-process machines; ones in which the entire part is etched at the same time. The primary advantage of these approaches is speed, because they work on the entire part rather than a small portion of it, they are potentially much faster than the serial etchers. We investigated three different 2-D etcher processes; thermal, electrochemical, and printing.

The printing process isn't purely a parallel process, but had the potential of providing a rapid method for 2-D smoothing. The concept was to use printing technology (specifically commercial inkjets) to write a liquid etchant directly onto glass sheets. As with other methods, this was intended as a multi-pass approach; after each printing/etching cycle, the thickness profile would be remeasured by a 2-D interferometer and a new cycle would be performed. The attractions of this approach are that the resolution needed for smoothing (as opposed to patterning) is coarse, well within the capabilities of commercial printers. No low-level control software or motion systems would have to be developed, as one would simply use commercial print engines. The intent was to use Rushford's 2-D interferometer to measure the 2-D thickness profile of the glass sheet, convert this to a gray-scale image, and then print it using commercial printers and software. The two key requirements were the need for a wide-format printer, capable of printing 80 cm wide sheets (of 0.7 mm thick glass instead of paper), and the ability to print an etchant rather than standard commercial inks.

Finding printers capable of writing on wide, relatively thick media was relatively straightforward. Wide format printers routinely deal with 36-60 inch media. Most commercial units have curved paper paths and can not handle glass sheets, but there are units (one being the Epson Stylus 10000 series) which do have straight feed paths and can handle thick media. While we did not progress far enough with this approach to buy one of these full-sized printers, we did use a letter-sized printer (the Allegro model from Affex) for process development work.

The largest challenge to this approach, which we never adequately solved, was finding a suitable etchant-ink, and using it to write onto glass. There are two main concerns, damage to the printer if we directly print (even dilute) HF acid, and excessive spreading of the "ink" caused by printing onto glass instead of paper. The ink question depends a lot on our budget and our choice of printer. One extreme is to build-up a custom printer. The attraction here is that one can use specialized, rugged, print-heads, which are capable of printing acid. The drawback, is that developing the motion hardware and software for a custom printer was well beyond our resources; hence our decision to stick with commercial, already developed, printers. Doing so, however, greatly restricted our choice of "inks". The fact that a full-scale etcher would have to use an Epson wide-format printer, also meant that the inks we developed would have to be Epson compatible; in accordance with this, the letter-sized Affex printer we used for developmental work also used Epson print-heads and cartridges.

We were unable to develop Epson-compatible inks containing HF acid. Accordingly, we investigated a two-step approach. This was based on printing a F-donor and then activating it later. For instance we were able to mix NH_4F into a clear Epson-compatible ink-base, and to print it. This printable F-donor did not damage or clog the printer. Later, we could introduce another acid, such as HCl, form HF in-place, and do the etching. When printing the “ink” directly onto glass, we still had the problem of the ink either spreading or beading-up. We used two methods to deal with this; one was to precoat the glass with a thin cellulose layer and print onto this. Another, superior, technique was to print the F-donor ink directly onto paper, and to then place this paper onto the glass sheet, flood it with HCl, and etch the glass. This latter method did allow us to etch patterns onto glass. However, we found that the presence of the paper layer prevented the removal of etch biproducts from the glass, and hence capped the etch depths we could achieve. While we believe that developing a etchant-ink is feasible, it was beyond our expertise and resources; future efforts here should probably use outside ink-developers.

The other 2-D smoothing approaches we investigated were truly parallel processes, in which the glass sheet is uniformly immersed in etchant, and the rate of the etching is spatially controlled. This technique offers two potential advantages, speed because the whole sheet is processed at once, and reliability because there are no moving parts. We worked on two methods, one based on electrochemistry and the other on thermal control; other possibilities to modulate the etch rate exist, but were not actively investigated.

The electrochemical process we investigated was based upon immersing the glass sheet in an inactive etching solution with a grid of electrodes. The etchant (a solution containing NH_4F) is normally inert, having no effect on the glass until the electrodes are charged. Once this is done, F is liberated and etches the glass near the electrodes; high (~ 100 nm/min) etch rates are achievable. This process offers a number of advantages; there is no background etch, and low power requirements.

During our research, we did face several challenges however. One concern is the nature of the electrodes; inert ones don't work and so we had to use active ones, typically containing nickel and chromium. During the course of the etching these erode, injecting metal ions into the solution near the glass, which then deposit on it. Accordingly we had to develop multi-component etchant solutions which not only introduced F to etch the glass, but which also prevented precipitation of the metal ions. Another challenge was the short range of the etching; we needed to maintain very close proximity of the electrodes to the glass. This resulted in the imposition of sharp tool-marks into the etched glass, and the need to mechanically keep the electrode face near the glass as it was eroded.

We built and operated sub-scale (~ 100 electrode) etchers during this effort, and were able to etch μm -deep patterns into glass. Ultimately, however, we tabled this effort. The primary difficulty was the complexity of the etch process and the difficulty in balancing the need to produce F etchant with that of preventing precipitation of the metals released from the electrodes. This electrochemical challenge was accompanied by implementational ones caused by the fabricational complexity of a multi-thousand pixel etcher, and the sharp tool-marks produced by this process.

2-D PCB-Based Thermal-etcher

The final 2-D etch process we concentrated upon, and the one we selected for large-scale implementation was a thermally-based parallel process. This was based on the fact that HF-glass etch rates are strongly temperature dependent. Typically, a 10°C increase in temperature can as much as double the etch rate. This offers a scenario in which the glass sheet is uniformly immersed in etchant and then has a 2-D lateral temperature profile applied across its surface. The temperature dependence of the etch-rate then causes a 2-D variation of etching. If the thickness is monitored by a 2-D interferometer, and this data is used to recompute the heating profile, then the sheet can be controllably smoothed (or profiled). In addition to parallel-induced speed and the lack of moving parts, this approach has advantages over the electrochemical approach of physical simplicity and diffuse tool-marks.

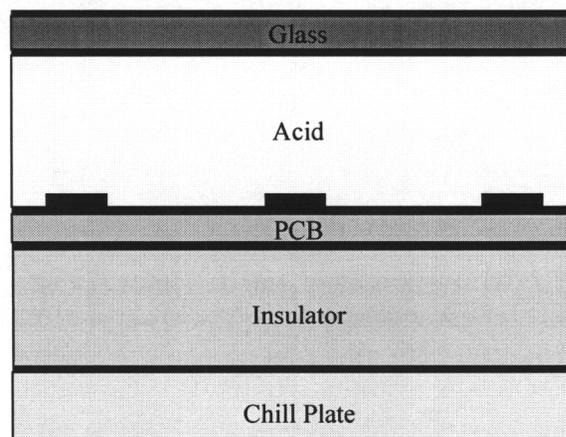
There are potentially many ways to apply a controllable 2-D temperature profile across the surface of the glass sheet. All face a common challenge, namely that it is difficult to maintain sharp thermal gradients given the presence of lateral conduction through the glass sheets and the etchant solution. This causes a direct relationship between the etch-rate (hence process-time) and the required thermal power; furthermore the power requirements increase as the inverse square of the lateral feature sizes desired. This type of approach benefits from an efficient heat source and from the relatively coarse, cm-scale, features sufficient for glass smoothing.

One approach, previously explored by Charles Thorsness and Jerry Britten, is to scan a laser across the surface, applying heat as desired. Thorsness has published extensive calculational modeling of this process, and some limited experimentation was done before our project. The challenges of this approach are the fact that it requires an active scanning system, that it is serial requiring frequent revisitation of each site to maintain thermal gradients, and the inefficiency of using high-grade laser photons to deliver low-grade heat. Another, more parallel, method is 2-D writeable light source such as a digital projector, to heat the surface. This has the advantage that it provides an already-developed, 2-D, rapidly-addressable, way to heat the surface. The difficulty is that commercially available projectors are relatively low-powered and expensive. Even top-of-the-line, movie-theater-grade, units do not deliver enough thermal power for rapid smoothing of glass sheets. Another approach is to combine the printing approach discussed above, with thermal etching. Here one would write a mask (on paper or transparency) with a commercial printer and then use it, when floodlit with a uniform light source, to impose a 2-D heating profile on the glass sheet. Since printers permit the rapid writing of such masks, using a pattern coupled to 2-D interferograms of the glass as it etches, this process could permit controllable 2-D etching. This has the advantage of using a lower-grade, more economical, light source than the other methods, but is human intensive in the continual need to print and mount masks.

We chose to pursue a different technique, based upon electrical heating, rather than using light sources. In this approach we use a 2-D grid of resistors to deliver a computer controlled heating profile to the glass/etchant system. To implement the resistor grid, we

use a custom designed printed-circuit board (PCB), containing thousands of tiny surface-mounted resistors with associated switching circuitry. This method leverages the fact that the nation already has a large, rapid-turn-around-time, PCB fabrication community; we do not have to build or hand-assemble this circuitry ourselves. The other advantages of electrical heating are simply that it is a much more robust and economical source of heating than are the optical methods we considered.

Before designing the PCB for this 2-D etcher, we first performed extensive numerical simulations to validate this etching approach, and to determine proper machine parameters. A number of layouts were modeled; the baseline configuration is shown below.



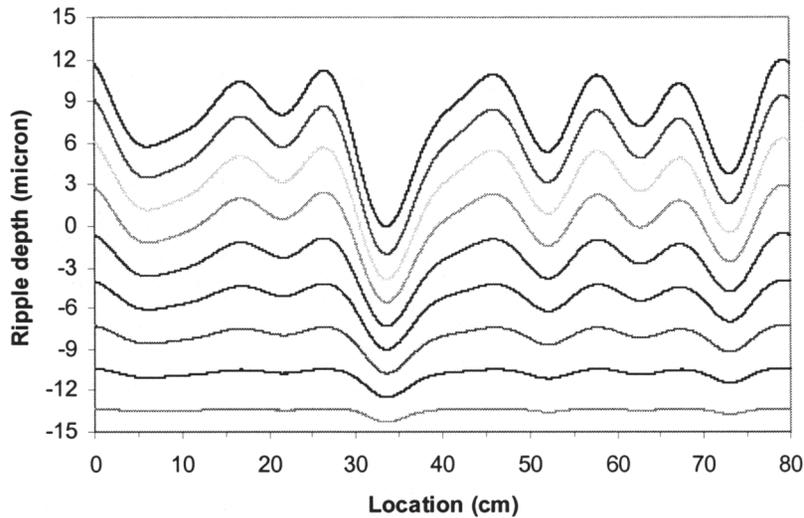
Here the PCB is placed at the bottom of a shallow (typically 6 mm thick) etchant bath (containing dilute HF and HCl). The glass sheet to-be-etched rests on top of the acid. In addition to doing the etching, the acid solution acts as the thermal coupling between the PCB heater and the glass sheet. It provides an efficient, but diffuse coupler, transporting heat to the glass, but also spreading it out laterally just enough to prevent sharp “tool-marks” from the discrete array of resistors. While some lateral conductivity is desirable, too much is not; hence the PCB is separated from the chill-plate (which of course is needed to preserve thermal balance) in order to prevent it from “shorting-out” the lateral gradients necessary for spatially-controlled etching. In this configuration, the glass sheet sits on top, allowing a clear view from-above for the interferometer, and also allowing straightforward insertion and removal of glass sheets. Since there is a potential drawback to this orientation, i.e., the possibility of structured convection cells, we have also designed and modeled inverted configurations.

We built a number of simulation codes to model the performance of this type of etcher. Scoping calculations are done with a time-dependent, 1-D code, modeling only the vertical flow of heat; this helps in setting power levels and thicknesses for the fluid and insulator. Most calculations are done with a time-dependent, 2-D code, allowing us to model one lateral dimension as well as the vertical one. While some sub-scale simulations were done with a full 3-D code, a large number of spatial zones would be required to model a full-scale (80 cm wide) machine; the 2-D code provides much faster results and captures most of the relevant physics.

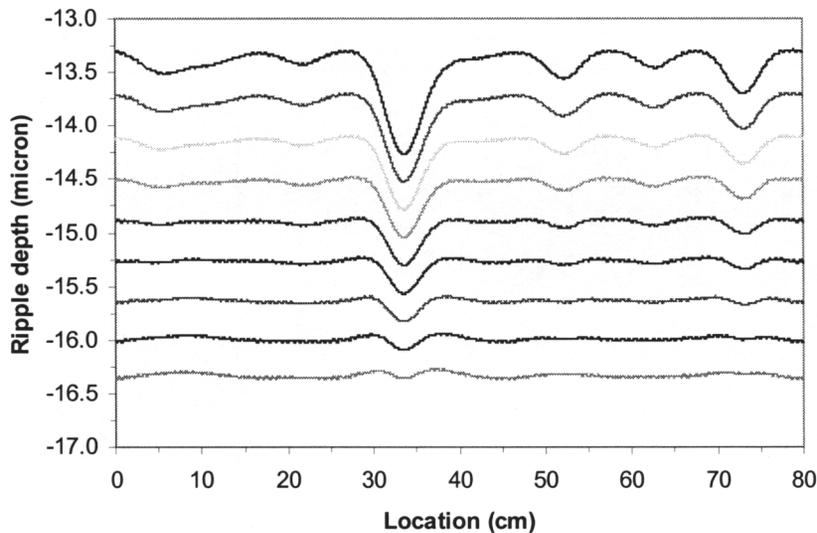
We used two types of 2-D simulations, open-loop ones to model the effects of a given resistive heating profile, and closed-loop ones to model our ability to smooth glass. A major question answered by the open-loop runs was the size of tool-marks; i.e., if we impose equal power on an array of discrete resistors, what spatial ripple is etched into the glass? Since our preferred, static etcher, has no way of removing such tool-marks, we want to avoid introducing deep ones. The ripples in our as-purchased glass have typical depths of 10-12 microns with spatial-periods (primarily in one lateral direction) of about 10 cm. We want to smooth the glass to an optical accuracy of $\lambda/10$, i.e., a physical ripple of about 100 nm. Unlike the electrochemical etcher, thermal ones etch everywhere; the lateral profile we intend to imprint (or remove) is accompanied by a background etch as the glass is immersed in a heated bath of etchant. This background etch does not present optical problems, but does mean that in order to remove a 12 micron ripple, the total etch depth will be considerably larger. Results of a typical tool-mark calculation are shown below. This used an array of resistors with a 7 mm pitch, separated from the glass sheet by a 6 mm acid layer; after 2 hours of heating, 28 mm of glass were removed. The removal was spatially quite uniform; the tool-marks show up as a sinusoidal ripple having a 7 mm period (matching that of the resistor array) and a P-V depth of only 24 nm, considerably less than our desired 100 nm smoothing accuracy. While the periodicity of this tool-mark ripple is set by that of the resistor grid, the depth depends strongly on both the thickness of the acid layer and the configuration of the etcher; ripple can be greatly reduced by a thicker acid bath and/or a configuration in which the resistors are underneath the PCB, not above it.

The most significant simulations are those which model the full closed-loop smoothing process. Here we start with a pre-existing thickness ripple on the glass-sheet, and try to smooth it. To perform this simulation we first take a “snapshot” of the glass thickness profile, feed the result into a custom control-law which generates a power command for each resistor. This discretized heating profile is applied to the system for some period of simulation-time (typically 5 minutes), at which time a new “snapshot” is taken, a new set of power commands is calculated, and the process continues. These calculations are run with the full width system (typically 112 resistors spanning 80 cm), and appropriate material thicknesses and properties for the glass, acid, PCB, and insulator. Below we show results for a typical simulation (run TE087) having a 1% HF/HCl etchant, a 125 mW/cm² peak power level, and using 5 minute “snap-shot” updates feeding a customized closed-loop control law.

Ripple : $0 < t < 6800$ sec

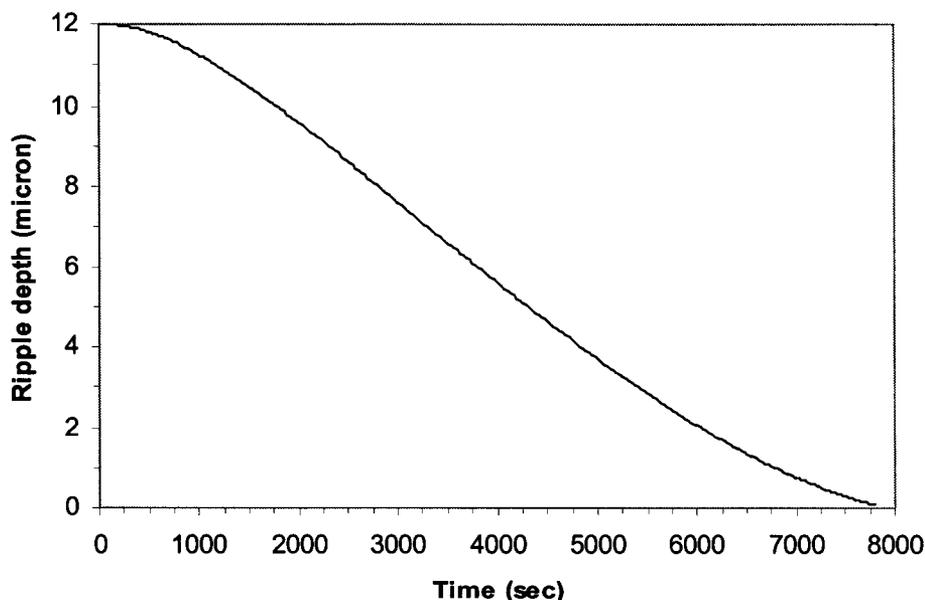


Ripple : $6800 < t < 7800$ sec



As can be seen, the ripples start at about 12 microns and are gradually reduced until, after a little over 2 hours, they drop below our target of 100 nm. The evolution of the peak ripple magnitude is plotted below. This monotonic reduction in ripple height is, of course, desirable, but not essential. A well designed control law achieves this monotonic performance, but less efficient ones generally “bounce” first; the ripple is monotonically reduced to 300-500 nm, then climbs again as the correction overshoots, before resuming a smooth reduction to the 100 nm target. Such overshoots lengthen the smoothing time (from 2 hours to 2 ½ or 3 hours), but do not interfere with successfully smoothing the glass; such, of course, are the virtues of having a densely-controlled closed-loop control process.

Ripple Reduction



This type of simulation shows that a resistor-based thermal etcher should be able to rapidly and controllably remove ripples from commercially purchased glass sheets.

After simulating the performance of this form of thermal etcher, we proceeded to design and implement it. This etcher has 3 main components; the PCB-based heater to deliver spatially controlled heating profiles, a wide-view interferometer to measure the thickness profile throughout the etching, and the rest of the system (the acid bath, glass holder, chill-plate, etc.) We did not complete this full 2-D etcher by the conclusion of this Strategic Initiative. We did, however, produce the key enabling components, the spatially controllable PCB-based heater and the wide-view interferometer. Their integration and testing in a wet-etcher is proceeding now, after the end of the current Strategic Initiative, in an effort led by Mike Rushford.

The design of a PCB-based spatially-controlled heater is governed by a number of factors. The first is the size we need to address and the spatial granularity we want to control heating over. Our glass sheets are 80 cm wide; to avoid boundary effects we want a somewhat wider heating region, say 90 cm. Now, 90 cm circuit boards can be built, but only by a handful of firms, as there is very limited demand for such sizes. At the same time, the smoothing capability we were developing was meant for two applications. It is needed not just for the panels of our 5 meter lens (calling for 80 cm sizes), but also for rapid smoothing of Disposable Debris Shields (DSSs) for NIF (calling for 40 cm optics). Because of these facts, we designed a modular heater system in which each PCB had a 45 cm active heating region; unlike the larger, 90 cm size, this sized PCB can be made by a large number of contractors. Single PCBs are capable of smoothing NIF-sized sheets (for DSSs), while 4 PCBs can be placed side-by-side to smooth the Eyeglass lens panels. To

insure gapless tiling for the latter application, we designed the boards so that the 45 cm square heating region extended up to 2 edges of the PCB, leaving the other 2 edges free for control circuitry. Our simulations showed that, in order to smooth out the 10 cm-wide ripples in the as-purchased sheets, we needed resistor periodicity of less than 1 cm. Once we began the actual electronic design of these PCBs, the advantages of a 2^N number of pixels in each direction became clear. As a result, our design led to each PCB having a 64 x 64 grid of resistors spanning 45 cm, i.e., 4096 controllable heating-pixels having 7 mm dimensions. The entire suite of 4 PCBs needed for smoothing lens panels hence has 16384 controllable resistors.

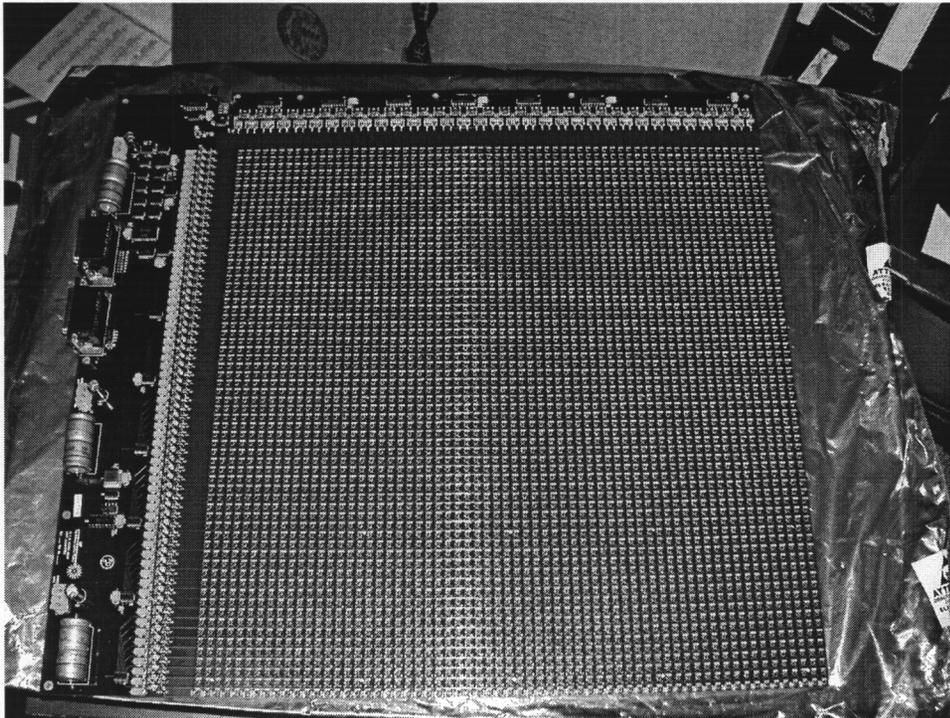
Designing the PCB-based heaters with this number of controllable heating pixels presents a number of challenges. One is simply fitting a resistor, and whatever pixel-level control circuitry is required, into the small 7 mm pixels. Another is determining the proper method to control the heating delivered by the large number of resistors. Our control system was based upon a number of key decisions. One is that we decided to use digital control; individual resistors are either on or off, but their power levels cannot be adjusted between these two extremes. This digital versus analog decision greatly simplifies the hardware implementation of the PCB, but does require a way to nonetheless deliver precise intermediate levels of heating, rather than the full-on or full-off values. Another choice was to keep low-level PCB control on the board itself, rather than in an external computer. While this leads to a more complex PCB design, keeping as much control as possible in on-board hardware rather than off-board software does permit more rapid control and simpler interfaces. A final, rather obvious, choice was to use a grid-based system to address and control the 4096 pixels, rather than individual commands to each; this lets us use 128 control points instead of 4096 ones.

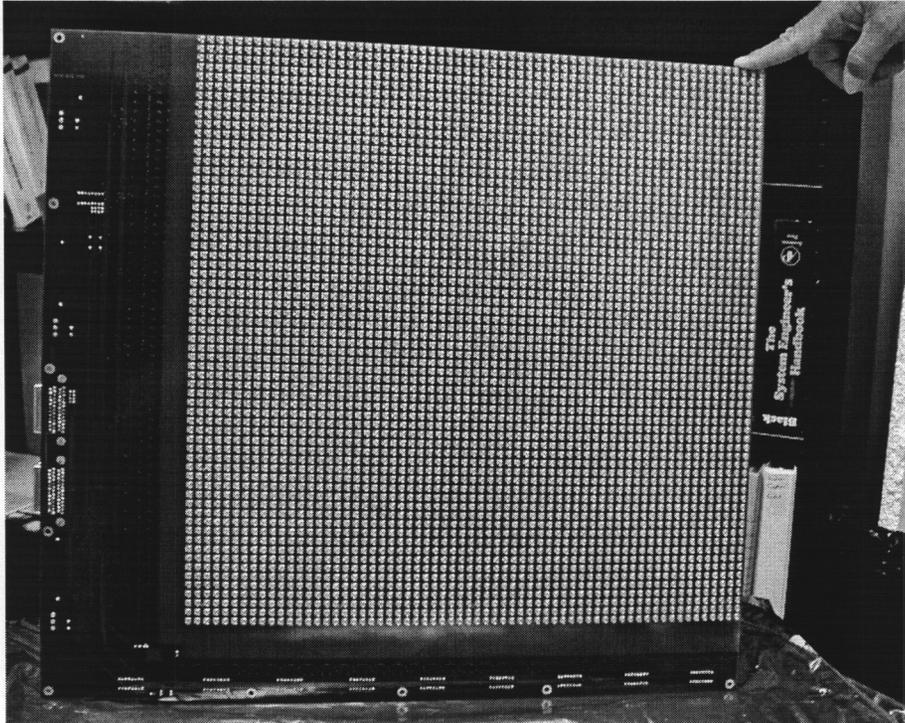
The logical and electrical design of the PCBs was carried out by an external consultant, Jordin Kare. The basic approach was to achieve the analog-like heating precision by digitization. Instead of a single 64x64 array of analog power levels, we used a 1024-level digitization, employing a rapid time sequence of 1024 separate 64x64 binary arrays. A command cycle starts with the wide-view interferometer taking a “snap-shot” of the glass. The computer then converts the interferometer’s phase map to a 2-D thickness profile. This is fed to the control law (developed and tested in our simulations) which determines the 2-D array of power levels which should be delivered by the resistors. Instead of directly applying these analog values until new interferometric data becomes available, we digitize the data into 1024 granularity and generate 1024 separate binary heating arrays (64x64 for each of the 4 PCBs). This control recipe consists of 512K 8-bit bytes, and is transmitted to the PCBs using a fast parallel port. It is stored on the board in a 512K RAM, and will command the board’s heating profile until new interferometric data is available and new data is sent to the board (nominally every 300 seconds).

The use of digitized heating is, given thermal inertia and diffusion times, indistinguishable from fully-analog control as long as it is cycled rapidly enough. Our PCBs have jumper-settable clock cycles, but are nominally designed to apply each 64x64 binary heating pattern for 0.5 msec, running through the entire 1024 sequence twice each

second, repeating this process until new data is sent. The actual low-level application of a 64x64 binary heating pattern from its stored recipe in the RAM to actual resistor commands is handled in a process designed to be simple and to involve as few elements-per-pixel as possible. The method used requires only 2 elements for each pixel, the resistor itself and one control transistor. Command occurs by a row and column method in which only a single column of resistors is heatable at any given time; for a period of 8 μ s. Each column contains 64 resistors, so requires 64 bits of control information. This data is clocked out of the RAM as eight 8-bit bytes and latched into row-registers. Once the column is activated a gate signal is sent to each transistor in the column; only the ones which also have row signals will switch on and feed power to their resistors. After 8 μ s, this column is done and the next is activated; after 512 μ s all the columns have been addressed and this 64x64 binary pattern has been applied. The next binary array is then read out and applied in the same fashion; after 0.5 seconds, all 1024 binary arrays have been applied and the whole sequence repeats. The thermal time-scales of the etcher insure that the temporal discretizations of this approach do not affect the performance of the etcher.

The PCB system is designed for peak heating values of 0.2 W/cm², so that each PCB must deliver 400 W of peak power (200 W on average). This requires only 0.1 W per resistor. The PCBs were designed by Ron Shaw using Kare's electrical design; they are 4-level boards and were metalized and stuffed by an external contractor. As mentioned earlier, they consist of two regions, a 45x45 cm active area containing the heating pixels, and the control circuitry which is consolidated along 2 of the 4 board edges. Each pixel contains only two elements, a resistor on one side of the board and a transistor on the other. Below we show photographs of the front and back of one of these circuit boards.





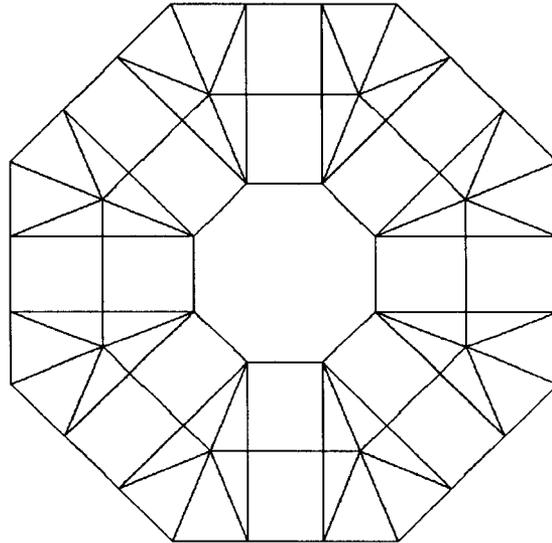
Because the development of this approach and the design of the 2-D PCB heaters occurred later than planned, we could not incorporate 2-D smoothing into the 72 lens panels needed for our 5 meter lens. Accordingly, we decided not to rush ahead by building all 4 PCBs, but instead to build and debug one PCB; after this is operational the others can be completed. We have purchased all the electrical components for the 4 PCBs, and have written (metalized) 5 of these PCBs; however only one of these boards has been stuffed with the electrical components.

At present we have one PCB fully assembled and operating in Mike Rushford's etcher. It is, after the replacement of a few defective components, working electrically as designed. Closed-loop etching experiments (performed by Rushford after the conclusion of this Strategic Initiative) are underway and (while still at limited power and etchant concentration) have achieved 2-D smoothing down to $1 \sim 2 \mu\text{m}$.

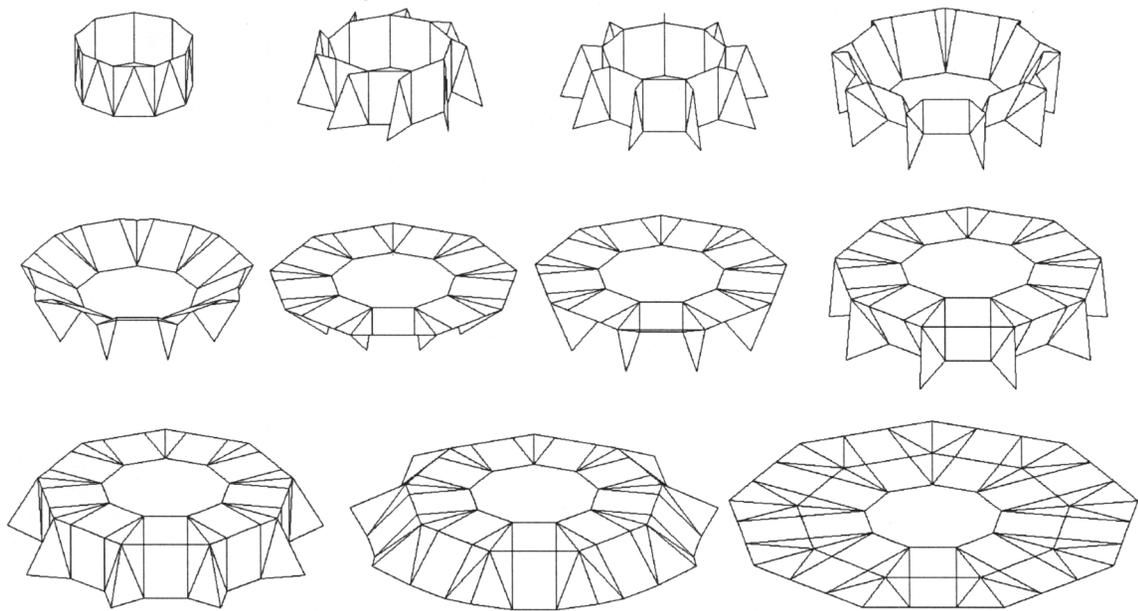
We believe that when the operational development of these 2-D thermal etchers is complete, that they will provide a rapid, effective, and economical way of smoothing large glass sheets, either for Eyeglass lens panels or for NIF disposable debris shields.

Design and Patterning of Diffractive Lens Panels

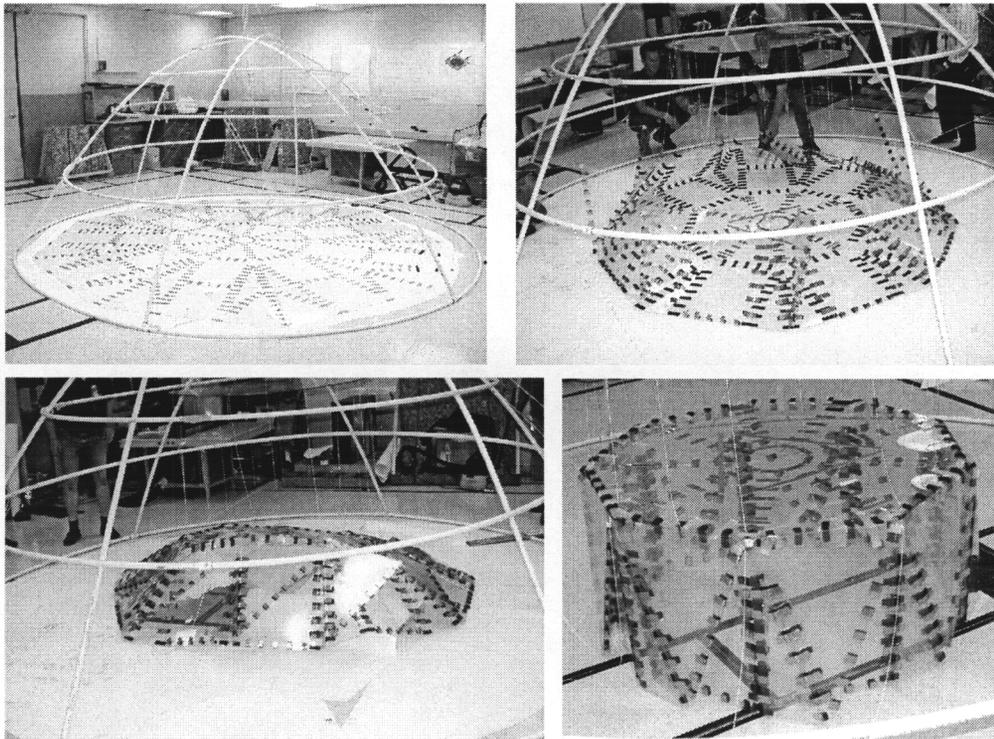
The 5-m Fresnel lens design we have chosen has an origami like pattern to enable it to be compactly packaged for space launch. This pattern (with 8-fold symmetry) was designed for us by an external origami-expert, Robert Lang.



The folding sequence, taking the lens between its flat unfolded state and a hatbox-like, folded configuration, is shown below. This sequence is for a lens pattern with 9-fold symmetry, but our adopted 8-fold pattern folds in the same way.



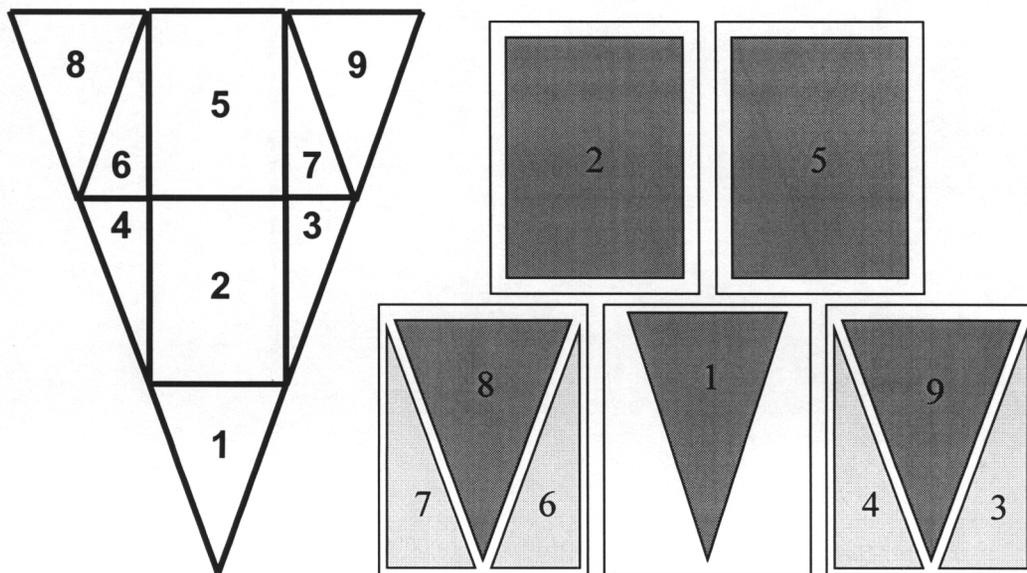
John Toeppen built a one-foot-to-one-meter scale model (i.e., 5 foot diameter) of this lens and repeatedly folded and unfolded it. The folding/unfolding process is simple and symmetrical, lending itself to mechanical assistance. He then built a larger, 2/3 scale, polymer version; photographs of its folding sequence are shown below.



Our 8-fold symmetry lens design consists of a total of 72 pieces made up of 16 rectangles, 32 right angled triangles and 24 isosceles triangles. The table below lists the sizes of these shapes.

	Number	Size
Rectangles	16	654 x 790 mm
Right angle triangles	32	327 x 790 mm
Isosceles triangles	24	654 x 790 mm

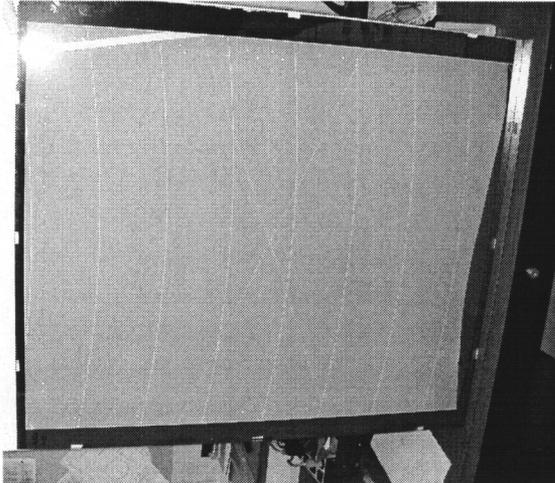
The high degree of symmetry in this pattern can be exploited to minimize the different types of tooling necessary for handling during the patterning and assembly steps. In addition to the 8-fold symmetry (reducing mask-count from 72 to 9), we can further reduce the number of masks to 5, by grouping some triangles into rectangular groups. This consolidation is shown below. It has the further advantage, that all lithography can be performed on identical (654 x 790 mm) rectangles, requiring only a single set of glass-handling tools. All mask writing and lithographic patterning is performed on these identically-sized rectangles; 5 masks are required, and 40 rectangular sheets are produced. Only after patterning is completed, are the 40 sheets cut apart into the 72 separate lens panels.



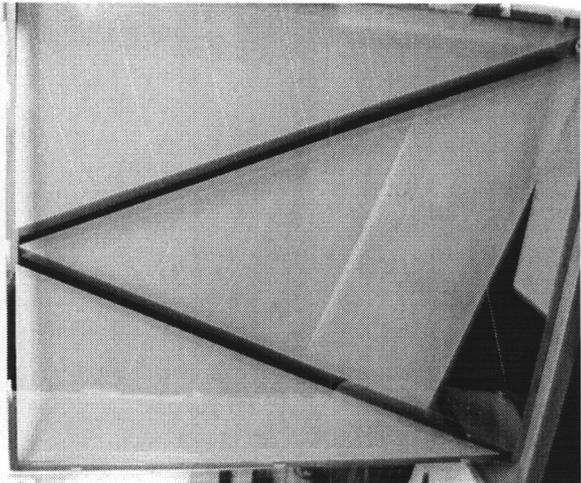
The diffractive pattern for the lens was designed to give a 250 meter focal length for 0.60 μm light. The basic parabolic paraxial profile was adjusted slightly with quartic terms specialized for point-to-point focusing, rather than plane-to-point. We adopted a simple binary phase-profile diffractive pattern, rather than more efficient multi-level ones. This was done because such a profile is perfectly capable of demonstrating full-precision diffractive focusing. Increasing efficiency from the binary 40% value up to the 80-95% values attainable with multi-layer designs was not crucial for us, and not worth the extra masks and patterning required.

The 5 rectangular masks needed to define our lens's diffractive profile were written for us by an external contractor, Advance Reproductions Corporation; they have the capability of writing the large, high-precision masks we need. Photographs of 2 of the 5 masks are shown below.

Mask for Panel 5



Mask for Panels 6, 7, 8



Apparent on these masks photographs is not just the circularly symmetric lens profile, but also a number of alignment features. The most crucial of these, are a large number of small fiducials along the borders of the panels. These fiducials are designed to be machine-read by the quad-looker of our assembly machine. They act together in groups of four (two on the border of one panel and two on the border of its neighbor) to allow precise, sub- μm , measurement of the panel-to-panel alignment. Another, unplanned, feature is apparent from the mask for panel 8; portions of its pattern are in error, following straight lines rather than circular arcs. This error, caused by offsite software errors, showed up only on this panel. While it was subsequently fixed, and a correct mask was produced, this occurred late in our fabrication sequence; accordingly the outer portion of eight (due to symmetry) of the panels along the periphery of our lens had improper diffractive patterns. These errors caused a small reduction in the areal efficiency of the lens, but are not of major significance.

After the masks were written and received by LLNL, they were used to lithographically pattern 50 rectangular panel-sheets (the 40 required for the lens and 10 spares). This work was done by a team led by Sham Dixit and including Paula Smith, Matthew Schmitt, Leslie Summers, and Mike Aasen. This patterning was done using procedures previously developed by the Diffractive Optics Group; the 1737 glass sheets were coated with photoresist, and exposed through the mask to a large UV light source. The undeveloped portions of the photoresist were then washed off, leaving the desired diffractive phase profile imprinted in photoresist on the surface of each sheet. Generally, at this point, such patterns are then etched into the underlying glass (or silica) by wet-etch using dilute HF acid. We did not do so for these lens panels because of the limited service

life required of our lens; accordingly its diffractive phase profile is formed by the proper depth of photoresist on top of the glass sheets, not by grooves etched into them.

After the lithographic patterning was completed, the final step in fabrication of the lens panels was cutting the 40 rectangular sheets into 72 rectangular and triangular lens panels. Given the fragility of these sheets (and their high value), we used an external contractor, Abrisa, expert in precision glass cutting, to perform this operation. They successfully cut apart the sheets, and delivered us 90 finished lens panels (the 72 needed for the lens plus 18 spares).

Precision Assembly of Large Diffractive Lenses

The Alignment Problem

Mechanics of Membranes

Viewgraphs and mental images routinely simplify Fresnel optics into planar disks whose thickness is negligible. Our ability to actually build large Fresnel lenses depends strongly on how thin lens parts can be made, and on other components that depart from a constant thickness. The optical performance that makes large Fresnel lenses attractive demands fabrication of nearly-sawtooth-shaped grooves in at least one surface of nearly-constant thickness components that fill almost all of a lens aperture. By themselves, these purely geometric considerations are insufficient to build and operate actual large lenses. Smaller Fresnel lenses can gloss over mechanical and chemical effects, but the full potential of large lenses imposes non-geometric requirements with few precedents.

This section of this report assumes that others have solved the problems of implementing the optically active area of large Fresnel lenses. Those solutions are nontrivial and not currently available commercially (for apertures larger than ~200 mm that can approach diffraction-limited performance). Therefore mechanical solutions capable of implementing large Fresnel lenses of optical quality are constrained in the choice of materials and thicknesses that will constitute most of lens area. Optical performance also exports a figure of merit for mechanical solutions – they must occupy a small fraction of the optically active area.

Despite these departures from the mental image of an aperture filled with a thin sheet, the solutions discussed in this section are perturbations on the mechanical model of a nearly-flat membrane with uniform thickness and materials properties, without gaps or holes. Under tension, this idealization obeys LaPlace's equation for small excursions away from planarity, and Poisson's equation for response to loads. Actual materials will add damping, but the only materials available for implementing 5 meter aperture lenses within the scope of this LDRD effort are glasses – which have essentially no damping and experience brittle failure under low stresses.

Finite thickness modifies this membrane analysis by including bending strength, which turned out to be crucial for handling large, delicate glass panels during fabrication. Although bending strength is otherwise irrelevant to the mechanics of the solutions discussed in this section, it will be crucial for applications of large Fresnel lenses and offers future prospects of a very different class of mechanical solutions.

Problems in Space and Air

A large Fresnel optic in space poses a plethora of mechanical and optical problems that could not be solved within the scope of an LDRD effort. Several deployment solutions appear adequate for converting a compact configuration capable of surviving launch conditions into a planar membrane. Folding which localizes out-of-plane

deflections along pre-determined lines of hinges remained the baseline solution throughout the LDRD effort, although brainstorming and some analysis was applied to crumpling, rolling, and snap-together-some-but-not-all-joints alternatives.

Engineering plans and some prototyping were applied to mechanisms that could fold 5 meter lenses in the folding geometry illustrated elsewhere throughout this report. Hinges adequate to make 5 meter prototypes foldable were developed, but their anchorage to optically active glass panels was not developed.

Hinges suitable for folding in space remain a costly engineering challenge. Their requirements were analyzed, and include thermal matching, local loading of the optically active membrane, nesting geometries for packing at launch, and shadow reduction to maximize the useable solid angle wherein a lens could function with sun angles close to its plane. Both shadowing and point loading extend beyond the active area of a lens into a support structure that can provide both stiffness and damping to a large Fresnel space optic. Baseline tensioning designs for space deployment rely on centrifugal 'forces' to keep the planar element planar, but service in orbits other than geosynchronous appear likely to impose slewing maneuver time constraints (to point a lens in different directions) which favor external structures. All of these space deployment problems are so specific to differing sponsors' potential applications that it makes sense to defer their solution until someone appears willing to pay for their costly solution.

Deployment in air provides several options to fully validate large Fresnel optics at the 5 meter aperture scale. Outdoor testing was used to validate 5 meter lens focusing using laser sources, while indoor testing inside the Nova bay was used to validate monochromatic diffraction limited optical performance as well as the preservation of that performance after a folding/unfolding cycle (performed with metal hinges glued to glass panels). Both of these validation tests were performed in air with a 0.75 meter aperture foldable glass Fresnel lens built with NIF large-area microfabrication processes.

Similar validation tests at the 5 meter scale were costed out and would have exhausted LDRD funds, but remain perhaps two orders of magnitude lower than the cost of space testing at this scale. Index of refraction nonuniformity requirements leave few sites capable of validating diffraction limited performance of a 5 meter lens. Of these, a tunnel at NTS provides the windless, 0.5 C temperature variability environment that was most affordable when such sites were evaluated in mid-2001.

Ultimately we demonstrated a 5 meter Fresnel optic in the open air at LLNL. A beampath more than 200 meters long was found that made experimental validation of optical performance possible at the 5 meter scale. Although this validation was unable to approach diffraction limits, it went a long way towards demonstrating that large Fresnel lenses were feasible. The outdoor locale imposed thermal variations in index of refraction along the beampath that make wavefront diagnosis nearly impossible, although speckle interferometry may well be able to overcome these difficulties in future experiments (especially with the aid of a big tent).

Wind loading provided the most severe challenge to experimenting at this scale in open air. Rain and bird strikes were also environmental challenges that were overcome with the deployment apparatus discussed elsewhere in this report, while large temperature excursions (esp. freezing of condensed water) was not addressed and proved to be serious but avoidable environmental effect.

Joints in a Membrane

Eyeglass research and development efforts at the 5 meter scale demanded so much engineering of both the optical elements and their assembly technologies that the LDRD effort assumed there would be not time for a second iteration of the overall design. A baseline optical design was frozen early, so that photomask design could begin. Mask completion enabled the months of processing of glass panels necessary to produce the optically active elements that were subsequently assembled into the first 5 meter lens.

The Eyeglass LDRD lens design picked a panel size at the extreme limits of previously proven fabrication processes, which combined with a preferred folding geometry to impose the exact shapes of the parts that required assembly. Given such parts (72 panels), the problem of how to assemble them required a specific solution chosen from among several general forms considered prior to 2002. In order to design the masks that built the panels, a nominal width for the gaps between panels was chosen early, based on preliminary hinge designs.

A cobalt alloy, Elgiloy, was chosen for its high elastic strain rating, which makes it an ideal hinge material. Available thickness of foils then set the hinge bend radius. Two different hinge geometries provided for hinges that could fold gaps in the lens through 180 degrees and through excursions below 90 degrees. These hinge geometries combined with the constant width gaps between panels to consume less than 5% of lens area.

Preliminary adhesion tests found glues adequate to attach this alloy to glass, so that all portions of the lens design were consistent before mask designs were frozen. These designs were being implemented when external funding was provided for the outdoor demonstration. The masks already in production (at a total cost near \$50,000) then rendered the semi-automated assembly approach, and the hinge design suboptimal.

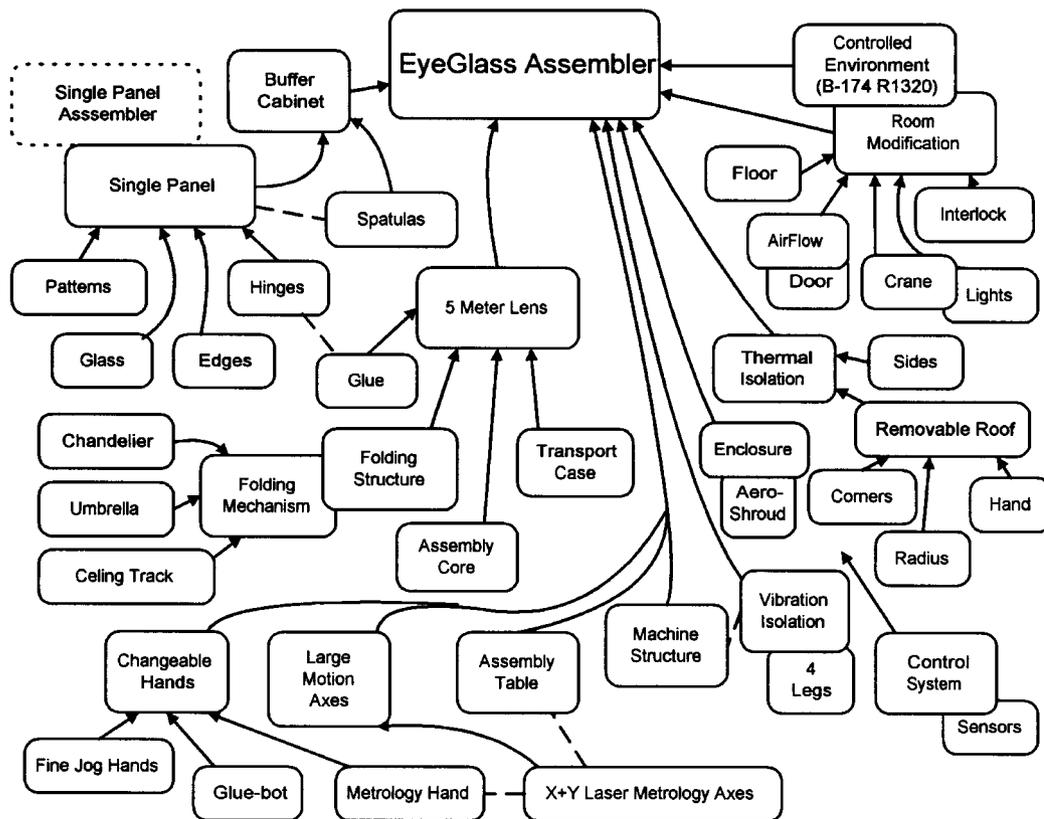
The lens that was actually built was manually assembled, from thicker, less flexible, aluminum hinges. Initial assembly of un-smoothed panels that were cut to the appropriate shapes was performed in April 2002, on top of the partially completed table that forms the working surface of the automated assembly apparatus then under construction. Simple handling of the cut glass panels broke 2 of the 18 cut spares, and convinced the assembly team that frames would be required to prevent crack propagation inward from almost invisible defects in the panel edges.

The incorporation of 0.020" thick aluminum frames around the edges of panels doubled the number of glue bonds in the membrane tension load path, altering the adhesive choice from hinge alloy to glass. The original folding hinge alloy failed to meet

glue shear strength requirements for all adhesives tested. These hastily obtained negative pull test results combined with the complexities of wind-driven flexure, the cost of the Elgiloy hinges, and the sponsor's schedule exigency to force the temporary abandonment of flexible hinges in favor of thicker (0.010") aluminum hinges.

Although the 5 meter lens that was fielded outdoors did not solve the majority of the assembly problems of a diffraction limited lens, a full solution was found and partially implemented before the decision for outdoor fielding in June of 2002. That solution met all the assembly accuracy requirements of a diffraction limited 5 meter glass lens, as well as the materials handling problems of the delicate panels. It remains a good enough solution to build diffraction limited Fresnel lenses from various glass materials in arbitrary panel geometries as thin as 0.003".

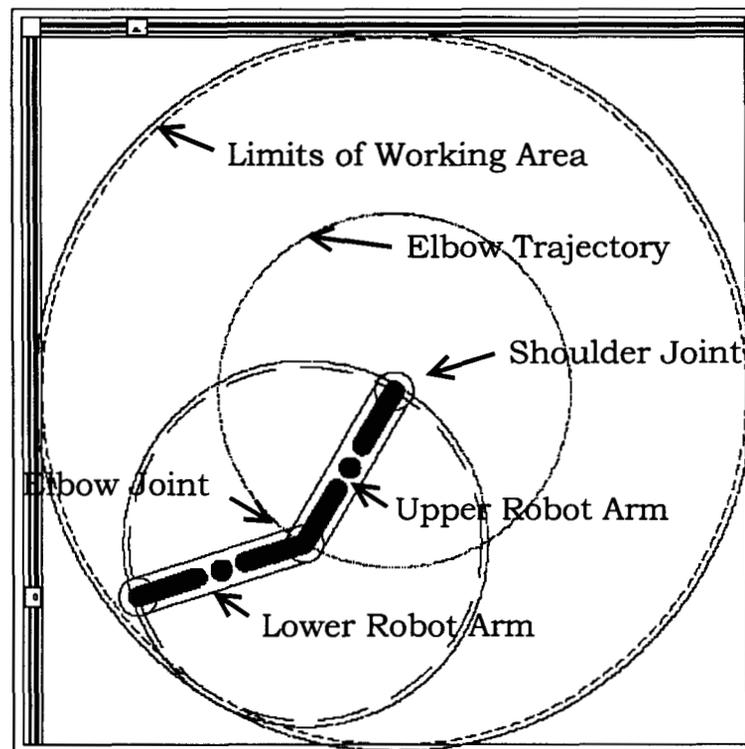
The Figure below shows a bubble diagram representation of the Eyeglass assembly team's parts list for the apparatus that has been partially implemented. Essentially all of its (nearly 800 kinds of different high-technology) components were procured (costing roughly \$400,000), the largest (structural and facilities) subsystems components were approaching construction completion, and other subsystems critical to its accuracy were being debugged when our effort shifted towards the fielding and testing the 5 meter lens outdoors.



Errors in Assembly

Most of the difficulty in building a diffraction limited 5 meter lens results from the radial location accuracy requirements (~6 microns global maximum location error) of the light bending grooves imposed on one or both sides of its optically active panels. Circumferential errors can be roughly an order of magnitude larger, decreasing as the angle subtended by a panel grows.

The first iteration of 5 meter lens design (performed in the Autumn of 2001) presumed 9-fold symmetry, but facilities revamping issues soon altered this to 8-fold (in order to preserve the option of folding the lens in half to get it out of the room where it would be assembled). Uncertainties in the geometry of the 5 meter lens panel locations within an assembled lens imposed excessive risk on any assembly apparatus that was tied to a specific panel geometry, so a solution that can assemble arbitrary panels within a 228x230" rectangle was adopted. The Figure below shows schematically the circular area of assembly within this rectangle that a 2 joint robot can reach in the horizontal plane. This solution extends the 6 micron global accuracy requirement to both axes of a nearly planar assembly process.



If the lens being assembled from processed and cut glass panels is going to achieve these accuracies, and the assembly team can't afford a second try, too many operations are involved to allow human error. Probabilities of error for human operations generally

run above 1%, and several hundred hinges were anticipated. The actual lens that was built by hand in the Summer of 2002 relied on very skilled technicians (Phil Stephan assembled it), had no hard accuracy specifications, broke one panel in its corner (which stayed put outside the tension load path due to its aluminum frame), and drooled glue over two others. The entire manual assembly operation would not have been possible within better than ~500 micron accuracy due to the thermal perturbations.

The automated assembly apparatus (discussed in the next few sections) was designed to hold 0.030 degree C temperature control during its precision assembly phase to restrict thermal expansion within roughly 25% of a nominal error budget that is expected to outperform the 6 micron global accuracy goal. In hindsight, considerably more manual intervention in assembly operations for a diffraction limited lens could save development costs, without implying the potential to remove and rework panels, as long as the operator remotely controls precision manipulation.

Limits of the Planar Approximation

Several flavors of potentially disastrous wishful thinking were discovered during apparatus design. Mental images that are purely geometric fail to capture the complexities of realistic diffraction limited assembly. Of these, the departures from planarity were most surprising.

With roughly a dozen terms in the error budget, and a margin of two for optimism, each error term in a “root-sum-of-squares” analysis must be held within roughly a tenth of 5 microns over 5 meters. The one-minus-cosine (or the angular deviation from flatness) error term that arises from excursions away from planarity matters when slopes are allowed to approach 100 microradians. This precludes assembling on surfaces that have not been smoothed with uncommon surveying and leveling techniques.

Since it was not possible to procure single pieces of suitable materials 5 meters across, a working surface for assembly must be constructed from multiple ‘tiles’ whose edges don’t exactly line up. If with tile flatness can routinely be obtained within .002”, the spatial period of this error term is probably shorter than 10” and threatens to exceed allowable slope error. Much worse are the steps at the interfaces between tiles, which must be smoothed so that the departures from planarity are spread over several inches away from the step.

Intermediate structures that are useful to reduce or eliminate final position errors due to friction, called ‘spatulas’ by the Eyeglass team, were thrown out of the design due to the cost of roughly a hundred panel-sized subassemblies (\$100 each would have made them an expensive, panel-geometry-specific subassembly). A design using spatulas can pick where the spatula is supported, preferably on three points as far apart as possible, and would preclude the need for smoothing seams. The team’s baseline solution involved specialized laser-based surveying (a service available at LLNL) and mechanical reworking of the table surface to impose global flatness in the .001” vicinity, although curing a 19’ square sea of low viscosity epoxy was gaining on this baseline.

Solutions for Hidden Variables

In-plane Stresses and Strains

Extensive brainstorming on assembly approaches that the LDRD team performed in 2001 was routinely plagued by the illusion that solid artifacts have fixed positions of their constituents. Position errors in the 100 ppb range result from strains in the 100 ppb range. These strains are trivial to exceed with a one degree temperature excursion in almost all materials, and preclude any panel joining technology that exchanges a few watts of heat. Glasses have relatively low thermal expansion, so modest temperature control can stabilize their position, but the accuracy of temperature control available in entire rooms isn't good enough. Active cooling of any residual heat imposed by attaching panels to hinges is routine to millidegrees, but intimate thermal contact between the work and a thermal control subsystem is a major requirement of any adequate assembly solution.

The stresses that result in 100 ppb strains in glasses are in the 1 psi range, and can result from slight forces applied locally to thin parts. Because of the delicacy of parts being assembled, and the desire to assemble even thinner ones without abandoning a costly apparatus, the team's initial solution held the pieces to be assembled in place while they were being manipulated into their final positions as components of a lens. They still had to move with a smoothness that precluded static friction, so any clamping solution had to rest atop a 'spatula' with some sort of 'bearing' underneath.

Air bearings were developed that could levitate or clamp panels, but their air stream was in intimate contact with the work and their heat transfer coefficients were very sensitive to local variations in their air gap. A baseline spatula design with electrostatic clamping was partially proven with a 2 square inch prototype. The prototype expected to evolve into cheap circuit boards in spatulas, whose cost was hoped to be less than \$200). This approach was demolished in a design review with outside critics (largely from NIF and LODTM) in February of 2001.

Besides the technical risk of electrostatic clamping (vacuum clamping didn't work without covering the assembly table with tubes or with predetermined vacuum connection locations, that require a fixed panel geometry), the 'clamped' approach ran into trouble with differential thermal expansion. Stresses necessary to keep glass clamped in one place on spatulas or tables made of different materials would break the glass or cause it to slip locally unless thermal expansions were matched within roughly a par-per-million-per-degree-C. Therefore clamped designs were predicated on the availability of Kovar (an exotic nickel alloy that can be expansion match to glass), which took too long to procure in thin sheets sufficient to cover the required area (6 months for a furnace run). After the design review, the Eyeglass LDRD team emerged with the mantra "Free the Lens", and henceforth pursued designs that approached zero in-plane stress instead of zero in-plane strain.

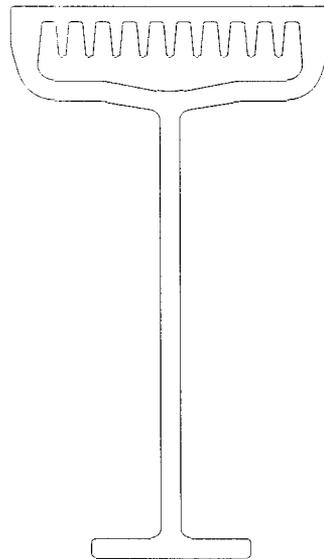
Thermal Requirements

In-plane strain errors result from forces applied (via elastic constants) and from thermal excursions (via coefficients of thermal expansion). If the panels are free to move during manipulation, before they are secured to hinges and thence to their neighbors, the only forces applied to them come from their suspension and from a manipulator that shoves them sideways. That suspension must provide “intimate thermal contact” with a thermal control subsystem below an assembly table, because the volume above that table will be occupied by manipulation and position diagnostic equipment, which must be free to move across the reachable disk illustrated previously if that costly equipment isn’t to be replicated for every panel. The Figure below shows that table under construction, when its skin was first put in place, before thermal control components have been installed underneath.

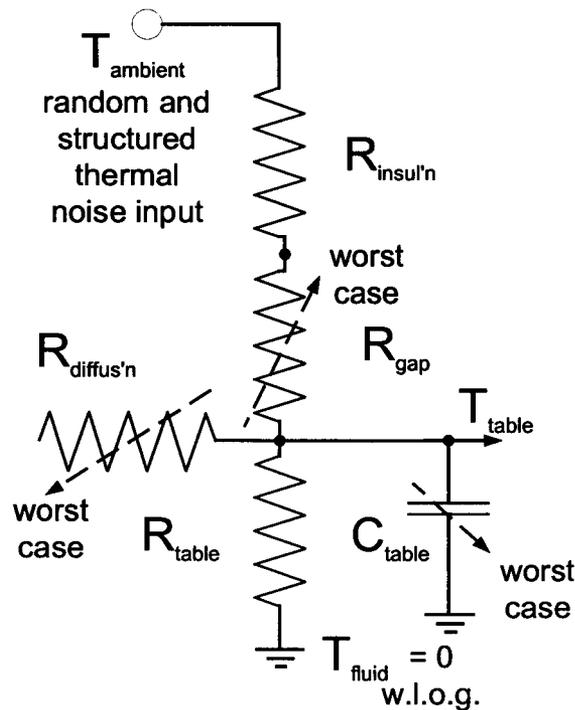


The thermal control systems that are easiest to develop rely on procurable centralized components, as opposed to distributed thermal actuation (e.g. an array of controlled resistive heaters). Distributed thermal sensing is still required to prove that sufficient thermal uniformity is being achieved for accurate assembly, but compact thermal control components can be procured that can routinely control a single location’s temperature within millidegrees. This thermal accuracy must then be distributed by a working fluid, and delivered to the horizontally spread out panels through heat exchangers. Below we show the cross section of the custom heat exchanger designed to provide both sufficient surface area in contact with a flowing 50/50 water+ethylene glycol and strength in

bending sufficient for the table to support a hypothetical 300 pound technician crawling across the table on knee pads.



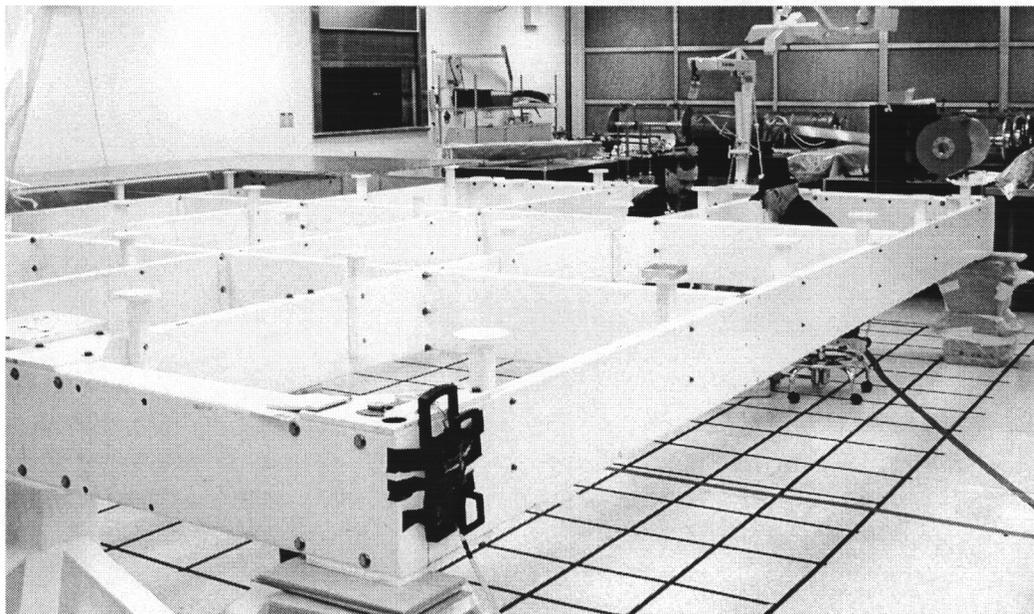
The 0.030 degree C thermal control requirement may be routine to achieve in a central box, but several error terms enter its distribution across the table to panels. Random heat sources can come in from above and below the table. Reducing these random effects puts constraints on thermal conductivity along several heat flow paths. All four known sources of random heating were engineered down to acceptable levels in the solution whose block diagram is shown below, for which all components were procured.



The four sources of random heating that thermal engineering anticipated are motors in the coarse manipulation drive subsystem, heat coupling through the table support structure to the floor, conduction to the room air (the B174 r1320 facility was chosen and qualified for its 1 degree C temperature regulation), and spatial variation along the heat exchangers under the table. To control the input room heat within its quarter of this error budget, 6 inches of still air will be necessary in an insulating enclosure around the machine, and at most a millimeter of air can be allowed in the heat flow path between a centrally regulated fluid and the work.

Over 15 different options for air and liquid cooled heat exchange were explored in the thermal subsystem design. Many of these were motivated by the thermal expansion matching requirements of a 'clamped' approach, and became superfluous in a 'free' suspension. It remains possible to build tables or spatulas out of glass or Kovar that also act as heat exchangers, but considerable cost savings (compared to the original apparatus hardware budget) were achieved using custom aluminum heat exchangers under an aluminum honeycomb table.

The high conductivity of aluminum in both heat exchangers and table surface allows a serpentine heat flow path to be followed by heat flowing to or from the fluid to the work. An inexpensive grade of aluminum honeycomb 1/2" thick (intended for clean room partitions, and made in Oakland to LLNL specifications with two weeks delivery time) became the most cost effective construction material for the entire apparatus. The huge table surface shown previously was supported on adjustable-height post spaced every 38x47" on a steel substructure shown at an earlier phase of construction below.



The honeycomb assembly table surface's low mass per unit area of 1 pound per square foot give this design a low heat capacity, which brings its thermal equilibration time (assuming thermal conductivities obey the thermal error budget) below 2 days. An

e-fold decrease in random temperature is predicted to take 4 hours for, while the entire four component random heat flux has been calculated at below 30 watts.

Panel-to-Panel Departures from Planarity

Besides the table flattening requirements mentioned three sections above, other assembly requirements were solved in the gap between panels. The panels themselves can be at different levels, and even tilted within the allowable 100 microradian slope error, but this leads to significant variation in the level of hinges on either side of a gap. Hinge and fine manipulation contacts with the panel must be very elastic in the out-of-plane direction to allow for these variations, which do not affect optical performance of the assembled lens. The suspension of the panels themselves also has a height variability specification, which is only a serious concern for the ‘velvet’ option.

The baseline suspension design maintains tight vertical position control of an air gap between table and work with regularly spaced and highly uniform 1 mm ball bearings, but the actual operation of this design is expected to be quite sensitive to small departures from planarity for a supposedly flat table. Even 0.030” thick glass is not stiff enough in bending to be suspended just on 3 balls (a fully-constrained level solution) without imposing excessive slope errors. A more closely spaced array of balls allows some balls to drop out of their suspension role. This was experimentally demonstrated on a granite flat. Air suspensions were proven to function splendidly despite vertical position errors in their air gap, but those errors dominate the variations in thermal conductivity that makes moving air next to the work problematic. (Air bearing solutions also complicated the focus control of fine positioning sensors.)

Validation Constraints

Residual Risks without Optical Test

The LDRD team did not have the resources to test a 5 meter lens as a diffraction limited optical element, even if the resources had been sufficient to build a lens worth such a test. Temperature variations in air along the beampath have been mentioned above as the reason for the least costly selection of an optical test site underground at NTS. Even indoor LLNL facilities suitable for testing a portion of the aperture, which might have been sufficiently thermally controlled to stitch together 2 meter wavefront tests, could not accommodate a 200 meter long beampath.

Without testing optical functionality to prove that assembly had been performed adequately, the Eyeglass LDRD team’s fallback position was to make global position measurements at fiducial patterns dispersed across the lens, which would determine whether assembly had placed the optically active patterns in the right places within the 5 meter aperture. Unknown unknowns remain the residual technical risk that purely metrological testing cannot resolve. No such unpredicted effects were encountered in the diffraction-limited indoor testing of a foldable/folded-once 0.75 meter lens in 2001. This earlier success at smaller scale doesn’t not prove such risks don’t exist, because of the

much tighter thermal control and residual strain tolerances that a 5 meter assembly process must solve pose challenges that smaller lenses in temperature controlled rooms would not have detected.

Although the magnitude of unknown risks can't be assessed, non-optical tests can be designed to eliminate those we could imagine. If the position of lens patterns is measured on the same temperature controlled table the lens is built upon, this will not eliminate the chance that unmeasured thermal nonuniformities would degrade optical performance. If the ability to measure global positions is built into the assembly apparatus, it too could have spatially nonuniform errors that metrology wouldn't detect. Both of these risks can be eliminated by rotating a finished lens before putting it back on the assembly table to qualify with global metrology. Because the lens is almost the size of the table, this will not detect radially symmetric unknown errors, although translating the lens by several inches would reduce the possible magnitude of such hypothetical errors (which don't affect optical performance).

Lack of Plane Wave Sources

If a plane wave was available over the full aperture of a 5 meter lens under test, it could shorten the beam path by a factor of four. Without such a lens, smaller apertures illuminated by a plane wave can test the coherence of two or more adjacent panels. That departure from coherence could be computationally extended to the entire lens, although errors would accumulate analogous to making 5 meter measurements with a considerably smaller ruler. Such error accumulation makes partial-aperture optical testing unattractive compared to testing with a point source. Large wavefront errors due to temperature variations, even in a sealed tunnel such as the one located at NTS, leave little room for errors to accumulate. On the other hand, partial aperture testing indoors (especially of lenses with smaller $f\#$'s), could be very effective for removing the possibility of unknown unknowns.

Point sources for testing avoid the cost and complexity of large aperture optics, which would have to be moved across large apertures if they were smaller than the lens under test. Point source testing of smaller apertures appears attractive for an alternative approach to fabricating large, broadband Fresnel optics. The color correction approach embraced and proven by the Eyeglass LDRD requires a second Fresnel optic, much smaller than the primary Eyeglass optic. Building a corrector that matches and cancels the misalignment of patterns in the primary Fresnel lens moves the precision alignment requirement task to a size regime below 1 meter, in which existing precision motion stages with submicron accuracies (made by Kensington or Fox) could suffice.

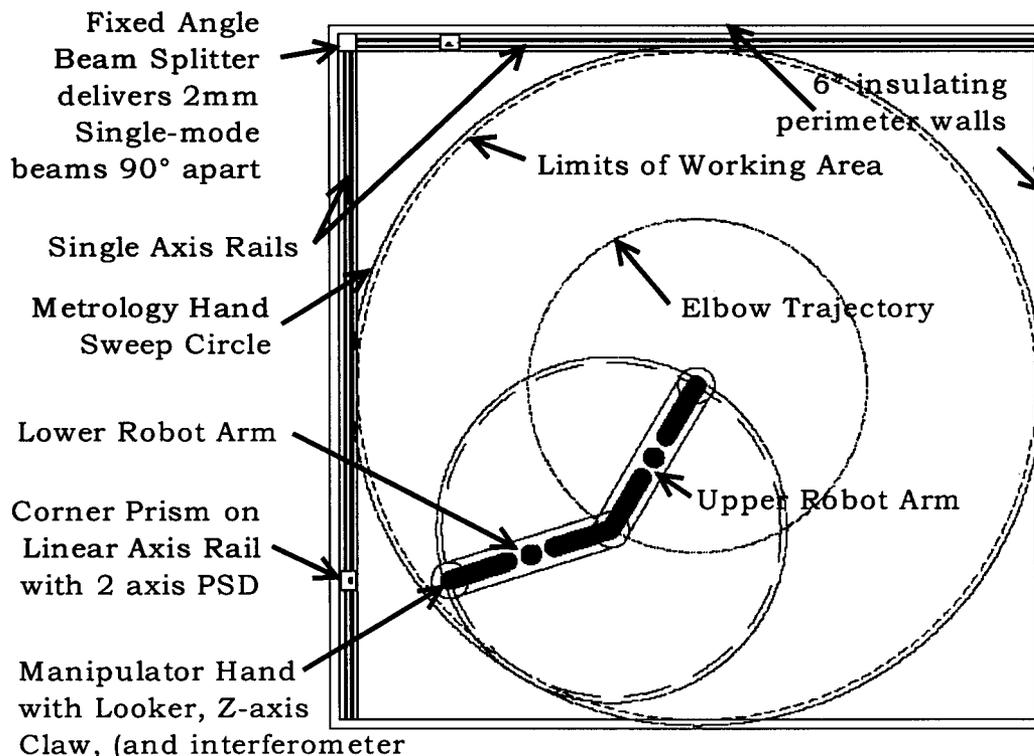
Practicing with sub-aperture point source optical diagnostics that can capture wavefront errors near a focus may provide the evidence that a low accuracy lens assembly alternative that deals with assembly errors in a much smaller corrector could eliminate the necessity of precision alignment across large apertures. This alternative would be particularly valuable for affordable construction of diffraction limited Fresnel lenses with apertures even larger than 5 meters. This alternative approach remains

infeasible in the near term, because the processes required to develop corrector Fresnel lenses (which require much more closely spaced blazing) will be very costly to develop, and many candidate processes would be made considerably more difficult if they had to cancel wavefront errors based on more than one wavefront 'image'.

Global Metrology Solution Postponed

Given LDRD resource limitations, plans for optical qualification of diffraction limited lenses were deferred with the rationale that even a risk-averse potential sponsor could afford this final measure of qualification, if all the other, already known risks had been eliminated in a 5 meter lens LLNL had already built. To eliminate all the risks we already knew of, a full global metrology capability that could locate any microscopically observable feature within the 5 meter lens was designed into the apparatus presented at the February 2001 design review.

The Figure below schematically illustrates a global metrology subsystem design that relies on three axes of submicron-accurate measurement. Measurements with accuracies better than 100 ppb would be provided by an \$18K Zygo interferometer. This state-of-the-art metrology instrument can reach across the air gap in the few inches above the work with a 2 mm diameter, RF modulated beam of collimated laser light. The subsystem that applied it to the Eyeglass assembly apparatus was expected to cost over \$150K to develop, and was deferred indefinitely when FY02 funding for the Eyeglass LDRD Strategic Initiative was cut in half. This proposed global metrology subsystem remains the most cost effective way to prove that assembled 5 meter optics have been assembled without any of the known error sources exceeding the 5 micron radial error tolerance.



The postponed global metrology subsystem uses two of the three interferometer axes to measure the lengths of the two legs of a right triangle, locating a microscope beampath perpendicular to the work placed at the right angle vertex. The third measurement axis of the interferometer actually cost less than nothing (in Zygo's price list of 2001) and has some utility when used to measure a fixed length of the same air inside the temperature-controlled apparatus (to cancel out temperature, pressure, and humidity errors in the speed of light). Besides this third leg, two other beampaths and two inexpensive, 20'-travel motion axes instrumented with PSD's (position-sensitive detectors) accurate to $\sim 1/2$ micron are required to measure X and Y relative to the table.

The nearly-right angle of the measured triangle in this scheme is held constant over temperature and motion by gluing together two of Zygo's beam delivery prisms, in a symmetric configuration immediately adjacent to the fiducial observation microscope's objective. This requires disassembly or re-implementation of Zygo's prisms, which are sold with aluminum enclosures that would produce thermal variations of the angle between the been if not removed before gluing. The actual angle will not be 90 degrees, but can be calibrated to high accuracy with the interferometers and a special target panel made with the same glass processes used to make lens panels.

Instead of costly two-axis global metrology, a form of adequate global metrology was developed sufficient to measure radial excursions of a series of circular fiducial rings built into the lens itself. When the necessity of folding a 9-fold symmetric lens was first considered, a variety of folding mechanisms that could perform the sequence of coordinated motions required were envisioned. One of these was actually tested on a 2/3 scale plastic model, so that folding requirements could be understood. In order to place such a mechanism over the work table, and to lift a completed, folded lens off the table, arrangements were built into the apparatus that could move a specialized crane over the center of the lens at will, before and after assembly operations. If this crane were used to mount a rigid, lightweight arm that spans the table's diameter, which has been balanced and instrumented with a series of tele-microscope beampaths at different radial stations, so that it can be spun up to rotate on a precision air bearing spindle (a \$2,000 item in the Dover catalog, that has not been procured yet), submicron measurements of the radial position (as well as ~ 2 micron accurate measurements of circumferential position) can be made with photodiodes.

Local Metrology Errors Accumulate

Without global metrology functions available during assembly, accuracy requirements for positioning individual panels correctly are tightened to ~ 3 microns. Instead of knowing where a panel is located within the lens, local metrology can only figure out where a panel is located with respect to its immediate neighbors. Local metrology errors add up to predict the errors that global metrology could measure. If these errors are uncorrelated, their magnitude can be RSS'ed (square root of the sum of squares). Otherwise, correlated error magnitudes add, and accuracy requirements would be squeezed below ~ 1.7 microns. Little difficulty and many options accompany sensing

position errors to this level of accuracy, but actually altering position errors of delicate panels without imposing strains that would ruin their optical performance is nontrivial. The mechanical control system that can do this job is nontrivial, and most it has been implemented.

Because the difficult precision manipulation function conflicts spatially with other assembly functions, it makes sense to decouple the sensor and actuator in a precision motion subsystem'. The actuation components are collected in a subassembly known to their development team as the "Tweaker", while the sensor cluster that detects relative position errors between adjacent panels is known as the "Looker". Fiducials built into the light bending diffractive patterns on the top surface of glass panels are observed by the Looker, measurements are combined to compute a least-squared error estimate, and commands are issued to the Tweaker that should zero out one panel's position errors. Multiple iterations could of this closed loop position control subsystem would further reduce position errors, although this fine motion subsystem has been design to achieve at least threefold better accuracy than required in a single iteration. Although all sensors and actuators required to implement this critical subsystem have been implemented and debugged, the computational loop has not been closed, nor have error sources been statistically characterized, so adequate assembly accuracy capability remains unconfirmed.

Coupling with Optics Design

Decoupling Hinges and Patterns

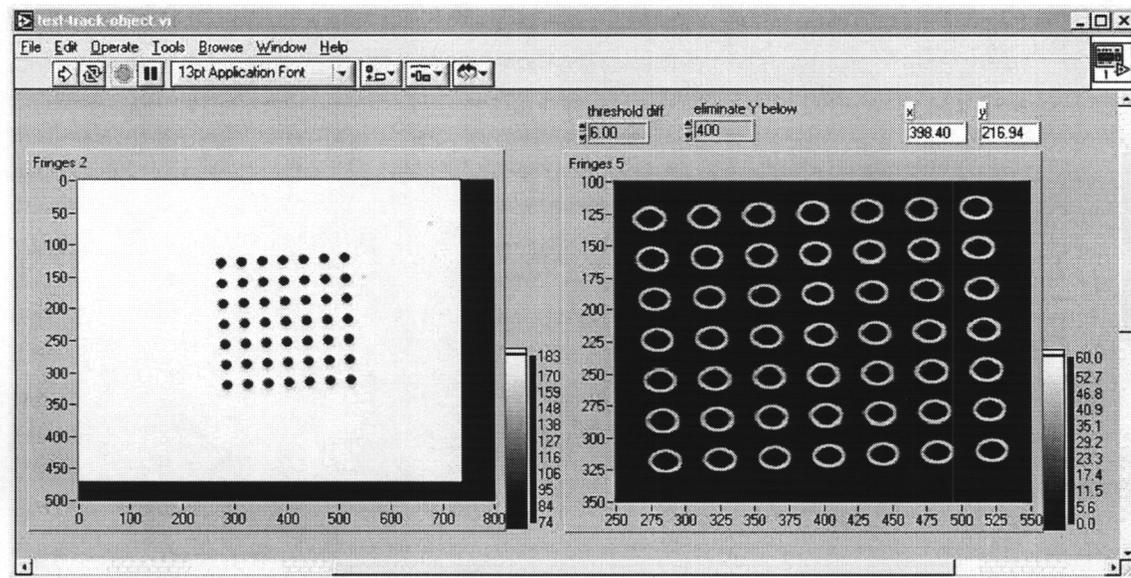
Optical function of an assembled lens does not depend on exactly where the cut edges of glass panels are located; it only depends on where the diffraction patterns are located globally. By building fiducials into those patterns that the Looker can locate, the tight accuracy requirements of assembly need not apply to panel cutting or hinge placement operations. Fiducials and the diffraction patterns they locate can be more than a millimeter away from their intended position on individual panels without consequences.

Panels fabricated to date were cut into their final shape before the patterns were imposed on one surface, using guidelines that are also part of the pattern (in order to minimize the cost of broken panels, and avoid the delay of producing a second batch if any patterned panels were broken in transit to the cutting subcontractor). Those patterns were built with a photolithography process developed at LLNL (in B298, for NIF), which has no commercial equivalent because semiconductors currently have no need for patterning across more than 120 millimeters. That patterning process requires costly photomasks, and the software that generates the masks is thereby the first step in the entire lens fabrication sequence.

Because the Looker design determined what kind of fiducials belonged in which positions on those photolithography masks, mask design produced a chicken-and-egg problem that coupled optics fabrication with assembly subsystems design. Fortunately, the overall configuration of an assembly apparatus was in place by late 2001, along with

a nominal hinge design that could adequately determine the gap widths. Many aspects of the assembly apparatus design could be decoupled from mask design. For example, hinge positions and quantities were only determined a week before assembly began because Fresnel patterns on the panels could be blocked by frames or hinges in any location around the edges of the panels.

Fiducials that would provide the location information for accurate assembly were intrinsically part of both mask design and Looker design. In order to decouple all of the Looker design from mask design, a single microscope beampath was implemented and plugged into image acquisition hardware and software that was in use for other (NIF B298) research. That transitional prototype was pointed at an array of possible fiducial geometries built into a hastily generated 6" mask, and the images acquired were processed with two different software packages to select the best fiducial design. The Figure below shows a typical result of this machine vision prototyping

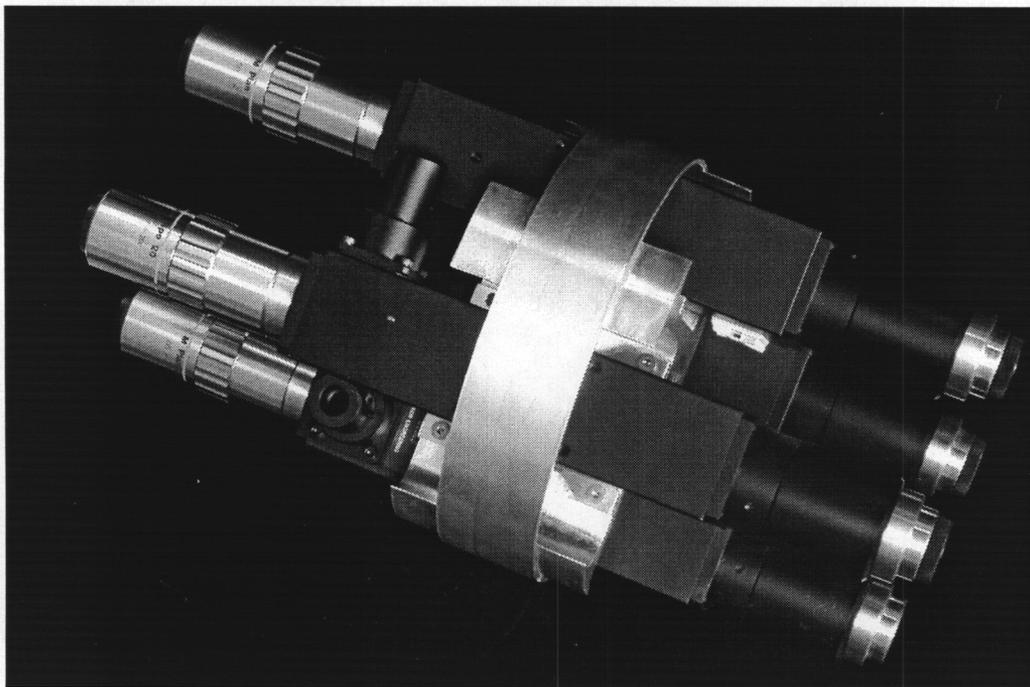


The best fiducial locations were guessed that would accommodate either of the top two contending Looker configurations. One of these contenders used two beampaths built into one instrument, and the other used four. Because the two beampath design could work with fiducials located on the square pattern required by the four beampath design, a square array of fiducials was selected to instrument every gap between panels. Since hinge locations weren't known exactly, and hinges would block the view of fiducials, a second square array was interspersed with the first along the gap.

The spacing in the square array of fiducials was set to an easy multiple of mask pixel size by picking a round number of 75 millimeters. That spacing was nearly midway between an upper constraint set on Looker package diameter and a lower constraint set by individual beampath component diameters. The upper constraint came from the difference between the size of the table whose substructure was already drafted and being procured and the size of the lens whose outermost points should fall within range of the

Looker. The lower constraint was emerging from lengthy shopping discussion with microscope optics experts.

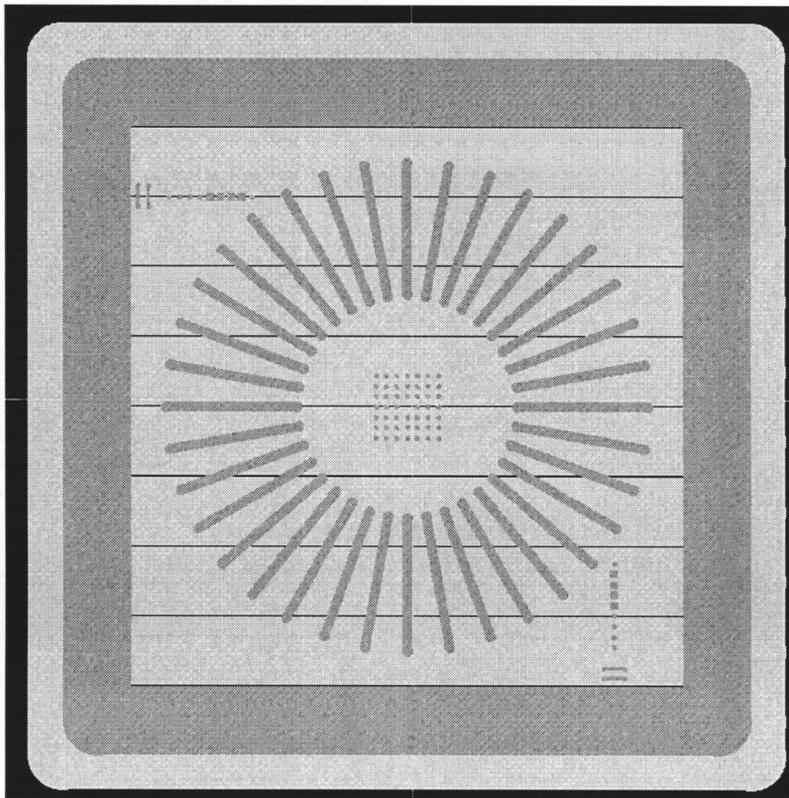
The 75 mm fiducial array spacing superseded an earlier 'guestimate' of 100 mm diameter for the circle that the four beampath design would place its beampaths on (which would have an inscribed square 70.707 on a side). The 100 mm specification came from the then-current best guess of the closest that two Nikon apochromatic microscope beampaths could be squeezed next to one another inside the still-to-be-designed instrument. That instrument is shown in its nearly-completed state in the Figure below, with one expensive beampath removed for qualification in a precision machine vision testbed. Besides algorithms and optics, the single beampath testbed was required long before the Looker was needed to evaluate the adequacy of processing lens patterns and the imaging consequences of several illumination source candidate components when applied to form images of actual processed fiducial pattern.



The Nikon microscope parts were found to be the best alternative for this machine vision application, and the new beampath circle diameter specification slightly increased the diagonal of their square pattern. Looker design then proceeded to match a fiducial pattern already being built into masks and later into panels. Although several manual recognition and alignment features were added to the masks just in case they were needed, a full array of fiducials covering the entire lens (which would turn it into a giant encoder) was forgotten in the Eyeglass team's rush to hold schedule. An array of fiducials covering the entire lens, rather than just lining its gaps, might have been highly effective subsequently to interpret mask production problems that forced the mask maker (who imported LLNL data files) to send the first mask fabricated to Sweden in order to figure out why its patterns were defective a few microns off.

Fiducial Patterns on Glass

The fiducial patterns themselves were designed for both machine and human location determination. The Figure below shows the idealized machine vision fiducial in the form used to check mask software. When actual microfabricated fiducials on lens panels are viewed through a microscope and television camera, their images contain two axes of position information that locates the fiducial relative to the beampath optical axis. The location and orientation of that optical axis isn't known very well with regard to the rest of the assembly apparatus, and it doesn't have to be, since at least one more beampath is rigidly built into the same instrument. It makes sense to maximize the instrumented area of the table and reduce any residual thermal sensitivity of the Looker to place adjacent beampaths as close to each other as possible, but far enough apart to straddle the gap between panels.



The relative locations of multiple beampaths can be calibrated within nanometers by observing test masks calibrated elsewhere. The lateral distances between beampaths can be made insensitive to thermal and vibration errors (within dozens of nanometers) by rugged, symmetric packaging built from a single material with fasteners. The Looker design which we implemented optimizes this insensitivity. Even several degrees C variations experienced during testbed debugging in a poorly thermally conditioned room doesn't significantly corrupt these relative distances, which add in to the machine vision observations of the relative locations between two adjacent panels.

Two beampaths instrumented with machine vision algorithms can measure four simultaneous dimensions, while four beampaths make eight measurements, and three don't convey more panel-to-panel relative location information than two. Four beampaths were ultimately selected in the Looker design (despite doubling the Looker cost to over \$25,000 in Nikon parts) because the extra information they generate precludes a variety of subtle error sources.

The core of the fiducial pattern was designed to minimize the errors from machine vision, also known as image processing, when applied to the task of extracting the fiducial's position within the microscope's field of view. The higher the magnification, the lower the error, as long as the machine recognition pattern's image remains entirely within the field of view. That field of view was chosen to keep the recognition pattern within the bounds predicted by the inaccuracy expected from initial panel placement added to Looker location uncertainty (1000 and 75 microns respectively). A square array of circles occupies the middle of each microscope camera's field of view. This array must remain entirely within the field of view for 1200 micron translations in either direction along the narrower ("vertical") axis.

A blank region surrounding the square array of small circles precludes other patterns from being erroneously interpreted as part of the machine recognition pattern. Magnification choices were limited by the microscope objectives available in the Nikon product line (5X, 10X, 20X, etc.), which is used almost exclusively in most similar automated alignment applications (e.g. for semiconductor manufacturing). Of the three camera image plane formats (1/4", 1/3", 1/2"), the most expensive was calculated to offer few advantages in the Looker error analysis (pixel noise errors are dwarfed by lithography process errors), and the least expensive didn't save enough to offset the expensive, distortion-free field of view it would have discarded.

Initially a checkerboard was considered the optimal machine vision pattern for this lateral location measurement application because it maximizes the number of pixels contributing information to the estimate of lateral position. When the Looker was being designed, however, the process that would be used to fabricate panels was uncertain. Under-developing the photoresist in a checkerboard pattern prevented the use of image processing algorithms that segment the digitized camera image. Square edges were also considered more prone to systematic errors that do not get averaged to lower variance by adding uncorrelated noise from the many pixels that capture location information (which comes only from edges in the image).

These systematic errors that do not average to lower variance are quite likely to arise from the uncollimated light that has been used to date when large masks have been contact printed onto photoresist. Given illumination that almost saturates the camera to maximize signal strength, the Looker error analysis was dominated by edge location noise in the image due to the random scalloping of the actual edge topography microfabricated into or onto the panel. Such systematic, process-induced errors were expected to be minimized by a chromium etch process (that has yet to be attempted on panels of the size required for the 5 meter lens due to its cost). This kind of error would

have been maximized by glass etching process (that was not used on the 5 meter lens panels due to its additional cost compared to simply leaving photoresist in place to form the simplest, 2-level blaze). All error sources combined to an estimated variance below 250 nanometers.

Surrounding the blank area that surrounds the array of circles, a starburst shaped pattern completes the fiducial design to allow both human and machine observers to find the fiducials. Machine recognition of these starburst patterns would only be necessary if initial panel location errors exceed their expected values by up to 2.5 millimeters. Beyond that level of location uncertainty, human observers can see the starburst well against the diffracting field of Fresnel grooves, allowing them to steer a microscope into the vicinity. Otherwise, without a correctly placed panel, it is very easy to get lost and very hard to find the fiducials looking only through a microscope.

The Starburst features have turned out to be very handy for debugging machine vision without motion stages. The starburst patterns may provide the only hope of locating severely misplaced panels in an assembly apparatus where human operators can't look directly. Other fiducials derived from an older generation of human-aligned semiconductor patterning equipment that rely on an inverse-patterned mask movable in two axes close above the imaging target (to "put the peg in the slot") were placed on the mask to enable advanced manual 'tweaking'. Other patterns that microscopically resemble long railroad tracks cross the entire lens in various directions succeeded in assisting human operators to coarsely align panels, but were never used for their intended function as vernier encoders.

Design of an Assembly Station

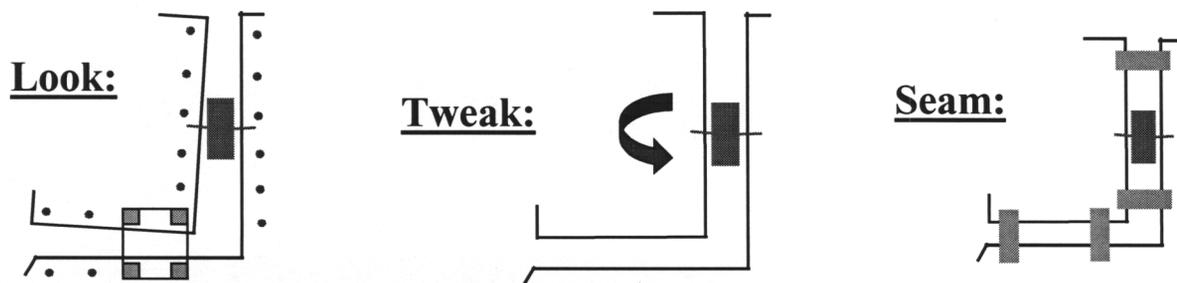
Panel-by-Panel Assembly Approach

Our first step in assembling the lens will be the coarse-placement of all the lens panels. This operation does not require high precision, simply that fiducials along the borders between neighboring panels be close enough together so that their precise separations can be read by our local-metrology tool. These loose precision requirements permit us to tolerate the arm's inherent accuracy (several mils) and to forego tight thermal control during this operation. Panels are picked-up, held, and released by a top-gripping placement-tool. This tool basically consists of a vertical motion stage (for lift-up and set-down) and a vacuum-chuck (for gripping the panel).

After all of the lens panels have been placed near their proper positions, we are ready to align and seam them. This is done, one border at a time, by specialized panel-joiner tools. Since this stage of the assembly demands high-precision, it calls upon our thermal control system. Accordingly, before beginning these attachment operations, we must bring the working-surface and lens-panels under active thermal control.

There are 3 panel-joiner tools; a Looker, a Tweaker, and a GluBot; we use the same tools for all panel-to-panel joints. The sequence of operations used to attach a new panel

to the already assembled portion of the lens begins with delivery of the Tweaker. This is delivered to the border between the new panel and the previously assembled ones, released from the arm and attached to both panels, holding the new panel in place. Then the Looker is delivered by the arm, and is carried along the border(s) of the new panel; where it inspects, 4 fiducials at a time, the alignment of the new panel. After all the fiducial sets are examined, a least-squares filter is used to determine the misalignment (x , y , and ϕ) of the panel. The Tweaker is then commanded to perform the appropriate x , y , and ϕ shifts to the new panel. The Looker then reinspects the border fiducials and, if necessary, the Tweaker can perform a new adjustment. After proper alignment is reached, the GluBot attaches metal foil seams along the interface, firmly attaching the new panel to the previously assembled ones. This panel-by-panel alignment/attachment process continues until the lens is fully assembled. The schematics below, illustrate this assembly process.



Coarse and Fine Panel Manipulation

Many aspects of the 5 meter lens assembly mission are similar to other technical solutions already proven, yet differ by an order of magnitude in one physical parameter or another. The large travel required basically exhausts the prior art in precision motion stages. If just 3 meter travel were required, a few vendors could attempt stages with pricetags that would approach \$100,000 per axis, and delivery times that would approach one year. Those stages would be one to two orders of magnitude heavier than the solution the Eyeglass team designed and procured. Although the weight of a more conventional apparatus by itself would not be a problem, thermal inertia would likely have required an unprecedented thermal control scheme (perhaps with kilowatts of authority) in order to keep settling time below a week.

Rotary rather than linear motions can easily be stretched to control precise travel over longer distances. Air bearing spindles capable of supporting radial arms that weigh several hundred pounds were available, but with pricetags and delivery times similar to the 3 meter travel linear stages. Even the best roller bearings ruined any rotational motion approach to 100 ppb position accuracy that tried to avoid air bearings. Magnetic and jewel rotational bearings were incapable of handling the loads in even the lightest apparatus. One rotary encoder with sufficient precision to implement a 100 ppb rotational motion control degree of freedom was procured for \$7,000, but buying at least one more of such an expensive sensor would have taken its normal delivery time of 13 weeks, and a better approach was found.

Partitioning the lens panel manipulation task into two subsystems resulted in an affordable design that could be procured and debugged within the time and budget available to the Eyeglass LDRD effort. In order to make sure that coarse motion would not put panels in place with too much error for fine motion to correct, almost an order of magnitude of overlap was specified into this partition of motion control functionality. Instead of 7 orders of magnitude, routinely available mechanisms deliver nearly 5 orders of magnitude in dynamic precision of coarse motion control (covering 5 meter travel with an estimated 75 micron accuracy), then hand off to fine motion controlled by the Looker and Tweaker subassemblies.

Fine motion was driven by New Focus micropositioners intended for laser alignment and microscope translation. These linear actuators make repeatable 2 nm steps along a single axis. Although their piezoelectric steps can accumulate to travel over ½” (and 4” versions are available), robust Looker and Tweaker designs limit their travel to roughly a millimeter. As discussed above, such limitations are important for correct machine vision in the Looker and to preclude getting entirely lost so that microscopes see only a seemingly unbroken striped pattern of diffraction grating.

Even with their travel limited, the chosen New Focus actuators takes more than a minute to traverse the full 1 mm range. The miniscule 2 nm step was far below the Looker noise floor, it made sense to use smooth elastic hinges to amplify their output motion five fold. An elastic motion was developed that performs this amplification without the ‘crunchiness’ of friction, while constraining the 5 other rigid body degrees of freedom that the actor itself leaves free. One of these little motion stages was built and debugged, and another three were constructed, which provide ~3.5 mm of limited travel with nearly constant 5.3-5.5 to 1 amplification to deliver more than that order of magnitude overlap in motion authority, which may well be necessary in case the coarse motion subsystem’s accuracy is considerably worse than predicted.

This partitioning allows coarse motion to occur without the thermal control system engaged, and hence without an insulating shroud around the assembly apparatus. Any travel that could go far enough to put a panel in the wrong place can be observed without the sort of teleoperation necessary to maintain 30 millidegree thermal control. This expedient also disposes of the requirement for some sort of loading station, since technicians can place panels on the assembly table themselves. Such manual panel placement can also be accomplished with very low accuracy (i.e. 1/8”), which a human can correct by moving a low magnification beampath over starburst fiducials he can see with his/her naked eyes.

Removable insulating covers could have been a design feature of a more expensive apparatus with a single, high dynamic precision motion stage, but it would not be a good idea to let people bump into such a mechanism. Removable covers would have been necessary whatever design was adopted for motion control, simply to get the assembled lens out of its fabrication location, but a more dynamically precise machine would probably have retained its precision better over time with the smallest possible access

opening, as compared to the configuration adopted which exposes the work table in almost all directions.

Suspension Options and the Fine Motion ‘Tweaker’

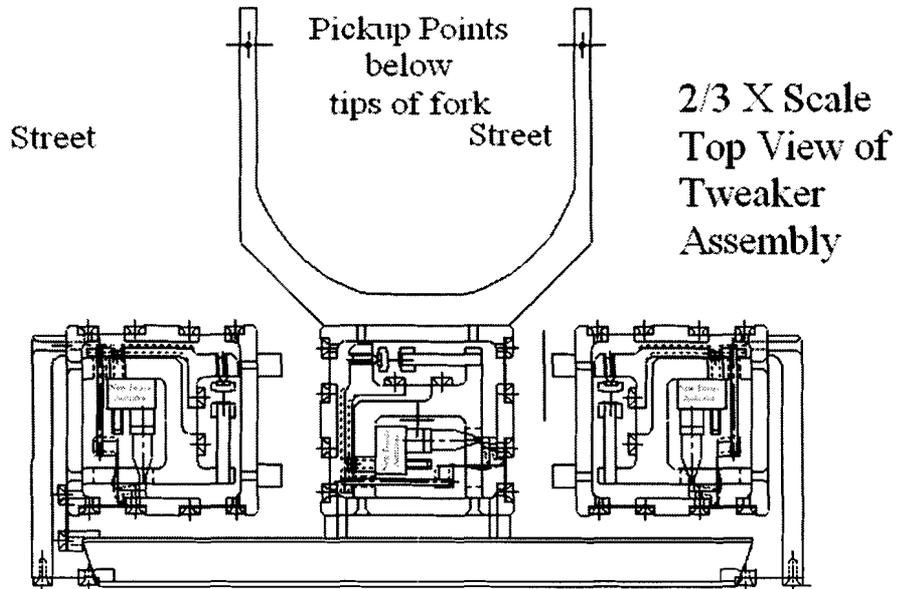
Fine motion was designed to deliver at least three orders of dynamic precision, from a ~1 millimeter travel down to a submicron accuracy. In order to actuate fine motion, the ‘crunchiness’ of ordinary bearings must be eliminated. Air bearings are the obvious candidates, although they generally provide more freedom of motion than fine motion requires, and will need other components to prevent panels from “running away”. Limited travel without ‘crunching’, and at least 6 orders of magnitude of repeatability is available from elastic suspensions similar to those used in some electron microscope stages.

Proven elastic suspensions can be used in the Tweaker, and in spatula designs that interpose a mechanical translation stage between the panels and table. But a population of nearly a hundred spatulas, each of which must have a shape that is specific to a particular lens design, (which would have required four shape variants for the design that was built), turned out to be an extremely costly suspension option. Velvet fabrics have a small enough air gap to enable satisfactory thermal control with 6” insulation around the apparatus, although they seem likely to impose out-of-plane random height errors that exceed the allowable slope tolerance discussed above. Custom molded silicon sheets with little stubby conical fingers may be a likely elastic panel suspension alternative, and would be the high-tech equivalent velvet whose compliance could be engineered in three directions. Although samples of possible relevant cast sheets were obtained, the relatively affordable development of an engineered ‘velvet’ still costs far more than a visit to the local fabric store.

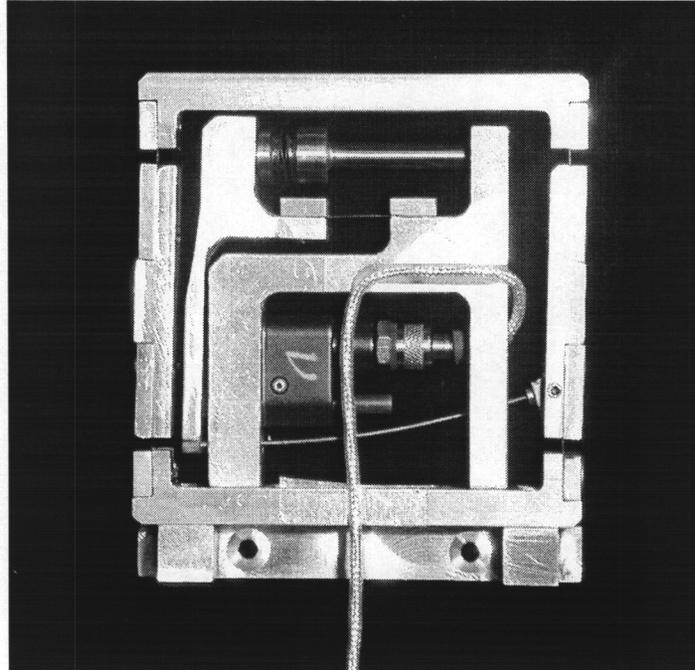
The baseline suspension option is ball bearings just 1 mm in diameter. To keep them from wandering around the table, a custom gasket needs to be cut in four different shapes to provide an array of little jails within which the bearings can roll. Without such a gasket, which is easy for the coarse motion system to drag harmless across the table, ball “run away” with the slightest tilt. A 3 mm circular jail would allow a 1 mm ball to translate a panel in any direction from its original position. Initial experiments showed that a ball bearing is needed every 6” or so to support the .032” thick glass sheets chosen to implement the first lens. More balls would be needed to support thinner lenses, and below some undetermined glass thickness, a ball loader would be necessary to save technician time loading the gaskets with balls.

The considerably harder problem of making sure that the balls are all placed in the middle of their jails before any tweaking has been solved in principle by using the vibration isolation ‘legs’ of the apparatus to deliberately tilt the machine. Departures from planarity must be constrained to build a diffraction limited lens, but that plane need not be horizontal. With a coefficient of rolling friction below .005 (measured by tilting experiments with stainless steel and sapphire balls on polished aluminum), balls are guaranteed to roll when tilted by 5 milliradians, which is within the excursion of “two

legs up, two down". Sub-milliradian tilts keep the balls from rolling with static friction, so panels delivered when all balls have been rolled to the East can be shifted to the centers of their jails by a deliberate 1 mm coarse motion translation to the West. Static friction of the gaskets on the table has been shown to keep them in place during these limited tilts and ball excursions.



The Tweaker itself is an unprecedented mechanism with three actuated degrees of freedom. Another three complaint degrees of freedom are engineered to make contact with the panel being tweaked, touching the glass of the finely manipulated panel through two soft silicone pads on the end of a long, Y-shaped leaf spring. The Figure above shows a top view of the Tweaker subassembly, which includes three nearly square amplified actuators. Its mass is designed (and on track) to be below 1 pound, which is supported on the adjacent panel across the gap. A large 8x15" soft, clean 'foot' distributes that weight over many square inches of suspension, and would effectively clamp the panel that supports it to the table with either ball bearing or air suspension options. Below we show one of the three Tweaker stages ready for debugging.



The Tweaker's foot is intended to rest on the portion of the lens that has already been assembled, or on the first panel before any panels are joined. The Tweaker is picked up and put down in its intended location by the same robot arm that deliver coarse motion, which also carries the Looker to arbitrary positions and orientations around the lens. The Tweaker subassembly has been deliberately design to be less than 0.6" thick from bottom of foot to top of actuators, allowing it to operate entirely within the relatively long working distance of the Looker's microscope objectives. This will prevent the arm for crashing into it despite software errors and operator inattention. Because the power demands of the three New Focus actuators is so low, four thin wires can power it on a ~36" arm that flexes in the plane of the table, but remains constrained vertically in the microscope working distance gap.

Design of a Custom 4-Beampath 'Looker'

Although the initial rationale motivating Looker design was discussed in the Decoupling and Fiducial subsections above, its preferred form was the result of 4 months of collaboration with silicon valley microscope experts. Technical Instruments is the local firm representing most microscope vendors (except Zeiss and American Optical) to firms in silicon valley. The majority of microscope applications that call for high performance optics are very different from Eyeglass Looker requirements. Human observers don't notice a few percent of distortion, as long as straight lines appear straight. Machine vision calls for very flat fields in the target and image (CCD face), as well as low distortion and relatively wide fields of view.

Dimensional stability of microscope hardware is very important when the beampath is built into equipment that moves more often and more violently than its human operators generally do, and that may need to hold its pointing calibration over years of

service. Because Technical Instruments was supplying the understanding that could minimize Looker errors and development risks that LLNL staff hadn't thought about, the high price of the solutions they supplied was still a bargain given the higher cost if LLNL effort.

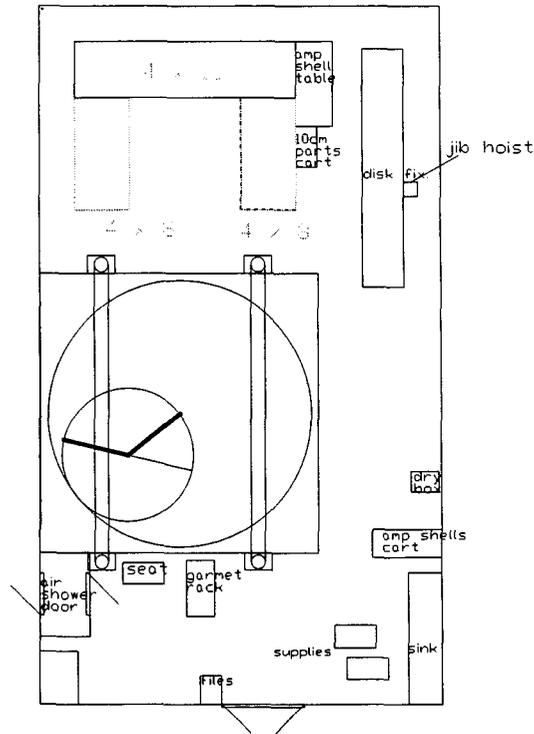
Considerable discussion of expensive and relatively unproven autofocus equipment resulted in the conclusion that the Looker would have to provide this function with a three degree of freedom New Focus actuated mounting that could lift and tilt all four beampaths at once, which had no viable commercial equivalent. Implementing autofocus in the LLNL apparatus called for some development, although the algorithms were in hand to determine best focus by image processing, but offers the significant fringe benefit that the Looker can be used to determine actual vertical positions of fiducials and hence of panels so that out-of-plane excursions are detectable.

The Nikon microscope body and objectives that the Looker design was built around solves machine vision problems similar to but more severe than the majority of recent automated alignment equipment being built in and for silicon valley. Although the microscope objectives chosen (apochromatic, long working distance, 20X) would have exceeded all the optical performance specifications LLNL staff knew of, similar applications to ours had benefited from these excesses. Our design turned out to benefit from long working distance in avoiding potential collisions between Looker and Tweaker, and from the low distortion that eliminated a potential 300 nm error term.

The high pricetag of the Nikon solution could have been reduced 40% by building our own microscope bodies, but LLNL machine shops couldn't align the axis of a tap wrenches with enough confidence to meet the Lookers specification for 100 microradian parallelism between beampaths (which limits an error term coupling lateral image position errors to focus errors). The Nikon body provided most of an illuminator beampath and eliminated the analysis LLNL would otherwise have to perform on thermal flexure (which causes image position errors on the camera face when the microscope body curves). The LLNL design completed this beampath with an inexpensive super-bright LED and a \$500, 1/3" Sony CCD camera that saved \$1200 per beampath compared to Technical Instruments' solution, which did not allow beampaths to be packaged as closely.

Packaging of the Looker was nontrivial because it affected many other design dimensions throughout the apparatus. It is the heaviest part of the coarse motion robot arm's payload. Its diameter reduces the table working area inside the apparatus' working envelope, which was not a free variable because the facility where the apparatus was being constructed had to continue to serve as a cleanroom capable of refurbishing Janus laser modules. The floorplan of that facility shown below was the result of four iterations of negotiation and considerable cleverness, which left a maximum East-West dimension of the apparatus constrained within inches of 20'.

B-174 R-1320 room layout



The vertical extent of the Looker determined the height of the air layer underneath the removable insulating roof of the apparatus, since it was taller than the structural element of a robot arm segment. The Nikon bodies with shortened illuminators and cameras stripped of their enclosures barely fit into the initial 100 mm beampath circle diameter specification. That specification minimized the looker diameter given hinge and frame widths, which earlier had been used to determine the minimum fiducial spacing across the panel-to-panel gaps. The revision specification thereafter matches the 1.414 X 75 mm specification built into the lens panels via the masks.

Four beampaths won out over the less expensive two beampath solution due to possible tilt sensitivity errors that could not be eliminated in a two beampath design. As a two beampath design was moved along a gap between panels, it would have to rotate with respect to the panels to stitch together the diagonals in a square array of fiducials. Otherwise a two beampath design would be insensitive to along-the-gap relative position information (one of the quantities the Looker must measure to inform successful tweaking). Because illumination of fiducial patterns with half a micron to several microns of surface topography produces systematic position errors in fiducial locations when that illumination isn't axially symmetric and the target is tilted, it made sense not to vary the target tilt by rotating Looker beampaths with respect to the panels.

Four beampaths with simultaneous image capture can also make relative panel position measurements that are totally insensitive to the motion of the Looker, whereas two beampaths must ensure that the Looker rotates about an axis that doesn't move with respect to the work. Although small limit cycles could be suppressed in the coarse

motion control system that positions the body of the Looker, it seemed prudent to eliminate any additional motion control system development at the cost of an additional \$12,000. When four measurements of X and Y position errors (of fiducials with respect to the camera fields of view) are made simultaneously, these can be processed into the three relative degrees of freedom between the two panels (that must be found to close the loop on fine motion), the three rigid body degrees of freedom of the Looker itself (in the plane, which are irrelevant), and two intra-panel distances that serve as a measure of the fiducial quality (and can thereby detect processing defects in the fiducials).

Thermal Control Options Downselect

If the entire lens could be built at one temperature, thermal expansion wouldn't degrade its diffraction limited performance, just shift its focus very slightly. Because even the fastest automated assembly might take two days, any lack of thermal stability is equivalent to a temperature nonuniformity. (Human supervision makes sense to call a halt in case of any obvious malfunction, which precludes running all night in case the humans are lulled by repetitive movements on the television and computer screens that provide their only window on the apparatus). Observation of YSI thermistors with state of the art, low-drift amplification suggests a temperature measurement noise floor can be achieved below 300 microdegrees.

Covering the table with an array of such low-drift sensors would make a very expensive distributed thermal control system, costing perhaps \$300 per instrument. The alternative distributed sensor approach would place a single sensor in a vacuum insulated enclosure behind a germanium far-infrared tele-microscope on the end of the coarse motion robot arm, which could carry this remote temperature sensor to observe any spot the robot can reach on the assembly table. Lateral thermal errors can be much more short lived compared to the vertical errors whose settling time depends on thick external insulation, so robot-arm-mounted distributed temperature sensing would likely be inadequate for distributed real time temperature control. Without measurements as frequent or more frequent than errors could arise or propagate, a distributed temperature control system was problematic.

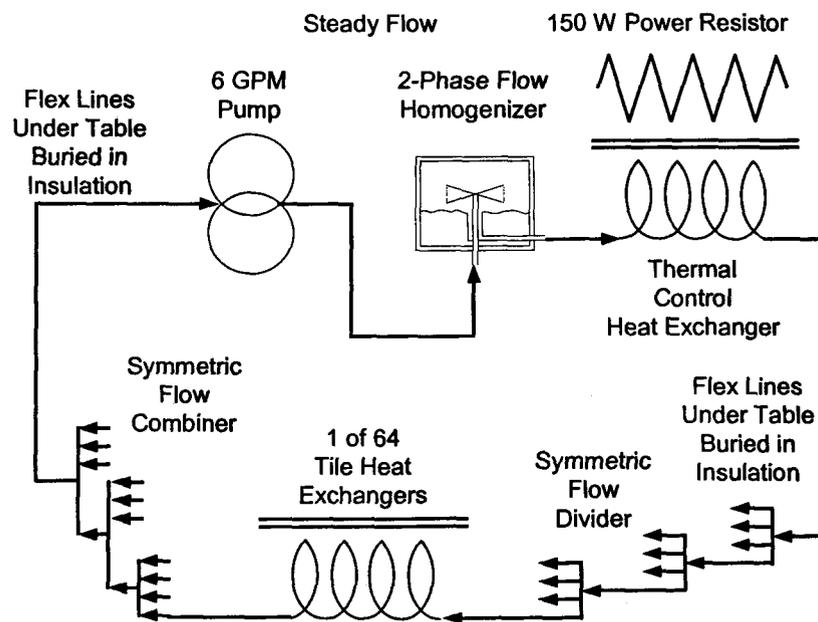
In comparison, a centralized temperature control system that imposes a stable temperature at a single sensor has much lower complexity than a distributed system, and can be built from available components. A chiller rejects the heat from resistors controlled by feedback derived from single point, low-drift-amplified thermistors. The control of frequency doubling crystals uses this scheme to routinely keep a single point within a few millidegrees of a set point that doesn't drift. In the Eyeglass assembly apparatus, four such control points are designed in. The most important of these regulates the temperature of a heat exchange fluid that circulates beneath the table, at the point immediately before it enters heat exchangers that spread the fluid's temperature across the table with adequately low nonuniformity.

The other three points where thermal control is required accompany each of three motors in the coarse motion robot arm, which could dissipate as much as 1200 watts per

motor when driven at peak torque. The motors' heat dissipation is rendered almost constant by driving a resistor in series with the fluid inlet of each motor's cooling jacket to deliver a computed estimate of that constant minus the power being supplied to the motor.

Dissipation in television cameras, microscope illuminators, temperature sensor amplifiers, and Tweaker power circuitry is already constant, and small enough that an extra few inches of the copper tubing that jackets the motors will balance it out. Microscope illuminator LEDs are pulsed for only the duration of a single video frame, so their contribution remains three orders of magnitude below the total thermal noise. The robot arm with its potential constant heat sources must be balanced to impart zero net heat to the air surrounding it, otherwise it will cause spatial variations in temperature due to its movement.

The constant temperature fluid entering multiple heat exchanges beneath the table exits those exchangers with a temperature change due to random heat inputs. The pattern that that flow must take is illustrated schematically in the Figure below. Sufficient heat capacity must flow beneath the table to limit these random fluid temperature changes to a fraction of the allowable random temperature fluctuation. In the worst case, some of the exchangers have exit temperatures at opposite ends of the random variation, resulting in spatial thermal variations that could be constant in time. To reduce the likelihood of this worst case, heat exchanger flow is directed parallel to the airflow in the surrounding clean room, so that spatial variations (due to the cycling of room air between its air conditioner's high and low temperature set points) over the insulating roof affect all heat exchangers almost identically.



The other random heating effect analyzed in its worst case for thermal control design is spatially invariant, but time varying (as the limit cycling of the room's air condition

and the variations of floor temperature have been observed to be). Averaged to be spatially constant, this worst case thermal error is sized by the ratio of thermal resistances between the fluid and the lens and between the lens and the random environmental temperature (air or floor). This error term sizes insulation requirements, limits the air gap allowable between lens and table, and distributes the heat exchangers with sufficiently small lateral spacing to restrict lateral conductivity nonuniformity.

Of the many fluids and heat exchanger alternatives studied, water mixed with ethylene glycol flowing through custom aluminum extrusions won the design tradeoff that included low development risk, high conductivity between fluid and table, short thermal settling time, availability of components, ease of installation, interference with other subsystems, and risk to damage the apparatus. Besides the Kovar and glass heat exchangers mentioned above, whose utility disappeared for 'free' assembly, various air and liquid options were traded off. The air options posed no risk of leakage, but occupied considerable volume beneath the table that liquids left unobstructed. The air distribution tree would have blocked the only possible volume sufficient for storage of the apparatus' bulky insulating roof. Vibration due to turbulent air flow might pose a nontrivial risk to planarity of the leveled table and integrity of connections (electricity, vacuum, and cooling fluids) over the long term.

Liquid systems had better thermal conductivities and equilibrated faster. Their smooth laminar flow and quiet, available pumps demanded an unusual mixer (a lawn sprinkler inside a pressurizable canister that sprays the liquid through air) to make sure that the random variations in outlet temperature from one heat exchanger don't get back into the same exchanger despite the temperature control loop which keeps only the average inlet fluid temperature constant. The availability of custom aluminum extrusions and the sealing of them against a maximum pump pressure of 60 psi (that could be encountered if flow were blocked) turned out to be nontrivial, although each issue was solved with less than two man weeks of effort and a month's delay.

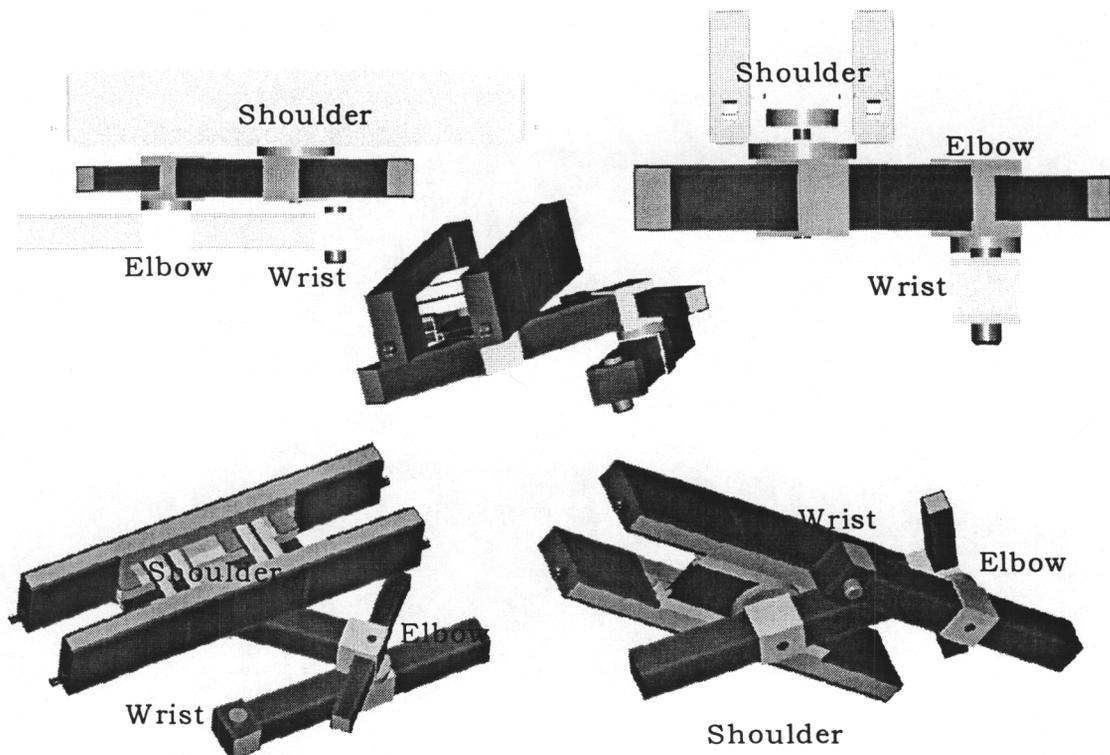
Robot Arm to Transport and Join Panels

The coarse motion system whose precision requirements were degraded into the vicinity of 75 microns (with a factor of three in margin before the job of fine motion control becomes much more difficult) also considered numerous options. The potential for collisions in the working volume during debugging, by faulty human command, or due to software error, made two motion systems operating in the same volume a bad idea. The functions of the Looker are improved by looking a numerous fiducial 'quads' along more than one edge of a panel, if more than one adjacent panel is already joined to the lens. But tweaking by pushing the suspension-mounted from a single coarse motion platform wasn't going to work because of the unknown and relatively inaccurate position stability of that platform.

A compact Tweaker that by pushes against the table near the panel being moved offers stiff dynamics anchored to a position fixed relative to the portion of the lens already assembled. This is a far superior solution, given a Tweaker that could operate at

a location independent of coarse motion being applied to the Looker. The solution eventually adopted was even slightly better, insuring that the lens portion already built was clamped to the table with the weight of the Tweaker. Unlike previous, more complicated contenders, this solution calls for a Tweaker that is not autonomous and can't collide with a robot arm that puts it in place, then departs to steer the Looker over every relevant fiducial 'quad', and later returns to pick the Tweaker back up.

The robot arm that travels over the table in this solution to coarse motion delivers panels to their initial position with one 'hand', then changes hands to deploy three separate instruments sufficient to precisely join the panels. Precision motion must occur after thermal equilibration and hand changes, and includes a third custom instrument not mentioned above, called the "Gluebot". Both Gluebot and Looker are built into the arm, whereas the Tweaker and panel pickup hand come and go off locations just above the table surface. In order to arrive at and depart from the table in a gentle, controllable way, another New Focus fine motion actuator combines (by adding displacements in series) with three pneumatic cylinders (that can reach just four screwdriver-adjustable positions in second rather than minutes), a 2 pounds-force force sensor, and a pneumatic claw (developed to load lathes automatically) to furnish one vertical degree of freedom and a binary clamping function that can hold a wide range of diameters (from 1 to 2") securely.



The robot arm itself has three degrees of freedom, the full range of motion for rigid bodies in a plane. Several views of the robot arm rendered by the ProEngineer drafting package are shown in the Figure above. Two of these degrees of freedom sweep this arm

to arbitrary position within a disc shaped working volume. These angular motions are powered with the maximum power conveniently available from the most capable motion control product (a controller available from Delta Tau). A third degree of freedom powered by the same motor and controller rotates the hand, including the Looker, so that arbitrary positions and angles are accessible.

The three degrees of freedom on the robot arm, which are driven by the Delta Tau controller, which in turn is commanded by a PC, are all rotational degrees of freedom that have a finite range of angular travel. Both 'shoulder' and 'elbow' joints move through 400 degrees to avoid any gap in reachability. The elbow is not restricted to a range closer to 180, although that choice could still reach arbitrary positions, because some conceivable operations might require the ability to chose between a clockwise and counterclockwise approach to a final position. One of these conditions I likely when the Gluebot is called up to apply hinges directly adjacent to the Looker. The 'wrist' joint which rotates just the hand must move through slightly more than 180 degrees, despite the 4-fold symmetry of the Looker, in order to cross-calibrate the distances between non-adjacent pairs of Looker beampaths, to deliver the Tweaker above arbitrary inter-panel gaps directions (from either side, thanks to the elbow), and to place the grasper which clamps the panel holding hand off center from the wrist in arbitrary locations.

Many other geometries were considered, including the X-Y gantries that come to mind first when robotic solutions are considers to traverse a rectangular table. However crossed roller bearings with adequate precision and stiffness were available for rotary motion, while neither linear bearings nor stiff rotary-to-linear drive components (such as ball screws, acme screws, or racks) were available in 5 meter travel. An exceptionally stiff drive train was designed that pushes a ball-screw-driven blade against an optimal diameter capstan to drive the two low-mechanical-impedance degrees of freedom (shoulder and elbow), while the wrist is driven through a convention drum-to-drum capstan. Capstan drives are the only alternative to unavailable custom hydraulics that can avoid the 'crunchy' rolling contact of gear teeth (even in harmonic drives) that would otherwise ruin coarse motion accuracy. Spreadsheet design were used to mechanically impedance match a single style of motor to the predicted angular inertias suspended form each joint, delivering minimum slewing times.

Structure Provides Support and Isolation

Considerable complexity was required from large structural elements that could not be procured out of any catalog. Structural elements roughly 20 and 30 feet long were designed, drafted, procured, received, and assembled which could handle the weight and size of the apparatus while providing it with more than the requisite stiffness. Although the assembly apparatus is only slightly larger than the first 5 meter lens it was designed to build, it is far lighter than conventional machines of its size. The parts to be assembled into lenses weigh much less than the table they are built on, at most 70 pounds of glass and hinge alloy and aluminum frames compared to nearly 400 pounds of honeycomb table surface, 700 pounds of heat exchanger, and nearly 4000 pounds of substructure.

The custom steel substructure was designed to provide the stiffness in bending that $\frac{1}{2}$ " thick honeycomb can't. That substructure mounts posts with provisions for leveling the table top, connects to four vibration isolation legs that can hold up to 12,000 pounds before bottoming out, and mounts massive L-brackets that support columns extending towards the ceiling. These steel components were designed in two batches, and procured from March Metalfab (who won a competitive bid) for a total cost of roughly \$50,000. (Each structural order was received in roughly 5 weeks after ordering.)

Although finite elements were not used for this steel design, stiffness analyses based on the beam equation found that its steel box beam sections were considerably stiffer and stronger than the necessary. There was no point in reducing weight, however, because that would have reduced the vibration isolating dynamics of the supporting Melles Griot laser table legs. Those legs had relatively low load capacity and were scavenged off a salvaged optical table.

Atop the columns, a much lighter aluminum structure was designed, drafted, procured, and partially fabricated. This superstructure was designed to mount the robot arm, which also holds up the insulating roof when it is not being stored under the table. Aluminum structures adapted to lifting moving loads stand above the robot's shoulder joint, which can roll North-South along rails built into the 30' long beams of this superstructure. Two trolleys roll along these rails and cannot pass one another, although both can be dismounted from the North end without risking heavy components dropping onto the table. One of the trolleys (the South) is 24" wide and mounts the robot arm. The other is 9" wide and mounts a crane sufficient to install and remove the 3 pounds-per-cubic-foot roof, or a folding mechanism, or a global metrology 'flycutter' arm on an air bearing. The superstructure is best illustrated with fully dimensioned and detailed ProEngineer solid models shown below.



Both trolleys are driven North-and South independently along the superstructure rails by manually operated winches of the sort found on racing sailboats. Another pair of winches feed through the beams and trolleys to provide 1500 pounds of manually controllable hoist functionality when either trolley is clamped to the beams. This allows the entire robot, roof segments, a folding mechanism, a flycutter, or the lens itself to be lifted safely off the table and over its edge to the North. Installations and service requirements were envisioned that would allow the robot arm to be worked on without any risk to the heat exchanger and table surface (incapable of supporting its weight), or assistance for LLNL Riggers.

Current weight calculations make the robot the heaviest load on trolley or beams at 600 pounds, while the superstructure itself is designed, and its beam structure has been tested, to safely mount 2000 pound rolling loads. The rolling trucks that slide along the rails have very little accuracy requirements, but should operate without shedding particles. They were selected in a relatively inexpensive a grade suitable for food processing, rather than much more expensive grades that may not be available in variants that can travel 30' for clean rooms or machine tools. Because these rolling trucks were specifically not rated to carry humans or have humans underneath them, structural details have been built into beams and trolleys that prevent trolleys from descending more than $\frac{1}{4}$ ".

A lighter grade of aluminum beam with construction details somewhat similar to the superstructure was developed through a first generation to implement the segments of the robot arm. A much heavier grade of steel structure and a happy coincidence combined to complete eyeglass structural subsystems. Earthquake safety for this large apparatus called for Class II tolerance (no humans at direct risk from a collapse) to 0.45 gees of lateral ground motion. The lightweight superstructure helped reduce earthquake loads, particularly the bending moments due to lateral acceleration from mass close to the ceiling. But roughly 13,000 pounds-force of could still slam sideways in the forecast worst case of Class II earthquake.

A separate set of heavy earthquake legs were designed, drafted, procured, redesigned, remanufactured, and then adequately installed which could stop the apparatus from moving laterally or vertically during an earthquake, yet did not touch it during normal operation. These addition, extremely strong in bending legs preserve the vibration isolation functionality of support atop the Melles Griot legs by remaining out of the load path unless an earthquake slams the substructure into flanges that surround the substructure in all six directions.

In between the first and second earthquake leg design, at least four designs based on the weight of lead, then granite, and then concrete were attempted. These transitory earthquake structure designs covered significant portions of the floor under the apparatus to shed the large bending moments that threatened to break the "duck's foot" webs of the first earthquake leg design, and the 5" thick concrete floor. Earthquake bending loads on the legs clearly exceeded the bending strength of this paltry floor, which derived little

benefit from rebar in its middle. Because this apparatus was being built a clean room performing a vital role for 15 other labs worth of laser experiments, cutting holes through the floor (even inside a tent) was not an option.

It took months to find the blueprints showing the construction of the correct portion of B174. At the same time, in the Summer of 2001, Plant Engineering's Structural Engineer did not have time to begin the requisite finite element analysis of possible bolt patterns that might spread bending loads through the inadequate floor. But the blueprints, once discovered, showed an earlier floor roughly 12" below the current floor. Blake Myers, the Eyeglass team structural designer who solved this problem, recalled (from the days before he was a ME Division Leader) that floor as the portion of the former B154 where the Pluto nuclear rocket was tested!

Core sampling found that floor beneath the current floor, and a drill/vacuum/mortar/drill/glue sequence recommended by the structural epoxy's sales representative was performed to sink 1" threaded rod through both floors, forming a compound beam more than three times thicker than the topmost floor. The first threaded rod installations were tested to above 50,000 pounds force of pullout force. The test with the best epoxy pegged the dial on a custom-kluged, hydraulic-press that used I-beam leverage. This result was nearly four times the pullout strength required for the second iteration earthquake legs to withstand a Class II earthquake's bending moment. The bolt pattern was deemed adequate in a service contract with a university professor who teaches Civil Engineering graduate students how to anchor in concrete, completing a chain of proof that this apparatus was installed with sufficient strength.

Partial-completion of automated assembly station

This closed-loop, automated assembly station was not fully completed by the conclusion of our Strategic Initiative. The station was fully designed, all hardware was purchased, and some subsystems were build and operational, but other systems were not completed and debugged, hence the complete station was only partially operational.

The reason for this, as will be discussed in the next section of this report, is that in late Spring of 2002 we decided to field and optically test the lens outdoors. The schedule required to mount, field, and test the lens greatly shortened the time available for its assembly. We did not have enough time to complete the assembly station early enough to use it for assembly of the lens, and did not have enough resources to continue building and debugging the station in parallel with fielding the lens. Accordingly, we suspended work on completing the automated assembly station.

Fielding and Testing of the 5 meter Lens

During the summer of 2002 we built, mounted outside, and optically tested a 5 meter diffractive lens. While this lens used the design and optical elements developed during the Strategic Initiative, it was assembled, mounted, and tested using external funding.

While the goal of this Strategic Initiative was always to fabricate a lightweight 5 meter diffractive lens as a demonstration that such large lenses could be built, our plans of how to achieve this changed over the course of the project. This evolution was a natural response both to our technical progress in exploring various implementation routes and to budgetary pressures. Our initial goal was to build a diffraction-limited lens and to optically demonstrate this in a controlled environment at NTS. However, as the project progressed, funds for NTS testing were denied, so our ability to optically demonstrate diffraction-limited performance of the lens disappeared. We still planned to build a diffraction-limited lens, but performance would have to be inferred from metrology rather than being explicitly demonstrated. As our technical work proceeded, it became clear that optical smoothing of glass substrates was a considerably harder problem than we had anticipated, and that we would be able to develop smoothing processes, but would not have the time to incorporate them into the actual panels of the 5 meter lens. This, of course, further eroded our ability to build and demonstrate a diffraction-limited lens.

In spring of 2002, with roughly 5 months left in the project, our plan was to build the full 5 meter lens, but to do so with unsmoothed glass and to leave the lens indoors on its assembly table. The assembly precision necessary for diffraction-limited performance would still be achieved through use of the semi-automated, closed-loop, assembly station. While this plan was the natural response to the preceding developments, we were not comfortable with it for a number of reasons. The first, of course, was that it would not result in a truly compelling lens demonstration; there is a big difference between an outdoor, optically functional lens, and an untestable hanger-queen, no matter how precisely assembled. The second worry was that we were facing a tight schedule and budget in activating the closed-loop assembly station; if debugging took more time than expected, assembly precision would likely be sacrificed.

At this time, we became aware of interest in Eyeglass from an external agency. They were interested in large diffractive lenses for future applications but needed a rapid demonstration (before Oct. 1, 2002) that we could actually build, field, and test such a lens. An actual full-up demonstration of a 5 meter lens was, of course, the intent of our project, and much more desirable than the untestable, assembly-table, lens we were currently planning. Accordingly, we accepted this challenge.

We realized, of course, that performing a full-up demonstration this rapidly would be very challenging. There is a major difference between assembling a thin-glass lens in the controlled and well supported environment of our assembly table, and taking this large, fragile, object outside in normal wind and gravity (which, of course, it would never face

in space). In addition to these risks, the need to field and test the lens by Oct. 1, greatly compressed the time available for patterning and assembling it. Accordingly, we were forced to adopt a simpler, manual open-loop, assembly method than the semi-automated closed-loop one we had been developing. This simpler assembly method could not, of course, achieve the few-micron levels of precision of the one we had been developing; but this full precision would not be observable in any event, being completely masked by the effects of unsmoothed glass and atmospheric testing.

The first step in our lens fabrication process was the creation of 72 diffractive lens panels. This was performed as discussed earlier, as part of this Strategic Initiative.

Given the fact that our lens would now be fielded outside, rather than left on the assembly machine, we did decide to add “window frames” to each of the lens panels. This was done because our handling experience with these thin-glass sheets showed that breakage almost always occurred because of edge cracks and/or loads. The placement of “window frames” along the panel edges was considered a crucial step in protecting the panels during the real-life fielding of the 5 meter lens. The “window frames” were made from 1/16th inch thick Al sheets, cut to precisely match the outer border of the various lens panels and extending ½ to 1 inch inside this perimeter. The frames were aligned to the lens panels and glued onto them using UV-cured Norland adhesive. The opacity of the frames led to a small (~ 5%) loss in areal-transmissivity of the lens.

The biggest risk to our plan to field and test the lens was mounting, transporting, and erecting it without breaking the very thin (700 μm) glass panels. Since our previous plan was to not do any of this, leaving the lens resting horizontally on its assembly table, the new plan to field a vertically-mounted lens outdoors in the wind and rain imposed major technical and scheduling challenges.

We devised the following fielding plan: As discussed before, the lens would be assembled manually, by open-loop alignment. The 72 panels would be seamed together with metal hinges, just as planned before for the assembly machine. At this point plans entered new grounds. We needed a way to structurally mount the lens, to transition from the assembly table to this mount, to transport the lens to our test site, to mount it vertically there, to protect it from wind during operation, and to protect it from wind and rain between operations.

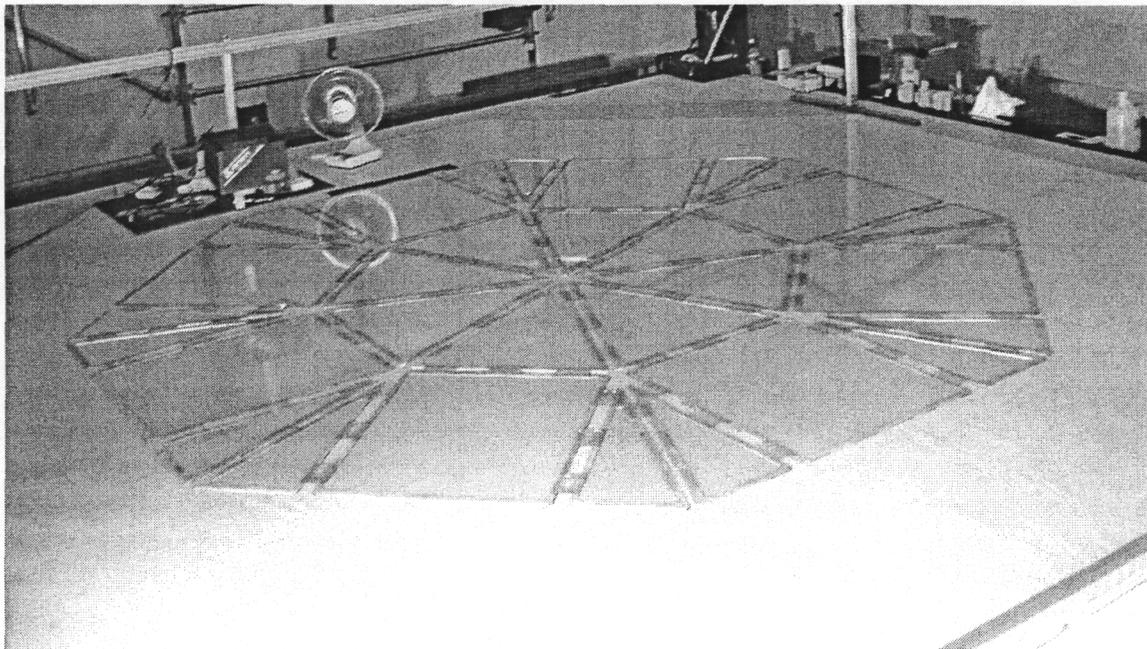
This planning, and its implementation demanded the invaluable assistance of a new team of engineers and technicians (most of whom had not previously taken part in this Strategic Initiative). Key members of this team, led by Andrew Weisberg, were Blake Myers (lead engineer), Phil Stephan (lead technician), Dean Urone (lead designer), Jim Peterson (fabrication and fielding expert), Steve Mooney (lead machinist), lab carpenters (led by Lynn Camp) and the NIF riggers (led by Gerald Walker).

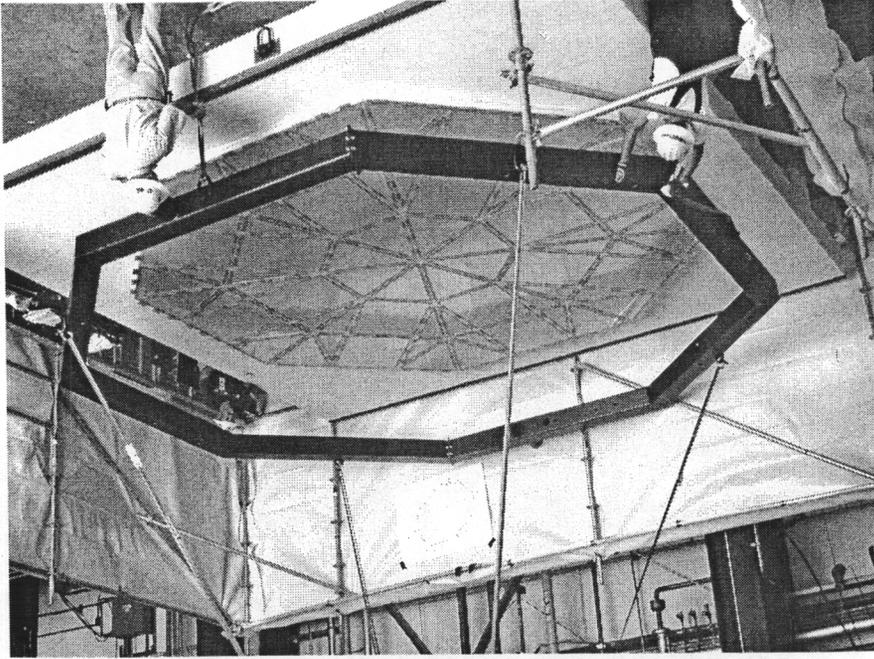
The plan we decided upon called for the lens to be assembled on the surface of a flat, stiff, octagonal platform. When fielded, the lens would be mounted by radial suspension within a stiff steel rim, similar to a drum-head. However, the radial loads required to hang

the lens from this outer rim depend critically on whether it must be supported horizontally (requiring large tension levels) or just vertically (needing much smaller tension). Given the vagaries of brittle failure in thin glass sheets, we did not want to apply enough tension to enable the ring to support the lens horizontally, or under heavy wind loading. For this reason, we designed a support grid to be placed on both the front and back of the lens. This would provide load and stiffness support for out-of-plane forces (such as wind when the lens was vertical and gravity when it was horizontal), but was sparse enough so that it occluded only a small fraction of the lens' optical area. The lens would be supported at its outside test side from two pillars, and have a cover to protect it when not in operation.

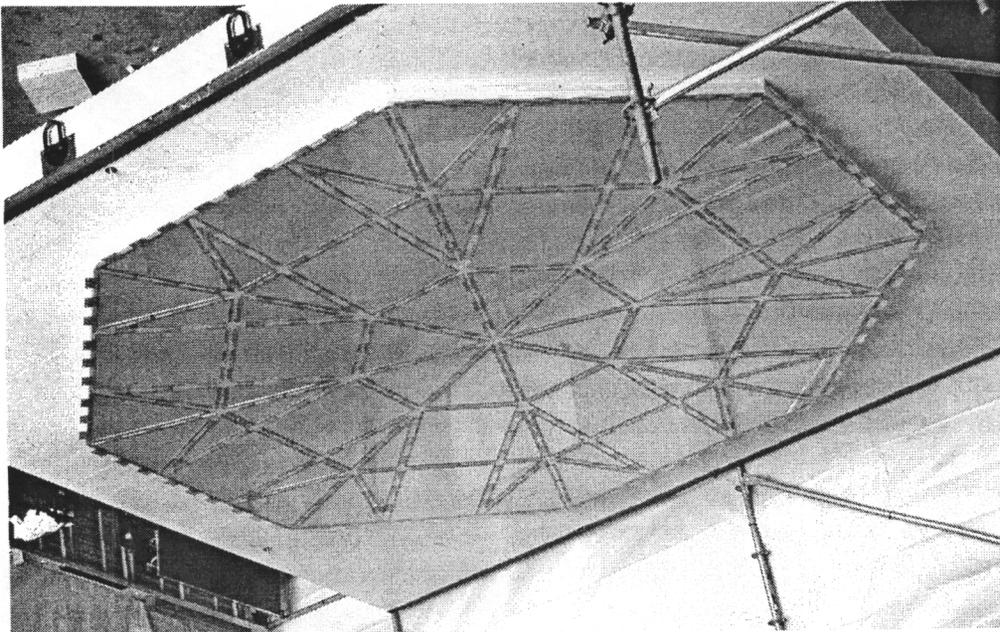
All assembly operations took place in B445, which offered both a high-capacity overhead crane and large exterior doors; both were essential for our work.

The lens was manually assembled from its 72 panels (mostly by Phil Stephan) on a stiff wooden octagonal platform matched to the shape and size of the lens. Alignment accuracy was insured by placing precise mechanical spacers between the lens panels; this process produced alignment accuracy of ~ 0.5 mm over the 5 meter span of the lens. This accuracy was close to that expected from the open-loop phase of the semi-automated assembly machine, and was more than sufficient for our fielded lens (given the unsmoothed glass and atmospheric turbulence). Once panels were properly aligned, they were seamed together by thin metal foils as previously planned. Due to budgetary and delivery-time constraints, we used 10 mil aluminum seams instead of the 2 mil Elgiloy ones that had been previously planned for and prototyped. Below we show two photographs of this assembly process, one partway through and the other at completion.





After the lens was successfully assembled, it had to be mounted. This was done by connecting it, via springs, to a steel support ring. This 18 ft wide octagonal support ring was constructed for us from 8 inch box beam. Note in the above photographs, that the octagonal portion of the platform where the lens was assembled, was raised above that of the outer part of the platform. This 4 inch offset was designed so that once the steel support ring was lowered into position around the lens, its midplane would be level with the lens.

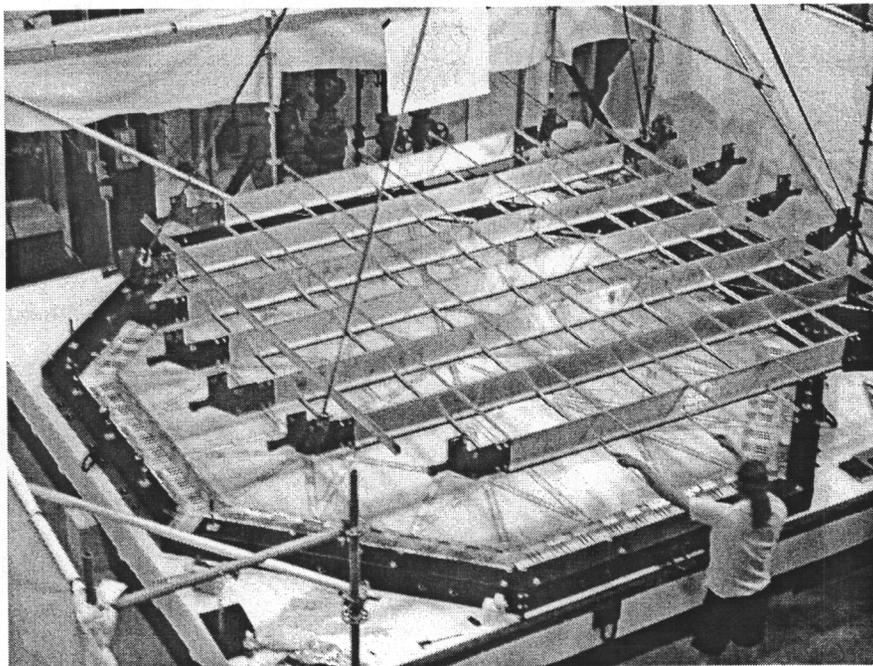


Below we show the first step of this operation, the positioning of the wind-support grid over the top of the lens/ring/platform system. The grid was constructed from orthogonal aluminum bars, ½ inch wide, spaced on 27 inch centers; it occluded only 4% of the lens area. The bottom surface of the grid, where it touched the lens panels, was covered by foam strips (like high-quality weather stripping) to avoid damaging the thin-glass lens. The grid was placed over attachment bolts on the rim-support ring, carefully lowered into contact with the lens, and then locked in place.

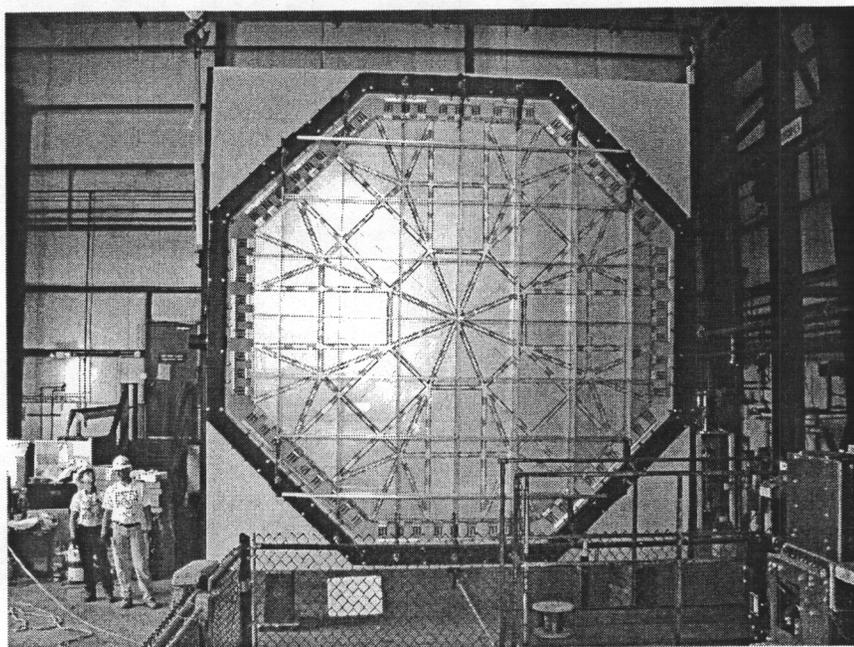
Once the lens was assembled and tension-mounted to the outer ring, it was time to remove it from the assembly platform. However, to avoid breaking the glass panels, the rim tension was kept low; the outer ring would only support the lens when vertical, not horizontally. So removing the mounted lens from the assembly platform was not a simple matter of lifting the ring and lens off the platform. Our plan involved three steps; first we mated one of the two wind-support grids to the top of the lens (attached to the ring, not the lens). Next we flipped the entire system (platform, lens, rim-support ring, and wind-support grid) 180° over, like a pancake. At this point we could lift the platform off from above, as the weight of the lens was safely supported by the, now underneath, wind-support grid.



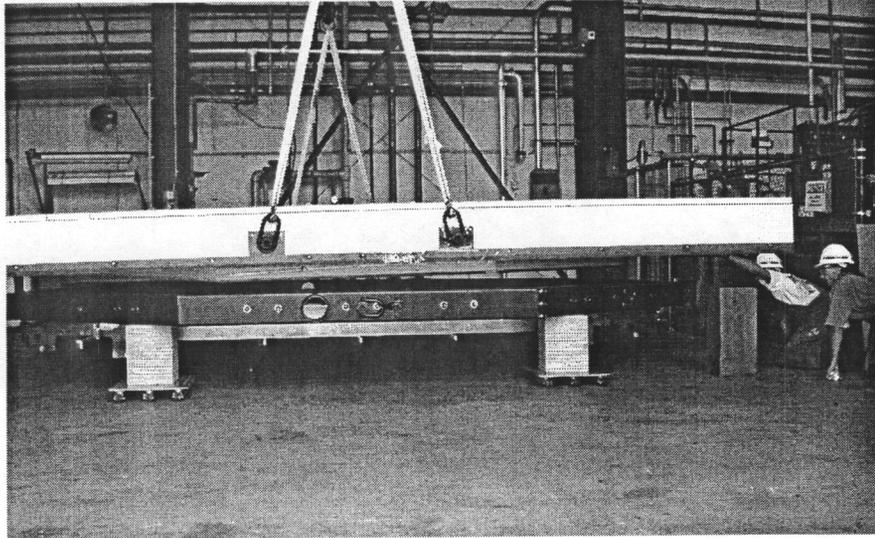
Once the ring was precisely implaced around the lens, it was securely bolted to the assembly platform. At this point the lens was attached to the ring by a set of 216 springs and carefully calibrated tension was applied. The lens-to-ring attachment details are shown below.



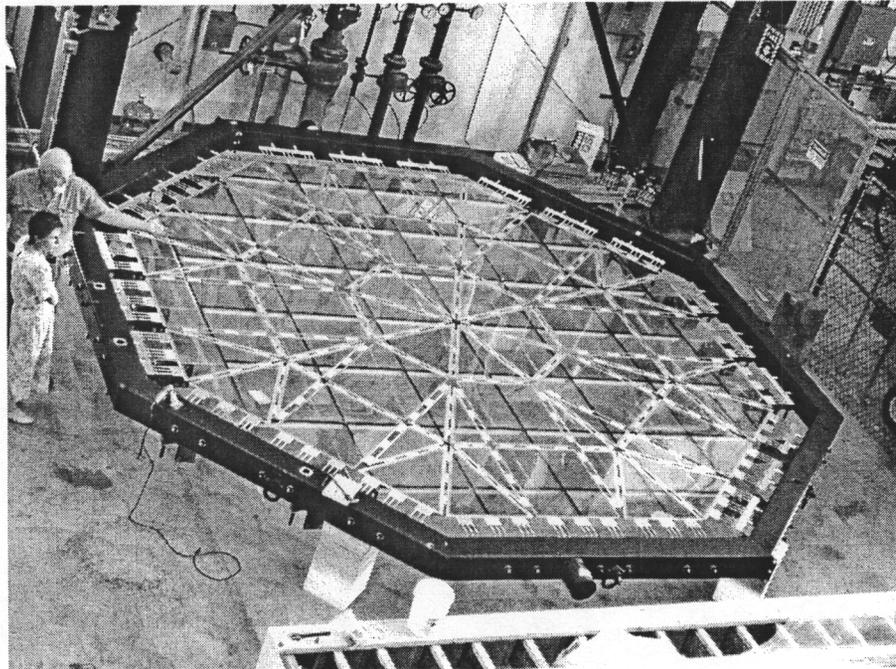
After the wind-support grid was securely attached, we proceeded with the 180° flip of the entire system. This was considered the single most risky operation of the project; insufficient stiffness in the 20' platform, the wind-support grid, or their coupling via the rim-support ring would immediately shatter the thin, 700 μm thick, glass lens. The photograph below shows the assembly in mid-flip, at the 90° point.



The 180° flip was completed without major incident. We then removed the bolts attaching the platform to the rim-support ring, and lifted the platform vertically off the lens/ring/grid system. The initial separation, where we could see the lens had remained intact, is shown below.

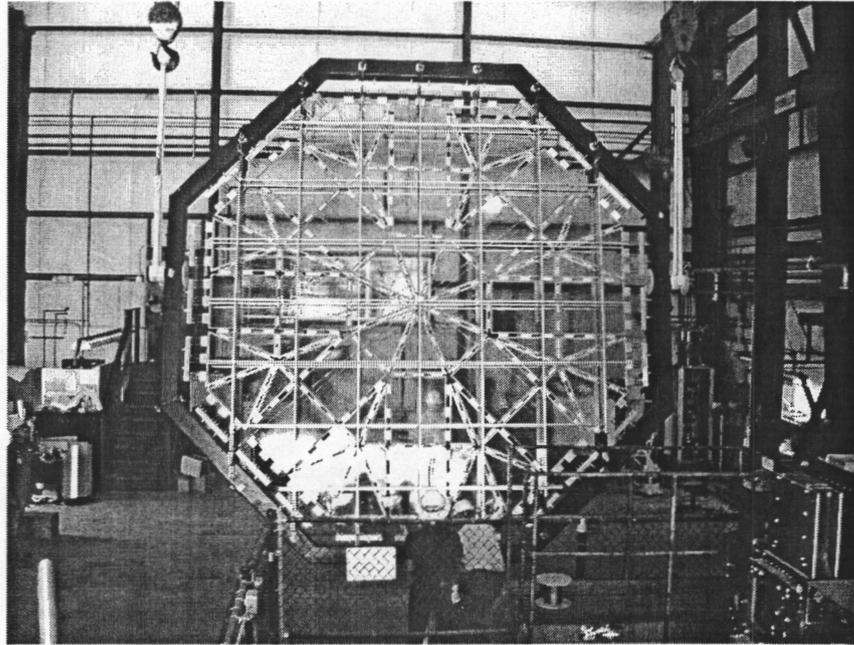


After the assembly platform was completely removed, we inspected the lens and found it to be in excellent condition.

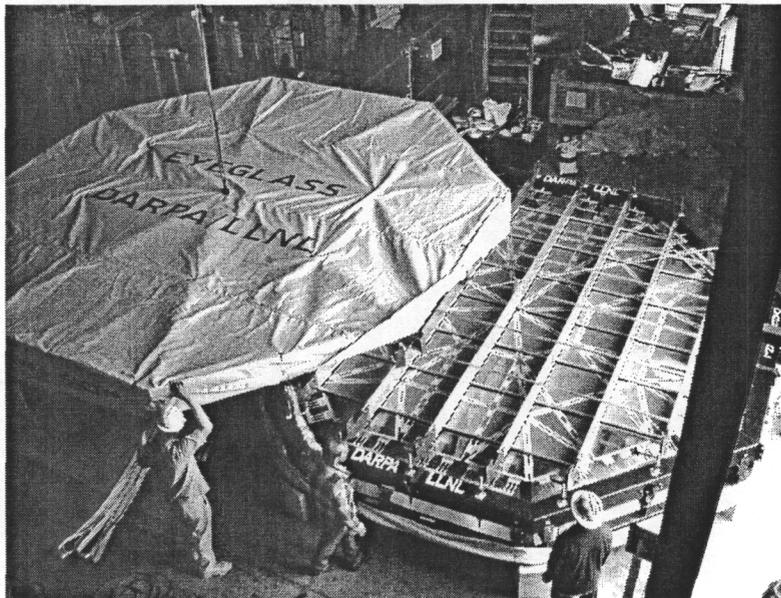


At this point we added the second wind-support grid to the upper surface of the lens. This was done in the same way as the first one (now on the bottom of the lens). The two grids were carefully aligned to one another so that the

obscurations due to the second grid matched, and hence did not add to, those of the first grid. We then flipped the lens assembly through 180° again to put it in position for packaging and transport.

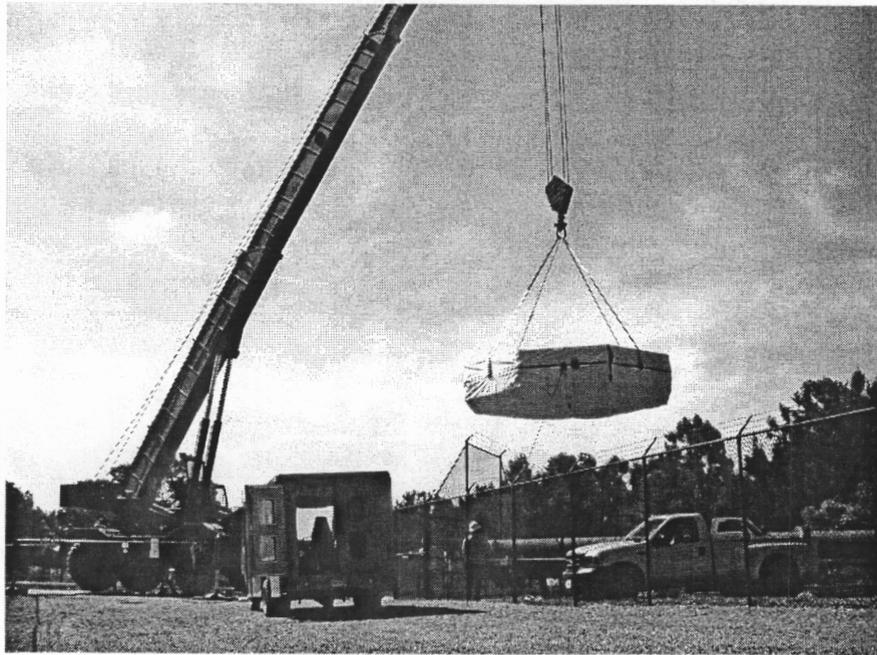


Before removing the lens from B445 and transporting it to its test-site, we had to package it for protection, both during its journey and from on-site weather. Our first step was a temporary one; we inserted thick foam blocks into the cells of the lower wind-support grid. These were pressed up against the lens panels and formed a firm, broad-area “floor” to protect the lens from vibrations during its truck journey from B445 to the test-site. The second measure was a more permanent one; we enclosed the lens, top and bottom, within canvas weather-covers. These consisted of octagonal steel frames covered by custom-made (and insulated) canvas “boat-covers”. The photograph below shows the upper cover being moved into position over the lens after the lower one has already been attached. The foam block “floor” helping to support the lens can also be seen in this picture.

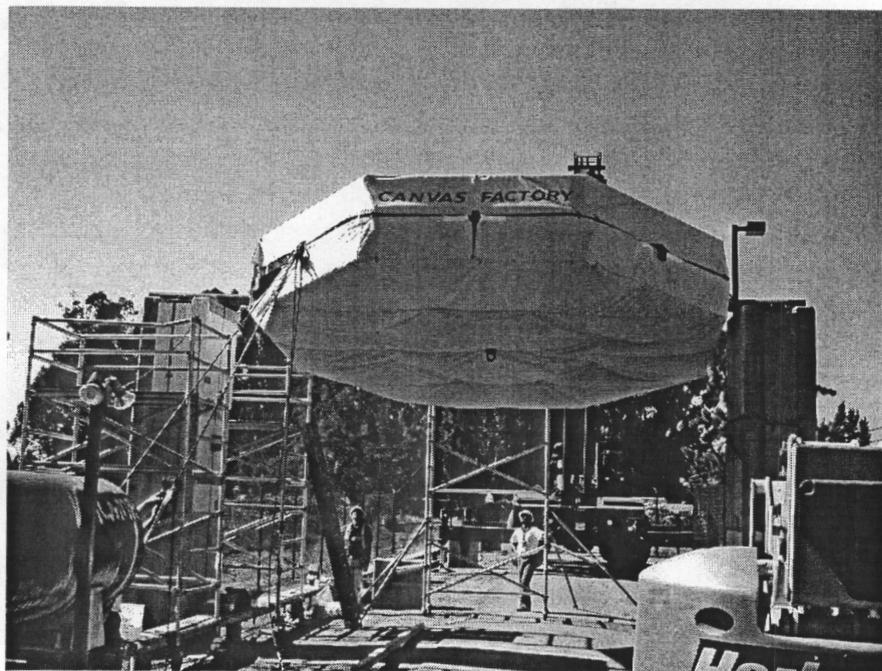


Once the covers were attached to the lens, it was ready to be transported to its test site, mounted there, and optically tested.

Our lens was designed to have a 250 meter focal length for 600 nm light. Since we had no 5 meter plane wave with which to test it, we adopted the previously discussed point-to-point optical configuration. This, of course, stretched the required test range to 1 km. We did not have any convenient clear-path this long on-site, so decided to exploit the multi-harmonic focussing capability of diffractive lenses in order to shorten our range requirements. Selection of 5th harmonic focusing as our baseline allowed us to shrink the needed range to 200 meters. While this is still long, there are sites available at LLNL offering this length. We selected a site just south of B191 for our optical testing. This allowed a beam path slightly greater than 200 meters long, most of which was within a fenced security buffer zone associated with the HEAF facility. The lens was transported to this site on a flatbed truck and hoisted into position by a crane, as shown below.



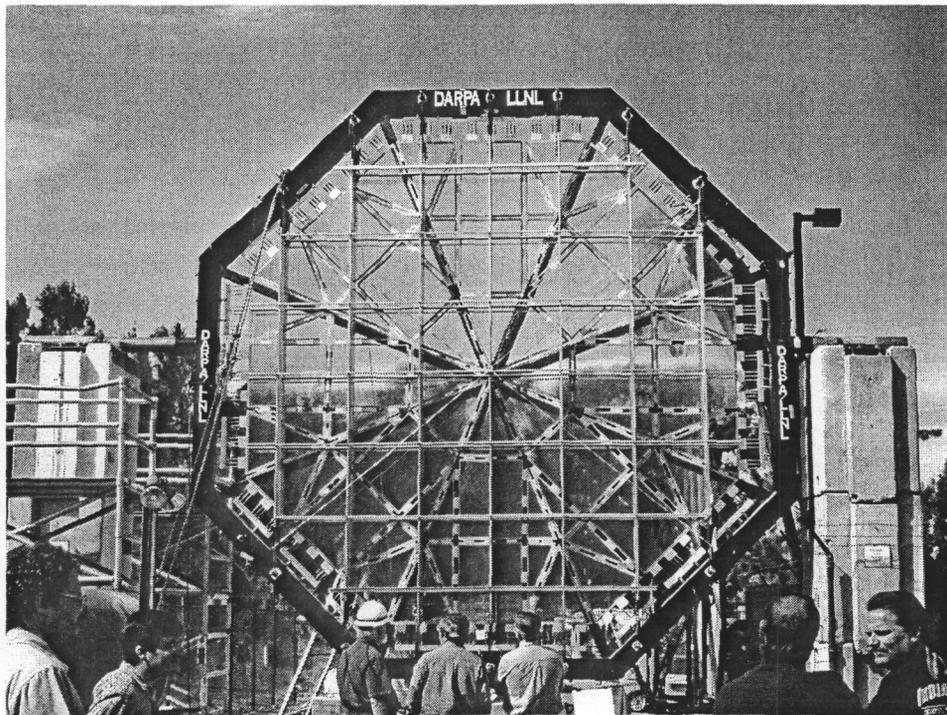
To position the lens for testing, we built two support pillars from stacks of large reinforced concrete blocks. The lens was held between these, at a height of 4.2 meters, by an off-center axle. When not being used, the lens is stowed in a horizontal position, held in this configuration by cables. Below, we show the stowed lens, in its covers, mounted between the two pillars.



To test the lens, of course, we had to remove the covers and rotate it into the vertical position, allowing point-to-point testing with a horizontal beam path. Our original intent had been to fully remove both covers for each test, lowering

the first to the ground, rotating the lens through 180° , lowering the second cover to the ground, and then rotating the lens back 90° into a vertical position for testing. This turned out to be rather cumbersome, so we adopted a simpler operational procedure. First, the lower cover was permanently removed; the lens is protected from weather when stowed only by an upper cover. When testing is required, the lens is rotated 90° into the vertical position, and the canvas portion of the cover is removed, exposing the lens for operation. The metal frame for this upper cover remains attached to the lens; because its octagonal spoke beams are aligned with interpanel gaps, this frame adds little optical occlusion.

A daylight photograph of the uncovered lens in its operational configuration illustrates the size and features of our 5 meter diffractive lens.



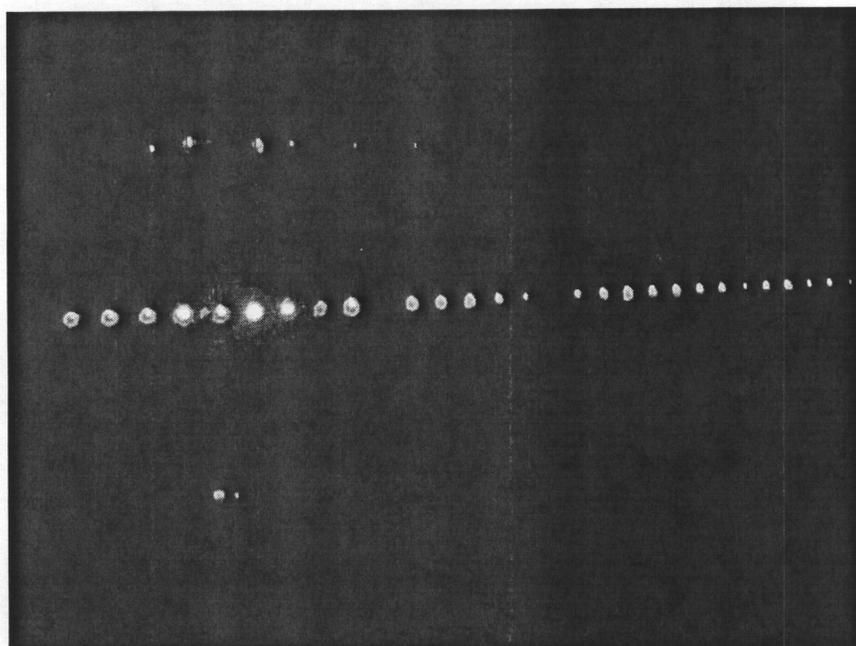
The main purpose, of course, of fielding the lens outside, is to optically test it, convincingly demonstrating that such a flat, thin, lightweight object provides an effective means of focusing light.

Because the 5 meter lens is a diffractive element (and is not accompanied by a color-correcting Fresnel Corrector), it is most effective when focusing monochromatic, i.e., laser, light. Since the lens is based outside in an open environment, safety considerations led us to use low power laser sources. This, in turn, meant that optical testing was done at night, when the focused light was most observable.

We tested the lens with three different laser wavelengths, blue at 405 nm, green at 532 nm, and red at 670 nm. Since our green laser had the most power

(20 mW) and the most efficient eye response, green light tests were the most visually impressive. The laser(s) were mounted 123 meters from the lens and elevated above its center, so that their focal spots were at a convenient viewing level. They were projected toward the lens through a beam expander so that their light filled the 5 meter aperture of the lens.

Our lens had a simple binary phase profile, hence we expected to see light focused into multiple focal spots, rather than being concentrated into just one spot. This was indeed the case, as can be seen in the following photograph. When viewed by eye, the line of focal spots was even more impressive; the nonlinearity of the eye's response made an entire chain of spots appear bright and intense.



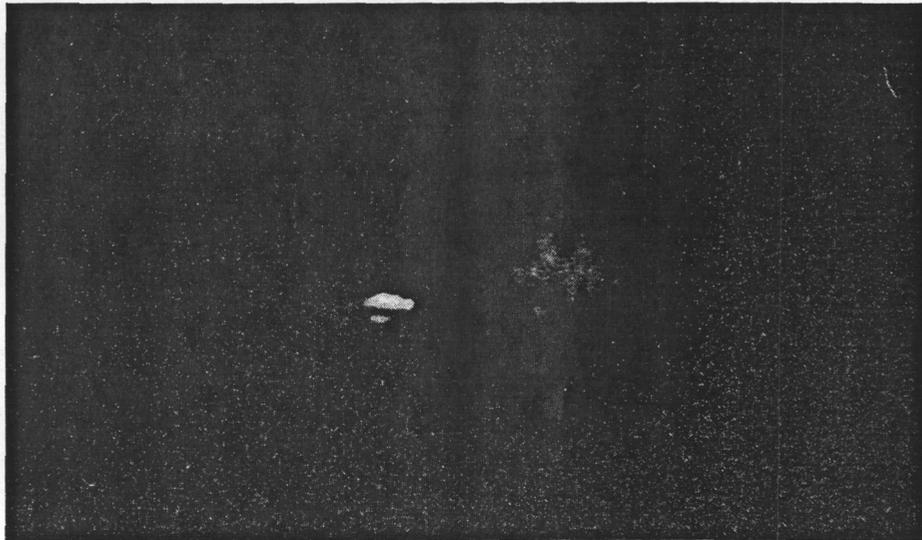
We characterized the lens' optical performance by measuring the position and quality of individual focal spots. The table below, shows the position (meters from the lens) of the lower-order focal spots at each of the laser colors.

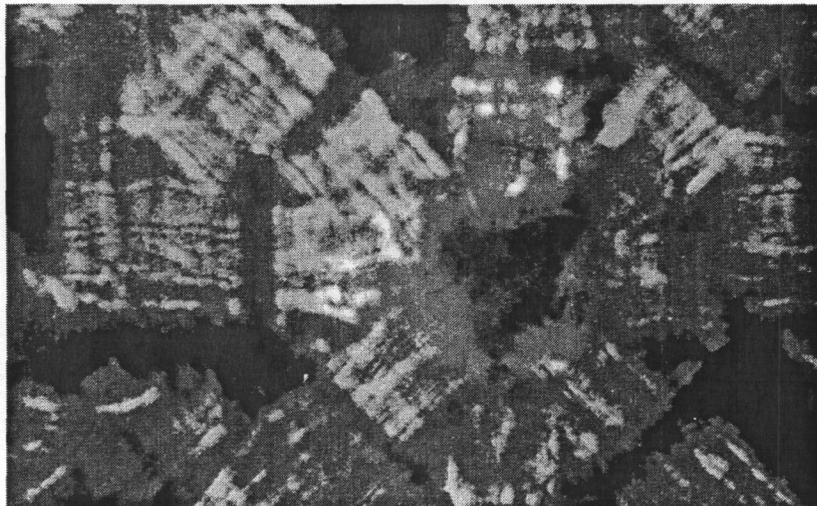
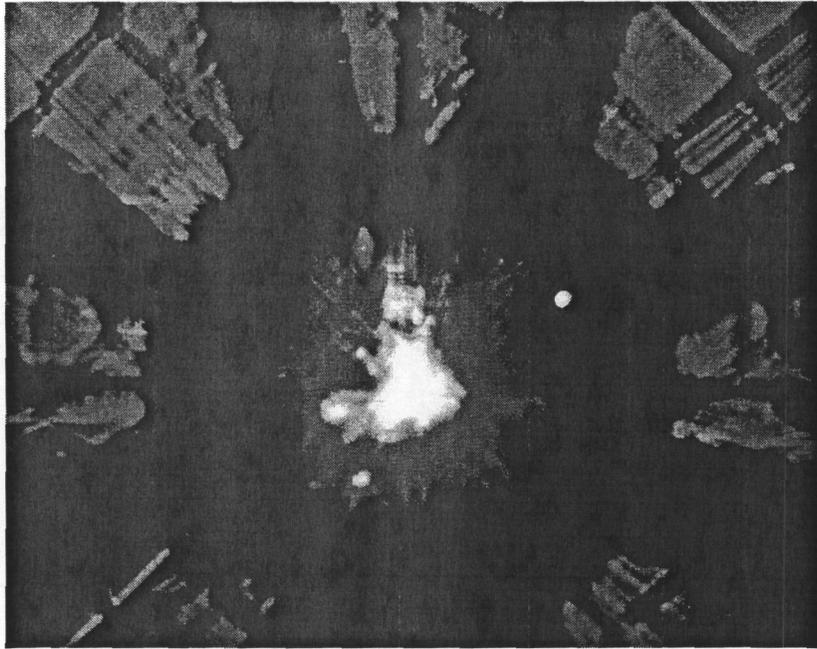
Harmonic	405 nm	532 nm	670 nm
3	8047	368	181
4	347	158	99.1
5	177	101	68.3
6	119	73.8	52.1
7	89.6	58.2	42.1
8	71.9	48.2	35.4
9	60.0	40.8	30.5

A diffractive lens with an "ideal" binary phase profile has only odd-harmonic foci. Our lens delivered focal spots at all the harmonics since its phase depth was set to be half-wave deep only at a single wavelength (600 nm); we saw both odd and even focal spots with (to the eye) little difference in intensity. Also, we saw, particularly for the green laser, many of the higher-order focal spots as well; a long chain of spots existed, stretching outward from the lens.

The length of our beam-path was limited, extending only ~ 150 meters north-westward of the lens; it was cut off by a treeline and the fact that the optical axis (purposely slanted downwards) ran into the ground. The furthest spot we identified was the 4th-order green one, tracked down at shin-height in the trees. Most of our spot measurements occurred at ~ 100 meter range from the lens.

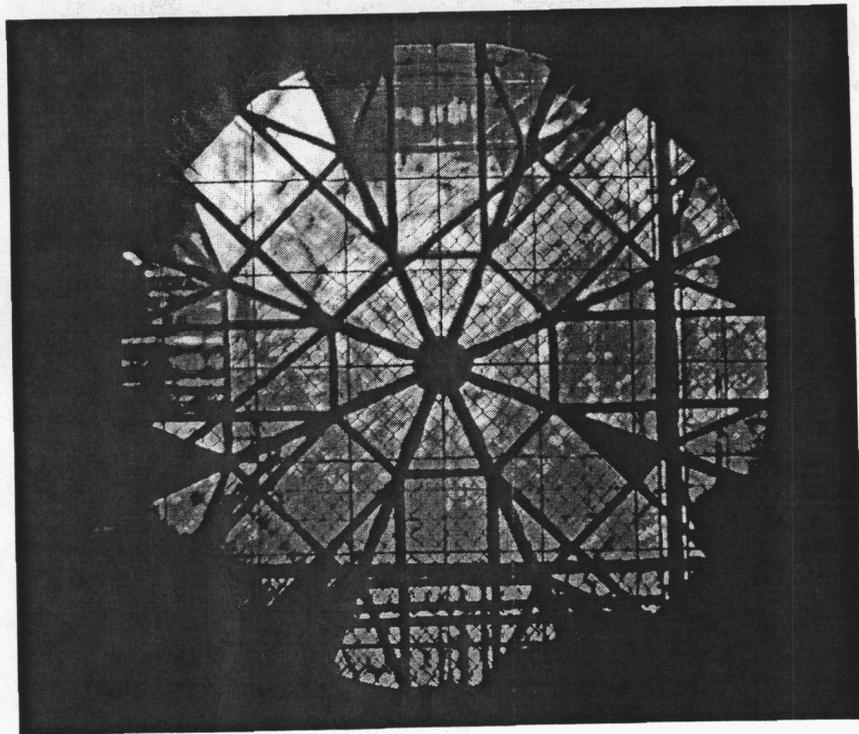
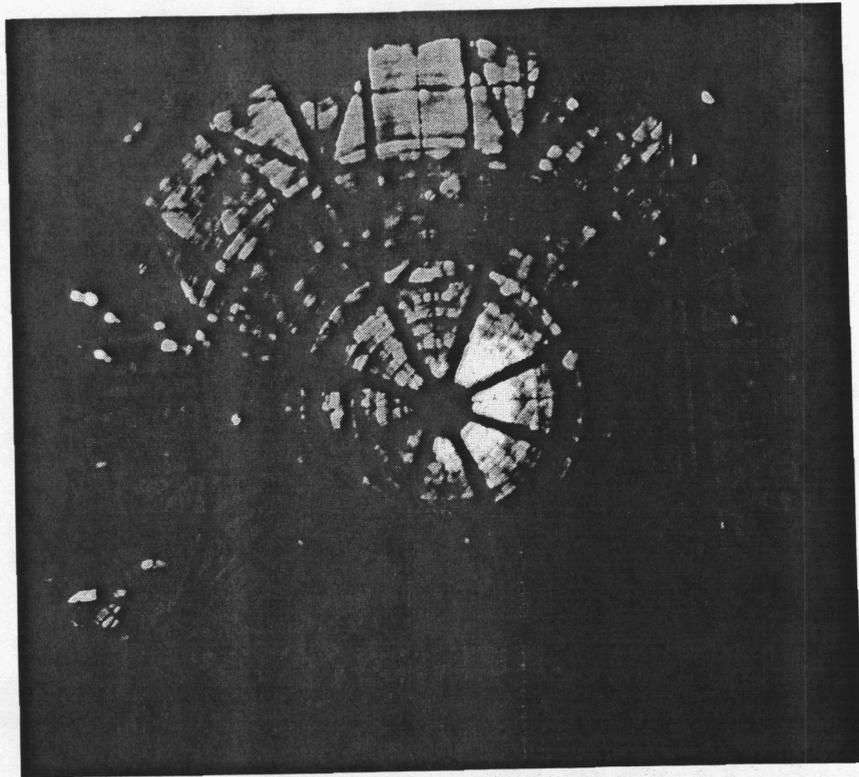
We took two types of spot measurements; images of the spot shape (by placing an opaque card or translucent diffuser at the spot) and images of the lens (by placing the camera at the spot location and imaging backwards to the lens). Spot images were used to determine the size and shape of the focal spot, while lens images showed focal contributions from across the lens. The three images below show the red (7th-order), green (5th-order), and blue (4th-order) focal spots. The red spot was imaged in reflection, while the green and blue spots were photographed in transmission; the fact that these green and blue foci are close together is readily apparent.

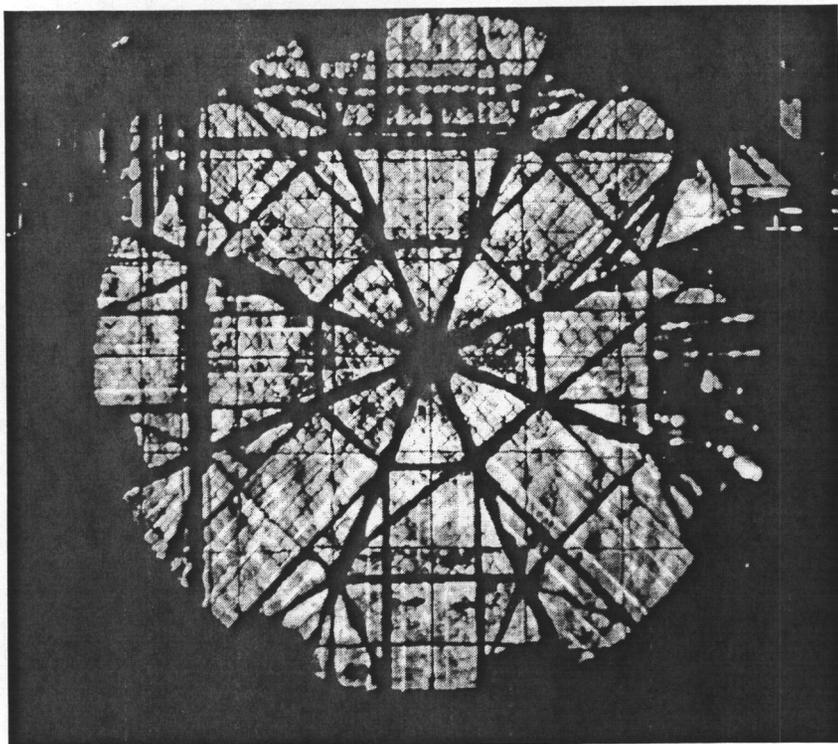




Scales are not shown on these photographs, but spots were typically 1 – 2 cm in diameter. These spot sizes are close to what we expected; they are dominated by the fact that the lens was constructed with as-purchased, unsmoothed glass sheets. The ripples are random, but typically 10 microns deep with 10 cm periods, at 100 meters range they will cause 1 cm spots. The two other factors which prevent diffraction-limited focussing cause less spreading; air turbulence and panel misalignments should contribute a few millimeters of spot broadening. While our 5 meter lens is not diffraction-limited, it does deliver tight focal spots at 100 meter range.

We also took images of the lens aperture, looking backwards from spot locations; these show focal contributions across the aperture. Below we show red, green, and blue lens images.





These images clearly show the various imperfections and obscurations of the lens and the beam-path (a chain-link security fence was located between the lens and focal sites, adding some obscurations of its own). The gaps between the panels and the eight octagonal spokes from the lens cover are readily apparent. One can also see the effects of the mask errors on panel #8. As discussed previously, a portion near the rim of this panel was mispatterned and shows up as 8-fold symmetric missing aperture.