LIGO Laboratory / LIGO Scientific Collaboration

LIGO-T1300756-v1 *LIGO* 10 September 2013

Status of UF Ring Heater Prototype development and testing

Mark Barbet, Giacomo Ciani

Distribution of this document:
LIGO Scientific Collaboration

This is an internal working note
of the LIGO Laboratory.

California Institute of Technology
LIGO Project

Massachusetts Institute of Technology
LIGO Project

LIGO Hanford Observatory

LIGO Livingston Observatory

http://www.ligo.caltech.edu/

Contents

1 Scope and purpose ... 3

2 Motivation for improvement over baseline design ... 3

2.1 Performance ... 3

2.2 Mechanical design ... 4

2.3 Production ... 5

3 UF RH ... 7

3.1 Mechanical design ... 7

3.2 Performance ... 10

3.2.1 Current prototype ... 10

3.2.2 Reduced gaps ... 13

3.2.3 Masking ... 15

4 Conclusions ... 17

List of Nomenclature

Element	The heating assembly that constitutes half a RH, usually excluding the shield.
Retainer	Part providing mechanical support and electrical connection to the two ends of an element
RH	Ring Heater
TM	Test Mass
UF	University of Florida

1 Scope and purpose

Although the baseline RH design is expected to meet specifications, production and testing have shown that it could benefit from improved temperature uniformity, production repeatability and mechanical robustness. The alternative RH design developed at UF has been adapted to fit in the same shields developed for the baseline RH and solves many of these issues.

While the base design and performance have been demonstrated, repeatability assessment, final tweaking and transition to bulk production requires further testing and funding. This note is meant to provide the information needed to decide if to continue development of the UF RH design, and on which timescale and scope (post-project?).

2 Motivation for improvement over baseline design

LIGO-T1300176, §2.10, points out the need of a RH with improved uniformity. In addition, based on production, testing and installation experience, other possible areas of improvement have been identified and are described in this section.

2.1 Performance

Based on the nominal design of the TCS system, the RH should only take care of correcting the HR surface of the TM, while the CO_2 projector illuminating the compensation plate will take care of compensating any residual thermal lensing inside the TM.

Requirements on the uniformity of the TM HR face correction are expressed in LIGO-T000092 in terms of round trip losses. However, the more immediate accepted reference is that the distortion of the HR surface should not exceed the polishing requirements expressed in LIGO-E080511 in terms of the Zernike coefficients:

Astigmatism: < 3 nm Amplitude of the Zernike coefficient $Z_{2,2}$ as defined in Born and Wolf pp. 523-525

According to FEA simulations (the only available tool to date), the baseline RH design easily meets these specifications, with a typical residual astigmatism $\lesssim 0.25$ nm.

Two main factors, indeed, can be expected to contribute to low-pass filtering the irregularities seen in the temperature profiles:

- The temperature pattern gets smoother and smoother as it diffuses inside the TM
- The HR face is on the opposite side of the TM with respect to the RH, and the thermo-elastic deformation induced by the thermal gradients is low-passed because of the TM rigidity.

However, only the first applies to the thermal lens induced by the thermal gradient. The result is that the thermal lens is much more influenced by any irregularity in the thermal profile.

The spatial variations observed in the baseline RH temperature pattern are much more likely to result in a "complicated" (i.e. non spherical) thermal lens that will make the job of the CO_2 projector harder, in terms of both creating the suitable illumination pattern and, possibly, maximum power required.

Reducing the higher order modes content of the thermal lens (and thus of the RH temperature profile patters) will thus serve a twofold purpose:

- Ease the job of the CO_2 projector
- Buy some margin over the (already good) results predicted by the COMSOL simulations for the HR face correction

2.2 Mechanical design

There are a few weak points in the baseline RH mechanical design:

- Glass rod weakness: glass rods have a fairly extreme aspect ratio (~.5 m long by 5 mm in diameter) and cannot be currently formed with high precision (±1 mm deviation from nominal radius of curvature are routinely observed). When trying to fit them in the tight RH geometry, they are subject to stresses that are not easy to control. Thermal expansion and contraction of the entire structure during operation makes this stresses even larger. Out of the 40 glass rods selected for RH production (after preliminary inspection), about 15 failed with no apparent reason (i.e. not the result of an accident or erroneous installation). A spring-based central standoff is being redesign to ease the problem, but complexity and miniaturization are making the task challenging.
- Macor retainer: for a variety of reasons that have become clear during development, the initially very simple design of the Macor end retainer has become increasingly complicated (see Figure 1). Now it has a fairly complex geometry that makes it difficult to machine, expensive and relatively fragile. It also includes threaded holes in the Macor, which make this a single-use part (since per aLIGO policy threaded holes in ceramic materials are not allowed to be used more than once).

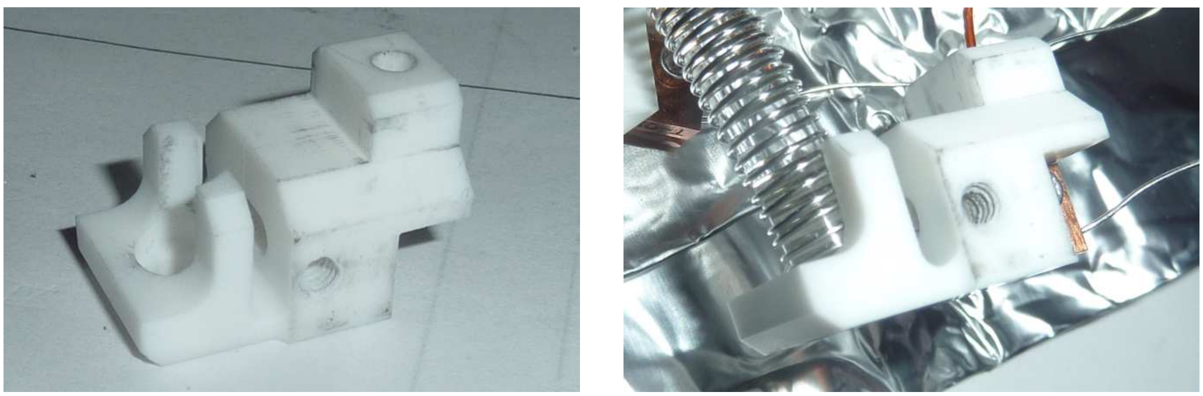

Figure 1. The baseline RH Macor end retainer. Modified over time to improve thermal isolation, provide better support for the glass rod and allow for electrical connections, it has now a fairly complicate and fragile shape. The threads in the Macor also make it a single-use part.

- NiCr winding: the NiCr winding is held in place only by friction against the glass rod, and by the fact that it is pre-formed in a coil that must be slightly enlarged to fit around the rod itself. While this seems to work very well in the center, it becomes increasingly weak close to the ends. Indeed, for the winding not to move in the final ~5 cm on either side, the NiCr wire has to be kept in tension; this can only be done when the retainers and the clamps are assembled. The result is that every time an "element" (glass rod+NiCr winding) is assembled/disassembled from the end retainers, or manipulated when the retainers are not installed, there is a risk of substantially modifying the winding distribution and the resulting profile. It's worth noting that the last few cm of winding on either side are particularly critical, as they are precisely tuned in production to compensate for the power lost though the retainers, effort that can be easily frustrated if they are allowed to shift around.

2.3 Production

- Handling: the problem of securing the NiCr winding described in the previous section is particularly risky with the current production cycle. In fact, because of facility limitations, the RHs currently go through these phases:

  Production -> Testing -> Disassembly -> Cleaning -> Re-assembly -> Installation

  Since no testing is done after the final re-assembly to verify that the thermal profile is still what measured after initial production, any accidental modification of the NiCr winding will go undetected. It is thus very important to reduce the risk of accidental modifications after the testing phase.
- Repeatability in production: despite the fact that the production technique has been refined throughout the development phase, the produced elements still show substantial differences in the measured thermal profiles. The causes of these differences have not been completely understood, although several likely candidates (probably all contributing to some extent) have been identified. Because of this, careful pairing of elements (again based on FEA simulations) has become critical to produce RHs with good performance. As an example, Figure 2 shows the result of such a paring performed on one of the production batches destined to be used to form two ITM RHs.

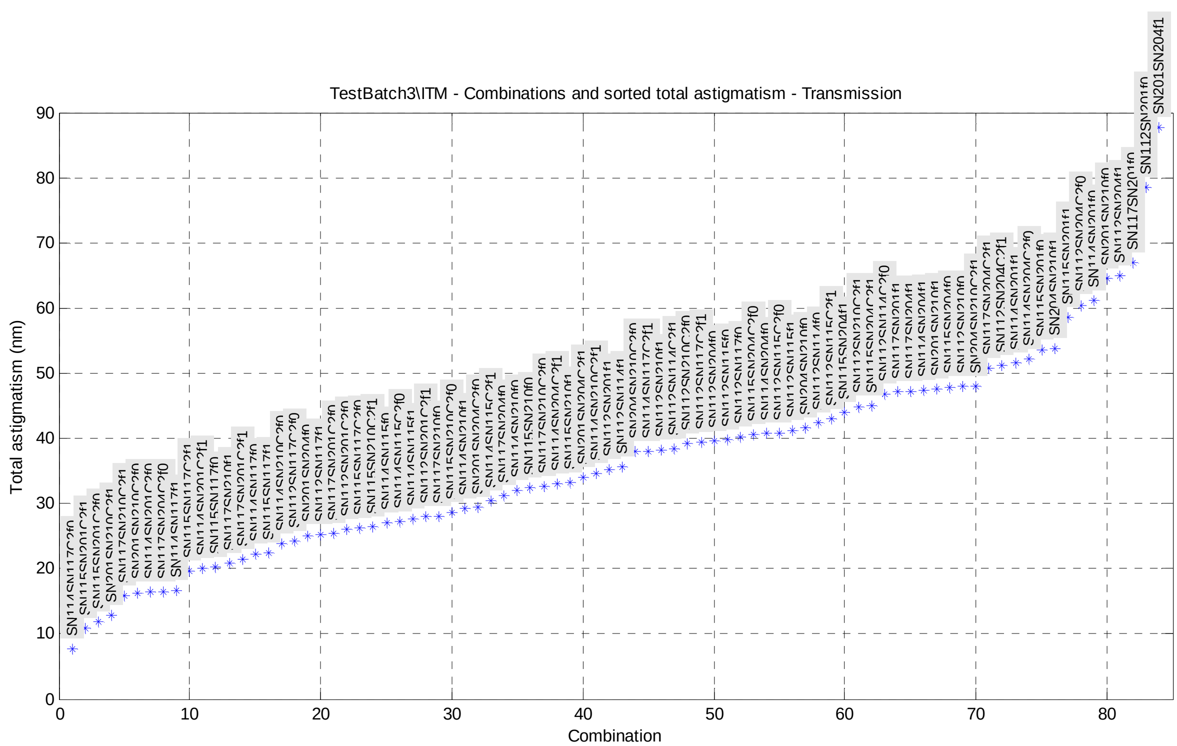

Figure 2. Simulated performance of all the possible RH that can be formed by pairing the elements tested as part of "Test Batch 3". Each point represent a RH formed using the elements indicated by serial number in the label (*SNxxxSNyyy*), including the relative flipping of one with respect to the other (indicated by *f0* or *f1* at the end of the label).

Among the all possible combinations that can be realized with a typical batch of elements, the best performing pairs usually get down to 10-20 nm residual astigmatism (but keep in mind that some of them reuse the same elements, and thus cannot be really built simultaneously), while the average is more around 40 nm, with a long tale going all the way to about 100 nm. This has several disadvantages:

- It is necessary to produce elements in excess to have a "population" from which to chose the best pairing
- Assessment of relative performances heavily relies on FEA modeling, and we might end up choosing the wrong pairing if the FEA model is not accurate.
- Good performance of a specific pair could be the result of fairly big irregularities in the thermal profiles of each element, which by chance compensate each other in that particular configuration. In this case, the performance of the RH would be even more dependent on maintaining the pairing and thermal profiles unchanged.
- Failure of an element is likely to require replacing the entire RH, going though another pairing process and probably accepting a degradation in performance (since the failed pair was already chosen as "the best one"), unless new elements have been produced in the meanwhile.

3 UF RH

3.1 Mechanical design

Details on the mechanical design of the UF RH were given in a presentation at the September 2010 LVC meeting (LIGO-G1000945). It is basically made of two concentric aluminum half-rings that sandwich NiCr wire in between. Electrical insulation and improved emissivity are both realized using a fire-sprayed alumina coating.

This design has undergone no modification since then, except for the redesign of the end connectors. These have now been adapted to be fully compatible with the current RH shields.

Some close-ups of the connector can be seen in Figure 3. It is comprised of three main elements:

- *Stainless steel plate* to connect the aluminum RH to the retainer. Thanks to the slim geometry and relatively low thermal conductivity it provides some thermal insulation.
- *Retainer*: built out of Peek for fast prototyping, it can be easily made out of Macor. It features a much simpler and more solid geometry with respect to the baseline RH retainer. Also, it does not have any threaded hole.
- *Clamp*: a simple block used to clamp the retainer to the shield. This is made out of Peek and could be kept this way, since it’s pretty far from any heating element and in good contact with the shield (assumed at room temperature). However, it’s an easy enough geometry to machine out of Macor if desired.

The stainless steel plate is attached to the RH body using two screws. Stainless steel plate, retainer and clamp are kept together using a single bolt inserted in through holes, with a nut that is tightened to clamp the entire assembly to the shield.

The connectors on both sides are identical. However, since the wire runs back and forth in each half, both electrical connections are on the same side. On the other side, a short piece of copper wire serves as a bridge between the two NiCr wires (leaving the NiCr as a single continuous wire throughout the retainer is possible, but has been found to release too much heat there, creating an imbalance in the temperature profile).

Connection of the NiCr wire with the copper bridge and with the conducting leads is realized by crimping the wires in small copper tubes.

If necessary, the clamp on the side of the electrical connection can also be used (directly or with the addition of another similar clamp on top) to provide an anchoring point to relieve mechanical stress from the conducting leads.

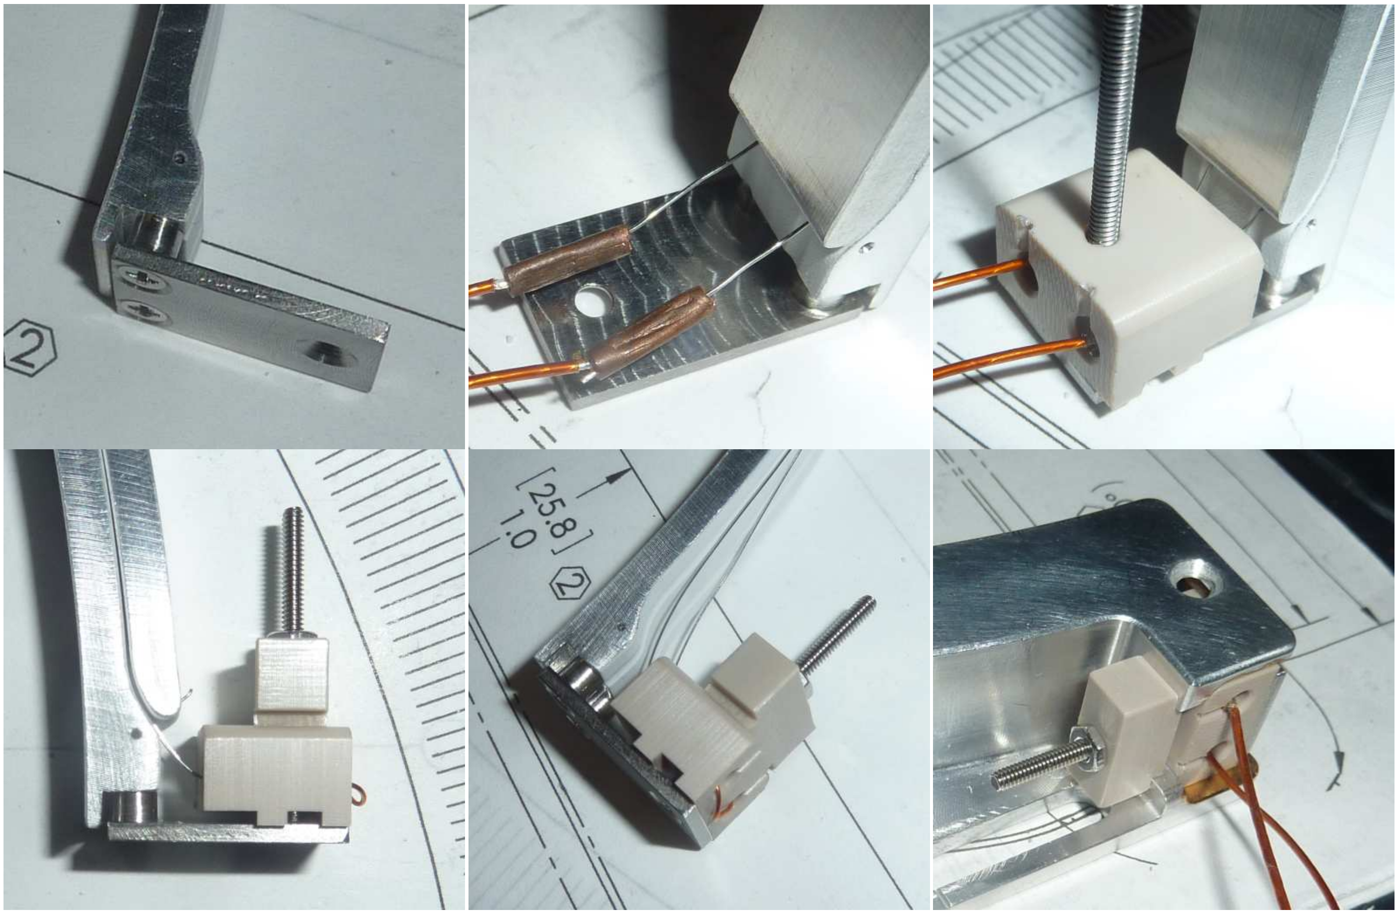

Figure 3. Details of the new UF RH mechanical retainer designed to fit into the current RH shields. Left to right, top row: *stainless steel plate,* NiCr wire crimped to electrical wires using copper tubes, *retainer* (built using PEEK for convenience of prototyping, but can easily be made out of Macor); left to right, bottom row: *retainer* with *clamp* on top an bolt+nut keeping attaching them to the SS plate, from two different angles (note that this end is the one with the copper bridge); whole assembly clamped to the shield.

Together with the redesign of the end connectors/retainers, a small modification of the aluminum RH would be required to achieve optimal performances. Indeed, the original design featured fairly big gaps at the conjunctions between the two RH halves (~6 mm), as shown in Figure 4.

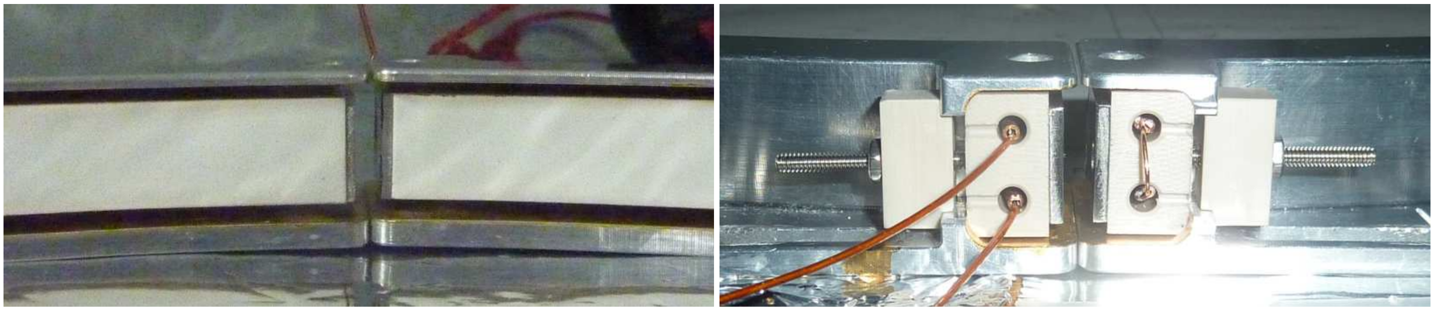

Figure 4. A close-up of the junction between the two RH halves, from inside the RH (left) and from outside (right). The fairly big gap has no reason to exist, and can be easily eliminated by a simple modification to the RH design. However, at present we don't have the resources to build and coat a modified prototype.

These gaps are the result of a design error and are not necessary in any way, since there is no risk of electrical contact, and the impact of mechanical and thermal contact, if any, can only be of further homogenizing the temperature across the entire RH. This design modification is illustrated in Figure 5 and is trivial even without modifying the connectors at all, but it has been left behind for now due to time and budget constraints (it would require to machine and re-coat two new RH halves). Nonetheless, an estimate of the effect is shown in the section 3.2.2.

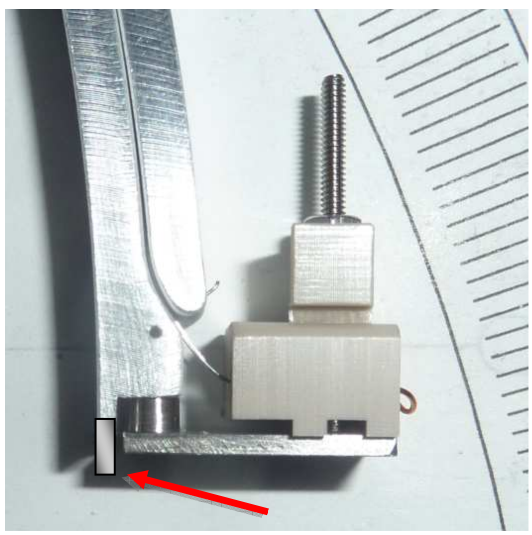

Figure 5. An illustration of the simple modification needed to extend the RH inner surface and close the gaps now present between the two halves when the RH is assembled.

3.2 Performance

3.2.1 Current prototype

The last report on the UF RH performance has been given at the March 2012 LVC meeting (LIGO-G1200216). At that time, the RH showed a very uniform temperature profile compared to the baseline RH. However the average temperature in the center of each half was a few percent higher than close to the end connectors, due to the lack of heat sinks in that region. This added to the dip at the end connectors (reduced with respect to the baseline RH, but still present) to create an overall astigmatism because of the higher amount of heat released along the vertical axis compared to the horizontal one.

The profile measured with the redesigned connectors (Figure 6) is very similar to the one taken with the old ones. The transmission phase map obtained from FEA simulations is also shown: the residual astigmatism is about 40 nm, equivalent to the average of all the pairings simulated for the baseline design.

When the UF RH temperature profile (Figure 6) is compared with a typical baseline RH profile (Figure 7), a few key differences can be noted:

- The radiation efficiency seems higher for the UF RH. The two profiles shown are taken with the same power sent to the RH, but the effective temperature measured by the thermopile sensor (i.e. the amount of IR radiation collected) is much higher for the UF RH. Higher emissivity of the alumina coating with respect to the glass+NiCr assembly, bigger effective radiating surface (the UF RH exposes a flat radiating surface filling almost the whole shield, versus the smaller and round surface of the glass rod) and better thermal insulation at the end retainers can all contribute to this, but the exact role of each effect has not been investigated.
- The "small scale" (<1/4 of the entire RH circumference) temperature variations along the RH profile (both those that are understood in nature and the one that are "random") are very much reduced. This is mainly due to the high thermal conductivity of the aluminum structure, and the fact that the wire is not exposed. While this is not a huge advantage due to the low-pass filtering effect of the TM, it can lead to a reduction of higher order modes distortion (in the specific examples shown in Figure 6 and Figure 7, bottom right corner, the distortions due to terms of order >2 are reduced by about a factor 2).
- The UF RH profile shows a very high regularity, and the astigmatic axis is very well aligned with the vertical direction, as expected. This is not true for the baseline design, where the residual astigmatism is due to uncompensated deviations from the expected profile and is thus basically random.
- The above observation and the small difference between the two UF RH halves suggest that the consistency and repeatability of performance of the UF RH elements can be expected to be much higher than those observed for the baseline design elements.

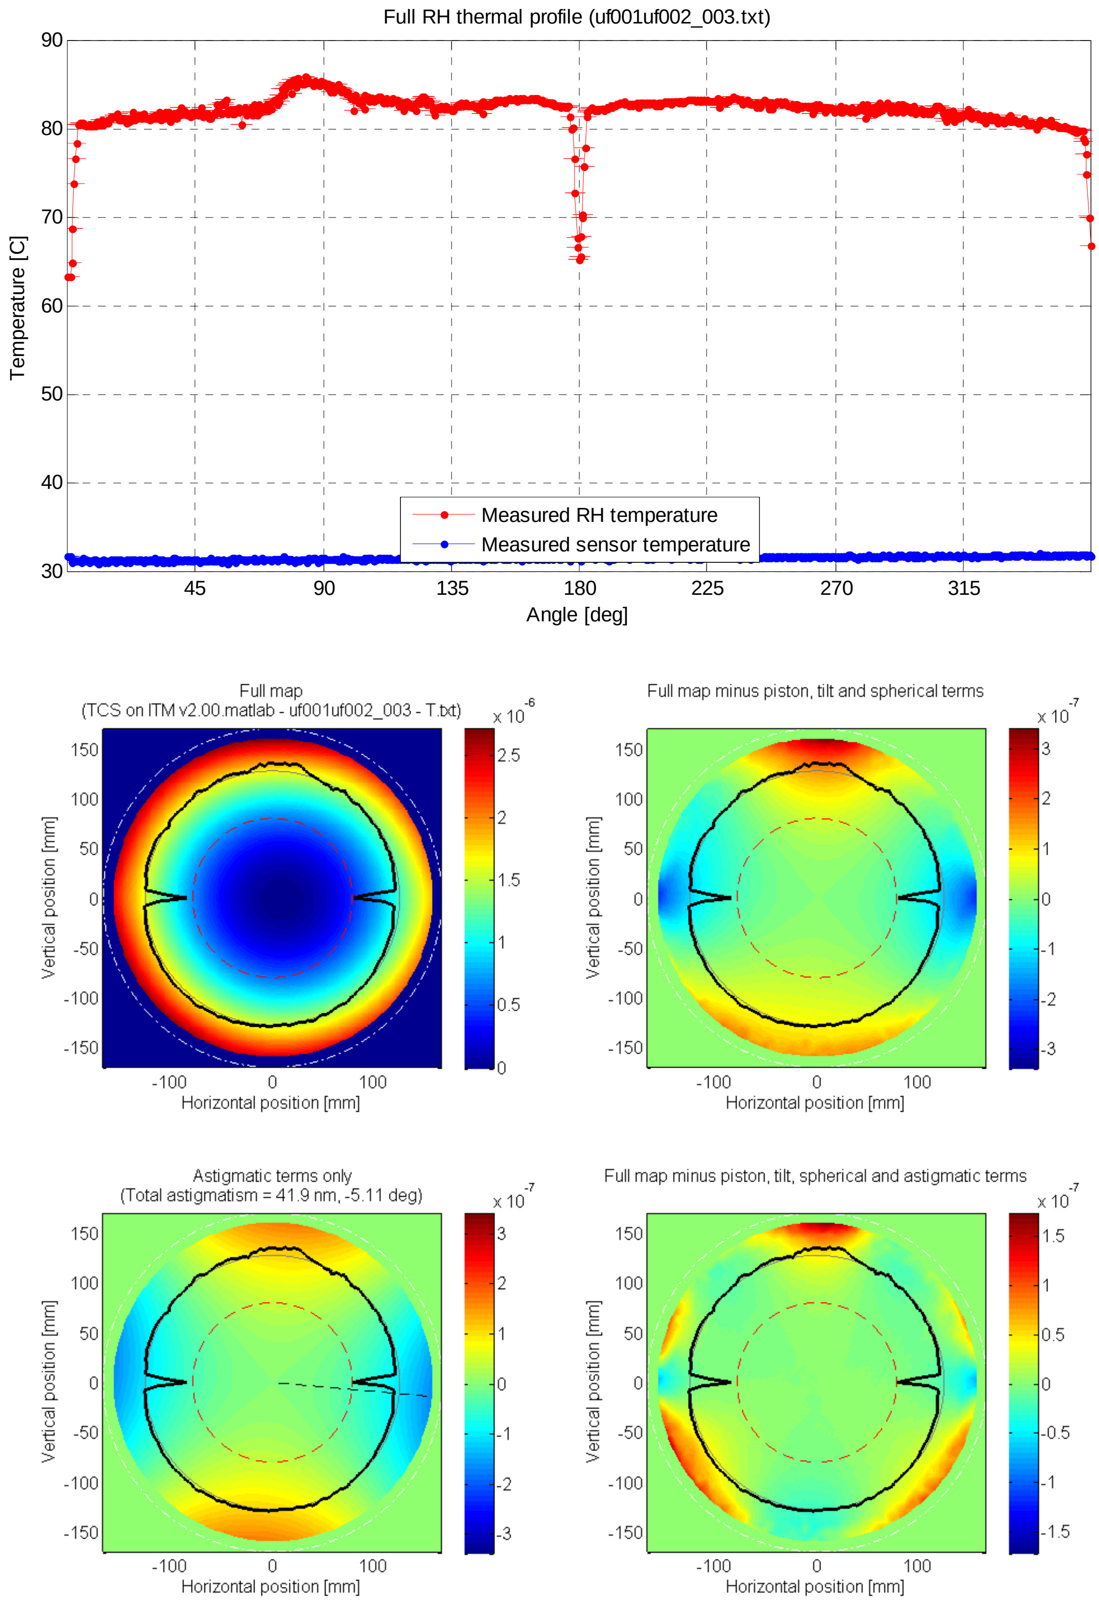

Figure 6. Top: measured temperature profile of the UF RH prototype with 5 W of electrical power per half-RH. Bottom: different visualizations of the simulated transmission phase map (in nm)

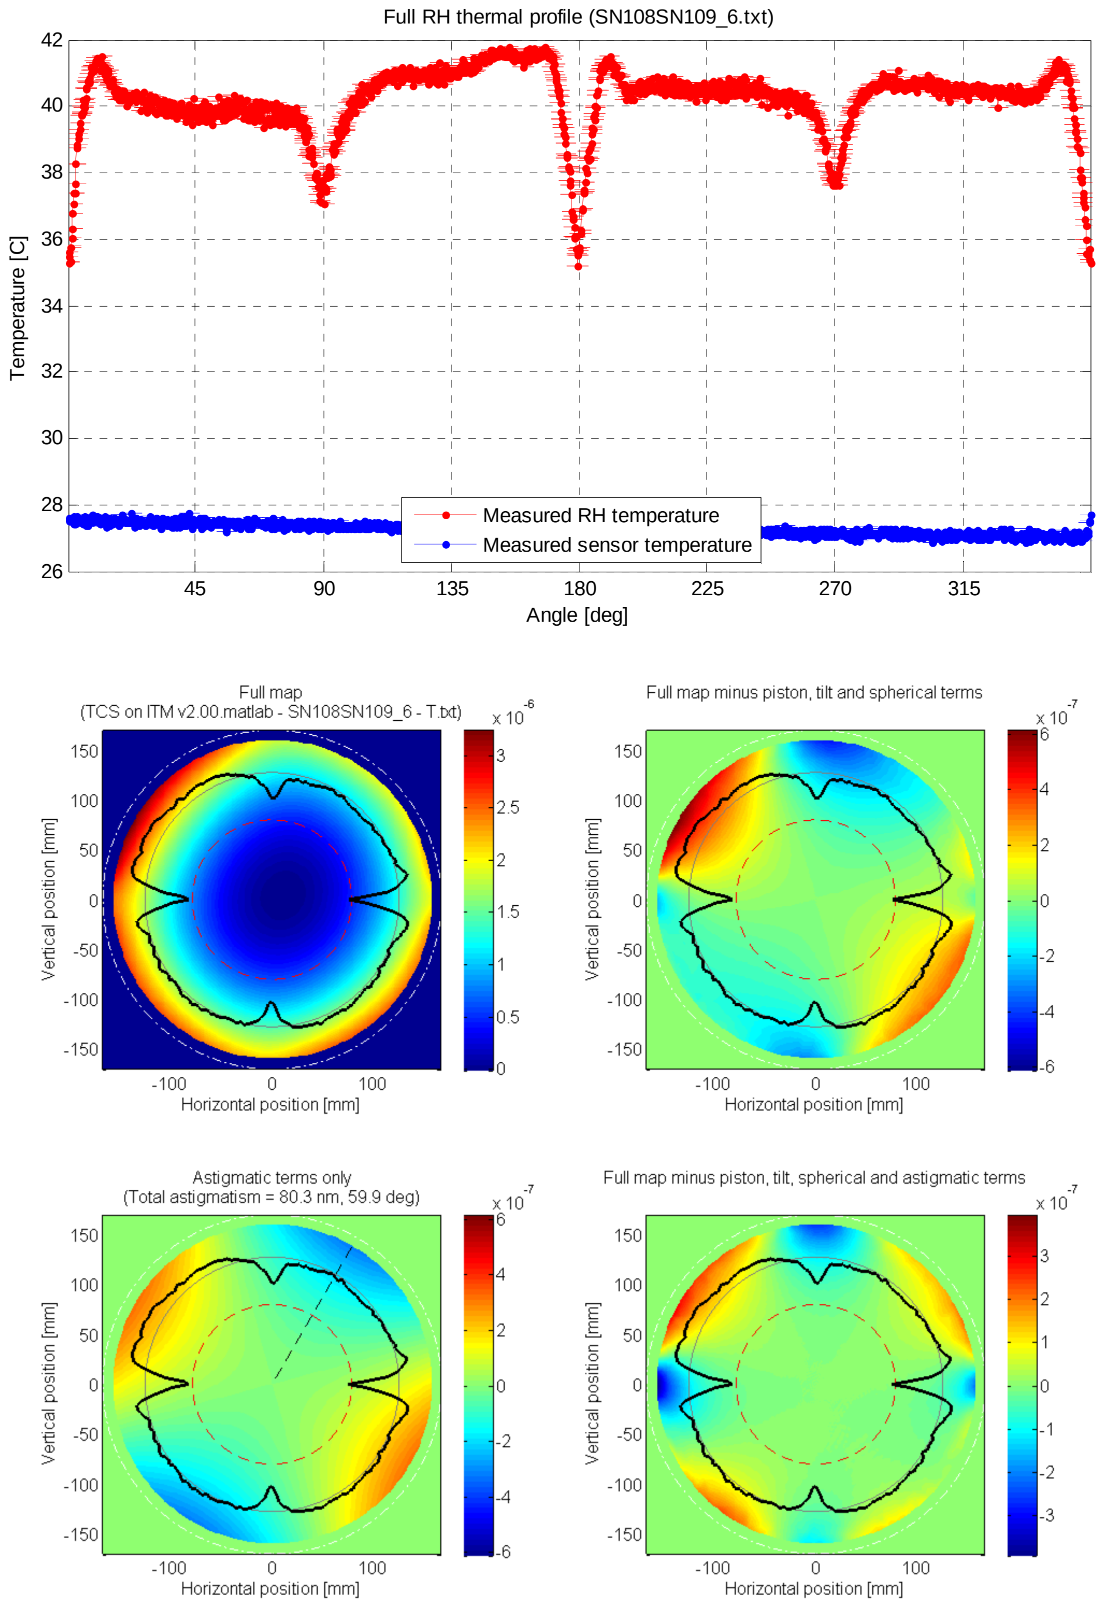

Figure 7. Top: measured temperature profile of a baseline RH with about 5 W of electrical power per half-RH. Bottom: different visualizations of the simulated transmission phase map (in nm).

3.2.2 Reduced gaps

The profiles of Figure 6 and Figure 7 show very similar temperature dips (in proportion to the average temperature of the RH above ambient) at the end retainers (0 and 180 deg). However, it is important to realize that while this is basically a limit of the baseline design (and in fact the tweaking of the NiCr wire close to the ends has been introduced to compensate for it), in the case of the UF RH it is the result of a design imperfection and can be easily corrected, as explained at the end of section 3.1 and illustrated in Figure 5.

While we did not have the resources to machine and coat a new RH, we tried to estimate the effect of such a modification by using temporary end connectors so that the RH halves would come closer together (at the expense of reducing the distance between the RH and the TM along the vertical direction, since due to the – wrong – design of the RH halves, now the RH becomes sort of an ellipse, as illustrated in Figure 8).

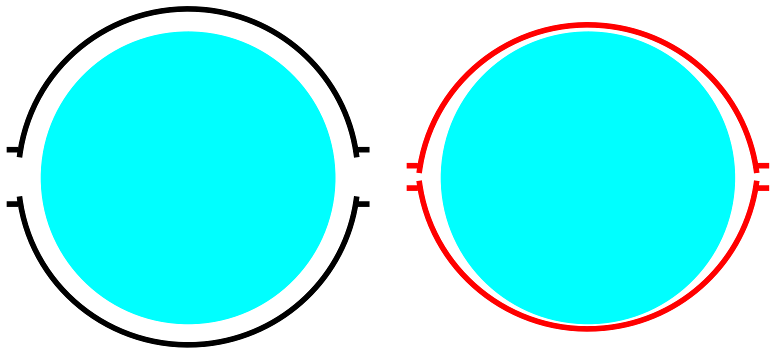

Figure 8. A scheme illustrating the expedient used to evaluate the effect of reducing the gaps at the retainers of the UF RH. Left: the RH correctly installed, but with large gaps due to an incorrect design. Right: the RH halves have been moved closer together by modifying the end retainers. While this yields an "elliptical" RH, it can help estimating the effect of closing the gaps.

While in the modified configuration the RH halves are still a few (2-3) mm apart (we can't do better without redesigning the RH itself), the dips at the ends are noticeably reduced, and so is the residual astigmatism that goes from 40 nm to about 30 nm, as can be seen in Figure 9. This can be further improved by modifying the design so that the RH halves touch each other.

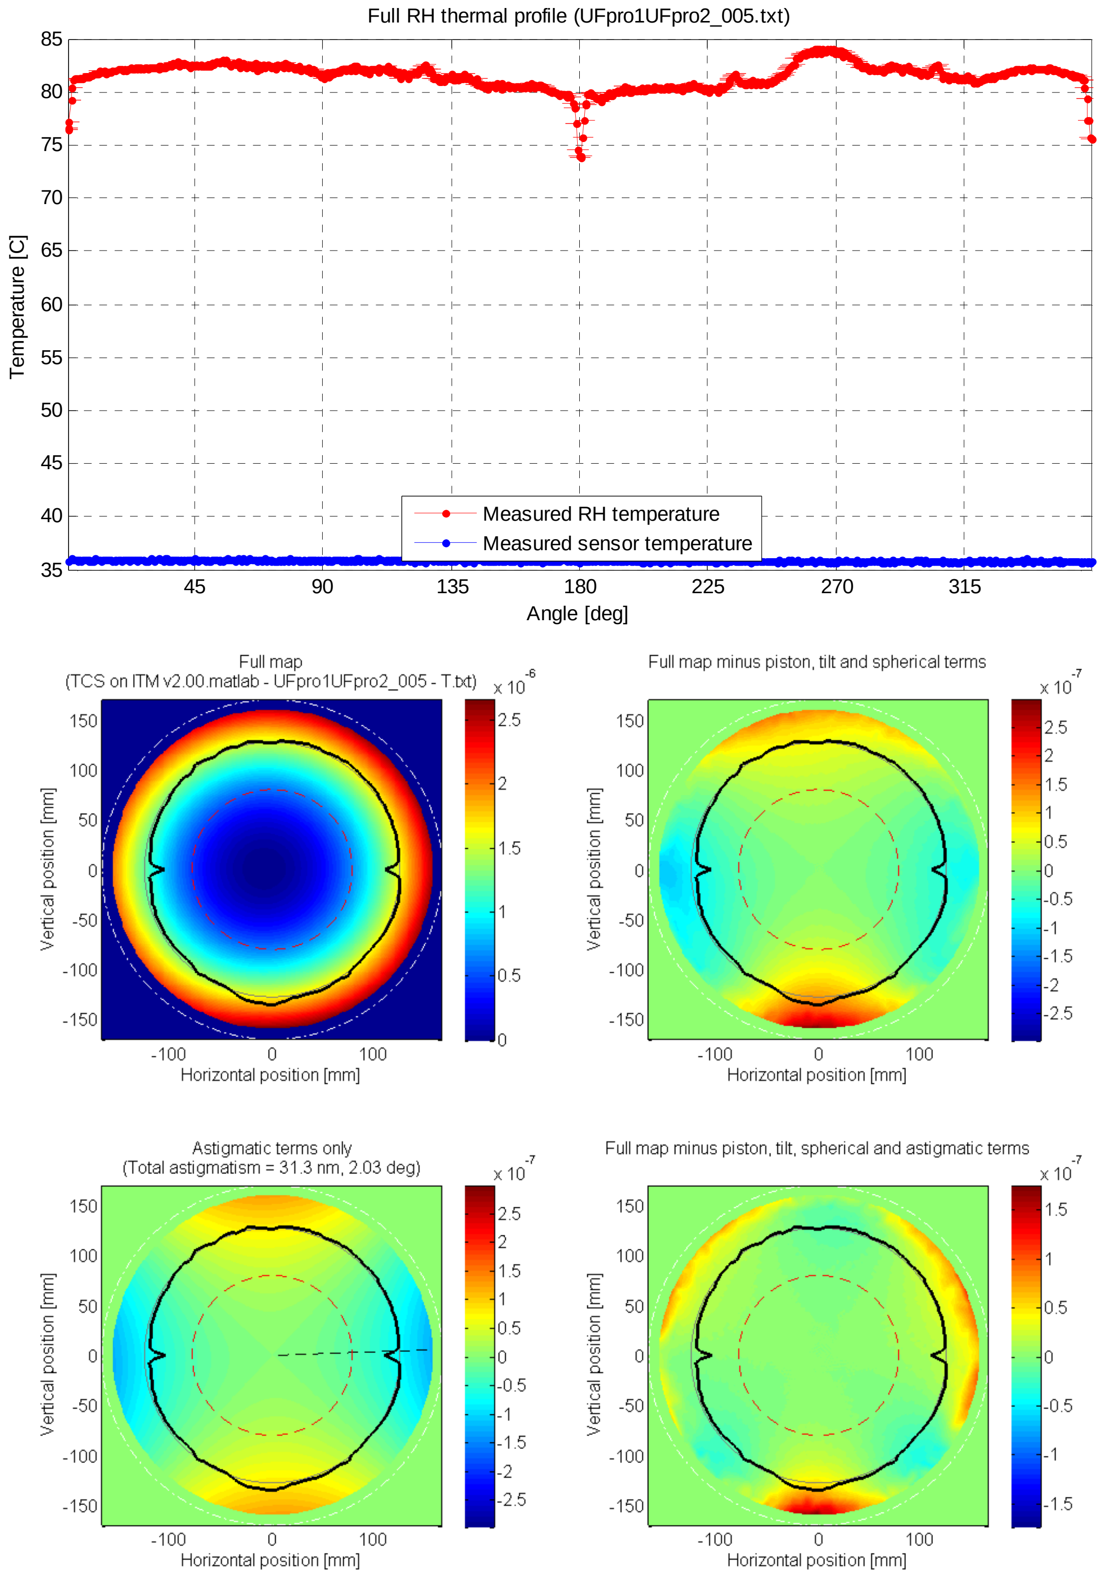

Figure 9. Measured temperature profile and phase maps of the UF RH installed so to reduce the unwanted gaps at the end retainers. The temperature dips are much smaller (compare Figure 6), and the residual astigmatism is reduced from ~40 nm to ~ 30 nm.

3.2.3 Masking

Contrary to the baseline design, the UF design does not allow for an easy way of controlling the heat released in different parts of the RH (e.g. adjusting the spacing of the NiCr windings). Moreover, the high thermal conductivity of the aluminum support will likely strongly reduce the effectiveness of such efforts.

On the other hand, since what matters is the infrared emission, reducing the emissivity of the surface is a very efficient and well controllable way of adjusting the effective temperature profile seen by the TM.

With adequate resources and time one can implement several solutions that allow for a precise and very well modulated control of the temperature profile (e.g. masking the deposition of the alumina coating so to reduce the width of the high emissivity surface in specific areas). However, we have found that little C-shaped metal strips clipped onto the RH represent a cheap, convenient and flexible way of locally reducing the emissivity (alumina is expected to have an emissivity of ~0.9 versus ~0.1 of the bare metal).

We have hand-made a few of this clips out of metal paper binders to try the concept. Figure 10 shows an example of such clips applied to the RH.

Figure 10. Hand-made metal clips used to locally reduce the emissivity of the UF RH. They turned out to be a convenient way of correcting the heating profile and reduce the residual astigmatism.

The effect of the strips is that of creating very narrow depressions in the effective temperature profile, which gets easily averaged out as the heat diffuses inside the TM. This method has several advantages over other possible alternatives:

- Repeatable, as it relies on the size (and emissivity) of the strips rather than on thermal contact (difficult to replicate on different assembly)
- Controllable, as the strips can be placed with high precision, and easily re-checked for unwanted movement during handling.
- Efficient, as modification of the thermal profile is achieved by reducing the emissivity rather than dumping the excess heat in a heat sink.

Figure 11 shows a test conducted by applying three clips per element, in “standard” positions (i.e. not optimized for the specific profile measured for each element). As can be seen, the residual astigmatism has been reduced to about 13 nm (from ~40 nm, since in this case we used the proper connectors that install the RH in the right position, but with the large gaps) that is at the level of the best baseline RH produced.

We expect to be able to further reduce this figure by tweaking the location of the clips for each element.

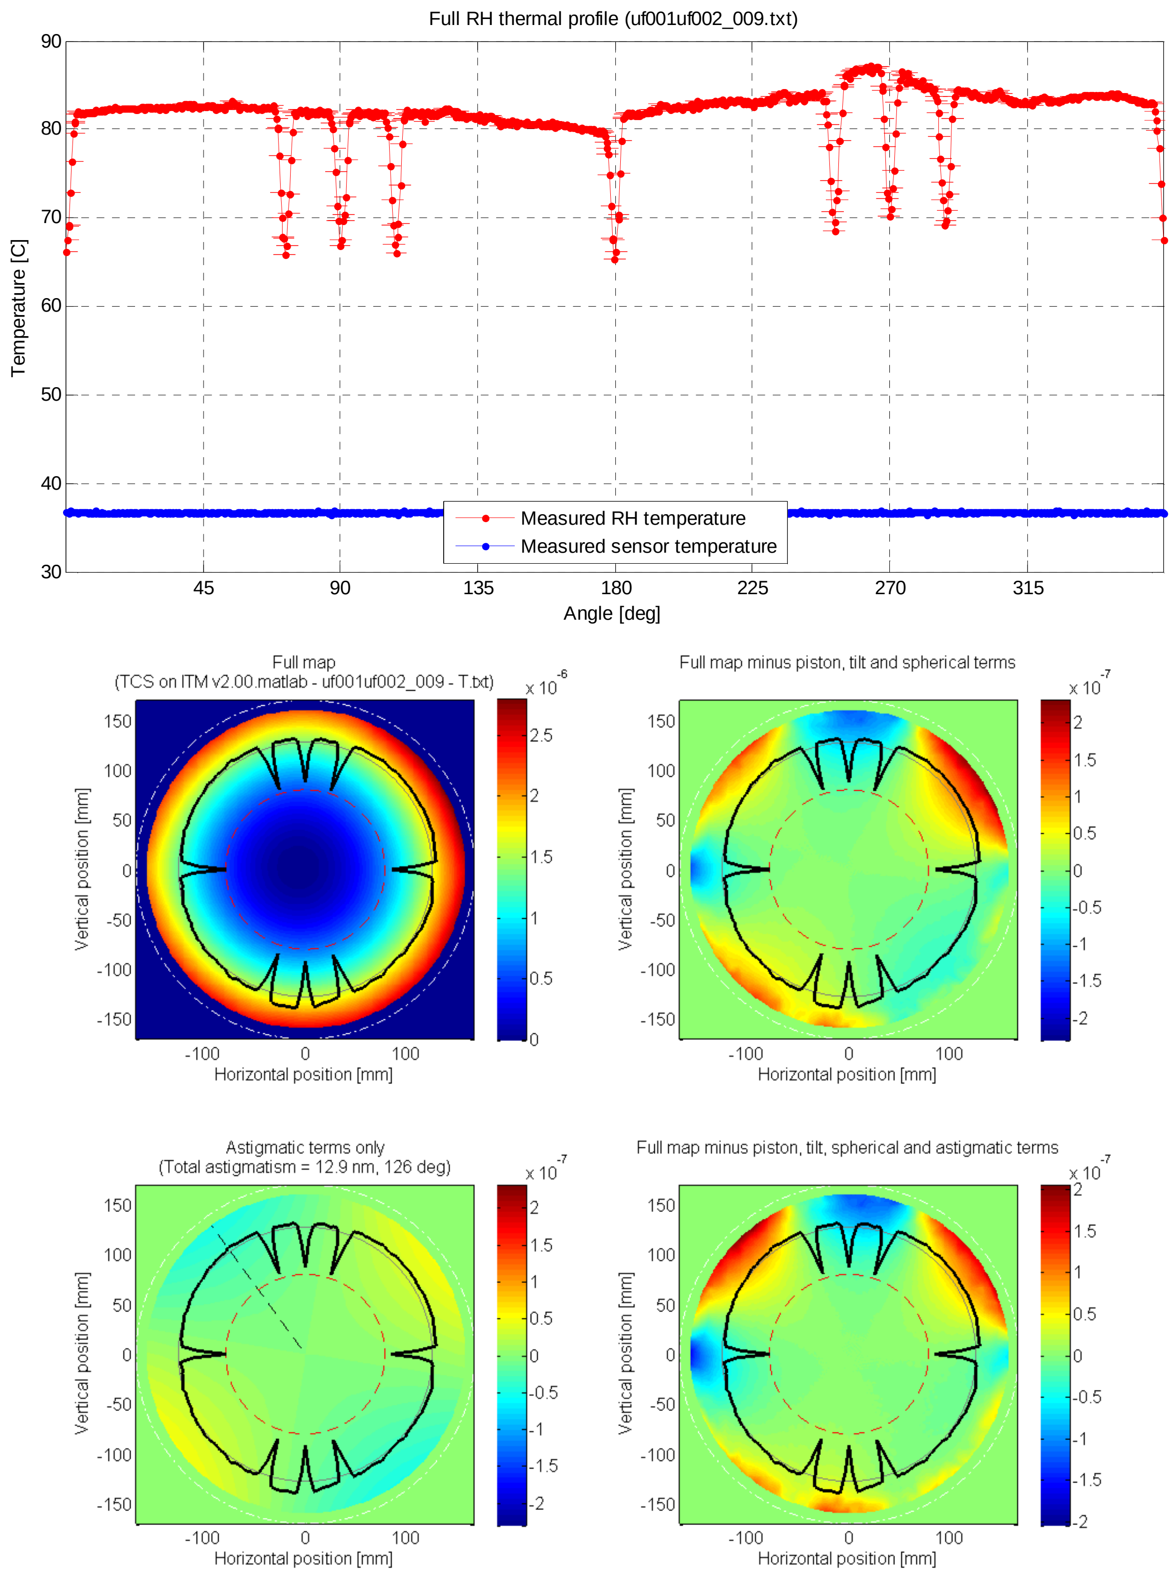

Figure 11. Measured temperature profile and phase maps of the UF RH, masked to compensate for the higher temperature in the center of each half. The six narrow dips created by the clips are clearly visible. The astigmatism is reduced from ~40 to ~ 30 nm (compare Figure 6).

4 Conclusions

We have shown that the current UF RH prototype, with redesigned end retainers to fit inside the current RH shield, can solve many of the issues that afflict the baseline RH design. In particular:

- Greater mechanical robustness
- Better production repeatability
- Much reduced sensitivity to handling
- Improved temperature uniformity

With a trivial design modification it has been shown to perform within a factor ~2 from the baseline design RH (~30 nm residual astigmatism), but with much improved reliability.

With the adoption of a “standard” masking (equivalent to the practice of tweaking the NiCr winding density, but much more robust) the performance difference has been eliminated, with the tangible possibility of even further improvement with a more sophisticated mask.

While there is no doubt about the viability of aluminum and NiCr as materials, the big surface of alumina coating remains to be discussed. Alumina is accepted as a vacuum material for use in aLIGO, and preliminary RGA test at UF (that do not reach the sensitivity need for aLIGO qualification) have shown no excess outgassing from the UF RH even when heated at ~200 C. However, the close proximity to the TM requires caution in evaluating its usability.

With the exception of this last issue, which is not up to us to tackle, the development of the UF RH design is from our point of view now mature. Further work is only needed and justified in the perspective of production for installation in aLIGO, and would require further funding.