

711-44-CR

183170

p. 107

(NASA-CR-194226) A LARGE
DEPLOYABLE SOLAR CONCENTRATOR WITH
RECEIVER AND HEAT STORAGE, PHASE 2
Final Report, 18 Aug. 1986 - 18
Jan. 1989 (Energy Science Labs.)
107 p

N94-70597

Unclas

Z9/44 0183170

SBIR-10.04-7039
RELEASE DATE 12-3-90

Final Report on SBIR 1984 Phase II
Phase II Contract NAS3-24882:
A Large Deployable Solar Concentrator
with Receiver and Heat Storage

submitted to NASA Lewis Research Center
Donald Chubb, Contract Technical Monitor

September 1990

WAYNE GIRAARD - ADMINISTRATOR
3-2799

Report written by:
Joseph A. Carroll
J. Kevin Cross
Andrew H. Cutler
-George E. Henschke-

Energy Science Laboratories, Inc.
PO Box 85608, San Diego CA 92138-5608



PREFACE AND INTRODUCTION

This report describes the work done under NASA Contract NAS3-24882. This contract was a Small Business Innovation Research (SBIR) Phase II research and development contract between Energy Science Laboratories, Inc., and NASA Lewis Research Center. The purpose of the contract was to further develop and to test solar concentrator and receiver design concepts conceived during the preceding SBIR Phase I study contract, NAS3-24397.

The contract performance period was from August 18, 1986 to January 18, 1989. Mr. Donald Chubb of NASA Lewis was the contract monitor.

The key personnel on the project and their major responsibilities were:

Joseph A. Carroll	Principal investigator
J. Kevin Cross	Concentrator fabrication and testing
Andrew H. Cutler	Receiver design and testing
George E. Henschke	Testing and test data reduction

Chapter 1 reviews our work on the concentrator. We developed design and fabrication concepts that seem quite attractive for fabricating and assembling full-scale concentrators. Significant errors were introduced during a critical fabrication step, so the deliverable 20% scale concentrator had much lower performance than expected. We believe these errors are avoidable, and describe methods to test and verify this in future work. Developing suitable optical testing techniques was a significant part of our effort. With further refinements along lines suggested in the report, this test method should be useful for testing a wide range of concentrators, including the prototype now being fabricated for space station.

Chapter 2 reviews our work on receiver design and component testing. The results show that beryllium is fairly sensitive to "thermal ratcheting." This results from small external loads combined with large internal stresses caused by anisotropic thermal expansion within each grain. This problem plus the potentially better performance and lower cost of a higher-temperature system using graphite as the storage medium led us to generalize our receiver design concept to be relevant for either beryllium or graphite.

The testing phases of tasks 1 and 2 required considerably greater effort than expected. This forced us to shift some effort away from tasks 3 and 4. Chapter 3 details the design work on an insulating receiver door plus hinge and actuator that we accomplished under task 3. Chapter 4 describes our systems effort, which was focused on refining our simulation and analysis programs, and on using those programs to identify key features of a sensible-storage receiver using graphite, and key issues that need to be resolved in developing it.

Readers who have questions, comments, or suggestions about this report or about the design concepts and test methods described in it, are encouraged to contact the principal investigator in San Diego, at 619/455-6125.



TABLE OF CONTENTS

1. CONCENTRATOR DEVELOPMENT, FABRICATION, AND TESTING	1
1.1 Membrane Design, Assembly, and Testing	2
1.1.1 Adhesive bonding on paraboloidal jigs	5
1.1.2 Self-aligning mechanical fasteners	9
1.2 Strip Design and Fabrication Details	14
1.2.1 Material selection	14
1.2.2 Alignment/fastening holes	17
1.2.3 Anisotropic stiffeners	23
1.2.4 Reflective and emissive surface coatings	27
1.3 Rim Design, Fabrication and Testing	31
1.3.1 Integral rims	31
1.3.2 Separate rims	38
1.4 Optical Testing	41
1.4.1 Hardware	43
1.4.2 Test procedures	47
1.4.3 Software	52
1.5 Optical Test Results	56
1.5.1 Early test results	56
1.5.2 Test results on final 20% scale concentrator	58
2. RECEIVER DEVELOPMENT	75
2.1 Beryllium Receiver Design	75
2.2 Fabrication Techniques Suitable for Beryllium	80
2.3 Beryllium Testing for Creep and Thermal Ratcheting	83
2.3.1 Background	83
2.3.2 Test hardware	86
2.3.3 Test procedures	92
2.3.4 Test results	95
2.4 Receiver Efficiency and Mass Calculations	103
3. RECEIVER DOOR, HINGE, AND ACTUATOR DESIGN	105
3.1 Strawman Door Design	105
3.2 Strawman Hinge Location and Design Options	107
3.3 Strawman Actuator Design	109
4. SYSTEMS AND SAFETY ANALYSIS	111
4.1 Strawman Graphite Receiver Design	112
4.2 Preliminary Thermal Analysis	115
4.3 Materials Issues and Selection	118

LIST OF TABLES

Table I	Characteristics of 10% scale concentrators	5
Table II	Characteristics of 16% central section concentrators 6-8	7
Table III	Characteristics of 2.8 meter diameter concentrators	9
Table IV	Characteristics of integral-type bi-stable rims	31
Table V	Rim deflection as a function of disturbing force	35
Table VI	Characteristics of our three optical test methods	42
Table VII	Number of pixels lit in different numbers of images	58
Table VIII	Number of pixels with different error variances	63
Table IX	Comparison of selected tantalum alloys and superalloys	77
Table X	Creep test results	95
Table XI	Thermal cycling test results	99

LIST OF FIGURES

Figure 1	10% scale concentrator #5 (deployed but with rim flat)	3
Figure 2	10% scale concentrator #5 (fully deployed)	3
Figure 3	Concentrator #10 (20% scale: 2.8 m dia, 1.2 m focal length)	4
Figure 4	Depth gauge for measuring changes in concentrator shape	6
Figure 5	Riveted 2.8 meter central section and assembly/stowage tube	10
Figure 6	Sample riveted concentrator section - four 5.5" wide strips	11
Figure 7	Mill and stacked-strip holder used for drilling strip holes	17
Figure 8	Hole pattern for joint-error samples	18
Figure 9	Pattern of edge shapes used in final 2.8 meter concentrator	20
Figure 10	Operator drilling holes for final 2.8 meter concentrator	21
Figure 11	Cross-section of initial corrugating tool	23
Figure 12	Concentrator #8 on its fabrication jig	24
Figure 13	Longitudinal corrugating tool for reinforcing strips	25
Figure 14	Reflectance fixture for Varian 634 spectrophotometer	27
Figure 15	ASTM D-3359-83: measuring adhesion by tape test	28
Figure 16	Specular reflectance of 5657 aluminum alloy	30
Figure 17	Cross section of rim #4	32
Figure 18	30% scale rim (5 mil/10 mil) supported at 6 locations	33
Figure 19	Rim reinforcement corrugating tool	34
Figure 20	Equilateral triangle rim cross-section	35
Figure 21	Finger/pocket support for optical tests on 2.8 m central section	39
Figure 22	20% scale thick rim - horizontal and vertical positions	40
Figure 23	Optical testing methods	41
Figure 24	Light source scan pattern for "point" and "bar" sources	45
Figure 25	Optical test setup - method 3	48
Figure 26	20% concentrator rim position changes, after stowage vs before	51

Figure 27	MacPaint file made using method 1	52
Figure 28	Focal plane errors of an ideal paraboloid under our test conditions	54
Figure 29	Rudimentary optical test of concentrator #4	56
Figure 30	Pixels lit in 20% Scale Concentrator Tests 1 and 2	59
Figure 31	Pixels lit in different numbers of images in scan 1H	60
Figure 32	Pixels lit in different numbers of images in scan 1V	61
Figure 33	Horizontal errors in tests 1 and 2	64
Figure 34	Vertical errors in tests 1 and 2	64
Figure 35	Total errors in tests 1 and 2	64
Figure 36	Flux vs total error in tests 1 and 2	65
Figure 37	Error maps of 20% scale concentrator	66
Figure 38	Changes in horizontal and vertical error, test 2 - test 1	68
Figure 39	"Saturating" focal plane images for 20% concentrator, tests 1 and 2	70
Figure 40	Computed flux map for 20% scale concentrator test 1, in suns	71
Figure 41	Computed flux map for 20% scale concentrator test 2, in suns	72
Figure 42	Pixels lighting apertures w/CR of 600, 200, and 60 in tests 1 and 2	73
Figure 43	Schematic of receiver design	78
Figure 44	Beryllium heat storage tube and its end connections	79
Figure 45	Individual heat storage tubes and receiver structure	80
Figure 46	Hot-end shield and storage tube assembly prior to receiver assembly	80
Figure 47	Cross-section showing tube assembly supported by removable structure	81
Figure 48	final assembly of receiver	81
Figure 49	Change from launch to on-orbit configuration	82
Figure 50	Initial beryllium testing concept	83
Figure 51	Beryllium creep test setup	86
Figure 52	3-zone Lindberg tube furnace - open	87
Figure 53	Micrometer slide and telemicroscope	88
Figure 54	Beryllium test bar	90
Figure 55	Beryllium test bar inside quartz tube	90
Figure 56	Initial test bar support scheme (caused bar to bend)	91
Figure 57	Final test bar support scheme (no bending)	91
Figure 58	Distance between marks vs. time - test 1, LR	96
Figure 59	Distance between marks vs. time - test 1, RL	96
Figure 60	Distance between marks vs. time - test 2, LR	97
Figure 61	Distance between marks vs. time - test 2, RL	97
Figure 62	Distance between marks vs. time - test 3, LR	98
Figure 63	Distance between marks vs. time - test 3, RL	98
Figure 64	Distance between marks vs. time - test 4, LR	100
Figure 65	Distance between marks vs. time - test 4, RL	100
Figure 66	Distance between marks vs. time - test 5, LR	101
Figure 67	Distance between marks vs. time - test 5, RL	101
Figure 68	Distance between marks vs. time - test 6, LR	102
Figure 69	Distance between marks vs. time - test 6, RL	102
Figure 70	Strawman door design	106
Figure 71	Two hinge design concepts using refractory hinges	107
Figure 72	Hinge location options and door positions, open and closed	108
Figure 73	Redundant Nitinol actuators and refractory tensile bands	109



1. CONCENTRATOR DEVELOPMENT, FABRICATION, AND TESTING

Our concentrator development effort flowed from our previous work under the SBIR Phase I contract, and is best introduced by a brief review of that work.

The SBIR Phase I effort initially studied an umbrella-like solar concentrator that could stow on the side of the shuttle payload bay, in a space reserved (but never used) for a 2nd remote manipulator arm. Since shuttle payload is often volume-limited rather than weight-limited, it seemed that payloads stowed in this volume might often be launched "for free," i.e., without displacing other payloads. However, a thermal analysis during the Phase I study indicated that when an umbrella-like concentrator passed from sun to shade, the thin membrane would cool off and shrink far faster than the frame would. This would induce loads radically larger than any experienced during launch. We studied ways to reduce this problem. The main options were high compliance between frame and membrane, good thermal coupling, or similar time constants. Concepts we developed that provided high compliance seemed impractical or at least very inelegant, so we considered options that provided good thermal coupling or similar time constants.

The key features that we wanted to retain were the long, thin stowed shape and a self-deploying capability. We realized that a paraboloidal membrane could be rolled up into a cigar shape without inducing in-plane compressive or tensile stresses, and that such a membrane might only need support at the rim.

We then studied bistable rim concepts that could be flexible enough to roll up in one state, but adequately rigid in the other. (We had originally looked at such concepts to help stiffen an umbrella frame.) This led us to a "diamond" rim integral with the concentrator. We fabricated two small concentrators about .42 meter across by laying barrel-stave-shaped strips side by side on a paraboloidal membrane and taping them together. This concept worked well enough for us to propose it for the Phase II effort.

During Phase II we fabricated five 10% scale concentrators, four central sections of a full scale concentrator, and one 20% scale concentrator. We identified or developed techniques intended to be suitable for fabricating full-scale concentrators, and fabricated a large number of small sections to test those techniques. We assembled the 20% concentrator using those techniques in key areas. This chapter of the report reviews our Phase II work on the concentrator, under the following major headings:

- Membrane design, assembly and testing
- Strip design and fabrication details
- Rim design, fabrication, and testing
- Optical testing
- Optical test results

1.1 Membrane Design, Assembly, and Testing

We made some key decisions early in the contract. One was to stay with an on-axis design. This provides the highest concentration ratio for any given RMS slope error. It also simplifies receiver design by making energy distribution within the cavity more symmetrical. We would have liked to be compatible with the off-axis baseline space station configuration, but we felt that the off-axis design was a result of too much emphasis on minimizing receiver moment of inertia and not enough on the direct and indirect costs of doing so.

We also decided to continue with the 6 meter focal length, 14 meter diameter design concept we baselined during Phase I. This diameter is small enough to allow stowing along with the Ku-band radar on the shuttle ledge. As discussed on page 17 of the Phase I report, f-numbers near .43 provide the highest concentration ratio for any given RMS error. We saw no net reason to bias the design above or below that value during Phase II.

All of our membrane designs used the "barrel stave" shaped strip concept conceived during the Phase I effort and proposed for the Phase II effort. Our initial Phase II work focused on membranes made by laying strips down on a paraboloidal jig and bonding them with adhesives. To keep the tape joint flexibility from dominating the stowage and deployment processes, we rolled the concentrators up with the joints perpendicular to the stowing axis.

We made five 10% scale concentrators. The last one is shown before and after rim deployment in Figures 1 and 2. The concentrator is suspended by multiple lines hanging from the ceiling. Each line has rubber bands in it, so it can provide comparable support even when the concentrator changes shape. We refined this soft-suspension technique as we went along, and found it adequate for all our deployment tests during this contract.

We then fabricated three 2.2 meter sections of a full-scale concentrator on a very shallow jig, to determine the handling properties of membranes with representative surface curvatures. We quickly found that the seams between strips needed to be rigid, or the membrane would buckle at the seams. We also developed an anisotropic stiffening concept suitable for our stowage geometry. To be effective, the stiffeners had to be continuous, and that practically meant that they had to run along the strip axis. This meant that stowage had to be done with the strips parallel to the stowing axis. This change allowed us to shift to simple lap joints between sections, because lap joint shear loads during stowing are modest if the lap joints run along the stowage axis, but high if they are perpendicular to that axis.

Our experience in working with the 2.2 meter sections led us to focus on the practical issues involving a full-scale jig and fabrication procedures. This led us to study concepts using self-aligning mechanical fasteners that could precisely position adjacent strips during assembly. This greatly simplified assembly, and it eliminated the need for a full-size paraboloidal jig. We made many small test sections to characterize the limitations of this process, and then fabricated a 20% central section of a full-scale concentrator using pop-riveted lap joints.

Further experiments on small sections showed us that eyeleting provided lower distortions than pop-riveting. We used eyelets to assemble the last concentrator fabricated under this contract: the 20% scale model delivered to NASA Lewis late in 1989. This concentrator is shown in Figure 3.




Figure 1a & b. 10% Scale Concentrator #5 (deployed but with rim flat)



Figure 2 10% Scale Concentrator #5 (fully deployed)



Figure 3 Concentrator #10 (20% scale: 2.8 m dia, 1.2 m focal length)

We videotaped the steps involved in fabricating, testing, and stowing the above concentrator. We have edited the tape and are providing the results to Don Chubb of NASA Lewis.

Inadequate controls on the step-and-repeat operation during hole drilling on our mill for the final concentrator reduced its performance well below the range expected. However, we did the drilling on groups of strips that were symmetrically disposed about the vertical axis. Drilling errors thus result in slope errors with high bilateral symmetry. As is quite clear from the test data shown in section 1.5, the slope errors have high bilateral symmetry, and are much larger than the inherent strip-related errors, which are bilaterally symmetric (and mostly horizontal).

Reducing drilling errors can increase concentrator performance significantly. In addition, scaling the concentrator up will make it more accurate, unless hole position errors scale up proportionately. We believe it is feasible to make full-scale strips with the critical "high frequency" hole position errors smaller (in absolute terms, not just relative terms) than those in our 20% scale model. This should have a quadratic effect on the achievable concentration ratio, because slope error varies roughly linearly with hole position error, and achievable concentration ratios vary with the inverse square of slope error. This can be verified in two steps: fabricate full-length strips and measure hole position errors, and then fabricate and test sections of full-scale concentrators using those strips. We believe that with refinement, our design concept should allow concentrators with adequate performance for space station and other applications.

The rest of section 1.1 goes into more detail on the ten concentrators, how we made and tested them, and what we learned. Section 1.1.1 covers concentrators 1-8, which were assembled adhesively, and section 1.1.2 covers concentrators 9 and 10, which were mechanically fastened.

1.1.1 Adhesive Bonding on Paraboloidal Jigs

We made a total of 8 concentrators/concentrator sections using adhesive fastening on jigs: five 10% scale models and three 16% central sections of a full-scale concentrator.

We began with a primarily experimental effort in order to rapidly gain some experience with concentrators larger than the 3% scale models we fabricated during Phase I. We found that we could obtain a suitable form for making 10% scale models from a company two blocks away that sold satellite receiving dishes in various sizes. The most appropriate dish happened to be perforated. This turned out to be useful later on, because it allowed us to check shape changes using a depth gauge through the holes. We designed and built tooling to cut the strips to size and fabricated five 1.4 meter diameter (10% scale) concentrators with the characteristics listed below in Table I.

Table I Characteristics of 10% Scale Concentrators

#	Alloy	Thick	Finish	Polished?	Tape	RimType	RimShape	RimThick
1	3003	.002"	mill	no	3M-#600	Integral	Diamond	.002/.0014"
2	3003	.002"	mill	no	3M-#600	Integral	Flange	.002"
3	1350	.005"	mill	no	3M-#600	Integral	Diamond	.005"
4	1350	.006"	Nameplate	yes	3M-#850	Separate	Diamond	.005"
5	1145	.001"	Bright	yes	3M-#853	Integral	Diamond	.005/.002"

All 10% concentrators were made from 30 barrel-stave-shaped strips of extra-hard (H19) foil 1.94" wide in the middle. Strips were polished with Simichrome paste, laid side by side on a .58 m focal length paraboloidal mandrel, and taped together. Rims were assembled in the same way, from curved strips of foil. Rims with two thicknesses used the thicker foil in front, for better buckling resistance during stowing.

Being made on the same mandrel, all five had the same focal length of .58 m. We used different gauges of material to gain experience with gauge-related effects.

We fabricated the first three concentrators simply to gain experience with various design and fabrication concepts, and with the resulting concentrator mechanical properties. We damaged the first concentrator slightly in stowing it. We improved our handling techniques with the next two concentrators, but we found that stowing and re-deployment still caused slight shape changes. We measured them by mounting the concentrator back on the satellite dish with a controlled standoff around the rim, and using a depth gauge through the holes in the dish to measure local variations in standoff distance. Since the concentrators were quite flimsy, we used electrical continuity to determine gauge contact with the concentrator. The assembly is shown in use in Figure 4.

We traced the changes in shape to the use of adhesives that were creep-sensitive under low load. This led us to make a brief review of suitable adhesives, as discussed below.




Figure 4 Depth Gauge Assembly for Measuring Changes in Concentrator Shape

Adhesives

A 3M representative told us that tapes such as #600 that have aggressive "tack" tend to have high creep, and that tapes #850 and #853 had much better creep resistance. We obtained samples of the #850 tape first and compared its creep rate under modest shear load to that of the tapes we had used on the first three concentrators. Tape #850 was about 6 times better than the #600, so we used it on concentrator #4. We polished the strips to obtain a specular enough finish to allow optical testing. The test results later indicated that stowing and re-deployment still caused some creep (see 1.5).

We did not receive the #853 tape until after concentrator #4 was under assembly. The #853 tape turned out to be about 24 times more creep resistant than the #850. We used #853 for concentrator #5, and for much of our remaining Phase II work (rim seams, etc.) because of its high creep resistance and the ease of use of tape-type adhesives.

We also did tests on non-tape adhesives, many of which have negligible creep. We found that Loctite Speedbonder 325 might be suitable for use in orbit. It is an activator-cured acrylate that retains remarkable strength, adhesion, and flexibility even at 77K. We made joints with it and tested them by cycling 50 times between 77K and 300K. It worked very well. It is not space-qualified, because of excessive outgassing, but a Loctite technical representative indicated that they could probably find variations on formulation and cure that would reduce the outgassing enough to allow use in the shuttle and on space station.

We eventually obtained six more adhesives and technical data on them. They included cyanoacrylates, methyl acrylates, and epoxies. We tested each adhesive as follows: we bonded together two pieces of foil and thermally cycled the specimen 50 times between a LN2 bath and room temperature. After each cycle, we wrapped the specimen around a 6" diameter cylinder. Based on these tests, lap shear strength data provided by manufacturers, and information in the NASA "Compilation of VCM Data of Nonmetallic Materials" handbook, two epoxies manufactured by Crest Products Corporation were the best structural adhesive candidates for use on flight concentrators. We eventually selected Crest epoxy 7344 as most suitable. We used it to bond corrugated stiffeners to our final concentrator.

16% Central Sections

We learned a great deal from the 10% scale concentrators, but uncertainties associated with compensating for both gauge and scale effects in scaling from 10% to full-scale led us to experiment with central sections of full-scale concentrators much earlier than originally planned. Our main interests for these experiments were shape stability through a stowage-deployment cycle, and membrane resistance to "dimple" deformations. These properties depend on local membrane characteristics, and can be adequately tested using only a small representatively curved section of a concentrator. Thus we did not need a large or accurate paraboloid, but merely a shallow dish that was stable enough in shape to allow us to check for membrane shape changes. After investigation, we found that Orange County Machine Works could power-hammer a large flat 3/16" steel plate to "dish" it to our specifications. (This process is used to form domes for gas storage tanks.) Limitations on door sizes at our facility constrained the form to a 2.4 m diameter. This made the dish only about 6 cm deep. We assembled concentrators #6-8 on this form. Their characteristics are listed in Table II:

Table II Characteristics of 16% Central Section Concentrators #6-8

<u>#</u>	<u>Alloy</u>	<u>Temper</u>	<u>Thickness</u>	<u>Finish</u>	<u>Bonding</u>
6	1350	H19	.005"	mill	3M #853 tape
7	1145	H19	.005"	Nameplate	Loctite 325
8	1145	H34	.010"	mill	Loctite 325

These sections were made by laying down twelve 7.5" wide strips on the form, with no overlap, and bonding a 1" wide strip of tape (#6) or aluminum sheet (#7-8) over the edges. No rims were used to stiffen them around the periphery.

We tested these concentrator sections as follows. First we suspended concentrator #6 vertically from a series of lines. It maintained its dish shape in the one-gee environment. Next we weighted it down with 18.5 lbs of weights suspended from the bottom. This caused the same loads that would be carried by the top central portion of a full-scale concentrator in one gee. The membrane collapsed into an accordion shape by folding at the tape joints.

We made concentrator #7 using a stiffer joining technique: a lap joint made with the same material as the concentrator itself. This guaranteed that the minimum stiffness at the joint was no less than that of the basic membrane.

When we weighed concentrator #7 down with 18.5 lbs, it behaved much better than #6: the accordion shape observed with #6 changed to a shape consisting of 2-3 "waves" with vertically oriented peaks and troughs. But concentrator #7 was not very resilient. Small dimples induced on its surface did not spring back automatically (either with or without added weight). Adding a second set of .005" bonding strips over the joints to locally stiffen them did not improve the dimple resistance or resilience significantly.

Concentrator #8 used the same design as #7, but it was twice as thick. This gave it eight times the stiffness, or four times the specific stiffness of #7. This concentrator performed much better. It was .010" thick, so we weighted it down with 38 lbs. It maintained its concave shape well both with and without added weight. Furthermore, when we induced dimples by pushing in a seam from the convex side, we found the dimples quickly vanished when we let go, both with and without the added weight. We also found that by pushing with two hands (rather than with one finger), we could push near the center of the convex side until the dish was over 20° from the vertical, without the dish inverting. We then stowed concentrator #8 around a 10" diameter cylinder, redeployed it, and placed it back on its form. We found a few new deviations of the concentrator from the form. These deviations were of order 1-10 mils in magnitude.

As described in section 1.2.3, we later added corrugated stiffeners to concentrator #8. This further increased its specific stiffness and hence its resistance to dimpling.

Conclusions on adhesive bonding on jigs

We learned several things from our work with adhesive fastening on jigs:

- o Suitable adhesives (i.e., ones with high enough creep resistance, etc.) are available, but it may be necessary to flight qualify them.
- o Good smooth seams are possible but difficult to obtain. A major difficulty, even on 10% scale models, is properly aligning the strips for fastening.
- o Keeping strips properly aligned during adhesive curing and further assembly is not entirely straightforward.
- o Assembly work will probably be simplest if the jig is vertical. This requires a high bay and a "cherry picker" for assembly work.
- o The jig needs to be collapsible for storage and transportation. This and its size and accuracy requirements would most likely make it expensive.

These difficulties with adhesive fastening and with jigs prompted us to investigate alternative alignment and fastening methods and eventually led us to develop the mechanical fastening techniques discussed in the next section.

1.1.2 Self-Aligning Mechanical Fasteners

The difficulties we encountered with adhesive fastening of strips on jigs led us to investigate mechanical fastening techniques. We started by considering drilling precisely positioned holes in the edges of adjoining strips as an aid to positioning. Then we realized we could use a lap joint, align holes in adjoining strips on top of each other, and use the holes to both position and clamp strips for adhesive bonding or even spot-welding. We then realized that pop-riveting strips together might provide adequate positioning and joining by itself.

As we envisioned this process, we would rivet strips together one by one, and roll up the concentrator as we assembled it. This eliminates any need for a precise full-scale jig, and allows all work on even a full-scale concentrator to be done at a convenient working height and position, in any room longer than the stowed length of the concentrator. We used two variations on this technique for the last two concentrators fabricated under this contract. Those concentrators are described below in Table III:

Table III Characteristics of 2.8 Meter Diameter Concentrators

#	Shape	Focal length	#Strips	MaxWidth	Fasteners	RimType
9	Spherical	6.0 m	22	14 cm	Pop-rivets	Separate
10	Paraboloid	1.2 m	40	9 cm	Eyelets	Integral, Triang.

These were both made from 10 mil 5657-H28 "Nameplate" finish aluminum sheet, by cutting and stacking barrel-stave-shaped strips, drilling rows of holes down each side, and aligning and joining the strips using fasteners through the holes. A separate strip of .002" thick tempered aluminum foil was corrugated axially and adhesively bonded to the back of each strip to stiffen the membrane against "dimple" type distortions.

This assembly process does not eliminate the need for precise jiggling, but it reduces the jiggling to that required to precisely drill holes on each strip. As discussed in 1.2.2, we used a step-and-repeat operation on a mill to do this.

We procured samples of aerospace grades of pop rivets and learned how to make good joints. We made tests with two closely spaced rivets at one edge of a pair of long coupons. The tests showed that the hole alignment errors were <0.001". We then fabricated an 11.5" x 26" paraboloidal panel by riveting together 6 strips of .010" aluminum. We were impressed by the simplicity and cleanliness of this method and decided to pursue it further. We then assembled a 20% central section of a full-scale concentrator with pop rivets. It is shown in Figure 5, along with the cylindrical mandrel onto which it was wound as it was assembled.

We made this concentrator section spherical rather than paraboloidal because it allowed us to use identical hole patterns on all strips. This allowed us to stack and drill all strips at the same time. The spherical shape also simplified optical testing in relatively close quarters. The increase in radius of curvature between the center and edge of this part of a concentrator is <2%, so using a sphere does not change handling properties significantly.




Figure 5 Riveted 2.8 meter Central Section and Assembly/Stowage Tube

Over 95% of the concentrator fabrication effort turned out to be in fabricating the individual reflector strips (see section 1.2). More elaborate jiggging would have considerably reduced our workload and would be used for full-scale concentrator fabrication, but it did not seem cost-effective to design and fabricate such jiggging for our limited needs during this contract. Final assembly of the elements took only a few hours and involved no unpleasant surprises. We were pleasantly surprised by how quick and trouble-free the assembly was.

Our fabrication and simple testing of this shallow central section provided three benefits. First, it gave us increased confidence that this basic concept and many design details were worth pursuing. For example, as described in section 1.2.3, we reinforced this shallow dish against "dimple" distortions by attaching corrugated foil to the back of the membrane. When we pushed on the back of the dish to create "dimples" in it, the dimples generally disappeared as soon as we released the pressure. This reassured us that our corrugated stiffeners were adequate for full-scale concentrators with similar radii of curvature. This dish also allowed us to fine-tune our design and techniques to improve performance and reduce labor on the final 20% scale concentrator. It allowed us to use narrower strips and corrugations, and much closer rivet spacing. Finally, we used this concentrator section to test techniques that might be used in orbit to attach a membrane to a separate rim. Our "finger-pocket" design allows more compact membrane stowage and decouples membrane and rim thermal response. (See section 1.3.2, and the videotape sent with the report.)

On the other hand, it was clear that our riveted seams were not yet of high enough quality for our deliverable concentrator. We knew that we would have to use far closer rivet spacing for a strong and well-defined seam, and we had to reduce local distortion around the rivet. So we began methodical experiments with pop-riveting and similar processes.

Rivet Tests

We began by comparing the pop-rivets used in the central section with another type of pop rivet from the same manufacturer, Cherry Textron. Both had a nominal diameter of $3/32$ ". The type used for concentrator #9, AAP-32, were made of 5052 aluminum and cost \$27.54 per thousand. The second type, CR9117-3-2, were made of 5056 alloy, had a pull-through mandrel, and cost \$190 per thousand. Their weights after installation were similar: .07 gram and .06 gram respectively. The rivets with the pull-through-mandrel could be installed with almost no shock to the material being riveted. They also had a neater appearance after installation than the first type. However, they deformed the surface around the holes more than the AAP-32's because of higher radial forces induced by the pull-through mandrel.

We experimented with a variety of methods of riveting .010" thick 5657 H28 1SB (One side bright) aluminum sheet. We fabricated six small concentrator sections measuring $20-1/2$ " x 23". Each consisted of four .010" thick, $5-1/2$ " wide 5657 H28 aluminum strips and three seams of AAP-32 rivets spaced $1/2$ " apart. One of these sections is shown in Figure 6.




Figure 6 Sample Concentrator Section - Four 5.5" Wide Al Strips Riveted Together

We found that a back-up washer was essential to prevent large deformations around each rivet. We used .025" thick, .250" dia steel washers, which weighed .13 gm. Aluminum washers which weigh only about 0.05 grams, or .12 grams with the rivet, are available but have to be special ordered. Placing the back-up washer on the material's shiny side (and the rivet head on the dull side) further reduced deformation. The rivet body and washer cover an area of .32 cm² and have a dull finish.

We estimated that a full-scale concentrator would have about 450 meters of seam. We selected a rivet spacing of 1.5 cm as a compromise between strength and distorted area. This means that a full-scale concentrator will have about 30,000 rivets and washers. They will weigh about 3.6 kg (2% of the concentrator mass) and cover about .6% of its area. (By comparison, the overlap of adjacent strips causes a higher mass penalty.)

Next we contacted a company called Avdel Corporation that makes a repetition riveting system. Mark Grenda, one of their representatives, visited us and demonstrated their Chobert repetition riveting system. Although the seams he produced during his visit were not nearly as good as those made with our hand-held rivet gun, the Chobert system was very attractive for a full-scale concentrator, because it was fast. Therefore, we gave samples of aluminum sheet to Mr. Grenda to see if he could improve on the Chobert seams. After two false starts (failing to use back-up washers, and using brass rivets), he was able to get satisfactory results. Meanwhile, we discovered that using a slightly oversized hole diameter (.100" instead of .098") and an undersized mandrel (.071"-.072") gives distortion-free rivet seams. It does so, however, at the cost of a small uncertainty in positioning.

Even with these improvements, we were not entirely pleased with either hand-held or automatic riveting techniques, because of the need for back-up washers and because of the sensitivity of performance to technique. We then came across the eyeleting process.

Eyelets

Eyelets are used in an very wide range of products, from clothing and shoes to printed-circuit boards and ram-air parachutes. An eyelet consists of a short length of ductile tubing, one end of which has been expanded and flattened into a flange. An eyelet is "set" by compressing it between two precisely shaped tools in a press. One tool holds the eyelet, and the other flares the unflanged end and rolls the flared region out and down until it clamps the sheets being joined, a short distance outboard of the holes. Increasing the setting force allows a pre-load to be applied between the joined sheets.

We obtained literature from several companies, obtained suitable eyelets and forming tools, and began experimenting. Our experiments involved making 3 small rectangular concentrator sections. Each section consisted of six 3.4" wide, 18" long, .010" thick 5657 H28 aluminum strips with five seams of brass eyelets spaced .6" apart. All three sections had a 2.4 meter radius spherical curvature, the same as the central section of a 1.2 meter focal length paraboloid. The first two sections had surface distortions and hole position errors as great as .002" to .003".

We traced the surface distortions to the fact that the holes used for these two sections had a nominal diameter only .001" larger than the eyelet body (.090" and .089", respectively) and to the fact that the hole drilling procedure we used resulted in holes of non-uniform diameter. We corrected the drilling procedure and drilled .004" oversized holes in the third section's strips. The surface of the third section was virtually distortion-free. We decided to use eyelets for the final 20% scale concentrator, with holes between .001" and .004" oversize, and set about doing further experiments to better understand the eyeleting process and its idiosyncrasies.

We did alignment tests similar to those we did with pop rivets. We eyeleted together pairs of sheets with three holes gang-drilled in them. Two holes were spaced .25" apart, and a third 2.25" away. The two close holes are eyeleted together. We found that with the proper eyelet, hole, and technique, the worst case displacement of the third hole was about .01". This indicates hole alignment errors in the eyeleted holes of order .0005".

We realized that two factors contribute to eyelets causing less distortion than pop rivets. First, the shock caused by the breaking of the central shaft of a pop rivet often bends the sheet around the hole. An eyelet is set by compression, and involves no shock load. The other difference is that in clamping up, a pop rivet exerts a significant radial force against the inside of the hole, particularly on the blind side. An eyelet exerts an axial clamping force, with some radial force at the end of the hole where the eyelet flares during setting.

We needed to use a reinforcing washer to avoid distortions with eyelets, but the washer is as effective on the dull side as on the bright side. (By comparison, pop rivets worked best with the washer on the bright side.) Hence blockage of the concentrator's reflective surface is far less with eyelets than with pop rivets.

We did our initial experiments with brass eyelets .089" in diameter because they were readily available and inexpensive, roughly \$15 per thousand. We used aluminum eyelets for the final concentrator to reduce the risk of corroding either the reflective surface or the eyelets. Timing constraints raised our costs. Stimpson Corporation in New York could have made a minimum order of 500,000 for us at \$1.08 per thousand, but their lead time was unacceptably long. We had to place a special order of 10,000 from Zeman Manufacturing at about \$100 per thousand (including tooling costs). Those eyelets have a flange diameter of .150", a body diameter of .094" and a body length under the flange of .100". They are made of 3003 quarter hard aluminum. The eyelets weigh .020 gm, and are used with .025 gm aluminum washers .200" in diameter.

We set the eyelets using an arbor press with a 4" throat, deep enough to accommodate the 3.4" wide strips used for the 20% scale concentrator. We placed this arbor press on a small dolly so that we could roll it along the seam to position it for each eyelet. We then attached an 8.25 lb lead weight to the arbor press arm. During setting the arm went over center, and it was nearly horizontal during final clamp-up. The weight provided a reproducible setting force. We moved the weight around, and found that a 7.5" arm gave the best results.

We spent roughly 70 hours joining together the 40 strips of the final 2.8 meter concentrator, using nearly 6,000 eyelets. We rolled the concentrator up on a 9" diameter, 126" long tube as we assembled it, just as we had with the preceding concentrator section. Our handling and stowing operations took advantage of the central holes in the eyelets. Kevlar thread looped through neighboring eyelets provides a convenient way to hold the concentrator, and the 39 rows of eyelets allow such loops to be attached nearly anywhere on the membrane.

Overall, we believe that eyelets are the method of choice for assembling concentrators made using our strip concept. They cause less distortion and blockage than pop rivets; weigh less; cost less (in quantity); have a lower profile (which allows tighter stowing); have smoother surfaces that are less likely to scratch adjacent layers during stowage, and provide a simple means of attachment for handling during assembly, testing, and stowing for launch.

1.2 Strip Design and Fabrication Details

Section 1.1 discussed our design and assembly concepts for concentrators, and noted that almost all of the effort went into making the individual strips that were then joined together. This section goes into more detail on the design and fabrication of those strips. It covers materials selection, hole drilling, stiffeners, and coatings.

1.2.1 Material Selection

Our baseline design at the beginning of the contract used barrel-stave-shaped strips of tempered aluminum alloy sheet. But we investigated other materials that might reduce weight or increase stiffness for our 2.8 meter concentrators and for full-scale concentrators. Our materials studies are discussed below.

Graphite Fiber Composite with Thermoplastic Matrix

We purchased a 1 lb sample of unidirectional carbon fiber composite manufactured by Phillips Petroleum Company. The fibers are embedded in PPS (poly-phenylene sulphide) high-temperature thermoplastic. This material is very strong and flexible and can be thermally formed and bonded. Its surface can be made quite smooth during forming. We laid up and bonded several laminates in experiments aimed at learning how to use this composite material. In one, we laid up a 0/90/90/0 laminate that showed (as expected) much greater bending stiffness in one direction than in another. In a second experiment, we were able to obtain a relatively specular final surface by heating and cooling the composite in contact with a glass surface. Our final experiment was to bond 4 plies on a compound-curved template. Our main conclusion from these experiments was that this material is more difficult to work with than we had expected, so we did not do any more work with it in this contract.

Foil-faced Sandwiches

We began with a paper core, with the paper simply serving as a low-cost surrogate for a thin layer of rigid foam. We glued aluminum foils of varying thicknesses to both sides of paper of varying thickness. The resulting sandwich was naturally significantly lighter than a monolithic metal sheet of the same stiffness. This led us to procure samples of 3 lb/ft³ and 18 lb/ft³ rigid polyurethane foam from General Plastics Manufacturing Company. We bonded 1 mil aluminum to both sides of foam sheets, using either a foil bonding adhesive made by Krasel Industries, or a double-sided acrylic adhesive sheet made by Betham Corporation. Betham's sheet adhesive was more convenient to use, and preliminary temperature tests indicated that it may be suitable for concentrators in low earth orbit: it survived exposure to 188 K, a typical "nightside" temperature. Extended cycling tests would be needed for more confidence.

Aluminum-Lithium (Alithalite) Alloys

Alithalite and other lithium-containing aluminum alloys have higher elastic moduli and lower densities than other aluminum alloys. This makes them attractive for use in our concentrator. However none of them were available on a commercial basis at the time we needed them. In addition, we do not know whether these alloys will be able to be rolled with an adequately specular finish, nor whether the corrosion resistance will be adequate.

Magnesium Alloy

Magnesium alloy AZ31B is an option, but it is expensive. In addition, as noted in work at NASA Lewis on magnesium reflectors, it is difficult to get a good specular surface on magnesium. Finally, since our corrugated stiffeners provided anisotropic stiffening, there was less need to reduce mass or increase stiffness by going to lower-density alloys.

Aluminum Lighting Sheet

Based on our success with anisotropic stiffening (see 1.2.3), we decided to use ordinary aluminum alloy sheet for our final concentrators. We obtained samples of Alcoa's Alzak and Coilzak, together with similar sheet from other suppliers such as Metalloxyd and Kingston Industries. These materials go by the generic name of "lighting sheet", and are used when high reflectivity is needed in the visible, such as in fluorescent lamp fixtures. They are made from 1000 series aluminum alloys (99% aluminum, minimum), coated or co-rolled with high purity aluminum (99.8-99.99%), and then hard-anodized about 2 microns thick. The yield strength of the base alloy is generally fairly modest, due to the low alloy content. The minimum thicknesses available from different suppliers were generally .012-.020". The low yield strength and the thickness were both of concern to us.

Final Material: 5657 Aluminum

Discussions with suppliers indicated that United Aluminum in Connecticut was known for being able to roll highly specular sheet. United's engineers said that they could produce sheet down to 8-10 mil that would have a highly specular finish on one side, if they used 5657 alloy and ran it right after polishing the work rolls. They explained that 5657 alloy is essentially a high-purity aluminum to which 0.8% magnesium is added. The magnesium provides a considerable increase in yield strength in the as-rolled condition, without affecting surface smoothness significantly. (They said that it is mainly iron and silicon that limit surface specularity.) United's minimum order was 250 lbs plus a mill allowance of up to 20%. The cost was under \$4/lb. We wanted 8" wide material for a 2.8 m central section, so we ordered 300 lbs of 8" wide, 10 mil thick 5657 1SB (one side bright) sheet in the H28 temper (rolled to extra-hard condition and partially annealed to "full-hard" condition.)

The material we received had an exceptionally good surface finish. Informal optical tests with a laser in a darkened room showed that the main non-specular reflection in the surface was from residual scratches aligned with the rolling direction. These scratches appeared to only involve a few percent of the light reflected.



1.2.2 Alignment/Fastening Holes

When we realized we could use pop rivets or eyelets to align and assemble concentrators, we also realized that precise location of the holes along the edges of each strip was the critical step in determining the accuracy of the overall membrane shape.

We knew that drilling was likely to be a major cost in a full scale concentrator fabrication effort, so we worked on ways of reducing those costs. We stacked 24 sheets of .010" sheet, clamped them between two plates, and drilled a pattern of holes through the stack using our mill. Comparisons of the top and bottom sheets with a dial caliper showed that drill drift within the stack was <1 mil. This gave us confidence that we could drill all strips for our 2.8 meter spherical central section at the same time. We stacked all 22 strips together, clamped them between two 1/4" aluminum plates, and drilled the stack on our mill, using a simple sinusoidal width variation. (To be precise, the width must vary sinusoidally along the length of the curved edge of the strip, not along the length of the center of the strip.) The mill and strip holder are shown below in Figure 7.




Figure 7 Mill and Stacked-strip Holder Used for Drilling Strip Holes

We learned from inspection of the assembled section that the differential longitudinal spacing of holes on strip edges that were riveted together were unacceptably large. It was often difficult to put a rivet through two "mating" holes. And after the rivets were installed, we could often see a small buckle on one side of the seam, between two adjoining rivets.

We decided to eliminate the buckles on our next concentrator by drilling all mating rows of holes at the same time. We did this by stacking adjoining strips face-to-face for drilling. We also realized that we needed to improve the accuracy of the transverse positions of the holes. We purchased an X-Y optical readout which read the mill table position to .0005".

The mill had a much shorter stroke than our strip length, so we knew that we would have to periodically loosen the vice, slide the stack of strips axially, and re-tighten the vice. We were not sure that the bottom plate which the vice gripped was accurately straight, so we developed and tested a technique for re-zeroing the optical readout after each shift.

We measured the hole positions by their offset from a thread stretched between two "toolmaker's buttons" attached to opposite ends of the top plate. This technique allowed us to periodically re-position the stacked strips without worrying about the straightness of the bottom plate. We drilled eight #42 holes through a stack of eight 5-1/2" x 7" pieces of 10 mil thick aluminum in a straight line using this method, moving the assembly in the vice jaws three times. When we riveted the samples together, flipping adjacent sheets over to double the transverse errors, no distortion was evident. After the thread broke, we replaced it by a .005" steel wire for our remaining work with the technique.

We determined the effects of transverse joint errors by riveting together two sets of two pieces of 16.4" x 5.5" x 10 mil aluminum using holes that were slightly offset with respect to each other, as shown in Figure 8, below.




Figure 8 Hole Pattern for Joint-Error Samples

We drilled 9 holes 1.6" apart along one straight line, and 10 holes 1.6" apart along another straight line that intersected the middle of the first at an angle of .020 degrees. Thus the holes at the ends of the second line were .0025" away from the line defined by the first set of holes. When two pieces were held together for riveting and one of them was flipped over, this "error" was doubled: if the first set of holes were lined up exactly, there would be a .005" offset between outer holes of the second set. After riveting was completed, we found that the seam was indistinguishable from those prepared using holes drilled along a single straight line. The holes for these tests were .004" oversize. This accounted for some of the results, but it is not clear how much, because pop rivets do expand somewhat when they are set. It is possible that for high-frequency errors up to several mils, most of the stresses are relieved by in-plane distortion rather than out-of-plane bending. We found this quite encouraging, because it indicated that the riveting operation would tend to "average" over local deviations from the desired profile.

These experiments indicated that the mill should allow us to drill holes accurately enough to fabricate a high quality paraboloid, so we decided to use this technique to drill the holes in the strips for the final deliverable 2.8 meter concentrator.

We then needed to calculate the exact hole positions required for our deliverable concentrator.

There are various ways to slice a paraboloid into strips. The one that appeared most useful was to define the rows of holes by geodesics parallel and uniformly spaced at their middle. The individual strips are then "inscribed" in the paraboloid. Using geodesics makes both sides convex, so the source material need be no wider than the center width of the strip. It also makes adjoining edges have the same edge curvature, so they can be stacked face to face and drilled at the same time. This eliminates any possibility of longitudinal offset between mating holes. But it does make transverse errors more important, because errors in mating holes correlate. Since the concentrator has bilateral symmetry, we can drill four strips at once: two strips for each of two identical edges equally spaced from the center.

Derivation of the edge shape went as follows. We knew from an analysis during Phase I that the local radius of curvature of the membrane surface (i.e., the sphere that fits best locally) is twice the distance to the focus. We also knew that the radius of curvature of the row of holes needs to be equal to the square of the local surface curvature radius, divided by the local half-width of the strip (center to hole). We wrote a program that "marched" geodesics along the paraboloid, to determine the distance to the focus as a function of length along each geodesic. We then marched the edge along each strip, determining local edge curvature from surface curvature and strip half-width, and double integrating edge curvature along edge length to get half-width. The program is written in Turbo-Pascal, and is called "StripWid." A copy is being provided to NASA Lewis on a floppy disk.

We then did a brief analysis of the effects of different types of fabrication errors. They are discussed below, in order of increasing severity:

Uniform displacement of an entire row of holes

If a strip is positioned wrong during drilling, the entire row of holes will be offset. If this is the only significant error, it is equivalent to slicing apart an ideal concentrator along that row and displacing the two pieces by the amount of that error. A transverse error makes a strip wider or narrower, and displaces the two pieces away or toward each other. Similar errors on all strips change the focal length and add a small amount of spherical aberration. Opposite errors in nearby edges nearly cancel out, except between the two offset edges.

Errors in radius of curvature of edge

If a mating row of holes has 10% less curvature than desired, the overall effect on the membrane is similar to having a 10% increase in one strip's width: the two halves of the concentrator are displaced .1 stripwidth from each other, causing the two halves to come to a focus .1 stripwidth apart. As with uniform displacement of a row of holes, opposite errors in nearby strips cancel out overall, except for a focus displacement of the intervening strips.

Rotation of a entire row of holes

If a strip is rotated while being stacked before drilling, its width error would vary linearly along the length. To a first approximation, it would rotate the two pieces of the concentrator, and not cause any problem. If this happened to all the strips, it would also shift the center of the paraboloid toward the narrow ends of the strips.

Higher-mode errors

Waviness along a row of holes will cause a washboard displacement of the row of holes out of the plane of an ideal paraboloid membrane. For a given linear error amplitude, the resulting slope error along the seam will vary with the frequency. However the effects of high-frequency modes decay away from the seam more rapidly than low-frequency modes do. Overall, we were far more worried about higher-mode errors than the other errors.

Averaging Edge Shapes

Ideally, there are 20 distinct hole patterns in a 40-strip concentrator: one for each seam having a different transverse offset from the center of the paraboloid. However, the hole patterns vary slowly and smoothly. (The edge curvature varies more quickly near the edge, but the shorter length of these strips cancels this out.) We realized that we could group nearby edges together and drill the entire group using an averaged hole pattern for that group. Opposite edges of one strip should not belong to the same group, because they cannot be stacked for drilling together. This meant that hole pattern groupings would have to be interlaced, and that strip groupings might be different for odd and even edges.

The key thing is to keep integrated errors small, so errors in adjacent seams cancel rather than adding up. We grouped the 20 ideal shapes into six groups. The six groups use median member shapes (averaged shapes should be better), with shapes 1 and 6 adjusted to cancel errors due to using shape 2 near the center and shape 5 near the edge. The six shapes are used in the pattern shown below in Figure 9. When used in this pattern, overall errors remain quite small: the overall effect is that the centroids of images from different strips do not coincide, but are displaced along a line a small fraction of a strip width long.




Figure 9 Pattern of Edge Shapes Used in Final 2.8 Meter Concentrator

Drilling procedure

We cut all strips to the same "barrel stave" shape, so different strips differ only in their hole patterns and their length. Cutting strips to shape and removing coil set (discussed in section 1.2.1) required nearly one hour per strip. After we cut out all 40 strips, we stacked a group of them between two 1/4" aluminum plates, bolted the plates together, and drilled one row of holes using hole pattern #1. We then unbolted the plates, regrouped strips, and drilled hole pattern #2, and so on. We drilled the holes .600" apart with a .096" diameter drill bit. To reduce the likelihood of errors, we printed out a computer listing for each row of holes, with .6" spacing between lines, and taped the list to the top plate so that the transverse position for each hole was listed close to where the hole was to be drilled. The mill setup for this drilling procedure is shown below in Figure 10.




Figure 10 Operator Drilling Holes for Joining Final 2.8 Meter Concentrator

After drilling 7-8 holes, we always re-checked the mill position relative to the wire. It frequently drifted by several thousandths of an inch. At first we thought this was due to irregularities in the mill table and possibly the readout itself. We realized later that most of the drift was probably due to the wire itself moving, due to the following mechanisms:

- 1) The ends of the bolted assembly plate were unevenly supported, so they twisted. This moved the tops of the wire support pins, where the wire was attached.
- 2) Variations in support height at the ends of the bolted plate assembly caused wire tension to vary. The wire was wrapped around the pins without pre-forming it to the pins, so tension variations altered the wrap angle and wire transverse position.

Each time the center line drifted, we re-zeroed the digital read-out and slightly enlarged two or more nearby holes before proceeding. Unfortunately, this seems to have introduced more errors than it cancelled, and may have been the largest source of error in drilling the curved row of holes.

This process can be drastically improved. One simple way is as follows. First make two wide strips with holes in a straight line. Use hole spacing of several inches, and mark the vise position on the supporting plate for each step-and-repeat position, but do not try to correct for any curvature of the supporting plate that is clamped in the vise. Then remove the two sheets, flip one over, and join them together at the two outermost holes. Lay the sheets down on a table, tap lightly until they are in an equilibrium position, and clamp the sheets together in this position. Then measure the offsets in the intervening holes and divide by two to obtain the local offset for each end of each step-and-repeat operation.

If the plates are positioned in the mill vise in the same position and with the same clamping force when the next set of strips is drilled, and if the plates are not mishandled in between, then those offset errors will recur. This allows them to be compensated for. The accuracy of the compensation can be determined by repeating the above calibration process, this time compensating for the measured offset. To speed drilling and reduce the possibility of error, the offsets for each end of each step-and-repeat position can be loaded into the computer program that prints out all the drilling positions for each group of strips.

Stacking strips, drilling a hole pattern, and unloading strips took an average of 12 hours, so we would have spent roughly 72 hours drilling eyelet holes if all had gone smoothly. However, we made an error in drilling the edge 3 hole pattern that caused a small delay. When we positioned the top plate over the stack to drill group 3, we dragged several concentrator strips out of position. We discovered this problem after we drilled about a third of the holes. We went back and re-drilled about a third of the row, by drilling holes offset .300" axially from existing holes, using a new computer hole position printout for the new axial positions. Hence, strips with an edge 3 hole pattern have an extra row of holes along part of their length. We prevented this problem in drilling subsequent groups by adding "inspection holes" in the top plate. These holes allow us to verify strip alignment visually, after the top and bottom plates are bolted together.

When we cut to length the strips for the final concentrator, we made another error which caused a few strips to be slightly short. We compensated for this after the membrane was assembled, by taping small triangular extensions to the affected strips to fill in the missing area. This locally reduced the stiffness of the membrane at four locations around the periphery. The small triangular regions had a different slope error than the strip they were attached to. This can be seen in the test results in section 1.5, between 4 and 5 "o'clock" on the periphery of the concentrator.

1.2.3 Anisotropic Stiffeners

A preliminary analysis of stiffening options for the membrane led us to experiment first with a very simple corrugation scheme. The concept uses sheets of 2 mil 1145 H19 aluminum foil corrugated along their length. The sheets are bonded to the back side of the concentrator strips. Bonding is done with the strips flat. When the strips are curved during concentrator assembly, the ridges of the corrugations go into tension. This helps protect the thin ridges against buckling during handling. We fabricated a "press brake" corrugating jig, sketched in Figure 11. After experimenting, we learned how to make uniform corrugations with flat enough bottoms to allow good adhesive joints to the concentrator strip.




Figure 11 Cross Section of Initial Corrugating Tool (Similar to a Press Brake)

After doing some informal handling tests on concentrator #8 (a 10-mil thick shallow central section of a full concentrator), we stiffened it with 6.5" wide, 2 mil thick corrugated strips. Figure 12 shows the stiffened concentrator, mounted on the 2.4 m diameter steel form on which it was constructed.

The corrugated stiffeners greatly improved the concentrator's handling properties. The corrugated stiffener has a very desirable anisotropy: it increase the transverse stiffness much less than it increases the longitudinal stiffness. As a result, a stiffened concentrator can be stowed almost as tightly as an unstiffened one. This leaves the integral rim as the constraining factor for stowage size. Based on this experiment, we decided to refine this stiffening concept for use in our remaining work.




Figure 12 Concentrator #8 on Its Fabrication Jig (note corrugated stiffeners)

We bonded corrugated sheets to several short strips and tested their properties. The measured increase in axial bending stiffness varied from a factor of 20 to 40, depending on the sample length. For a mass increase of <50% and a simple, easy to fabricate stiffener, a factor of 20-40 improvement in stiffness was quite attractive. However, preliminary calculations based on extension and compression only, neglecting web shear, indicated that the axial bending stiffness should increase by a factor of about 200. For the short sample lengths tested, web shear strain limits the amount of stiffening the corrugations provide. This means that the resistance of the concentrator to local "dimple" distortions is less than the larger-scale stiffening effect. This was of concern, because it is primarily the dimple resistance of the membrane that we want to improve. We knew we had to use narrower corrugations to increase the amount of web material and hence the dimpling resistance.

For concentrator #9 (the riveted 2.8 meter diameter central section), we used the same stiffener material (1145-H19 foil, .002" thick), but we refined the corrugation design. The strips are about 4.25" wide. They have 11 ridges .19" wide and .06" high, and ten .12" wide channels. The closer spacing increased the amount of web material, and increased the stiffness against dimpling. The lower corrugation height increased the amount of axial curvature the strip could tolerate without yielding or buckling. We corrugated the strips using the same type of "press brake" tool as before. We glued them to the reflector strips with Crest 7344, a two-part epoxy adhesive which meets NASA outgassing specifications. Short sections of stiffened reflector strips were about 50 times stiffer than unreinforced strip sections. This is a significant improvement over the 20-40 increase in stiffness achieved previously with slightly higher but much wider corrugations. The new strips also could tolerate more axial flexing than the first strips without taking a set.

We investigated the possibility of having the foil stiffeners corrugated commercially for the final concentrator. We found one aluminum roll-former (Woodside Engineering) that said it could corrugate tempered aluminum foil into the shape we needed, but they could not do it until after our contract ended. So we designed and made a corrugating tool ourselves. We spent several days making a small prototype tool. It uses two 2" wide rollers. A feed guide assembly pre-curves the foil transversely into a gentle sinusoidal shape to foreshorten it, so the foil does not tear when it reaches the forming rollers. We then spent about two weeks making a tool capable of producing a 5.9" wide strip having fifteen .05" high, .09" wide flat ridges spaced .375" apart. The tool is shown in Figure 13.




Figure 13 Longitudinal Corrugating Tool for Reinforcing Strips

The longitudinal corrugating tool has two mating steel rollers, the first of which drives the second through a set of gears. An array of 35 ramped guides forms the foil into a sinusoidal corrugation before it reaches the forming rollers. The guide array is in a chevron pattern, so it sinusoidally corrugates the foil first in the middle, and then gradually outward on each side. This minimizes transverse stresses and tearing during the roll-forming operation. It was difficult to fine-tune the rollers and guides to avoid slitting the 2 mil thick 5052 H38 aluminum it was designed to corrugate. We had to compromise by making the ridges lower than originally planned and their sides more gradually sloping. However, we were then able to produce reinforcing strips quickly and continuously.

For the final 20% concentrator, our intent was to mimic as much as we could the one-gee handling properties of a full-scale concentrator that could have deeper and much stiffer corrugations. To do this, we decreased the corrugation height from .05" to .025". The resulting 20% model has the same mass per unit area and comparable stowage diameter as a full-scale model, but the shallowness of the corrugations makes it axially less stiff.

Compact stowage with an integral rim requires that the membrane be thin and hence unstiffened near the periphery, where it constitutes one side of the integral triangular rim. The rim includes the outer 6" of the concentrator membrane. To leave adequate margin for rim joints, the corrugated stiffeners end at least 1" from the inner edge of the triangular rim. This means that at least 7" of the concentrator closest to the periphery (or over 20% of the overall concentrator area) is unstiffened.

The final strip/stiffener design on the final concentrator is as follows. The strips are all 3.4" wide in the center, and they all have the same barrel-stave taper toward a width of 3.0" at each end of a full length strip. (But not all strips are full length.) The rows of holes are all .25" inboard from each edge at the widest part of the strips, so the overlap there is .50". At other locations, the row of holes curves slightly differently for each group, so the overlap varies in the range of .45 to .55". The stiffeners have a fixed width of 2.4" over their length. They end when they come within 7" of the rim. The stiffeners are glued to each strip .1" off center. This gives enough unobstructed room on each side of the stiffener for the eyelet-fastened lap-joint overlap, even at the end of the stiffener. The stiffeners are made of .002" thick 5052 H38 aluminum foil. This alloy is much stronger and more resilient than the 1145-H19 foil used to stiffen concentrators #8 and #9. Each stiffener has five .025" high, .09" wide flat ridges spaced .375" apart on center. The stiffeners increase the strips' axial bending stiffness by a factor of 5-6.

The same logical error that caused us to cut some strips slightly short also led us to cut the ends of some corrugated stiffeners at the wrong angle. We left alone the side that was too short, but cut and removed the stiffener where it was too long. This gave us a smooth surface where we needed one for the inboard tape joint between the rim and the back of the membrane.

Cutting out, corrugating, and trimming the stiffening strips and gluing them to the back side of the strips for the 20% scale concentrator took approximately 75 hours. For a full-scale concentrator, fabrication could be done by Woodside Engineering or another commercial roll-former. The gluing process could definitely be sped up by using automated application equipment.

1.2.4 Reflective and Emissive Surface Coatings

We concentrated our coating efforts on improving the reflectance of the 5657 aluminum alloy sheet. We began by vacuum-depositing pure aluminum onto various aluminum alloy surfaces. We coated four 2" x 2" inch foil coupons with 99.997% pure aluminum. The results were unsatisfactory: the deposited material had a whitish finish and adhered poorly.

We then consulted L. Holland's standard reference work, Vacuum Deposition of Thin Films. This led us to believe that we should be able to increase reflectivity roughly .03 by rapidly depositing 500 angstroms of 99.997% pure aluminum in a high vacuum. We spent some time learning about evaporative and sputtering techniques, and a good deal more time cleaning, overhauling, and making various general improvements on our existing 6" VHS diffusion pump system, which was temporarily dedicated to this contract. We fabricated apparatus needed to do vapor deposition and ordered an ion-beam sputtering system.

We used a Varian Series 634 spectrophotometer fitted with reflectance jigs to compare the specular reflectivity of our aluminum coupons for wavelengths between 300 nm and 800 nm. The reflectance jigs for the test and reference beams are shown in Figure 14.




Figure 14 Schematic of Reflectance Fixture for Our Varian 634 Spectrophotometer

We determined acceptance angle using simple tests with a front-surface mirror shimmed to tilt it up, down, left, and right. Mirror tilts 2 degrees up or down dropped the measured reflectance by half. Comparable reductions in response for left and right tilts required angles of .7° and 2.3° respectively. We tested most of our 5657 samples with the sheet rolling direction (and resulting striations) vertical. This made the scatter from the striations horizontal. The .8 degree offset of an unshimmed sample eliminated most scatter on one side, but included more from the other side than is representative for a solar dynamic system. (Tests done with striations horizontal typically showed .05 higher reflectance.)

Because of our spectrophotometer's limited frequency range (55% of the extraterrestrial solar spectrum), we cannot confidently state the integrated solar specular reflectance obtained for any of our coupons. But we were able to compare coupons processed in various ways, and we believe the results are representative of relative performance over the entire solar spectrum. We found that 99.997% pure aluminum films from 750 to 1300 Angstroms thick, deposited at rates of 50 to 300 Angstroms per second at pressures of .7 to 3.0 microtorr all had roughly equivalent reflectances for wavelengths of 300 to 800 nm. Depending on the preparation method used, at a wavelength of 525 nm (near the peak of the solar spectrum), these reflectances were .11 to .18 higher than that of bare, clean 5657 aluminum alloy, which has a reflectance of about .67 at that wavelength.

We used ASTM Standard D-3359-83 to determine the adhesion of our coatings to the substrate. This test involves peeling off a strip of pressure-sensitive tape of known adhesion from a grid pattern cut on the sample coupon with a scalpel. 3M tape #850 had the most uniform adhesion of our various tapes, so we used it. Adhesion is rated from 0B to 5B, with 5B representing the best coating performance (see Figure 15).




Figure 15 Illustration of ASTM D-3359-83: Measuring Adhesion by Tape Test

To maximize the reflectance and adhesion of our vacuum-deposited coatings, we tried the following preparations:

- 1) Pre-cleaning with various organic solvents
- 2) Pre-cleaning with HF acid
- 3) Sputter cleaning immediately before deposition
- 4) Sputtering during deposition
- 5) Electropolishing before deposition

We also tried electropolishing by itself, and electropolishing after vacuum coating. Based on these experiments, we selected three techniques for improving our final concentrator's reflectivity. The two involving vacuum coatings both have adhesion ratings of 5B.

The first gives the best results and is the most complicated. We electropolish, strip, anodize, and seal the 5657 sample before placing it in vacuum for sputter cleaning and deposition of pure aluminum. This technique increases specular reflectance by .18 at 525 nm, with much larger increases at shorter wavelengths and moderately smaller increases at longer wavelengths. For actual production, the vacuum processing would be done using a continuous roll coater, before the individual strips are cut and drilled. Multi-station continuous roll vacuum coaters are available that can first sputter-clean and then sputter-deposit aluminum. It should be possible to also vapor-deposit an emissive overcoat at a subsequent station on one multi-station coater. The anodizing on the uncoated back surface can provide an emissive surface where it will show (i.e., between corrugated stiffeners).

The second method eliminates vacuum processing: simply electropolish, anodize, and seal the alloy. This technique increases the specular reflectance by .14 at 525 nm. As with method one, the reflectance improvements are higher at shorter wavelengths, since electropolishing appreciably smoothes the sample surface. This method is simpler and cheaper than the first one, but it does not yield as good a reflectance. Clear anodizing 1-2 microns thick can provide an emissive coating on both front and back surfaces at low cost, but at some cost in reduced concentrator reflectivity.

The final method involves vacuum processing only: sputter cleaning and vacuum coating. This increases specular reflectance by .11 at 525 nm. This has lower performance than the second method, and probably higher costs, so we did not pursue it further.

To determine the feasibility of electropolishing entire strips for the first or second methods, we contacted Fin-Tech, a firm capable of doing continuous strip electropolishing. We asked them to electropolish a small sample of our 5657 aluminum and to send us a price estimate for treating an entire roll. We received two samples back from them together with a price quote of about \$35/ft² for processing about 100 ft² of aluminum. The first electropolished sample had a reflectance of about .72 at a wavelength of 530 nm. This was a reflectance increase of .05 over untreated 5657 aluminum. The second sample was electropolished in a new batch of solution. Its specular reflectance was .80 at 525 nm. We also experimented with electropolishing in-house using small samples and achieved even better results: .85 reflectance at 525 nm. But we were unable to approach .85 for entire strips coated in a continuous process, and we do not know for sure whether Fin-Tech can also do as well with continuous processing as they did with a small sample. Figure 16 shows a comparison of reflectances obtained in-house, by Fin-Tech (their second sample) and for untreated 5657.

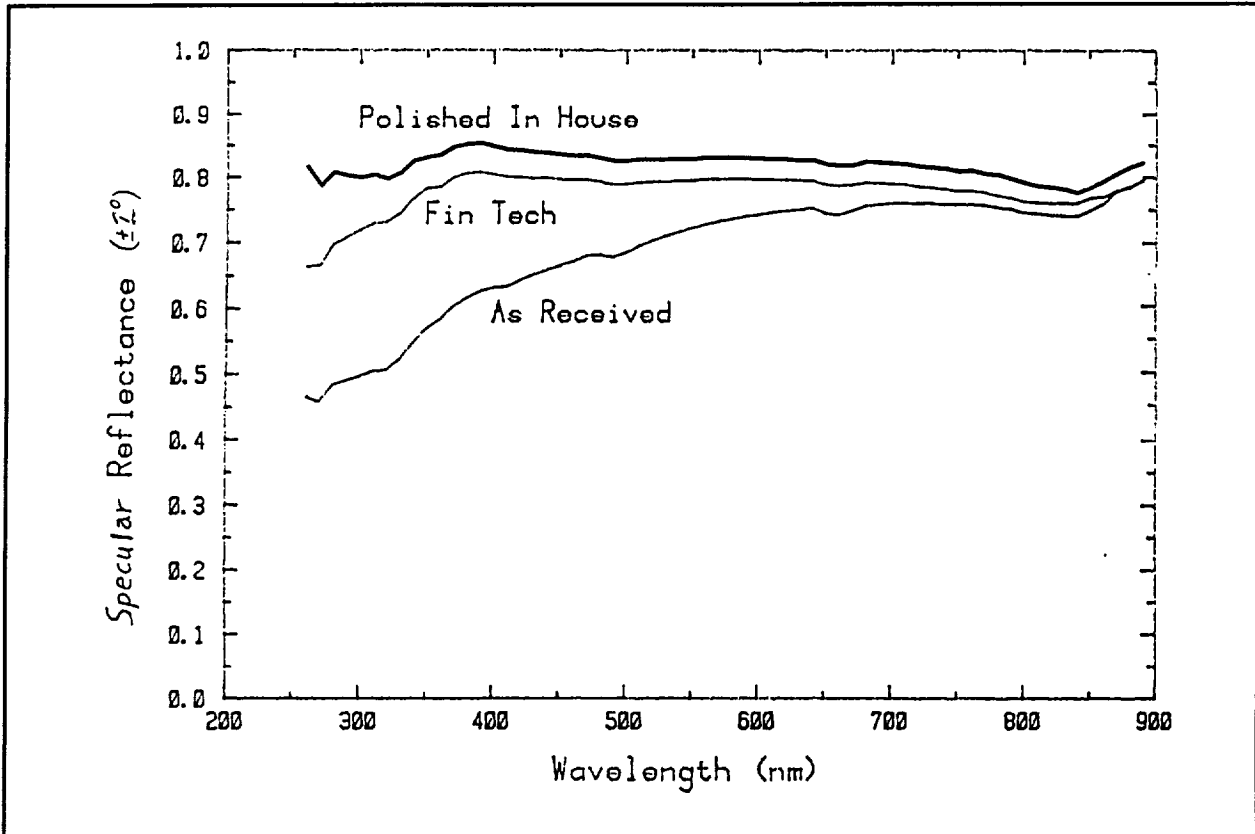


Figure 16 Specular Reflectance of 5657 Aluminum

We considered sending our 20% scale concentrator strips to Fin-Tech for electropolishing but decided against it due to time and budget constraints.

We made one final attempt to improve strip reflectance. We prepared seven "brightening dips" based on recipes in the Electroplating Engineering Handbook and experimented with them on small coupons of the 5657 aluminum. These dips consisted of varying amounts of phosphoric acid, nitric acid, acetic acid, sulfuric acid, and hydrogen peroxide. No electric current flow is involved. Unfortunately, none of these dips increased 5657's reflectivity.

In the end, we decided to use cleaned but uncoated 5657 aluminum for the 20% scale concentrator. Our specular reflectance data showed the 5657 to be comparable to Kingston Industries lighting sheet, and the 5657 was thinner and stronger. In addition, our emphasis in optical testing was mostly on accuracy and stability of concentrator shape after handling. The cleaned 5657 surface was more than reflective and specular enough for such tests.

We procured MgF and SiO sources for sputtering overcoats on the reflective surface, but did not have time to develop flight-like overcoats. Based on work for NASA Lewis by Harris Corporation, we believe that a vacuum-deposited MgF overcoat is probably best. This coating might be applied in a vacuum system after deposition of a reflective aluminum layer, at the end of the electropolish-anodize-seal-sputter-deposit process that gives the best specular reflectivity. The overcoat can be of relatively fixed thickness. The emittance variation that we need for controlled thermal distortion can be provided by ensuring that the back of the strip has a higher effective emittance than the corrugated stiffeners do.

1.3 Rim Design, Fabrication and Testing

The rim is an important part of the concentrator. It provides the structural attachment point between the membrane and the rest of the host spacecraft, and it constrains the periphery of the membrane to a plane so the membrane remains paraboloidal in shape.

The rim designs we developed fall into two broad categories: integral and separate. The integral rims have the advantage of being self-deploying, but they present difficulties in stowing and in differential thermal expansion/contraction between rim and membrane. Most of our effort was on integral rims, so they are discussed first.

1.3.1 Integral Rims

As listed below in Table IV, we fabricated 8 integral-type rims. Most of them were diamond-shaped, with each of the four sides being a frustum of a 60 degree cone that was fabricated from curved strips of aluminum foil or sheet. The inner side of the front formed an extension of the paraboloidal membrane. The back strips were often thinner than the front, because this reduced the tendency of the front to buckle and crease during stowing. The rims are deployed by inflating them. This makes them go "over center" from cone to diamond shape. Once they go over center, air pressure is no longer needed.

Table IV Characteristics of Integral-type Bi-stable Supporting Rims

#	Diameter	Section	Front	Back	Stiffening
1	1.4 m	Diamond	.002"	.0014"	None
2	1.4 m	Caret (^)	.002"		None
3	1.4 m	Diamond	.005"	.005"	None
4	1.4 m	Diamond	.005"	.002"	Integral circum. troughs
5	4.2 m	Diamond	.005"	.0014"	None
6	4.2 m	Diamond	.010"	.005"	None
7	4.2 m	Diamond	.010"	.005"	Radially corrugated foil
8	2.8 m	Triangle	.010"	.003"	Integral radial corrugations

The first four rims were made for the 1.4 meter test concentrators assembled on the perforated satellite dish. The first rim did not adequately support the concentrator, because the flat 1.4 mil foil in back was too flimsy. The second one, a simple caret shape, simply added a single strip to the periphery of the concentrator. This shape allows stowing around a much smaller diameter without rim buckling, but it also provides far less support than even rim #1. Rim #3 provided greatly increased stiffness because of its thickness, but the thick back made the front buckle when we stowed the concentrator.

We began to explore ways to stiffen the back against buckling when deployed, without causing buckling of the front during stowing. The first approach, used in Rim #4, was to use thinner foil in back, but form shallow circumferential troughs to stiffen the foil. The two back sides were embossed with two .1" wide "troughs." They increase the rim's rigidity when deployed, but the peripheral creases also help the rim go "over center" during transitions between stowage and deployment. This rim was integral, easily deployable, and strong enough to support the concentrator #5 in one gee when it was held either horizontally or vertically.




Figure 17 Cross Section of Rim #4

The success of rim #4 (shown in Figure 17) was encouraging. It seemed that we were closer to our initial goal of making an integral, bi-stable rim that could be handled easily in one gee and could be stowed within a 15" diameter cylinder for storage in the space intended for carrying a second RMS "arm".

But we realized that we needed to study scale effects. This led us to make rims #5-7. We chose 30% scale for these rims because this was as large as we could conveniently fabricate and handle in our work space, and it was about mid-way (geometrically) between 10% and full scale. Handling differences between 10% and 30% should be an indication of further differences to be expected in scaling from 30% to full scale. These three rims were not attached to any concentrators, but they had the same overall bi-stable diamond shape of the integral rims and could have been integral with a 30% concentrator, had we made any.

Rim #5, which used thin foil sections 8 cm wide, could be stowed on a 5" diameter cylinder without being damaged. However, we also found that it was too flimsy for handling in one gee: it buckled even when it was supported at 8 equally spaced points on its circumference. We concluded that scaling up this rim and increasing its thickness as needed would make it too thick to allow our concentrators to be stowed on the 2nd RMS ledge. At this point we began to reconsider separate rims, to allow the membrane to stow compactly.

At the same time, we had been trying to obtain data on the vibration environment of the ledge. A shuttle experiment involving a 50-foot heat pipe was planned to use this space, so we thought there should be data available regarding use of this volume. However, a NASA JSC representative told us that we had to work "through the system" to obtain the data. When our contract monitor contacted him, he found there was resistance to releasing the data unless there was an official NASA Headquarters policy statement permitting the use of that volume for other payloads. We were later able to obtain the vibration test report for the heat-pipe radiator indirectly, but this resistance may arise again in other ways.

Our contract monitor told us not to worry about trying to fit inside a 15" diameter cylindrical volume, if use of the 2nd RMS ledge was questionable. This eliminated any need to shift our focus from integral to separate rims. Furthermore, he noted that the simpler deployment process with an integral rim is a powerful advantage, particularly because solar dynamic power was beginning to be considered for unmanned spacecraft. We therefore continued our focus on integral self-deploying rims that can be handled in one gee.

We made rim #6 with front sections made of .010"/.005" tempered aluminum, with 8 cm wide sides. It is shown below in Figure 18.




Figure 18 30% Scale Rim (5 mil/10 mil) Supported at Six Equally Spaced Locations

Rim #6 was strong enough to allow support at only six locations. We found that we could stow the rim on a 10" diameter cylinder without damaging it. This led us to believe that scaled-up rims (.050" and .025" thick in front and rear) should enable us to perform optical tests on full-scale concentrators in one gee and to roll them up tightly enough to be stowed in the shuttle payload bay (but not the RMS ledge). But these thicknesses resulted in heavy rims and larger-than-desired stowage diameters, so we worked on increasing the deployed stiffness without losing flexibility for stowing. (This would allow us to use sheet thinner than .050"/.025" on full scale concentrators.)

To make rim #7, we corrugated strips of foil in the transverse direction and bonded them to the back (thin) sides of rim #6, using the tool shown in Figure 19:




Figure 19 Rim Reinforcement Corrugating Tool

The tool consists of two mating, tapered aluminum rollers 2" long with eight .10" deep channels in each roller parallel to its axis. The top roller is turned by a 3" diameter handwheel, and causes the lower roller to turn by means of gears attached to the axles of the two rollers. The transverse corrugating tool produces .12" wide, .06" high ridges spaced .25" apart. The tool only handles material up to 1.88" wide, so two corrugated strips are laid down side by side on each back face of the rim, with enough overlap to provide most of the stiffness that a single wider corrugated strip would give.

We made a total of 32 curved reinforcing strips with an inner radius of 99.2" out of 1350 and 3003 H19 .002" thick aluminum alloy. Each reinforcing strip was 1.62" wide.

We then disassembled the 30% scale rim and glued two reinforcing strips to each of the 16 .005" thick rim section arcs with Loctite "Depend" adhesive. Reinforcing strips overlapped each other by 1/4". We then put the rim back together with 3M #853 tape. The estimated weight of the 30% scale rim before the addition of the reinforcing strips was 2,380 grams. The reinforcing strips increased the estimated weight by 20% to 2,860 grams.

We inflated the rim using a small compressor, and found that when we lifted it at four points, it still buckled at the short unreinforced tape joints joining the .005" thick arcs. We glued pieces of corrugated aluminum to these eight weak regions and we again deployed the rim. It now required support at only four equally spaced points around its circumference. Hence the corrugations stiffened the rim enough to reduce the required number of support points from six to four. We also found that the stiffness up until buckling was better than expected. We measured it by hanging weights from the midpoint between two rim supports and measuring the rim's deflection. The results are shown in Table V.

Table V Rim Deflection as a Function of Disturbing Force

<u>Mass, gm</u>	<u>Deflection, in</u>
0	0
147.5	0.11
295	0.23

The linearity and the full recovery after removing the weights indicated purely linear elastic response up to 295 grams. (And no buckling was visible during the test.) Next we rolled the rim around a 9.6" diameter aluminum tube. This did not damage the rim in any way. But we thought we might be able to improve even on this design, so we began analyses and experiments on other bi-stable rim geometries that might be more flexible than a diamond when stowed, but stiffer when deployed. We soon realized that an equilateral triangle could stow the same way as a diamond, if its outer side could buckle outward. Our design is shown in Figure 20.




Figure 20 Equilateral Triangle Rim Cross-Section

Since the rim is round rather than straight, the middle of side #3 has to stretch about 5% when it buckles outwards during stowing. And during transitions between stowage and deployment, the outer edge of side #2 has to be able to stretch about 1.5%. We provide the required elasticities by corrugating sides #2 and #3 with gentle sinusoidal corrugations. In side #3 the corrugations are deeper near the middle of the side, because that is where the greatest length change occurs. The corrugations are gentle so that they can be mostly stretched out without yielding the material. That allows them to return when the side moves inward, shortening its peripheral length. The corrugations also stiffen the sides against local buckling in the deployed condition. The corrugations on side #3 are themselves bi-stable, because once the rim is deployed, they run continuously across the tape joint in the middle of side #3. The corrugations stiffen the joint, and hence the side, against collapse.

We then made test rim sections, varying material gauge, triangle size, and corrugation size. The initial prototypes were straight. Side #3 locked into shape in-plane. We then made several curved rim sections. When deployed, the mid-side joints did not stay "in-plane", but rather exhibited "super-corrugations" due to the large reduction in side #3 length when the rim was deployed. The super-corrugations look inelegant, but the rim was still stiffer, stronger, and lighter than diamond-shaped rims, so we accepted the super-corrugations.

The final triangular rim prototype was an arc 45" long. It had sides made of 6" wide, .003" thick 5052 H38 aluminum. The sinusoidal corrugations had an amplitude of .1" and a wavelength of 1". The super-corrugations were also roughly sinusoidal, with an average amplitude of 1" and average wavelength of 6". This prototype seemed to be stiff and strong enough for one-gee handling of the 20% scale membrane. We therefore decided to use a rim of this design for our 2.8 meter deliverable concentrator.

Our contract calls for development of fabrication concepts suitable for full-scale flight concentrators, so we investigated ways to attach rim sections to each other and to flight concentrators. Neither 3M nor CHR Industries sells a pressure-sensitive tape that met our requirements for low outgassing, flexibility, and adhesion in vacuum at low temperatures. So we developed a tape using Crest 7344 epoxy on a Kapton backing. Kapton has a coefficient of thermal expansion very close to that of aluminum, so thermal stresses are small. But Kapton needs to be aluminized or otherwise protected from atomic oxygen and ultraviolet radiation. Aluminized Kapton is available from Sheldahl at \$315 per 108' long, 1" wide roll (4 roll minimum order). We made several sample joints between .010" aluminum sheet and .003" foil using .003" and .005" thick samples of 1" wide Kapton.

The 7344 epoxy worked well with the Kapton, but it was far more expensive and time-consuming to use than 3M #853 tape. So for our deliverable concentrator we used #853 tape to attach rim sections to each other and to the membrane. The #853 tape has a .002" backing and the joints are more flexible than sample joints of .003" Kapton and 7344 epoxy, but thinner Kapton could compensate for this for joints on flight concentrators.

We spent about 65 hours making the final rim and attaching it to our final membrane. The rim was inflated through a 3/8" diameter grommet hole through a locally stiffened section of the membrane's front surface. The rim was quite rigid, but not rigid enough to overcome the astigmatism caused by the corrugated strip stiffeners. We reduced the astigmatism by using adjustable lines from four "mounting posts" to the rim. But another four posts (with lines extending only from one side) were necessary to eliminate the problem entirely.

Self-deploying self-supporting solar arrays

We felt that our collapsible equilateral triangle rim concept might have applicability to lightweight self-deploying and self-supporting solar arrays, which NASA Lewis was also increasingly interested in. We fabricated a six foot by two foot "mock solar panel" consisting of a piece of .003" thick tempered aluminum foil. We attached collapsible triangular booms along its two long edges. We also fabricated another panel with another somewhat simpler stiffening technique: the long edges were pre-stressed like a carpenter's tape measure, so that when the edges were nearly straight axially, they would curve transversely and stiffen the panel. In curving transversely, they also wrapped partially around two stub posts at the base of the array, cantilevering the array from the posts. We sent these two panels to our contract monitor for evaluation at NASA-Lewis in November, 1988.

1.3.2 Separate Rims

Separate rims appear to require EVA during deployment, but they solve other problems. They allow more compact stowage. They reduce concentrator fabrication cost, by eliminating many mutual constraints on rim and membrane design. They allow membrane replacement without rim replacement, if the membrane suffers damage or a loss of reflectivity due to aging or an episode of contamination.

The key issues are:

1. Interface design: how to attach the membrane to the rim
2. Rim structural design: compact stowage, light weight, and high rigidity

Finger-and-Pocket Interface

After considering a variety of interface design concepts, we realized that a simple pocket could be added around the periphery of the concentrator membrane, on the back side. Bent fingers projecting from a free-standing rim could tuck into the pocket and hold it in place. We used this rim concept with concentrator #4, built in the first 3 months of the contract. But we did not realize the full benefits of this concept until over a year later, when we began to wrestle in earnest with membrane/rim differential thermal response to sun-shade transients. Then we realized that with a proper geometry and springy enough fingers, this design could provide high stiffness against out-of-plane distortions at the periphery, while allowing a large amount of expansion and contraction of membrane relative to rim.

At the same time, we were worried that this concept would drastically increase the amount of EVA required to assemble a solar dynamic system. We decided to develop a prototype representative enough of a full-scale system that we could play with it to develop attachment procedures. By this time we were fabricating our 2.8 meter central section of a full-scale concentrator, and had not yet decided how to support it for optical tests. So we developed a refined finger-and-pocket interface. The 8 cm depth of this shallow dish allowed us to support it with a ring of short fingers projecting from the back wall of our lab. Since this concentrator was made using material gauges (strip and corrugation) representative of a full scale concentrator, and had the same membrane surface curvature, it was very suitable for tests of attachment procedures. The geometry is shown in cross-section in Figure 21.

The procedure we developed begins with hanging the membrane vertically from a series of spring-loaded lines just in front of the "crown" of fingers. This soft distributed support mimics some aspects of a zero-gee environment. Then a part of the membrane's edge is gently pushed against the bent fingers. A finger is pushed radially inward past the open edge of the pocket and released. This slips the bent tip of the finger into the pocket. Doing this in several places around the crown is enough to hold the membrane in place. The other fingers can then be pushed in and released without difficulty. (It takes so little dexterity that it can be done with an elbow.) Then the spring-loaded support lines can be removed.

To detach the membrane, the soft suspension lines are first re-attached. Then a string looped around the ring of fingers is tightened. This pulls the finger tips inward, out of the pocket, and frees the membrane from the fingers so it can be rolled up or otherwise handled.




Figure 21 Finger/pocket Support for Optical Tests on 2.8 Meter Central Section

This attachment technique appears entirely viable for use on orbit. Either EVA or pre-threaded guide lines between membrane and crown can be used to position the membrane in front of the fingers and insert several fingers in the pocket. Once they are in place, the remaining effort of pushing in and releasing fingers requires so little dexterity that a remote manipulator with a suitably shaped end-effector should be able to do it.

Our test support structure used 15 pieces of 3/4" thick high-density fiber board attached to the rear wall of our lab in the shape of a ring. Attached to this ring was a "crown" of 115 flexible steel "fingers" made from 0.020" feeler gauge material, 0.5" wide. The fingers projected 4" straight out from the wall, and were bent outward near the tip until the angle inside the bend was about 87°. For a full-scale flight concentrator, the materials and geometry would be refined, but the overall scale need not be radically different: fingers about 4" long can provide enough "stroke" to handle worst-case differential thermal distortion in a full-scale concentrator, without overstressing the rim, fingers, or membrane.

Rim Structural Design

The finger/pocket concept provides a good interface between rim and membrane, but a light, stiff, compactly stowable and easily deployable rim is still needed.

Because a separate rim does not stow with the concentrator, the rim can stow in a much larger diameter than 15" without taking up much room. This allows use of much thicker materials or deeper corrugations than are possible with an integral rim. Larger rim widths and depths are also possible without requiring that large areas of membrane near the periphery remain unstiffened. Finally, materials with much higher stiffness-to-weight than aluminum can be used.

We made two separate rims during the contract, both with a deep diamond cross-section. A deep diamond cannot go "over center" and collapse flat like a shallower diamond. However, if the rim has a joint somewhere around the periphery, the rim can be opened up and easily flattened into a circular arc. (And the rim is originally fabricated in that shape.) Once the rim is flat, it can be rolled up onto a slightly conical hub large enough to ensure that the inner side of the rim does not buckle during stowing.

The first separate deep diamond rim was designed for concentrator #4, the first polished 10% scale model. The rim was made by taping together curved strips of tempered aluminum foil .005" thick and 2.5" wide. We attached the membrane to the rim with 43 bent fingers. The assembly was rigid enough to support concentrator #4 vertically (with optical axis parallel to the floor) for optical testing.

The second separate rim was a by-product of the 30% scale corrugation-stiffened rim #7. We simply removed 1/3 of the rim's circumference and taped the remaining arc together at the ends. This made a much thicker 20% scale rim, with surfaces sloped 53 degrees to the plane of the rim rather than 30 degrees. This made the rim radically stiffer. As shown in Figure 22, we were able to support it in any orientation with only two supports, even outdoors in a mild breeze.




Figure 22 20% Scale Thick Rim - Horizontal and Vertical Positions (in mild breeze)

For operational separate rims, graphite-epoxy sheets with transverse stiffeners on all four sides can be used rather than aluminum. The hinges between sides could be thick Kapton-epoxy, or high-strength thin metal sheet bonded to the sides, or reduced-ply extensions of the graphite-epoxy sheets themselves. The diamond proportions are not critical, but a shape moderately higher than it is wide probably gives the best stiffness. The rim can be unwound by elastic strain energy, and lightly inflated to bring it to roughly the right shape. Then the two ends can be brought together and slipped over and latched to a rigid coupling by retracting guide lines using a small winch built into the coupling.

1.4 Optical Testing

Our concentrator design results in many discontinuities and local variations in slope error, so tests such as measuring distortions in an image of a checkerboard are clearly inadequate. Performance estimates require evaluating slope error at thousands of locations, and getting good diagnostics is even more demanding. We developed and used three optical test methods during the contract. They are all "inverse": a light source is moved around near the focus. Camera images taken at a distance each indicate what areas on the concentrator have a given error in the focal plane. We correct digitally for aberrations introduced by the camera being significantly closer than infinity. The test methods are shown in Figure 23. Key characteristics of the methods are listed in Table VI.

Method 1 uses a spot light source; methods #2 and #3 use bar-shaped light sources. The time required for testing and processing varies with quadratically in slope error resolution with a spot source, but only linearly with a bar source. This change radically reduced test and processing time, simplified the analysis, and allowed us to increase the slope error resolution. In addition, our concentrators have bilateral (rather than radial) symmetry in construction, so processing axial and transverse errors separately is more relevant.

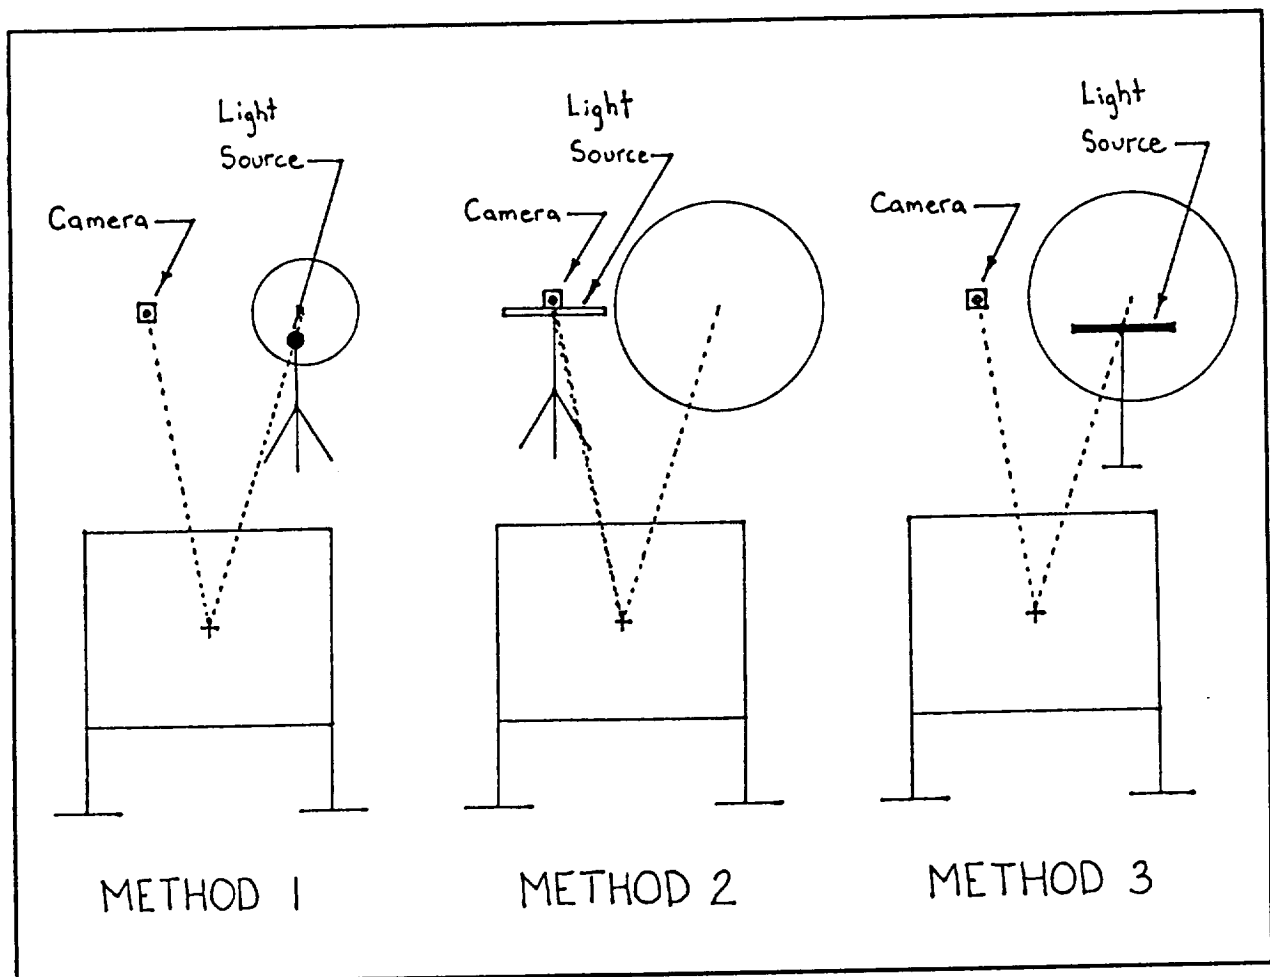


Figure 23 Optical Testing Methods

Table VI Characteristics of Our Three Optical Test Methods

<u>Method</u>	<u>#</u>	<u>Concent.</u> <u>fl,m</u>	<u>Light</u> <u>Source</u>	<u>Distance to:</u>		<u># points</u> <u>evaluated</u>	<u>Error res.</u> <u>millirad.</u>	<u>Time req'd, hours:</u>	
				<u>Light</u>	<u>Camera</u>			<u>Testing</u>	<u>Processing</u>
#1	4	0.53	25 spots	1.06f	20.0f	5,100	4-10	2- 4	12-16
#2	9	6.0	5x100cm	2.00f	2.0f	80,000	2	6-12	5-10
#3	10	1.2	2.5x50cm	1.14f	10.5f	93,000	2- 5	4- 8	3- 6

Method #2 has higher resolution of slope errors than method #3 simply because the light bar is much further away from the concentrator, due to the long focal length of the central section of a full-scale concentrator. (Its spherical shape made us co-locate the light source and camera at the center of curvature, 2 focal lengths from the membrane.) We could have improved the resolution of method #3 by using a narrower bar and digitizing more images.

Our methods have the camera significantly closer (in optical terms) than infinity. With the paraboloids tested in methods #1 and #3, this introduces a focus shift plus spherical aberrations. We compensate for the focus shift and the some of the spherical aberration by moving the light bar outward. (As noted in chapter 5, this adjustment should be larger, so as to minimize the worst absolute spherical aberration.) As described in 1.4.3, we compensate for the details of the spherical aberration digitally, when we process the data.

The image digitizer we used only determines whether a pixel is brighter than an adjustable threshold. The digitizer can create a gray-scale in the digitized image, but it does so with half-tone techniques: i.e., it varies the threshold locally. This process introduces noise, so we turned the digitizer contrast full up to fix the threshold. Hence our digitized images simply represent whether the light reflected in a given direction exceeded a fixed threshold. We set that threshold 1 stop below full image brightness with methods #1 and 2, and 2 stops below it with method #3. As a result, our tests indicate whether or not at least 1/2 (#1-2) or 1/4 (#3) of a pixel's area had a given slope error.

Our tests methods all characterize the slope errors of small regions, not of points. If a seam crosses a pixel more than 1/4 of the way from each side in method #3, then that pixel is represented on two images. The computed mean slope error of that pixel will necessarily be the average of the errors of the two sides. However, we also determine the variance of the specular reflection from each pixel, so we have a handle on sub-cell-size errors that involve at least 1/4 the area of that pixel. Note that for our method to allow derivation of an accurate flux map, the specular and "circumspecular" reflectance of the reflective surface must be known independently. Our main interest in these tests is to determine the accuracy and stability of our basic mechanical design concept, so this drawback is acceptable to us.

The remainder of this section describes the evolution of our test method in detail, under the headings of hardware, procedures, and software. The actual test results are in section 1.5. We believe that further refinements will make this basic test concept quite useful for evaluating a wide range of solar concentrator designs. The recommended refinements are described in chapter 5.

1.4.1 Hardware

Our optical tests began with the equipment listed below. Evolution from this to the final set of hardware is described component by component.

- GE VHS camcorder (1CVM6060E model) with f/1.2 9-54mm zoom lens
- MacVision video image digitizer
- Macintosh computer
- Multiple-spot light source on stand
- Position and brightness references
- 48" x 48" flat mirror
- Concentrator with support lines

Camera

We started with an ordinary VHS camcorder with a 6:1 zoom lens. The camcorder has automatic light control and automatic gain control (ALC and AGC) circuits which cannot be disabled. More seriously, the zoom lens has enough play in it that we were not able to guarantee that a sequence of images would all have the same size or position.

We used this camera with methods #1 and #2, but looked for a more suitable camera for our tests on the deliverable concentrator. We selected a Panasonic WV-1500X B&W TV camera. This camera has an internal switch that allowed us to disable the the AGC circuit, and it has a gamma control switch (.6 or 1.2). We set this to 1.2. We selected a 25 mm manual-aperture lens. This focal length allows a 2.8 meter concentrator to fill most of the image height at our test distance. The manual aperture means there is no ALC circuit.

We tested the camera for image stability and uniformity as follows. We found that line voltage variations over a typical range did not change the image size, and that the change in image size due to camera temperature variations was only about 0.01%/°C. Turning the camera on warmed it up about 14°C, mostly in the first 20 minutes. This is a much larger change than ambient temperature variations (typically 3°C) in the air-conditioned lab. We next tested the lens aperture control and found it very reproducible when turned in the same direction, and smoothly log-linear except near each end of the aperture range. We used this to test for photometric uniformity, by varying the lens setting with our digitizer set at high contrast. Over the image area used in our tests, with the 25 mm lens set at mid-range lens openings, the sensitivity varied over a total range of <10%, with the center being less sensitive than the periphery. The sensitivity at one location varied by 6% when the overall image brightness varied between the maximum and minimum we expected in our tests.

Monitor

The camcorder we started with had a built-in monitor. This was adequate with method #1. When we went to method #2, we did not have room to get behind the camera to view the image, so we used a monitor to display the image so that we could aim and focus the camera. When we shifted to the Panasonic camera, we used the same monitor. We turned the monitor off during image capture, to reduce stray light.

Digitizer

The MacVision video image digitizing box receives a video signal from the TV camera, digitizes it, and sends it to the Macintosh computer for storage in the "MacPaint" file format. This allows the image to be accessed and manipulated by a graphics software package called MacPaint. The digitizer allows brightness and contrast adjustments, but we found that for some reason these adjustments also have a small effect on the size of the digitized image. When we were setting up for our method #3 tests, we tested the digitizer for temperature sensitivity. We found it was about $0.15\%/^{\circ}\text{C}$. This was high enough that for the method #3 tests we put the digitizer in a plastic bag and then lowered it into an agitated ice/water bath to keep its temperature constant enough. This kept its temperature constant within $.3^{\circ}\text{C}$ during remaining tests.

Macintosh Computer

We started with a regular 512K Macintosh. We held our master processing file in memory. Hardware and software limitations required us to lump groups of pixels to reduce the number of discrete cells whose slope errors were stored. In addition, processing was slow. We bought a Macintosh SE for our later tests, and improved our programming techniques. This allowed us to process MacPaint files into a master file with single-pixel resolution.

Light Source and Support

We began with a fixed light source and several movable masks. In combination, they allowed 25 different equal-area regions to be exposed, one at a time. The source was quite light, and could be supported by a lightweight tripod. Tests using a light source of this design become very time-consuming as the desired slope error resolution grows. We realized that we could replace the "point" source with a "line" source. Holding a line source horizontal and moving it vertically determines vertical slope errors; holding it vertical and moved horizontally determines horizontal errors. A line source reduces the number of test images required from N^2 to $2N$, where N is the number of slope error "bins" desired.

For tests using method #2, we built an opaque box to hold a 48" fluorescent tube, and provided a 5 cm x 1 m rectangular opening covered by a diffuser. This source was far heavier, and large height scans were needed, so we used a large, heavy-duty tripod.

For method #3, the close proximity of the light source to the concentrator required a narrow holder and support (to minimize obstruction) and a narrow emitting region (for good slope error resolution). We used a 22" long 1" diameter fluorescent tube in a 25.5" long x 2" wide holder. It is shown in Figure 24. We used the tube without any diffuser. We masked 1" at each end of the tube because brightness varies in that region. We also masked the tube so only 180° of its circumference emitted light. (In the future we would reduce this to 60° , to make the source size and location precise even at large angles off-axis.) We moved the source in $1/2$ " steps in each direction. This caused most regions to be lit in 2 or 3 images, and provided better resolution and sub-pixel scatter data. The support used a pipe projecting from a flat baseplate on the floor. To adjust height, we put stacks of 1.000" and .500" thick shims under the baseplate. We took some care to ensure that the light source remained in the same plane and orientation as we moved through a scan.




Figure 24 Light Source Scan Pattern for "Point" and "Bar" Sources

In methods #2 and #3, we moved the light horizontally by sliding the support on the floor. We put guide rails on the floor to aid in properly positioning the support (see Figure 25). We estimate positioning errors with test method #3 at about .030" vertical and horizontal.

Position and brightness references

The geometric and photometric variability of our camcorder led us to include position and brightness references in each image. In method #1, we cut a hole in the back of the fixed light holder, added a diffuser, and put several sheets of paper over parts of the diffuser, so different areas varied in brightness. We then measured those areas with a light meter and selected one which was about 1/2 as bright as lit regions of the concentrator. Then for each test image, we adjusted the digitizer brightness control until this area straddled the threshold. Areas on the light source that were brighter than this served as the position reference, and the splotchy transition region showed that the brightness setting was correct.

For method #2, we attached 4 point sources of light to the wall around the 2.8 meter concentrators, in a square pattern. This pattern of lights was in each image and allowed us to document relative image position, size, and orientation. In method #3, we illuminated the top, bottom, left, and right edges of the concentrator to show position before the test, and did not touch the concentrator, mirror, or camera during the tests. (For future tests, we recommend using the 4 point sources to verify stability.) In method #2, we used as our brightness reference a small patch at the edge of the light bar. As with method #1, we used paper to reduce brightness to half that of lit regions of the concentrator. The geometry of method #2 allowed us to see this patch in the mirror. In method #3, the photometric stability and accuracy of the camera and lens allowed us to set the digitizer once at the beginning of the test, using fully lit regions on the concentrator as our brightness reference. We adjusted the digitizer until we could barely see lit regions, and then opened up the camera lens 2 stops, so the threshold was 1/4 full brightness. This resulted in the digitized image registering any sub-pixel-sized regions that had a lit area of at least 1/4 pixel.

Mirror

Our first optical tests were with 1.4 m concentrators. To minimize spherical aberration, we wanted the camera to be as far away as possible. This led us to "fold" our optical path using a 48" square mirror. For tests of 2.8 meter concentrators, we needed a larger mirror. We bought a 54" x 60" mirror for those tests and built a vertical stand for it. Increasing concern about accuracy and stability of our test setup led us to check this mirror for flatness by stretching thin lines across it. We ended up having to shim it in several places to flatten it. This mirror was barely large enough for test method #2. The mirror cracked due to loads at one of the support clips. We purchased a 72" x 78" mirror as a replacement, to give us size margin, and added rubber pads to the mirror support to reduce load concentrations. This mirror had noticeable astigmatism when first mounted. After shimming the base and the supports, the mirror surface was flat within 10 mils, with a fairly uniform concave deviation from flatness (i.e., it was nearly spherical). This distortion increased the effective test distance between concentrator and camera by about 5 cm. The mirror's deviations from flatness changed slowly over time, probably due to the use of wood in the support structure.

Concentrator Support

Our initial tests were on a light 1.4 m diameter concentrator. We suspended the concentrator from the ceiling and wall with thin Kevlar lines, and established a proper relationship with the light source with additional lines between them. The 2.8 meter central section required rigid support all around the periphery to establish a representative boundary condition for a part of a larger membrane. We used the "finger-pocket" concept shown earlier in Figure 21 for these tests.

The 20% scale concentrator was initially supported by lines running from its rim to 4 aluminum posts attached to the wall of our lab, with each post having two lines in front of the concentrator and two in back. This did not adequately restrain the rim to a plane: a 2nd order astigmatism remained. We added 4 more posts which only had lines in front of the concentrator. In an operational concentrator, these lines need not require posts, since they serve only to pull the rim forward toward the receiver. This support scheme is representative of a design that could be used in orbit with four support posts.

Final Test Hardware

After the evolution described above, the test hardware involved the following:

- Panasonic WV-1500X B&W TV camera
- f/1.4 25mm manual-aperture TV camera lens
- TV monitor
- MacVision video image digitizer (in ice/water bath)
- Macintosh SE computer
- 20" x 1" semi-cylindrical fluorescent light and adjustable support
- 4 lamps to light up concentrator at top, bottom, left, and right edges
- 72" x 78" flat mirror
- 2.8 m concentrator supported on adjustable wall-mounted support

1.4.2 Test Procedures

Our optical testing was simple at first, and got more complex as our need for accuracy grew. The evolution of the procedure is described below, under these headings:

- Position hardware
- Adjust instruments
- Take data
- Stow and redeploy concentrator (and check image size)
- Position and retest
- Measure test geometry

Position hardware

We mount the concentrator at a convenient working height, and adjust it until the rim and central seam are both vertical. (For the last tests, we used a plumb bob damped in water as a vertical reference.) Then we position the camera at the same height as the center of the concentrator. We then move the mirror into position at the other end of the room. We move the mirror transversely until its center is on a line perpendicular to the rim of the concentrator. For the last tests, we put squares of Post-it paper at the center of the mirror and concentrator to assist this process. For our last tests, we also shimmed the mirror supports until the mirror was vertical on each side, and shimmed the supports around the edge of the mirror to flatten it. We then turn the camera on, focus it, and adjust camera and mirror until the concentrator is centered both in the mirror and on the monitor screen.

We position and align the light source last. Absolute position requirements are modest, but accurate measurement is important for post-test interpretation, and stability in step sizes and alignment from image to image is important. We first put the light at the same height as the concentrator and camera. We then stand near the camera and move to either side. If the amount of lit area on the concentrator drops quicker on one side of the camera than the other, we move the light source. We repeat this until the dropoff is comparable on each side. We then align the light source parallel to the plane of the concentrator rim. (For the final tests, we then swiveled the light source around and tilted the camera until the raster lines were parallel to the light bar.)

We then move side to side again near the camera, and adjust the distance from the light source to the concentrator until there is little correlation between our side-to-side motion and side-to-side motion of the lit regions on the concentrator. This determines the best focal distance optically. For method #2, we adjusted distance by moving the camera and light source together.

In method #3, we also lit the top, bottom, left, and right edge of the concentrator with bright lamps, digitized ten images, and used those images to determine the height, width, and position of the concentrator for processing. Multiple digitizations differed in size by no more than 1 bit horizontally and vertically (except for one case in which the concentrator image was 2 pixels larger than the rest). The final test setup is shown below, in Figure 25.




Figure 25 Optical Test Setup - Method 3

Adjust instruments

We recheck the focus carefully after everything is in position. With the camcorder, we then adjust the image size with the zoom lens and re-check focus; with the Panasonic camera, we adjust the lens aperture until the image is mid-range in brightness, as indicated by a monitor previously adjusted for normal TV reception.

We then turn the MacVision contrast full up, and adjust the brightness control until a brightness reference light source in the image is just at the threshold. For tests with the camcorder, we did this by having the reference 1/2 as bright as a fully lit concentrator region. Because of the ALC and AGC circuitry, we had to adjust the digitizer threshold for each image, to keep the threshold constant despite variations in lit area. For tests with the Panasonic camera, we close the lens 2 stops, adjust the digitizer until the lit regions are at the brightness threshold, re-open the lens, and use those settings for all images.

For our final tests, we also attached small platinum-resistance thermometers (RTDs) on the camera and the digitizer, placed the digitizer in a plastic bag in an ice-water bath, and recorded temperatures periodically during a test. For those tests, we also checked and corrected the mirror's deviations from flatness before starting the test.

Take Data

This involves cycles of positioning the light source and capturing an image into a MacPaint file. For the first method, we sequentially uncovered each of the 25 regions shown on Figure 24, and made 25 files. For methods #2 and #3, we moved the light bar in small steps perpendicular to its long axis. We did two scans: a vertical scan with the light bar horizontal, and a horizontal scan with the bar vertical. For method #2 (with the 2" wide bar), we used 20 positions 1" apart, moving both the camera and the light source. For method #3, we used 40 positions in each direction, 0.5" apart, moving only the light bar.

For method #1, we videotaped each image. We played it back later into the digitizer, and adjusted the digitizer brightness threshold to cancel the effects of the ALC and AGC circuits. For method #2, we made these adjustments and digitized the images in real time. This eliminated image degradation due to the videotaping step. With method #3, the photometric stability of the Panasonic camera allowed us to digitize all images with one digitizer brightness setting.

In the first two methods, we took a single image for each position. In method #3, we took five images. The small amount of noise in digitization caused differences between images in areas where roughly 1/4 of a pixel was lit. We took five images and later used majority voting on each pixel to decide whether it was above or below the threshold brightness. (For future work, we recommend using a digitizer with a gray scale, to provide a linear response.) In method #3, we also recorded camera, digitizer, and room temperature every 1/2 hour.

Stow and redeploy concentrator

Our optical tests were intended partly to determine performance, and partly to determine whether the shape was stable through a stow-deploy cycle. To do this, we needed to test a concentrator, stow and redeploy it, and retest it with very similar test conditions.

The 1.4 m diameter concentrator was easy to stow and redeploy. We detached it from the lines to the wall and light source, leaving it softly suspended by long lines from the ceiling. Then we collapsed the rim by pulling a vacuum on it, and rolled the concentrator up.

To redeploy it, we detached a strip of tape and allowed the concentrator to unwind, and inflated the rim to make it go "over center" into a diamond shape. We then reattached the lines to the wall and to the light source that positioned it properly for testing.

With the 2.8 meter central section, we attached soft suspension lines to the membrane, tightened a string looped around the support fingers to pull them inward, out of the pocket at the periphery of the membrane, and let the membrane swing forward slightly, away from the fingers. To keep the pop rivets from scratching adjacent layers during stowing, we attached a thin polyethylene sheet to the membrane. We then rolled up the membrane around a cylinder 9" in diameter. The .010" material was stiff enough that significant tension (of order 10 lbs) was needed to tighten the membrane around the hub. To redeploy, we allowed the membrane to gradually loosen and unwind from the cylinder. This happens with a gradual change in winding radius, as in a spiral, rather than with a sudden transition between rolled and unrolled portions. We then orient the membrane and hold part of its periphery gently against the bent fingers, and push in and release a finger. This snaps it into the pocket. Doing this in several places holds the membrane in place. Then the remaining fingers can easily be snapped in, and the spring-loaded support lines removed.

Stowing and redeployment of the final 2.8 meter concentrator are shown on the videotape sent to NASA Lewis. We first attach the Kevlar & rubber band distributed soft suspension system to support the concentrator. Then we loosen and detach the 24 lines that attach the rim to the wall-mounted struts, and move the concentrator away from the wall to give us more room. Then we attach small tabs of adhesive tape to the middle of the hinged side of the rim, and pull outward to buckle the hinge. This is quite tricky to start, but once it starts, it is a little easier to "work around" the rim. Once the entire rim is collapsed, we pull a vacuum on it via the inflation hole in order to flatten it as much as possible.

We then cover the front surface of the concentrator with a thin sheet of polyethylene to protect it, and part of the edge is attached to a rim. To facilitate stowage we then push together the upper ends of the suspension lines, which are attached to sliders in a curtain rod. Next we attach the "3 o'clock" edge of the now flexible membrane to a long 9" diameter cylinder. (We used a 1' long x 2" wide piece of duct tape to hold it; in the future we would tie it down using loops of line threaded through eyelet holes.) We then turn the central cylinder to wind up the concentrator like a clockspring until it is tight.

We reverse the above procedure to redeploy the concentrator. The only changes have to do with the rim, which requires inflation and some local pushing to lock it into shape.

Position and retest

After we redeploy the concentrator, we adjust the support line lengths so the concentrator is as close as possible to its previous location and orientation. With test method #3, stowage of the rim caused changes that made the 3 and 9 o'clock edges move forward of their pre-stowage position. These rim positions only had support lines pulling forward, so the lines were slack and the rim remained forward of the first test position. Differences in rim position between first and second test are shown in Figure 26. "FB" and "F" indicate rim positions with lines pulling forward and backward, and forward only.

	-5	-8	0	-11		
			F		-3	
0 FB					FB 0	FB = location supported by lines in front and in back
+7					+4	
+18 F					F +3	F = location supported only by lines in front
-12					0	
0 FB					FB 0	
	-5		F		-7	
		+2	+2	-8		

Figure 26 20% Concentrator Rim Position Changes, After Stowage vs Before (in mm)

With test method #3, we also inspected the concentrator carefully. The inspection showed three small dimples, each involving 10-20 cm of one seam. Two were stable only in the dimple mode (they could be popped out, but returned when released), and one was bi-stable. We left it dimpled after checking its stability. We also found the 20% scale membrane slightly buckled at one location on the periphery.

After redeploing and repositioning the concentrator, we do our post-stowage optical test. When we completed our retest with method #2, we realized that the concentrator image size had changed. We therefore adjusted the zoom lens until the image size matched the pre-stowing test image size and performed a second retest. For test #3, we used the fixed-focal-length 25 mm lens so this was not a problem. We did check the image size before retesting and found that it was within 1 pixel of the first test.

We make final measurements of test geometry after the retest sequence is complete, to verify that nothing moved during the test.

1.4.3 Software

The software required for our image processing included:

- Macintosh system software (for 512K Mac; later version 4.0 for SE)
- MacVision software (ver. 1.4; used w/MacVision box to make MacPaint files)
- Lightspeed Pascal (ver. 1.0; for running Pascal on a Macintosh)
- In-house programs (written in Pascal)

The first three items are commercial, so this section covers only the in-house programs. They evolved with our test methods. We discuss the evolution of the software approach, to clarify the advantages of the final software approach, but the main focus is on the final software used with test method #3. Our data processing starts with a stack of floppy disks containing MacPaint files generated during an optical test. Each disk contains MacPaint files generated by the MacVision digitizer and its software. Each file has a name that indicates the test and light source position used for the image digitized in that file. A typical image is shown below in Figure 27:




Figure 27 MacPaint File Made Using Method 1. (Note shadow of light source & tripod)

We begin our processing by creating and zeroing a master array, and then digesting the MacPaint files into it, file by file. The data we store allows us to compute mean light source position and light source position variance for each pixel in the concentrator image. With test method #1, we held this master array in memory. To fit within memory allowed by the Mac and Lightspeed Pascal, we had to combine 16 pixels into one cell. Then for each cell, our program estimated the mean light source position by weighting light source position for each image by the number of pixels lit in that cell. We quickly realized that despite the polar configuration of light source #1, we had to use Cartesian coordinates to find average light source positions when different pixels in one cell were lit in different images.

This helped us decide to switch from a "spot" to a "bar" light source, because we realized this would cut our memory requirements for each cell (and hence allow more cells) by allowing us to evaluate horizontal and vertical errors separately. The equal steps in light position from image to image allowed us to process and store data in integer form in the master array. We then coded the data to compress two of the parameters into one two-byte integer. This further reduced memory needs. We then tried processing 1/4 of the image at a time, to further reduce memory needs and allow use of larger numbers of cells.

Soon thereafter, we realized that we could store the master array as a floppy disk file, and access the data for each cell individually. This relaxed the array size constraints enough to allow us to evaluate errors pixel by pixel over the whole image. This made it unnecessary to combine pixels into a cell, or to process the concentrator image 1/4 at a time. But the large array and the frequent disk accesses slowed the processing down considerably. Our final advance was to configure the master array as a file of smaller arrays, each of which contained data on a whole row of pixels. This reduced the number of disk accesses, while still allowing the master array to handle the full image, pixel by pixel.

Data processing for test method #3 starts as follows. We scan five MacPaint files, taking a majority vote on each pixel to determine whether it was lit or not with that light source position, and creating a new majority-vote MacPaint file. We do this for each of the 40 light source positions in a vertical or horizontal test. We then use the 40 new majority-vote images for a horizontal or vertical error test for subsequent processing.

For method #3, we automated our program to operate on a set of input files whose names are specified on a disk file. When the program is executed, it asks for the name of that file. The program then reads in the list of file names from that file, requests the operator to install each floppy disk in order; checks that the right disk is installed; creates the new majority-vote MacPaint files, and updates the master array pixel by pixel, for each pixel lit in each majority-vote image. This Lightspeed Pascal program for the Macintosh is called cMF.2 (create master file, version 2).

We found we could do subsequent processing much faster on an IBM-compatible 286 computer with hard disk, so we converted the master file into three files of arrays of integers and transferred them to a PC. In this format, each file entry covers one line of the image, and consists of an array of 400 integers, each concerning one pixel in that line. The array ordering is that of the MacPaint file, starting at the top, and starting each line at the left. Because the images were acquired through a mirror, left and right are reversed. Thus each line starts at the right side of the concentrator, as it is viewed from the front.

We digitally reversed the images shown in section 1.5, so that top and left on the image correspond to top and left on the concentrator, viewed from the front. Error data uses standard Cartesian coordinates: if one stands near the mirror and looks at the focal region directly (i.e., not through the mirror), then positive horizontal and vertical errors indicate light source positions to the right and upward.

Corresponding integers in the 3 files used in the 286 post-processing represent:

- N: the number of images in which a given pixel is lit
- Sum2X: sum of image #s (from -20 to +20, 0.5" apart) in which that pixel was lit
- SumSqr2X: sum of squares of image #s in which that pixel was lit (to compute variance)

We then computed the average position error for each pixel on the image, by dividing Sum2X by N and storing the result as an integer representing mils error (-10K to +10K). We stored the results in files named KAvgX (i.e., 1000 times the average error in inches), with suffixes 1H, 1V, 2H, and 2V representing horizontal and vertical errors in tests 1 and 2. For pixels with N=0 (unlit in any image in a given scan), we assigned an error of 30000, which is too large to be confused with real data. We also computed the pixel positional variance and stored it in files named, HVar.1H, etc. The individual integers represent 100 times the variance in square inches. For pixels with N=0, we set the variance to -1.

As mentioned earlier, our test conditions require digital correction for spherical aberration. We did this as follows. We first wrote a program, CloseAb, that computes the errors (all radial) that different locations on an ideal paraboloid would exhibit under identical test conditions: camera 10.5f from origin; light source plane 1.14f from origin. As shown in Figure 28, the error is modestly negative near the center (i.e., the rays pass through the axis before they reach the focal plane). Near the periphery, the errors become strongly positive, reaching 79 mm at the rim.




Figure 28 Focal Plane Errors of an Ideal Paraboloid Under Our Test Conditions

This pattern can be intuitively understood as follows. For testing at 10.5f, a prolate ellipsoid is needed for a distortion-free focus. This shape has a smaller variation of curvature with off-axis distance than a paraboloid does. Hence a paraboloid's curvature will be too tight near the radius, and too loose further out. This makes inboard rays cross the axis before they reach the "best-focus" plane, while outboard rays reach the best-focus plane before they reach the axis.

We stored the computed aberration data in a integer array file (CloseAb.out) of mils aberration for each .002f increment in off-axis distance. We then computed horizontal and vertical coordinates of each pixel, and from that the off-axis distance. We included an approximate correction for the fact that the concentrator has depth and is being viewed from finite distance. As a result, locations "within" the dish are further away from the camera than the rim plane is, and hence slightly farther from the axis than would be estimated if one simply scaled the planar image on the vidicon.

Once we had the off-axis distance, we retrieved the appropriate aberration from CloseAb.out. We then found the horizontal and vertical error components of that error, and adjusted the horizontal and vertical KAvgX data appropriately. We had already calculated mean light source errors for the whole concentrator, so we also subtracted the mean error.

The resulting data are stored in files with an "A" added to the filename suffix to signify that the data are adjusted: KAvgX.1HA, etc. The data represent the error of our concentrator compared to that of an ideal paraboloid under the same conditions, in the effective plane of the light source, which is 1.14f from the origin. Because the light source was a semi-cylinder and not a plane, we had to estimate the effective light source plane. It is 2 mm toward the concentrator from the centerline of the fluorescent tube. This position over-represents the effect of the tube's "bulge" in front of the mask plane for rays near the center, and under-represents it for rays near the periphery. (For future testing, we recommend that the source be masked so the emitting region is no more than 60° of the tube's periphery. This allows all regions of the concentrator to "see" both masked edges.)

The errors calculated for a plane 1.14f from the origin are not exactly the same as errors in an infinity/1f test, because the difference in geometry causes the error to have a scale factor different from 1.0. The exact error scale factor depends on location. Rays near the center have 1.14 times the error they would have in infinity/1f tests. For rays near the periphery, cosine effects and spherical aberration enter in. Transverse errors are 10% larger than they would be in infinity/1f tests, while radial errors are 13% smaller. The average scale factor is >1.0, so the results presented in section 1.5 are modestly pessimistic. For future tests, we recommend reducing the data to infinity/1f conditions, because that has direct physical meaning, is more directly useful, and is a more standard parameter to report.

Once we had the adjusted error files, we wrote over a dozen short post-processing programs in Borland Turbo-Pascal 5.0 to analyze and process them into the images given in section 1.5. Those programs are listed on a disk being provided to NASA Lewis. They are documented as needed in the text accompanying each figure in the next section.

1.5 Optical Test Results

This section discusses early test results briefly, to illustrate the need for improving our test setup, and then presents the results obtained on the final concentrator. We emphasize the final results because that concentrator and the test methods are both more relevant to future developments than our earlier test results were.

1.5.1 Early Test Results

We did our first tests on 10% concentrator #4. They included pre- and post-storage tests. We could not tell how much the changes in slope error and the differences in which cells were lit in each test were due to concentrator shape changes, and how much they were due to image position and size changes during and/or between tests. This made us more careful about verifying image position in later tests.

To get a crude flux map and a ballpark estimate of the concentration ratio of concentrator #4 we put a piece of tracing paper with concentric circles drawn on it at the focus to show the focal plane intensity distribution from a small light source 21 focal lengths away (see Figure 29).




Figure 29 Rudimentary Optical Test of Concentrator #4

At the best-focus distance, most of the energy seemed to fall within a 5 cm diameter circle. This corresponds to roughly an 800:1 concentration ratio.

The most significant finding on our tests on concentrator #9 (the 20% central section) was that only 59% of the pixels were lit in the vertical scan (with horizontal light bar), while 70% were lit in the horizontal scan. Remaining pixels either had horizontal or vertical slope errors above 20 milliradians, or had sufficient sub-pixel curvature that the intensity in any one image never reached the threshold required to register.

Differences between the average horizontal and average vertical slope errors should not affect the fraction of pixels lit significantly, because the light bar shape and scan size were roughly matched. Thus a vertical error large enough to be off-scale in a vertical scan will be off-scale "sideways" (i.e., vertically) in a horizontal scan, due to the finite length of the light source. So another explanation was needed for the 70%/59% discrepancy.

The brightness reference source varied in position in the two scans, so we checked the camcorder's photometric uniformity across the image. We found enough variation to explain most or all of the 70%/59% discrepancy. The differing threshold affected what fraction of pixels with sub-pixel slope error registered in one or more images. These problems convinced us that we needed a camera with better photometric stability and uniformity. It is likely that some of the discrepancy was due to a difference in the amount of sub-pixel-sized curvature in the two directions. The strip joints would tend to cause the vertical scan to have higher lit-pixel fractions. However, we joined the strips with pop rivets spaced several inches apart. This induced a noticeable periodicity in the vertical curvature of each strip. This could easily have caused enough added sub-pixel-sized slope variation to make "on-scale" errors fail to register, by spreading the light intensity over too many images.

Our inability to interpret our results helped us decide to use reduce the brightness threshold to 1/4 of full reflected intensity, and also to use a light source twice as wide as the step size. This makes it harder for modest sub-pixel dispersions to reduce the reflected light intensity below the threshold. As a result, this combination provides us far better data to allow us to evaluate sub-pixel slope error variability, something which the high and variable threshold prevented us from doing meaningfully with test method #2.

1.5.2 Test Results on Final 20% Scale Concentrator

The 20% concentrator membrane was fabricated while stowed, then deployed and attached to its integral rim. The complete concentrator was optically tested, stowed with rim, redeployed, and retested. Each test scan involved 40 light positions, with 3 of 5 image voting on each position. Each position overlapped 50% with the previous one. This section will present and discuss the data taken in these tests under the subheadings below:

- Lit pixel statistics
- Sub-pixel error variance
- Pixel error distributions
- Error maps
- Changes between first and second tests
- Flux maps


Lit Pixel Statistics

Table VII shows the number of pixels lit in different numbers of images in the four scans. Several factors account for the large number of unlit pixels. One is that the image area evaluated extends beyond the concentrator periphery. This accounts for 2,000-2,600 pixels (the exact number is hard to determine). A smaller factor is shadowing by the light source holder and support pipe, which together shadow about 1,000-1,400 pixels. (The shadowed region remains fixed through a scan only for the pipe, and only in the vertical scan. The other shadowing varies in position from image to image, but is responsible for a larger number of unlit pixels, though distributed over a much larger region.) A few thousand pixels are probably "in-range" but with large enough sub-pixel-size slope variation that they fail to register in any image. These factors together account for perhaps 6,000-8,000 of the unlit pixels. This leaves 10,000 to 20,000 pixels unlit in the various scans because they have large slope errors.


Table VII Number of Pixels Lit in Different Numbers of Images (20% Scale Concentrator)

Test \ NLit:	0	1	2	3	4	5	6	7	8	9	>9
1h	26585	10843	26138	23981	6793	1320	182	22	4	0	0
2h	27832	8817	21006	26317	9227	2125	444	88	11	1	0
1v	18255	8009	25501	30133	10589	2706	566	91	17	1	0
2v	18494	7206	22549	31146	12297	3211	789	148	21	7	0


Figure 30 shows the spatial distribution of lit pixels in the four scans (1H, 1V, 2H, and 2V). At the bottom, it also shows the logical "and" of the two scans in each test. Only the "and" pixels (those lit in both scans in a test) have a error defined in both coordinates, and only those pixels are can be mapped and analyzed in the concentration ratio and flux maps. Figure 31 and 32 show the spatial distribution of pixels lit in different numbers of images for the horizontal and vertical scans of test 1. The slightly elliptical plots are due to a printer limitation. The ellipses represent circular regions in the TV camera image.



69283 pixels lit in test 1H




68036 pixels lit in test 2H




77613 pixels lit in test 1U



77374 pixels lit in test 2U




64050 pixels lit in tests 1H and 1U




63064 pixels lit in tests 2H and 2U


Figure 30 Pixels Lit in 20% Scale Concentrator Tests 1 and 2 (lit = white)




Black = pixels lit in 0 images in test 1H




Black = pixels lit in 3 images in test 1H




Black = pixels lit in 1 image in test 1H



Black = pixels lit in 4 images in test 1H




Black = pixels lit in 2 images in test 1H




Black = pixels lit in >4 images in test 1H


Figure 31 Pixels Lit in Different Numbers of Images in Scan 1H




Black = pixels lit in 0 images in test 10




Black = pixels lit in 3 images in test 10




Black = pixels lit in 1 image in test 10



Black = pixels lit in 4 images in test 10



Black = pixels lit in 2 images in test 10



Black = pixels lit in >4 images in test 10

Figure 32 Pixels Lit in Different Numbers of Images in Scan 1V

To interpret Table VII and Figures 30-32, first consider the results of a test on a concentrator in which "high order errors" are small enough that errors do not change significantly within a pixel. Then a pixel with a given error will generally register in two images, because of the 1" wide source and 0.5" step between images. But pixels whose error is centered on one light source position will be marginally lit in the two adjacent ones. Those pixels will often be lit in one or three images, depending on details. For pixels which are a large distance transversely from the light position, the projection of the semicylindrical light source from the pixel location onto the $1.14f$ plane is significantly wider than 1.0". This means that more and more of the pixels will be lit in 3 images. At the edge, about 77% of the pixels will be lit in 3 images. Hence pixels in a "smooth" dish will be lit in 1-3 images.

Adding sub-pixel slope errors and an "area" threshold of $1/4$ pixel allows pixels to register in more images: at the limit, in 4 groups of 1-3 neighboring images. A small amount of image defocus (which causes mixing between neighboring regions) could allow some pixels to register in more images, but should not change the average much, because light "stolen" from a neighboring pixel reduces the likelihood of that pixel registering in extra images.

Consider now the actual test data. Of the pixels which are lit at all, most are lit in either two or three images, as expected. Fewer pixels are lit in 1 image. A few hundred 1-image pixels are lit in the last image at either end of the scan, indicating that their "other" image would have been off the end of the scan, but the errors of other 1-image pixels have distributions comparable to those of all lit pixels.

Most of the unlit pixels are near the periphery. Because of distance and more so because of cosine effects, pixels near the periphery need to have 25% smaller transverse errors and 63% smaller radial slope errors than near-axis pixels do, in order to register in the tests. This clearly contributes to the larger number of unlit pixel near the periphery. But the sharp increase in unlit pixels within about 15-20 cm of the periphery points to two other factors: the joint between the rim and membrane 15 cm inboard from the edge (along the slope), and the absence of stiffening corrugations in the outer 18-20 cm of the membrane. The increase in unlit pixels in test 2 is particularly strong in this region. Hence a separate rim, combined with stiffeners that reach as close as possible to the periphery, may be able to significantly increase concentrator performance.

Sub-pixel Error Variance

Pixels lit in more than one image have an apparent variance in slope error. The phrase "apparent variance" is used because even a perfectly curved pixel would show up in 2 or 3 adjacent images, and have an apparent error variance of .06 or .17 square inch respectively. This is for a variance computed as $\text{Avg}(\text{SqrX}) - \text{Sqr}(\text{AvgX})$. This logic provides population rather than "sample" weighting (N vs $N-1$), since population weighting appears more appropriate for the methodical scan conditions. Variance was computed in this way for all lit pixels in each scan. The results were binned in bins of different widths. The bin populations are listed in Table VIII, along with "exceedance" bin data expressed in standard deviation (the square root of variance).

Table VIII Number of Pixels with Different Error Variances (0-1 and 0-100 sq. inches)

Scan	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0 sq.in.
1H	36142	23266	534	6886	0	1578	102	318	67	17	
2H	28716	25319	756	9316	0	2350	91	582	152	43	
1V	31849	28498	1065	10699	0	3153	225	828	223	115	
2V	28249	29450	984	12219	0	3576	281	1076	257	124	

Scan	0	1	2	3	4	5	6	7	8	9	10	100
1H	68910	281	36	18	5	4	2	0	0	1	26	
2H	67325	535	65	27	8	4	2	0	0	2	68	
1V	76655	769	109	45	14	6	5	2	2	3	3	
2V	76216	894	152	52	26	9	10	5	1	1	8	

Percentage of Lit Pixels with Error Standard Deviations as Follows:



Scan	<0.63"	≥0.63"	≥1.00"	≥1.41"
1H	96.46	3.54	0.54	0.13
2H	94.23	5.77	1.04	0.26
1V	92.91	7.09	1.23	0.24
2V	91.64	8.36	1.50	0.34

Table VIII clearly shows that error variations within pixels (as measured by our tests) are very modest for the great majority of pixels. As a result, we decided to neglect error variations within each pixel and do our remaining data processing and analysis using the average slope error for each pixel.

Pixel Error Distributions

The first way we evaluated the error for each pixel was to plot the frequency of different horizontal and vertical errors. The data include adjustments for spherical aberration and a zero-shift resulting from the setup. They are plotted in SI units as are the remaining plots. Figure 33 shows the horizontal errors for tests 1 and 2 respectively, and Figure 34 shows the vertical errors for the same tests. Statistical data are listed with each plot. As seems reasonable, stowage and redeployment increased the standard deviation modestly (about 3%). A more obvious pattern is that the vertical errors are about 20% larger than the horizontal errors. This led us to wonder more about the spatial pattern of the errors, and led to the error maps in the next subsection.



We also computed the frequency distributions of total error for each pixel in tests 1 and 2. The total error R for each pixel is computed as $\text{Sqrt}(\text{Sqr}(\text{HorErr}) + \text{Sqr}(\text{VertErr}))$. The results are plotted in Figure 35. Dividing that distribution by R gives a circumferentially averaged intensity distribution. That is shown in Figure 36.



N=69017; MaxBinPop=1028; Avg= 0.052 cm; StDev= 6.087

N=67647; MaxBinPop=1034; Avg=-0.03; StDev=6.91 cm.



Figure 33 Horizontal Errors in Tests 1 and 2 (-25 to +25 cm)



N=77495; MaxBinPop=683; Avg=-0.048; StDev= 8.120 cm.

N=76937; MaxBinPop=739; Avg=-0.04; StDev=8.37 cm.


Figure 34 Vertical Errors in Tests 1 and 2 (-25 to +25 cm)



N=63670; MaxBinPop=673; Avg= 8.08; StDev=5.06 ca.

N=62179; MaxBinPop=602; Avg= 8.53; StDev=5.31 ca.

Figure 35 Total Errors in Tests 1 and 2 (-25 to +25 cm)



N=63670; MaxBinPop= 2.6E+0003; Avg= 8.08; StDev=5.06 cm. N=62179; MaxBinPop= 2.5E+0003; Avg= 8.53; StDev=5.31 cm.
Figure 36 Flux vs Total Error in Tests 1 and 2 (-25 to +25 cm)


Error Maps

The horizontal and vertical errors plotted above in Figures 33 and 34 are shown in shaded maps in Figure 37. All six plots represent errors by a dithered gray scale, with black representing an error < -6 cm and white representing an error $> +6$ cm. We create the gray scale by comparing the error in cm for a given pixel with a random number having a uniform distribution over the range -6 to $+6$. If the error is greater than the random number, then the pixel is printed white; if less, then it is black. The random threshold keeps one from having confidence about the error on a pixel-by-pixel level, but it provides a good indication of average error in regions more than a few pixels across. Hence it is good for intuitively evaluating the overall error pattern.

The three maps on the left indicate horizontal errors; those on the right indicate vertical errors. The gray scale logic results in positive horizontal or vertical errors appearing lighter than negative errors. Hence the horizontal error maps can be regarded as being "lit" from the right: areas tilted to the right then appear brighter. Vertical error maps are likewise "lit" from the top, so that areas tilted upward are brighter than those tilted downward. In both cases, pixels unlit in a given scan are printed white, because the resulting images seem slightly easier to interpret than when we printed unlit pixels black.

The top left and top right maps in Figure 37 indicate the "raw" error in test 1, before we compensated for spherical aberration and a mean offset. The middle maps show the adjusted data for test 1. The bottom maps show the adjusted data for test 2.

The clear message from the horizontal error maps is that the membrane has a large slope error component in the form of a triangular wave with a wavelength of 2 strips: one strip tilts right; the next one tilts left. The vertical extent of these waves is rather short. The errors result from strip width variations across a large number of adjacent strips. The vertical wavelength is equal to the pleat-to-pleat repeat distance. The membrane is pleated where it is because it is "cinched up" and slightly buckled by intervening unpleated regions where the strip width is less.




Hor errors, right lit: -6<gray+6 cm. File=KAUGX.1H




Vert errors, top lit: -6<gray+6 cm. File=KAUGX.1U



Hor errors, right lit: -6<gray+6cm. File=KAUGX.1HA



Vert errors, top lit: -6<gray+6 cm. File=KAUGX.1UA



Hor errors, right lit: -6<gray+6cm. File=KAUGX.2HA



Vert errors, top lit: -6<gray+6 cm. File=KAUGX.2UA

Figure 37 Error Maps of 20% Scale Concentrator (positive error = white)

The vertical errors are larger in amplitude and wavelength. They have a more obviously symmetrical pattern from side to side because the "lighting" is from the top, and hence symmetrical from side to side.

Looking at the horizontal and vertical errors together, it becomes clear that they are inter-related: the horizontal bands with significant pleating in the horizontal error map are all regions which are at the bottom of a "trough" in the vertical map, and hence dark on top and bright on the bottom, while unpleated regions correspond to crests, which are bright on top and dark on the bottom. Hence the narrower, unpleated regions tend to pucker and pull forward of the best fitting paraboloid, while the wider, pleated regions move back and buckle. The vertical wavelength of this error pattern is consistent with errors in our mill setup. (This is not presented as proof, but rather as the most straightforward explanation we can find for the error pattern shown in the sets of maps.)

Error Changes Between First and Second Tests

Figure 38 shows the changes in horizontal and vertical error between test 1 and test 2, pixel by pixel, using the same dithered gray-scale used in the error maps in Figure 37.

The changes in horizontal error are far larger on the left side than on the right. This is compatible with a reversal of the "pleating" on that side, so that forward buckles are now backward and vice versa. It seems unlikely to be due to image misregistration, given that strips are 8-10 pixels wide and the image sizes for the two tests matched within about 1 pixel. The horizontal image also is considerably darker on the right than on the left, as might be expected from the "underdeployed" rim distortion documented in Figure 26.

Likewise, the vertical error change map varies in density from top to bottom, because the underdeployed rim shape makes the radius of curvature in the vertical direction too large. This makes the top "better lit" on the second test than the first, and the bottom "better shaded." The reversal of this logic at the very top and bottom, so that the top is black and the bottom white, is due to localized effects of the top and bottom front-mounted lines, which pull the rim forward at top and bottom without having much effect elsewhere. (The secondary supports at left and right side need to pull the rim backwards, and cannot do that, because they only have lines on the front.)

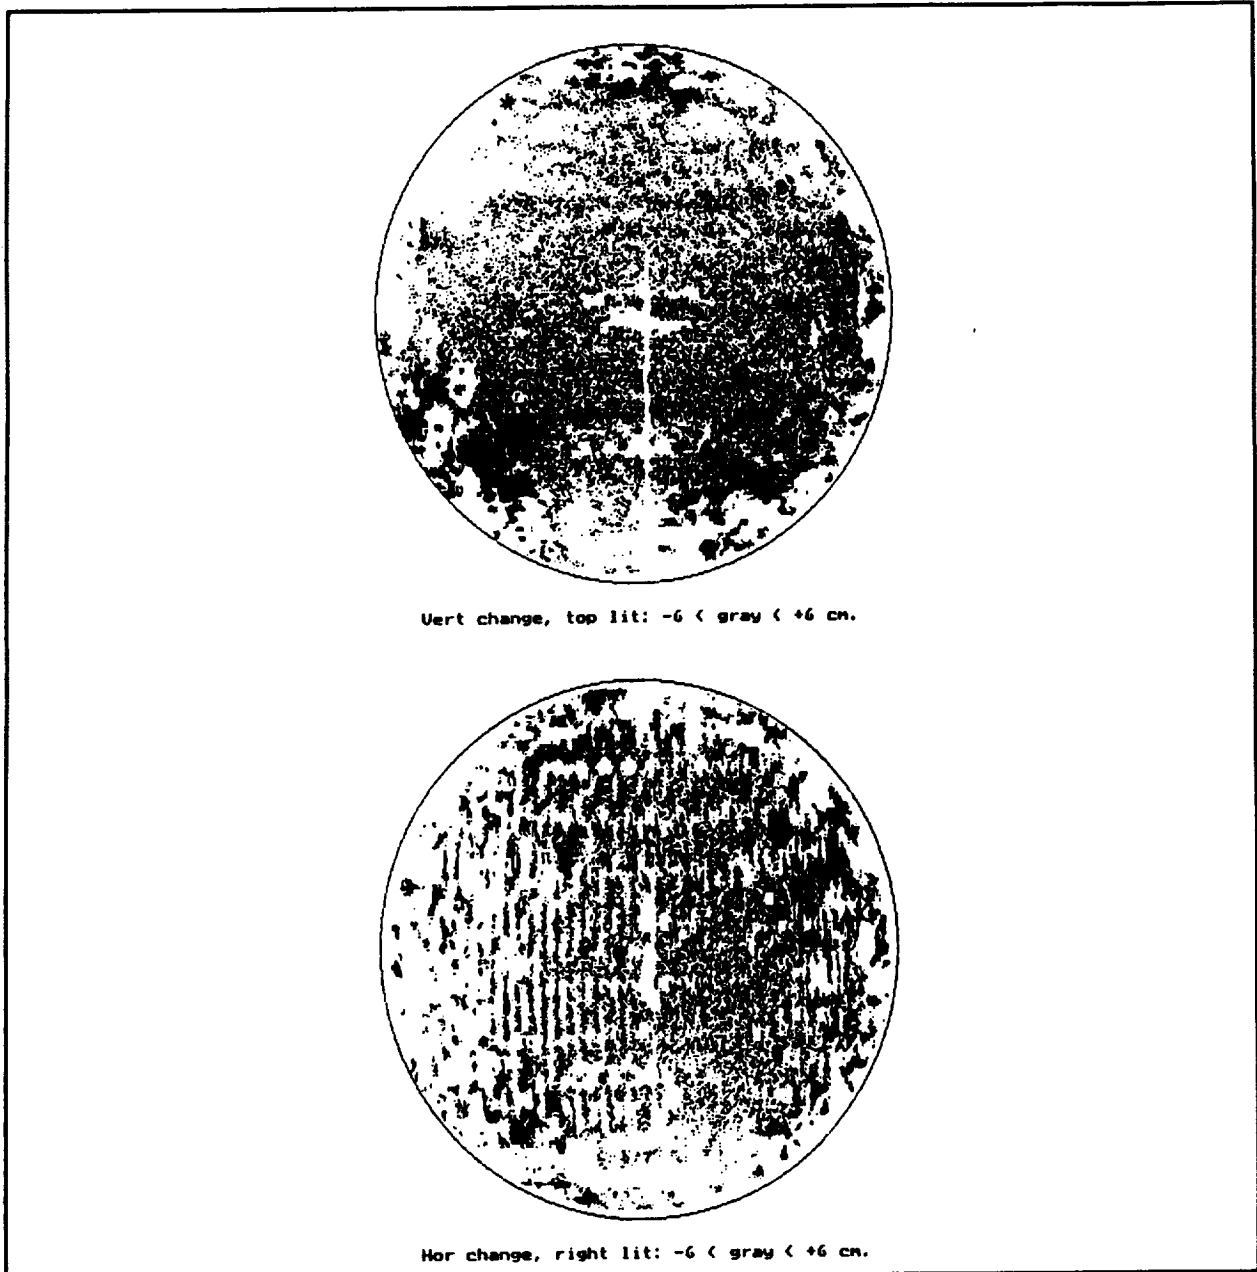


Figure 38 Changes in Horizontal and Vertical Error, Test 2 - Test 1 (white = positive)

Flux Maps

The position errors plotted above can be used to generate approximate flux maps at the focal plane. As noted at the end of 1.4.3, the errors are errors in a focal plane $1.14f$ from the origin, for our test conditions (rays diverging from a point $10.5f$ away), compared to the errors of an ideal paraboloid tested under the same conditions. The resulting errors are generally larger than they would be under infinity/ $1f$ conditions: 14% larger for rays near the center, 10% larger for transverse errors near the rim, and 13% smaller for radial errors near the rim. The flux maps below are not corrected to infinity/ $1f$ conditions, and hence are slightly pessimistic. If the average error scale factor for $1.14f$ test conditions is 1.08 (weighted for lit pixels, which tend not to be near the periphery), then fluxes and concentration ratios should be about 17% higher than indicated here.


Because the concentration ratios and flux densities are so much lower than the 46,000 theoretical maximum, the error associated with ignoring the finite diameter of the sun is of order .1-1%. This is so much smaller than the conservative error caused by not reducing the data to infinity/ $1f$ conditions that we felt justified in ignoring it.

Figure 39 shows a graphic flux map in a region defined by a concentration ratio of 60 (i.e., an area $1/60$ th of the concentrator area normal to the sun). That figure also includes circles that define the smaller apertures associated with concentration ratios of 200 and 600.


The gray scale on this map is not linear, but "saturating": each lit pixel on the concentrator "bleaches" one pixel on the flux map. If additional concentrator pixels illuminate the same flux map pixel, nothing happens. Thus two regions with 10% and 90% of their pixels remaining black will have a flux ratio of $\text{Ln}(0.1)/\text{Ln}(0.9)$ or 22, rather than 9.

Figures 40 and 41 give numerical flux maps, with a value for each 1 cm square region in a 50 cm square in the " $1.14f$ " focal plane. The flux map values are in suns, assuming a moderately conservative specular reflectance of 76%. Isoflux contours are drawn by hand at 10, 30, and 100 suns. To scale this map properly, note that the largest circular aperture that can be inscribed in that square flux map has a CR of 31. As with the concentrator maps, these two maps are not quite square, due to printer limitations. However they do map a 50 x 50 cm focal plane region.

Figure 42 shows which pixels on the concentrator have a mean errors small enough to pass light through apertures with CR of 600, 200, and 60.




Test 1: Flux map for CR of 60, 200, 600.



Test 2: Flux map for CR of 60, 200, 600.

Figure 39 "Saturating" Focal Plane Image for 20% Scale Concentrator Tests 1 and 2



31943 = total co2-suns. Map assumes 76% specular reflectance.

Figure 40 Computed Flux Map for 20% Scale Concentrator Test 1, in Suns




Figure 41 Computed Flux Map for 20% Scale Concentrator Test 2, in Suns

Vacuum Pump

The vacuum pump maintained vacuum in the system. It was equipped with filters to capture any beryllium particulate contamination which might have formed. The residual pressure with this vacuum was about 60 millitorr. This was not low enough to prevent some oxidation and migration of metal from the test bar and tensile load string to the inside of the quartz tube during testing. This made its surface milky, as shown later in Figure 55. The oxidation was not thick enough to significantly affect mechanical properties, but it made the scribe marks more and more difficult to see over time. We eventually had to interrupt the tests, take the vacuum assembly apart, and clean the tube. For future experiments, we recommend using a small (2") diffusion pump to improve the vacuum. These typically cost about \$600 used.

Quartz Tube and End Caps

The evacuated quartz tube protected the beryllium sample from excessive oxidation, controlled any beryllium contamination and allowed us to test the sample under typical service conditions. The end caps seal to the quartz tube to maintain vacuum. They are cooled, and there are cooled radiation barriers somewhat inboard of the end caps. This keeps the seal region cool enough to allow use of ordinary vacuum seals and vacuum grease.

Tensioning Apparatus

The spring-loaded tensioning device kept a constant stress on the specimen during creep testing (we were interested in constant stress rather than constant strain), allowed variation of stress as needed throughout the experiments, and allowed us to provide a wide range of tension with little change in the apparatus. The device was cooled to ensure accurate tensioning. The tensioning wire was made of 1/8" diameter alloy 230 welding rod.

Beryllium Samples

The test bars were used as purchased from Brush-Wellman. A drawing of the bar is shown in Figure 54. Brush was not very prompt with our beryllium test bar order. This raises questions about supplier responsiveness if there is a decision to actually use beryllium in a receiver. We made a request for a quote during the first week of December, 1987 and Brush responded in late January. The quote was 2.5 times higher than previous approximate quotes. Furthermore, they called back in mid-February and discussed changing the material and processing specified to something which they thought would better suit our needs, and which also would be a lot easier for them to make. It was not entirely clear which of those considerations drove them to ask us to change the order.

Figure 55 shows a photograph of the test bar inside the quartz tube. The photograph is a top view of the quartz tube inside the furnace. The furnace is open. The right side hole/alumina/refractory joint can be seen but the left joint is clouded by beryllium oxide on the quartz tube. Also shown in the photograph is the flat quartz plate (bottom center).

MISSING
73 — 88




Figure 54 Beryllium Test Bar




Figure 55 Beryllium Test Bar inside Quartz Tube

End Connections for Test Bars

We considered various end connection techniques. Two appeared most promising. One is to screw a multiply slotted sleeve onto the anodized threaded end of the test bar. This would work if the anodized coating is thick enough to form a good diffusion barrier and if there are enough slots to lower the modulus of the sleeve to keep it from overcompressing the end of the test bar and causing it to creep to failure. This transition method is not suitable for a gas-tight seal for an actual receiver. The other option is to drill holes through the ends of the test bar, put beryllia sleeves in them, and put superalloy or refractory pins through these. A similar scheme using beryllia discs hot-pressed into the tubes could provide a gas-tight seal in an actual receiver. Graded seals of beryllia to beryllium have been made and used at these temperatures in the past. This technique could be used for a gas-tight seal in the receiver. However, it would have been far too expensive and time-consuming to use in our creep tests. We finally decided to use the hole/sleeve/pin method.

It took several attempts before we found a proper solution. Our fourth control experiment verified that we had an appropriate tensioning method. The first three control experiments used "horseshoe" shaped connectors as shown in Figure 56. The bar gradually bent in the plane of the connectors, indicating uneven load distribution. The fourth test used the 3-piece load transfer device shown in Figure 57. This led to uniform creep, with no bending.




Figure 56 Initial Test Bar Support Scheme (Caused Bar to Bend)




Figure 57 Final Test Bar Support Scheme (No Bending)

Alumina spacers properly installed in the transverse holes at the ends of the test bar prevent bar-pin contact under tension. The spacers broke at the ends under certain conditions (usually during removal), but this was not a major problem. The spacers often could be reused. It is important, however, that the ends of the alumina tube not protrude beyond the edge of the hole in the test bar: this allows it to fracture, and the fracture propagates back into the hole and breaks out chips of alumina in areas where it is supported by the test bar.

2.3.3 Test Procedures

Prior to testing, we discussed appropriate safety precautions for handling of beryllium with the General Atomics (GA) safety office (GA, being our lessor, wished to ensure that we perform our work safely). GA felt that our precautions were adequate and that there would be no problem handling the beryllium safely. Since GA has extensive experience with high temperature beryllium testing, this both added to our confidence that we had chosen a reasonable and safe way to perform our experiments and ensured that there would be no unexpected safety problems or cleanup expenses.

After attending to the issue of safety, we performed an initial test using a piece of drill rod. We used drill rod initially, due to its availability and the relative ease of machining a test bar from it. Its composition was 0.20-0.40% Si; 0.40-0.60% Cr; 0.10-0.20% V; 0.40-0.60% W; balance Fe. We crept the bar at 710°C. Thermal expansivity was consistent with handbook values, and the apparent creep rate was consistent with a stress exponent of 7 or 8 and expected linear creep values. The metallurgical state and mechanical properties of the alloy 230 tensioning wire were unchanged during the test.

For the next three tests we used a 304 stainless steel bar. We crept the bar at temperatures of 750°C, 800°C and 995°C. At the first two temperatures the bar did not creep significantly, in accord with literature data, and the tensioning apparatus did not show any difficulty. However, the 995°C test led to unacceptable creep of the tensioning apparatus. Therefore, we thought that it was reasonable to expect that the apparatus would be adequate to test the beryllium bars at 813°C (1500°F).

Our method of installing the test bar in the quartz tube was to insert a wire through the quartz tube, attach the wire to the tension train, put the train under tension and adjust the supports as necessary, and then slide the quartz tube over the train. We then installed the end caps, keeping the whole tensile string under tension at all times while we installed it in the furnace. This worked about 2 times in 3. When it did not, we found it best to take the tube off and start over rather than to try to shake the tube or reach in with a rod to adjust the supports. The distances between the tensioning rods and the end of the test bar were adjusted with dial indicator calipers to be the same, usually about 140 mils. These preparations were common to both sets of tests.

Also common to both sets of tests was a limit on the allowable tension levels. This was set by the remainder of the tensile string, which crept too rapidly to allow us to maintain constant tension when we set the tension much above 10 lbs. Since the test results are limited to cases with low imposed load, they apply directly only to similar cases. The non-linear component of a temperature field within an object imposes mechanical stresses that cannot be relieved by bending, stresses of several MPa per degree departure from the best-fitting linear field. Hence our results are limited to cases where the maximum non-linear temperature variation is of order 1°C. In cases with higher non-linear variations, strain rates due to creep and thermal-cycling-induced ratcheting will be larger than measured here.

The test procedures themselves are presented together on the next page. This is followed by a brief summary of our error analysis.

Creep Tests

We conducted these tests at constant (or nearly constant) temperature and stress. Each time the distances between marks were read, the following procedure was used:

- Record temperature, tension and vacuum
- Read test bar mark positions from left to right
- Record temperature, tension and vacuum again
- Read test bar mark positions from right to left
- Record temperature, tension and vacuum again
- Reset tension to nominal value, to compensate for creep

We also recorded the temperature of the micrometer slide each time. We did this to support an error analysis, because the operator recording the positions of the marks altered the air flow through the hood enough to change the temperature near the micrometer slide by 1-2°C. This turned out to have a negligible effect on the readings.

Thermal Cycling Tests

We performed these tests by cycling the temperature between two nearly constant values and allowing the stress to vary as the test bar and the rest of the tensile string expanded and contracted during each cycle. We obtained the desired upper and lower temperatures using a timer to turn the furnace on and off at preset intervals. It usually took several hours to get the desired upper and lower temperatures at the beginning of each test. We then had to fine-tune the timing for a few days until the cycling stabilized. We used the following procedure to collect data:

- Record low-end temperature, tension and vacuum
- Set timer "on" (i.e., furnace on at all times) after recording low-end data
- Record high-end temperature, tension and vacuum
- Allow temperature to stabilize at a preset value ~15-25°C below high end
- Record temperature, tension and vacuum
- Read test bar mark positions from left to right
- Record temperature, tension and vacuum again
- Read test bar mark positions from right to left
- Record temperature, tension and vacuum again
- Reset tension to nominal value, to compensate for creep
- Set timer to "auto" so that it resumes cycling

We performed the thermal cycling tests after we determined that the temperature variations of the micrometer slide during recording were too small to have any effect. Hence we did not record this temperature for the cycling tests.

Error Analysis

We analyzed uncertainties contributed by the various factors listed below, and found them to be of the magnitude indicated:

Micrometer slide instrument error ($\pm 5 \times 10^{-6}$, fractional)
Furnace temperature controller instrument error ($\pm .005^\circ\text{C}$)
Timer instrument error ($\pm 1 \times 10^{-4}$ fractional)
Operator reading error of test bar mark positions ($\pm .002$ cm)
Adjustment for tensile strain uncertainty ($< \pm 1 \times 10^{-5}$ cm))
Adjustment for furnace temperature uncertainty ($\pm 1 \times 10^{-4}$ cm)
Adjustment for micrometer slide temperature uncertainty ($\pm 1 \times 10^{-5}$ cm)
Adjustment for creep uncertainty during reading ($< \pm 1 \times 10^{-5}$ cm)

This list shows that the largest error (by more than an order of magnitude) was the operator reading error. It totally dominated the final uncertainty in test bar mark position.

The reading procedure involved offsetting the microscope and then moving it until the cross-hair lined up with the center of the machined mark. We determined the reader error by doing this 20 consecutive times with the same mark position and calculating the standard deviation. Unfortunately, the 20 readings were done only moving the microscope from left to right. Doing the same thing from right to left will probably give slightly different results, partly due to play in the microscope slide, and partly due to perceptual differences in aligning the cross-hair with the center of the mark. Unfortunately, we did not at the same time do 20 right-to-left readings to characterize this source of error.

This would all be academic were it not for the fact that the right to left data sometimes give strange results (e.g., negative creep under tension) which fall outside of the uncertainty range. Equally strange results from the left to right readings were always within the uncertainty range. This is most likely due to the fact that we had to use the same operator error estimate for both left to right and right to left tests when in fact it should have been used only for the left to right tests.

These errors were much more of a problem for the creep tests, where the apparent creep was comparable to the uncertainty, than for the thermal cycling tests, where the results were an order of magnitude greater than the uncertainty.

2.3.4 Test Results

As mentioned in the introduction to section 2.3, we performed a total of 6 tests, not including the initial calibration tests. Three of these were creep tests and three were thermal cycling tests. The two sets will be presented separately.

Creep Test Results

The creep test results are summarized in Table X and presented graphically in Figures 58 - 63. We present results from both the left-right and right-left reading sequencess to show the level of uncertainty in the test results.

Table X Creep Test Results

Test #	Test Duration (seconds)	Temperature (°C)	Stress (psi)	Strain Rate ($\times 10^{-10} \text{ s}^{-1}$)
1 LR	982,057	695	134	4.7 ± 2
1 RL	982,057	695	134	-3.0 ± 3
2 LR	1,296,449	695	268	-4.3 ± 1.7 *
2 RL	1,296,449	695	268	1.0 ± 1.5
3 LR	2,611,410	803	90	1.1 ± 1.8
3 RL	2,611,410	803	90	-3.3 ± 2

Note: Residual vacuum pressure for all tests was about 60 millitorr.

* - This result is unexpected and not easily explained.

The results are within our previously-calculated allowable creep rates of 3 to $5 \times 10^{-10} \text{ s}^{-1}$. Therefore, from a creep standpoint beryllium could be used at temperatures of 695°C to 803°C (968 K to 1076 K) and constant stresses of 90 to 268 psi (.6 to 1.8 MPa). (As noted in 2.3.3, keeping the mechanical stresses this small requires good control of the non-linear component of the temperature field in the beryllium storage elements.)




Figure 58 Distance between Marks vs. Time - Test 1 LR




Figure 59 Distance between Marks vs. Time - Test 1 RL




Figure 60 Distance between Marks vs. Time - Test 2 LR




Figure 61 Distance between Marks vs. Time - Test 2 RL




Figure 62 Distance between Marks vs. Time - Test 3 LR




Figure 63 Distance between Marks vs. Time - Test 3 RL

Thermal Cycling Test Results

The thermal cycling test results are summarized in Table XI and presented graphically in Figures 64 - 69.

Table XI Thermal Cycling Test Results

Test #	Test Duration (seconds)	# of Cycles	Temperature (°C) Min/Max	Stress (psi) Min/Max	Strain Rate ($\times 10^{-10} \text{ s}^{-1}$)
4 LR	3,399,000	2266	724/817	83/120	$3.4 \pm .7$
4 RL	3,399,000	2266	724/817	83/120	$3.6 \pm .7$
5 LR	1,806,000	602	622/817	82/155	25.8 ± 2
5 RL	1,806,000	602	622/817	82/155	23.6 ± 2
6 LR	3,216,000	1072	632/826	47/115	$8.1 \pm .9$
6 RL	3,216,000	1072	632/826	47/115	6.8 ± 1.2

Note: Residual vacuum pressure for all tests was about 60 millitorr.

The results from test 4 are within our calculated allowable creep rates of 3 to $5 \times 10^{-10} \text{ s}^{-1}$, but they correspond to relatively modest temperature cycles, less than desirable in a sensible storage element.

Tests 5 and 6 use temperature cycles closer to those desirable along most of the length of a beryllium storage element. (The ends need smaller cycle amplitudes, both because of the connectors and for other reasons.) The results from tests 5 and 6 lie well outside the acceptable range. The cycle amplitudes in tests 5 and 6 are essentially identical (194 vs 195°C). Test 5 is at a slightly lower temperature, but has about 50% higher average stress. Since the resulting creep rates are over 3 times higher in test 5 than in test 6, the stress exponent appears to be roughly 3 (i.e., $\text{Ln}(3)/\text{Ln}(1.5)$). This indicates that non-linear temperature gradients will have to be severely controlled to keep strain rates under control.

It is also worth comparing tests 4 and 6. Test 4 has a comparable maximum temperature to test 6, and a higher average temperature than test 6. Test 4 also has a moderately higher stress. But test 6 has a thermal cycling amplitude 2.1 times higher than test 4. Since the strain rate is about twice as high in test 6 as in test 4, clearly cycling is important, even when it involves only extending the low end of the cycle downward, with no change in the maximum temperature.




Figure 64 Distance between Marks vs. Time - Test 4 LR




Figure 65 Distance between Marks vs. Time - Test 4 RL




Figure 66 Distance between Marks vs. Time - Test 5 LR




Figure 67 Distance between Marks vs. Time - Test 5 RL




Figure 68 Distance between Marks vs. Time - Test 6 LR




Figure 69 Distance between Marks vs. Time - Test 6 RL

2.4 Receiver Efficiency and Mass Calculations

During Phase I, we wrote a system performance and mass estimation program ("Brayton") to help us properly optimize an overall system using sensible storage. We revised and refined the program during this contract. The program predicts a very strong dependence of system mass on the peak cycle temperature: changes in peak temperature affect all other components (radiator, heat exchanger, engine pressure ratio, etc.), and have much larger effects on overall system mass than might at first seem likely. Detailed results of a sensitivity study are presented at the end of 4.1 as part of a comparison of beryllium and graphite receivers. Here we just present a rough breakdown of component masses for our strawman design. During this contract we also wrote a ray-tracing program to evaluate receiver optical efficiency as a function of cavity shape and storage element absorptivity. That program and the conclusions we reached after exercising it are described below.

Component Masses

Mass estimates on our receiver design showed that it involves about 20 kg of KBI 10 or 41 tubing to conduct gas flow through the receiver (though there is a slight possibility vitreous carbon would be used in a graphite receiver), 120 kg of platinized graphite insulation shell (including door), roughly 10 kg of door hinge mechanism, the required amount of sensible heat storage material (about 500 kg of beryllium or 850 kg of graphite - assuming 1088K operating temperature), 15 to 50 kg of material associated with end connectors for the sensible heat storage tubes, and a significant but unquantified amount of structure associated with stabilizing the receiver structure during launch.

Energy Distribution in Receiver

We favor an ogival receiver cavity, with fairly smooth cavity surfaces to reflect incident radiation (most of which is incident towards the aperture end of the receiver) towards the back of the receiver. This allows most radiation to be given several chances to be absorbed before it is reflected back out through the front. To quantify the feasibility of this concept, we wrote and used a ray-tracing program that estimates energy distribution inside a truncated ogival receiver. The program assumes specular reflectance, but does not make gray-body approximations, because such approximations can significantly overestimate the multiple-bounce absorptance of a cavity with non-gray walls. (Consider the extreme case of a surface which absorbs the solar UV and visible, and reflects the solar infrared. The multiple-bounce absorptance of a cavity made of such a material will be no higher than the single-bounce absorptance, because the infrared will keep reflecting until it escapes.)

The program approximates the reflectance of beryllium or platinized graphite by dividing the solar spectrum into 10 equal-energy bands, each with their own reflectance value. (For platinum, we used a range of 0.4 to 0.85 in 0.05 increments. This provided a reasonable fit to a reflectance curve obtained from an encyclopedia.) The program then pre-calculates the total energy absorbed at each of the first 10 bounces by summing over these 10 "bands." The sums are then used to accumulate energy into axial "bins" as each ray is traced.

We varied the intensity of incoming radiation versus off-axis angle in a way that is representative of a paraboloidal concentrator with a small receiver shadowing the center of the paraboloid. We assumed that concentrator slope errors were independent of location on the concentrator. We modeled those errors as having a rectangular distribution within selectable limits.

The program keeps track of how many times a given ray reflects off the receiver surfaces before it exits the receiver. Using the program showed us that an ogival cavity shape plus specular reflection can cause all rays except for those emanating from near the center of the concentrator (i.e., rays which hit the back of the cavity) to have the same number of reflections before they exit the cavity. For a typical ogive, the number is about 6. This is high enough to give a fairly high cavity absorptance even if the absorptance at any one bounce is fairly low.

3. RECEIVER DOOR, HINGE, AND ACTUATOR DESIGN

The effort on task 3 was more limited than originally expected, because much higher levels of effort were required to develop good testing techniques for the solar concentrator and the beryllium than we had expected. We were, however, able to develop preliminary designs for the door, hinge, and actuator that we believe should meet the criteria listed above. Sections 3.1-3.3 discuss those strawman door, hinge, and actuator design concepts.

To put this discussion in context, it is worth briefly reviewing our Phase I results. Our analysis and optimization efforts showed that heat loss from the receiver during eclipse was approximately 10 times as costly as equal heat loss per unit time during the day, because heat loss during the day could be compensated for with a slightly larger concentrator, while night losses require a somewhat larger proportional increase in the mass of the receiver, which is likely to be a much heavier component (particularly with our concentrator and receiver designs). We found that using a door to reduce night-time losses made a much larger aperture justified, and allowed equal overall performance with a concentrator having only half as large a concentration ratio. A review of work in the 1960s (and our own work under task 1) indicates that achievable concentration ratios are regularly well below those expected from the design, so an insulating door may be quite valuable.

But a door also has drawbacks: cost, complexity, weight, failure modes, etc. For a door to be worthwhile overall, it needs to have at least most of the characteristics listed below. We believe our strawman door, hinge, and actuator designs do this.

- heat losses through and around door should be much lower than with open aperture
- door should open and close rapidly ($< < 1$ minute), and should require little power
- surfaces should be dark and diffuse to prevent dazzling crew or damaging sensors
- door should tolerate concentrator focusing on it for a few days (open or closed)
- hinge and actuator should fail open if failure is possible
- door should open out of path of reflected rays that can reach the aperture
- door should open 90° or 270° to minimize door shadowing of concentrator.

3.1 Strawman Door Design

Our strawman door design uses several layers of graphite. The outer surfaces (both sides) are roughened enough to make them diffuse in the visible. They are coated with non-metallic coatings such as SiC or TiC to reduce atomic oxygen erosion rates (which would be several mm over a 10 year life otherwise). The inner layers are coated with iridium or platinum to reduce emittance. The door is assembled from flat sheets of appropriately coated Poco graphite sheet by using pitch as "glue" and then pyrolyzing the assembly. To make the door tolerant of large thermal gradients (through the thickness and across its surface), it is assembled as a series of overlapping shingles, with a high enough aspect ratio in the spaces between shingles to keep radiative heat leaks along the spaces adequately low. Figure 70 shows the design in cross-section.




Figure 70 Strawman Door Design

3.2 Strawman Hinge Location and Design Options

Our emphasis is on hinge designs which do not involve sliding, and it is also desirable to avoid rolling to the extent possible. Two design concepts are shown below in Figure 71:




Figure 71 Two Hinge Design Concepts Using Refractory Metal Springs as Hinge Elements

Torsion springs made of KBI 10 or KBI 41 tantalum alloy are attractive since these alloys have a low modulus and anneal at sufficiently high temperature that they can be used unprotected as a spring hinge with no concern about them being detempered. A torsion spring as the hinge element also provides an inherent opening force, so that actuation is required merely to close the door.

The hinge can either be located at the edge of the front face of the receiver, or displaced some distance down the side. The two positions are shown below in Figure 72. The proportions are appropriate for a graphite-storage receiver, which can be far smaller than a beryllium-storage receiver, as discussed in chapter 4.

If the hinge is at the edge of the front, then it is exposed to higher heating. It is desirable to have it open 270° to minimize shading of incident rays, but having it open 145° is enough to get it out of the reflected rays, which is more important. If it is located down the barrel, then heating is less, and it only has to open 90° , but the hinge has to be displaced far enough to keep the door out of the way of reflected rays that can enter the receiver aperture. In addition, the door is more subject to misalignment due to warping of the arm. Given the high heat resistance of the KBI 10 and 41 alloys, it is probably best to put the hinge at the edge of the front of the receiver, have the door open 145° , and make the concentrator diameter about 0.1% larger to compensate for the small loss of incident light.




Figure 72 Hinge Location Options and Door Positions, Open and Closed

3.3 Strawman Actuator Design

Given a hinge in a high-heat area which is spring-loaded toward an open position, there are four main requirements on the actuator:

- ability to tolerate high temperatures, or to exert force at some distance
- ability to overcome the spring force to close the door
- capability to be disabled if it fails in the door-shut position
- compatibility with redundant design.

The simplest design we have found that meets these criteria is a set of Nitinol springs located near the back of the receiver, with refractory metal bands to transmit a tensile force to the hinge. Resistively heating the Nitinol actuators causes them to shorten, which pulls on the metal bands, which are wrapped partly around the hinge. The same basic design is currently used by one Japanese car manufacturer, to open retractable headlights. The springs are in series with the headlights, so whenever the headlights are turned on, the springs shorten and raise the headlights. Nitinol memory alloys in these types of applications are known to survive tens of millions of cycles with no significant change in actuation properties. The transition temperature is selectable by controlling alloy composition. This concept is shown below in Figure 73:




Figure 73 Redundant Nitinol Actuators and Refractory Tensile Bands

The Nitinol needs to be located back far enough that it can be kept below its transition temperature. Equally important is arranging for the Nitinol to "see" a cool radiative environment: screened from direct and reflected sunlight, and from the warm outer barrel surface, but exposed to space (or the rather cool earth) otherwise. It is acceptable for episodes of tracking error to heat the Nitinol above its transition temperature, because the door should be closed then anyway. But it is important to keep the Nitinol below a much higher temperature at which it loses its memory. If necessary, the Nitinol can be in the vicinity of (and in thermal contact with) the cool side of the Brayton engine cycle, to keep maximum temperature down.

Keeping enough Nitinol springs hot to close the door should not take more than 10 watts. Initial heating at 20 watts will be required for about 10 seconds, to start the actuator moving. The Nitinol should thus be energized and de-energized shortly before the end of sunset and the beginning of sunset, so the actuators will close and open the door at about the time that the solar flux and the thermal loss flux through the aperture are equal.

Redundancy can be assured by having several actuators in parallel, each with their own band to the hinge. If an actuator heating circuit fails or is not used, that spring and band will become slack when another actuator closes the door. To take up the slack, each band can have a slight sinusoidal set near the cold end: strong enough to take up slack, but weak enough to keep from closing the door. The actuators and bands can be trapped in individual channels if necessary, so that if they break they cannot foul on others, or on the door itself. This will also reduce exposure of the bands to atomic oxygen. This allows the bands to be thinner and more flexible.

4. SYSTEMS AND SAFETY ANALYSIS

Under Task 4, we refined and exercised a preliminary system performance and mass estimation program we had developed under Phase I. The results show a high sensitivity of system mass to peak cycle temperature. A higher turbine inlet temperature raises efficiency, which reduces concentrator size (and cantilever loads). Lower heat rejection requirements and a slightly higher optimum rejection temperature significantly reduce required radiator area.

These results, plus our awareness of creep and thermal ratcheting as possibly imposing maximum use temperatures on a beryllium receiver, led us to look at other storage media. During Phase I we had looked at and discarded graphite, because its specific heat is 42% lower than beryllium's (at equal temperature). Specific heat rises with temperature, so a graphite receiver that runs hotter will have a higher specific heat. But this is a small effect compared to the cascading effects of system reoptimization for a higher turbine inlet temperature. We looked at graphite again, and found that if it allows receiver operation 120-150 K hotter than beryllium does, the resulting increase in efficiency can reduce overall system weight enough to compensate for its much lower specific heat. The results are shown at the end of section 4.1.

Peak cycle temperature can be limited by concentrator limits, the storage medium, the turbine, or other materials (toluene degradation, in the case of an organic Rankine system). Graphite itself is unlikely to impose any temperature limits, so the real limits in a graphite-storage Brayton system are likely to be the concentrator or the turbine. Hence if beryllium itself is the limiting factor in a beryllium-storage receiver, and using graphite allows the peak temperature to be raised at least 120-150K, then graphite results in a lower system mass.

In addition to reducing system weight, graphite offers other significant advantages: it is much cheaper to buy and far more convenient to work with; it is relatively easy to connect dissimilar materials to the ends of graphite tubes; there are no significant risks in handling it; it is not subject to creep, thermal shock or thermal fatigue in the receiver thermal environment; it is far less volatile than beryllium; and the little graphite that does sublime and condense on concentrator or other space station surfaces will oxidize to CO and leave. These advantages of graphite over beryllium (particularly the freedom from creep and thermal ratcheting) led us to study graphite receiver design issues in more detail under task 4, because of the diverse systems implications of a graphite-storage receiver. Sections 4.1-4.3 summarize our work on graphite receiver design, thermal analysis, and materials selection.

Our other systems analysis work under task 4 consisted of development of detailed simulation programs that model the optics, thermal transients, and coupled system behavior of a solar dynamic system using sensible storage (beryllium or graphite) and a Brayton engine. A disk containing those programs is being provided to NASA Lewis, along with typical output files from those runs. The results of other analyses and iterative design modifications made using these programs are incorporated into the beryllium receiver design described in chapter 2 and the graphite receiver concept described in this chapter.

4.1 Strawman Graphite Receiver Design

The graphite receiver design described below is the result of our work on a beryllium receiver (described in chapter 2), plus design iterations and analyses that take into account the interacting implications of the property differences between beryllium and graphite. Typical results of cycle and system optimization are listed at the end of the section.

Cavity Geometry

For a net thermal output of 65,000 W (day and night), the strawman receiver cavity is a cylinder 1.0 m long and 0.7 m in diameter. The storage elements are .15 m thick, so the outside diameter of the receiver cavity, about which insulation must be wrapped, is 1 m. These dimensions are baselines that need to be fine-tuned, but the ultimate values seem unlikely to change by more than about 10%, because overall losses go up when the receiver shape is changed much. For beryllium the cavity length and diameter should be at least 50% larger, and the storage elements no more than 1/3 as thick, to reduce thermal gradients and stresses.

Storage Element Geometry

The cylindrical wall of the cavity is formed from 16 graphite "billets." The billets are 1 m long and 0.15 m thick, and trapezoidal in cross section: 0.135 m wide on the inside and 0.194 m wide on the outside. These billet dimensions and a bulk density of 1.62 give a mass of 40 kg. For 16 billets the total mass is 640 kg. If fabrication is easier with other wider or narrower billements, the width and number of billets may be changed, but the total mass for a 25 kWe system is likely to be should remain between 600 and 700 kg.

Small holes are drilled through the length to serve as gas passages. Collecting manifolds are used at each end. The required flow area is only about 1% of the billet cross-section, so these holes need have only a negligible effect on billet mass and conductivity. If the holes are reasonably well distributed, then the ΔT from storage to hole-surface can be much smaller than the cavity-surface-to-storage ΔT , and thus not a major heat transfer issue. (But as discussed later, non-linear transverse thermal gradients and the resulting stresses may be an issue.)

Effects of a Platinum Coating on Billet Surfaces

The cavity and back surfaces of each billet are smoothed so that reflected radiation is directed further into the cavity rather than scattered back out the aperture. This raises overall absorptance modestly. We have considered putting a very thin coating of platinum on the billet. The reduced emittance of platinum-coated billets reduces the effective cavity emittance moderately, and hence reduces night losses if a door is not used or does not fit well. It also reduces radiative heat transfer between the inlet and outlet ends of the billet. Such thermal "short-circuiting" along the billets reduces the achievable gas outlet temperature, and hence slightly degrades system performance.

On the other hand, platinum significantly degrades optical efficiency. It also reduces the axial flux gradient below the desired range. It also decreases the peak local temperatures allowed, but on the other hand the lower absorptance also reduces local overheating.

Overall, platinum does not now appear justified in the baseline design. Adding it later does not change the basic design concept, but it does significantly change the analyses and affects design optimization. The key ingredients required for a more careful analysis appear to be as follows. Accurate spectral data for platinum and reflectance data for platinized graphite are needed, to determine whether platinized graphite can have a high enough multiple-bounce absorptance and specularity over enough of the solar spectrum. Average solar absorptance should not be used, because of non-linearities in multiple-bounce absorptance. Longer wavelengths are poorly absorbed, but also more specularly reflected. Most radiation will come in at a large angle to the axis. This increases the number of specular reflections required before the ray exits, but also decreases the number of off-specular or diffuse reflections required for the ray to exit. Parallel and perpendicular polarizations should be handled differently, since they have different reflectivities and (to a first approximation) keep the same polarization geometry upon successive bounces. It is possible that optical testing of models is the best way to measure solar absorptance.

The side walls of each billet should not be platinized. Then adjoining billets are better coupled thermally. This reduces billet-to-billet temperature variations and side-to-side distortions. On the other hand, there is increased absorption at the front of each crack. This probably is not an issue.

To increase the number of bounces, the surfaces might be grooved transversely, so as to direct radiation not across the cavity but to closer surfaces. This has diminishing returns, because it also increases emittance-area products.

Balance-of-System Mass

Because the receiver is very compact, the balance of system mass (structure, plumbing, door, and insulation) should be quite low. We assume for now that balance-of-system mass is <56% of the storage mass, for a total receiver mass of under 1000 kg. One way to make a fair comparison (and perhaps a reasonable estimate) is to use the same insulation weight per unit area, and the same plumbing weight per unit length, as in Boeing's receiver design.

Preliminary System Optimization

We used our system evaluation and optimization program, Brayton, to optimize radiator temperatures, gas pressure ratio, and heat exchanger NTU as a function of storage material (Be or C) and peak allowable cycle temperature (900-1300K). All runs assume the same total gas loop pressure drop outside the compressor and turbine: 4% of gas pressure. All runs use the same compressor, turbine, and generator efficiencies, and the same receiver mass effectiveness (the ratio of ideal sensible storage mass to total actual receiver mass.) Because graphite receivers are heavier but smaller, balance-of-receiver mass should be lower with a graphite receiver. Hence a fixed multiplier contains an implicit bias against graphite. The component weight scaling laws are listed with the output.

Fixed parameters:

0.82 Compressor efficiency
 0.86 Turbine efficiency
 0.98 Generator efficiency
 0.35 Receiver mass effectiveness
 0.04 Pressure drop around loop

Component mass scaling laws:

1.2 Kg/effective m² of concentrator
 5.0 Kg/Kw output for brayton + misc.
 7.0 Kg/effective m² of radiator
 20.0 Kg/(Kw/K) for heat exchanger

Optimized cases for different Tmax (K)

Case	NTU	CycEff	M2eCon	M2eRd	TotMass
Be900	7.0	0.3282	3.480	2.140	48.658
C1020	6.0	0.3767	3.032	1.752	48.652
C1100	6.0	0.3959	2.885	1.407	41.466
C1200	6.0	0.4112	2.778	1.048	34.916
C1300	6.0	0.4340	2.632	0.910	30.059

Optimized component inlet temperatures (K):

Case	Comp	HX	Rec	Turb	HX	Rad
Be900	250	421	604	900	630	447
C1020	245	439	648	1020	683	474
C1100	250	457	688	1100	726	496
C1200	260	486	738	1200	780	528
C1300	260	508	775	1300	819	553

Estimated component masses, Kg/Kw:

Case	Mconc +	Mrecv +	Mbray +	Mhx +	Mrad =	TotMass
Be900	3.053	24.185	5.000	1.440	14.979	48.658
C1020	2.660	27.871	5.000	0.857	12.264	48.652
C1100	2.531	23.351	5.000	0.735	9.850	41.466
C1200	2.437	19.509	5.000	0.632	7.338	34.916
C1300	2.309	15.853	5.000	0.526	6.371	30.059

Our overall conclusion is that heat storage in carbon rather than beryllium requires about a 13% higher turbine inlet temp for equal system mass. This drives mirror accuracy: for equal losses, the receiver aperture diameter must be 22% smaller, or 65% higher thermal losses must be accepted. But if concentrator performance or losses are not dominant, moderately higher losses may be affordable, and a graphite system is preferable overall.

4.2 Preliminary Thermal Analysis

This section discusses the implications of heat loss, radiative short-circuiting, and thermally induced stresses on design.

Heat losses occur in 3 main ways:

1. Transport through insulation around the "barrel" and back-wall
2. Radiation through the front wall and aperture-door assembly
3. Conduction through structure and engine plumbing

The outside of the barrel, and the back of the receiver, have a large total area, but the total heat loss can be modest because these areas are easy to insulate. We assume an equivalent overall emittance (including leakage & conduction) of 0.02. This corresponds to multiple layers with correspondingly higher emittances: 5 layers at .10, etc. With close-fitting insulation, the diameter at mid-thickness could be about 1.02 m. This makes the AE (area*emittance product) for the barrel and back surface about .08 m². Hence despite the area, the losses are modest.

The aperture is much smaller, but harder to insulate. Neglecting transparent insulation that can be used in the sun, the options are to leave the aperture uninsulated (and try to keep cavity area and perhaps longwave emittance low), or to close the aperture at night.

As shown in the Phase I report, closing the aperture at night can be quite effective. Night-time loss fluxes are roughly 10 times as expensive as day-time loss fluxes, because day losses just require a slightly larger concentrator, while night losses also require a larger fractional increase in a much heavier storage mass. Hence the more important heat loss condition to analyze is for the case when the door is closed. But "closing the door" does not eliminate losses entirely, and the remainder of the front wall, which must be tolerant of extreme heat loads, is hard to insulate very well. We have assumed an effective average emittance of .06. This gives an effective loss AE of 0.05 m².

At an average night-time billet MRT of 1000K, 0.13 m² of loss area has an average night-time loss of 7400 W. (Sensible storage allows the MRT to be only moderately above the average gas temperature in the receiver, and well below the exit temperature.)

There are structural leakage paths, but these can be minimal if launch loads are passed to structure that is removed or not in contact during on-orbit operations. Losses in engine plumbing can be minimized by keeping the piping and turbine assembly compact and their surface emittances low. We assume that the hot-side plumbing is 0.3 m long and 0.06 m in diameter, and that the hot-side turbine assembly is a 0.2m diameter cylinder 0.1m long with one end. The total area is about 0.15 m². An emittance of 0.1 should be easily achievable, for a loss of 900 W at 1020K.

Hence the night-time losses should be approximately 8,300 W, or about 13% as large as the net heat delivered. This is high enough to worry about reducing: eliminating it would allow an 11% reduction in storage mass, cavity area, and support structure. But it is not a dominant issue, and hence need not be a key topic during preliminary design.

Internal Short-circuiting

By "internal short-circuiting" we mean radiation and conduction from the hot exit end of the receiver to the cooler inlet end. It can be a significant source of entropy, because of the large axial thermal gradients in a sensible-storage receiver, and the high radiative transport that occurs at high temperatures. Minimizing it involves using long low-conductance billets, reducing clearances between billets and between the billets and barrel insulation, and reducing cavity area and/or billet emittance. The last item can be improved by platinizing. With low emittances, the cavity becomes a "hall of mirrors" and the detailed optimization of the geometry becomes less important than simply reducing area-emittance products.

With the door closed, the inside barrel area is 2.2 m². The low-mass front and back (0.77 m²) can be considered adiabatic for this purpose. If the billet surface emittance can be kept down to .2, the participating cavity AE product is .44 m². Clearances between billets should add under .01 m² (16 x 1m x <.6 mm), with near unit emittance.

Cascade short-circuiting between the billets (considered as a solid barrel) and surrounding insulation could exceed short-circuiting within the cavity itself unless the barrel insulation is close-fitting, because the barrel's outside area is 40% larger than its inside area. (The smaller of the AE product and the annular area constrains the annular short-circuiting flux.)

The axial radiative conductance of 16 billet-to-billet clearances of 0.6 mm by .15 m should be negligible since the total cross-sectional area is only .00144 m²; it appears to be <1% as large as true conduction effects.

Effects of Transverse Thermal Gradients in Billets

Finite conductance through the thickness limits billet thickness. The day-night average net flux is 65 KW, and the net loss flux through the barrel insulation is 4 Kw. If the flow paths are reasonably well distributed through the billet, so that the gas samples the average storage temperature, then the average surface overheat required (compared to the average billet temperature) is the same as having 1/3 of the heat (21.7 Kw) flow through the full thickness. This plus the barrel insulation loss of 4 KW (which must flow through the full thickness) is the equivalent average flux through the billets. With a flux of 25.7 Kw, a transverse graphite billet conductance (at temperature) of 70 W/MK, and thickness of .15 m, and an average (at .775 m diameter, 1/4 of the way out) cross-sectional area of 2.44 m², the average surface overheat should be 23K. This causes an average increase of about 9% in heat loss from the cavity.

The deltaT is highest during the day (and approaches 100K near the aperture), when added losses are cheap to replace. After night begins, it relaxes with a scale time of 900 seconds, 42% of the maximum night length. Averaging over the maximum time in eclipse, the average night-time deltaT should be 38% of the end-of-day deltaT.

The low fluxes and gradients near the back of the receiver make it possible to consider sculpting the billets so they are thicker (and possibly narrower) near the top. This allows a reduction in barrel length and/or circumference, reduces the adiabatic back wall area, and makes the hot-side manifold much shorter.

The billets might be made thicker overall, but the returns on this diminish quite rapidly. A 1% increase in thickness (and 1% reduction in heat flow area into the billet) increases the ΔT by 2%, and reduces heat loss area by only 1%. Things are actually a bit worse than this, because the ratio of barrel inner to outer diameter is already far from 1.0.

Thermal Distortion and Thermal Stresses

If expansion clearances are 0 at 1200K, at 900K with a 10 ppm/K coefficient they are 0.3%, and thus negligible. Required allowances for distortion should be comparably small. For example, let us assume that uneven flow within a billet induces a temperature difference of 20K from one side to the other. That shrinks one side 200 ppm relative to the other side, and makes it part of a circle with radius $.165\text{m}/200\text{ppm}$, or 825 m. Over a 1 m length this causes a displacement of center with respect to ends of 0.150 mm. From front to back a 60K gradient would cause bending of 0.5 mm.

Decreasing the cavity area and increasing the billet thickness has a first-order effect on the bending distortion, since the flux and hence the gradient (deg/meter) is first-order. But the non-linear component of thermal stress, which cannot be relieved by bending, grows with flux*thickness. The maximum non-linear stress for a parabolic temperature gradient occurs at the front and back surfaces and is 1/6 of the front-to-back ΔT . This stress is in compression, with the tensile stress being away from the surface and 1/2 as large.

Internal gradients from storage to the gas flow passages are mostly non-linear and thus cause thermal stresses which cannot be relieved by bending. Here the maximum stress is in tension, at the tube surface. Due to the radial inflow of heat from a large area to small holes, and the fact that bending cannot relieve the resulting thermally induced stresses, the maximum stress that cannot be relieved by bending is well over 2/3 of the total storage-to-hole ΔT , or 8 times larger than for an equal total ΔT between the front and the back. Hence it is probably worth focusing far more attention on minimizing mass-hole thermal gradients than on minimizing surface-mass thermal gradients.

Thermal Stresses: Graphite vs Beryllium

A comparison with beryllium is in order. At temperatures of interest, beryllium has a 90% higher volumetric heat capacity and a 50% higher conductivity than graphite, but its thermal expansion is 50% higher, its modulus 20 times higher, and its creep strength much lower.

For equal cavity area, beryllium storage billets can be 47% thinner. This gives them 65% lower thermal gradients through their thickness. But the non-linear thermal stresses are over 10 times higher than with graphite. This plus the thermal cycling "ratcheting" effect found in our tests described in chapter 2 limits the thickness of a beryllium receiver, and causes it to be much larger than a graphite receiver which has no creep or ratcheting problem at temperatures and stresses of interest. If the maximum stresses allowed are equal, then the beryllium cavity area must be 3.2 times larger. This leads to 3.2 times the structural and insulation mass, and (for equal insulation thickness), 3.2 times the heat loss. Even if the graphite receiver operates at a 15% higher temperature, its heat loss can be 46% lower (at 1/3 the insulation mass).

4.3 Materials Issues and Selection

This section briefly recounts the key issues we have discovered that affect materials selection, and lists materials options uncovered in literature reviews and discussions with suppliers and users of graphite in aerospace and other high-tech applications.

Type of Graphite

What is desired is low conductivity along the billets, and high transverse conductivity. Most processed graphite is anisotropic, with higher conductivity in the extrusion direction. Thus billets should probably be cut transversely from large blanks. Pyrolytic graphite would be ideal, but is probably unavailable in the required sizes. "Recrystallized graphite" (as described in the McGraw-Hill Encyclopedia of Technology), can have a density above 1.92, and controllable anisotropy in thermal conductivity. General Atomics has considerable expertise in high-tech forms and applications of graphite. Massoud Simnad (of General Atomics and of UCSD) is an appropriate contact.

Gas Leakage Through Graphite

Graphite heat storage tubes can be made impervious to gas leakage by coating them with metal carbides such as SiC or TiC. These coatings can be applied to the inside or outside of a tube using CVD by several vendors. General Atomics also has an amorphous carbon coating for graphite. They developed this coating to contain volatile fission products. Massoud Simnad says that those coatings can even be applied to the inside of long tubes. End connections to refractory metal alloy tubes are possible by using an intermediate piece of the metal carbide. This can be produced to net shape by hot pressing and is in equilibrium with both the graphite and the metal, so no reaction occurs. The refractory metals have thermal expansion coefficients near that of graphite, and most of them have low moduli, so thermal stresses will be fairly low.

Insulation

The main candidates here are:

1. Materials specified for the Boeing/Lewis receiver for space station
2. Advanced "tiles" for AFE and shuttle (contact: Howard Goldstein, NASA Ames)
3. Graphite felt (with protection from atomic oxygen)
4. Thin layers of platinized POCO graphite. Issue: launch vibrations

Materials for Balance of Receiver

Graphite sheet appears to be the material of choice for the inner liner, door and front plate of the receiver. POCO graphite is a high strength, fine grained graphite which has found many applications in aerospace. It is available as thin sheets which we believe could be laid up together with a pitch binder and graphitized into the required shapes.

We believe it is feasible to make formed shapes as thin as .4 mm in this manner. These sheets would weigh about 700 gms/m² (2 1/4 oz. per ft²). These sheets could be made oxygen resistant by overcoating with platinum or iridium (for reflective surfaces) or with an appropriate carbide (for emissive surfaces). Platinum and iridium can be vapor deposited in vacuum, have excellent adherence to graphite, and do not form stable oxides under the conditions obtaining in orbit. Coating graphite with various carbides (SiC and TiC seem the most interesting) is a commercially available service. On oxidation, these coatings form impermeable, temperature-stable metal oxide coatings.

Properly coated graphite sheets and constructs would not be susceptible to atomic oxygen attack. Micrometeoroids would damage the coating and allow local erosion, but proper component design should be able to tolerate local erosion everywhere but in the plumbing, which can be well protected. Properly coated graphite should not be damaged by short or long term exposure to stagnation temperatures at the focus of the concentrator. Carbide coatings can be dark and have diffuse reflectance, so accidentally focusing the concentrator on these parts will not damage instruments or dazzle astronauts.

Peak Temperature Capabilities

Carbon, platinum, iridium, titanium and silicon are not volatile enough at the stagnation temperature of our concentrator to cause significant contamination of the space station (or damage to the receiver), even if the tracking mechanism fails in such a way that it focuses the concentrator on one part of the exterior of the receiver and keeps this up for the 3 months until the next shuttle flight can bring tools to fix it.

The maximum allowable turbine inlet temperature (which will probably determine the maximum cycle temperature of a graphite receiver) is not clear from the literature or consultation with NASA Lewis personnel. The numbers typically quoted are 1300 to 1500F, or roughly 1000K to 1100K. A CBC turbine for space station use operates in a unique environment. The working fluid is a noble gas mixture, which is not chemically reactive in any way; and the system is surrounded by vacuum so air cannot leak in and react with system components. Steam and carbon dioxide are oxidizing with respect to carbon and refractory metals, so essentially all other turbines operate in an oxidizing atmosphere.

Graphite and refractory metals are very susceptible to oxidation at high temperatures, and thus are not normally considered as candidate turbine materials. In this unique case, graphite or refractory metal turbines are not subject to oxidation, and could be used at >>1100K. Several alloys are available which could operate at 1600K or greater, as could graphite. Since there is a lot of expertise at Lewis in long-duration use of refractory alloys and graphite in high temperature environments, this may be a fruitful area for further discussion, and perhaps some work at Lewis.

If a high temperature turbine (>>1100K) were available, it could lead to a solar dynamic power system substantially smaller and lighter than those already under development, with much smaller radiators and a relatively inexpensive graphite-storage receiver. Concentrator performance is likely to limit the overall optimum peak temperature to well below 1600K in the near term, but relaxing other temperature limits relaxes operating constraints and provides a useful driver for improvements in concentrator performance.

